

A Survey of Advance Reservation Routing and Wavelength Assignment in Wavelength-Routed WDM Networks

Neal Charbonneau, *Member, IEEE* and Vinod M. Vokkarane, *Senior Member, IEEE*

Abstract—Traditionally, research on routing and wavelength assignment over wavelength-routed WDM networks is concerned with *immediate reservation* (IR) demands. An IR demand typically does not specify a holding time for data transmission and the start time of the data transmission is assumed to be immediate (i.e. when the connection request arrives). The concept of *advance reservation* (AR) has recently been gaining attention for optical networks. An AR demand typically specifies information about the start of the data transmission or a deadline, as well as the holding time of the transmission. AR has several important applications for both wide-area networks and Grid networks. For example, AR can be used for adjusting virtual topologies to adapt to predictable peak hour traffic usage. It can be used to provide high-bandwidth services such as video conferencing and in Grid applications requiring the scheduled distribution of large files and for co-allocation of network and grid resources. AR can also be beneficial to the network by allowing the network operator to better plan resource usage and therefore increase utilization. Knowledge of the holding time can lead to more optimal decisions for resource allocation. This translates to better quality of service for users. In this paper we provide a comprehensive survey of the past and current work on advance reservation for optical networks. There have been many variations of the advance reservation concept proposed, so we will also provide a broad classification. In addition to the survey, we will discuss what we believe are important areas of future work and open challenges for advance reservation on optical networks.

Index Terms—Advance reservation, scheduled demands, WDM, survey, wavelength-routed, and RWA.

I. INTRODUCTION

OPTICAL wavelength-routed WDM [1] networks, or optical circuit switched (OCS) networks, are a potential candidate for future wide-area backbone networks as well as scientific Grid networks. In WDM networks, each fiber is partitioned into a number of wavelengths, each of which is capable of transmitting data. This allows each fiber to provide data transmission rates of terabits per second. An optical WDM network consists of fibers connected by switches, or optical cross connects (OXC). In order to transmit data over the network, a dedicated circuit is first established when a user submits a connection request. When a connection request

Manuscript received 7 April 2011; revised 1 November 2011 and 7 November 2011. This work was supported in part by the National Science Foundation (NSF) SOON project under grant CNS-0626798 and the Department of Energy (DOE) COMMON project under grant DE-SC0004909.

N. Charbonneau was a graduate student with the Department of Computer and Information Science, University of Massachusetts, Dartmouth, MA and is now with The MITRE corporation.

V. Vokkarane is an Associate Professor with the Department of Computer and Information Science, University of Massachusetts, Dartmouth, MA (e-mail: vvokkarane@ieee.org). (corresponding author)

Digital Object Identifier 10.1109/SURV.2011.111411.00054

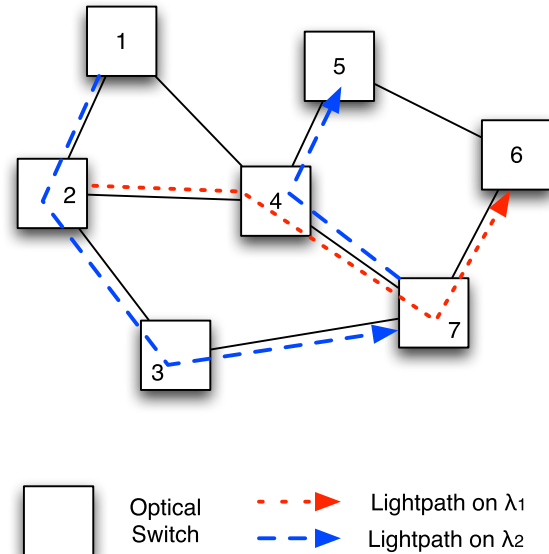


Fig. 1. Example of a wavelength-routed network. For each request, a lightpath is established in the network. The lightpath consists of a path as well as a wavelength. In this example, there are two wavelengths used in the network, λ_1 and λ_2 and three lightpaths shown.

arrives at the network, the request must be routed over the physical topology and also assigned a wavelength. This is known as the *routing and wavelength assignment* (RWA) problem [2]. The combination of a route and wavelength is known as a *lightpath* [3]. The RWA problem is NP-complete so heuristics are typically used [4]. The bandwidth granularity of the circuit does not necessarily have to be one wavelength. There is work on traffic grooming, which performs aggregation of multiple sub-wavelength traffic streams onto a single wavelength [5], [6]. An example of a wavelength-routed network is shown in Fig. 1 (with no traffic grooming). There are three lightpaths in the network using two different wavelengths. One lightpath is sourced at Node 1 with a destination on Node 7 using wavelength λ_2 . Another is sourced at Node 2 with destination of Node 6 on λ_1 . The final lightpath is sourced at Node 7 and destined for Node 5 with wavelength λ_2 . No two requests can use the same wavelength on the same link. If more requests arrive over time new lightpaths must be allocated as long as there are enough wavelengths to establish them.

In a single-hop, or all-optical, WDM system, the signal is transmitted all-optically through the network. There is no conversion of the signal back to electronics in the network. These

are also known as transparent optical networks. In multi-hop systems the signal may undergo optical/electronic/optical (O/E/O) conversion at some intermediate nodes. If O/E/O conversion occurs at every node, then the network is called an opaque network, whereas if only some nodes employ O/E/O the network is called a translucent network. In the absence of wavelength converters (which are expensive), a connection in a single-hop WDM system must use the same wavelength across all links. This is known as the *wavelength continuity constraint*. Multi-hop systems can use different wavelengths on different links because the signal may undergo O/E/O conversion at some intermediate nodes, allowing it to be retransmitted on a wavelength different from the received wavelength. This conversion process can be expensive, however, both in terms of cost of equipment and due to the dependence of the conversion process on the connection line rate and modulation format. The disadvantage of single-hop systems is that, in the absence of regenerators, the signal noise accumulates from physical layer impairments such as crosstalk, ASE noise, and nonlinear impairments like four-wave-mixing, cross phase modulation, and stimulated Brillouin and Raman scattering. To counter this, impairment-aware routing can be used to ensure the signal to noise ratio is at acceptable levels when the signal reaches the destination. There has recently been significant work in impairment-aware routing [7], [8].

Two traffic models are usually considered for wavelength-routed networks: static and dynamic [4]. A static traffic model gives all the traffic demands between source and destinations ahead of time. A traffic matrix is given and the goal is typically to find an RWA that can meet all the demands and minimize overall cost (e.g. using the least number of transmitters/receivers). Dynamic traffic requests arrive one-by-one according to some stochastic process and they are also released after some finite amount of time. When dynamic traffic is considered, the number of transmitters and receivers is fixed and the goal is to minimize request blocking. A request is said to be blocked if there are not enough resources available to route it. There is extensive work for these problems, see [2], [4], [9], [10], among others.

We can further classify the above traffic models as *immediate reservation* (IR) or *advance reservation* (AR) [11] requests. The data transmission of an IR demand starts immediately upon arrival of the request and the holding time is typically unknown for dynamic traffic or assumed to be infinite for static traffic. AR demands, in contrast, typically specify a data transmission start time that is sometime in the future and also specify a finite holding time. Fig. 2 shows the difference between an AR and IR request. We can see that in Fig. 2(a) the resource allocation occurs when the request arrives at the network. The duration of the request is unknown. In Fig. 2(b), the actual allocation of resources does not occur until a later time. The resources are reserved when the request arrives, but they can be used by other requests before the reservation time. The difference between the arrival of the request and beginning of the transmission is the *book-ahead time*, which is specified by the request. The duration of the request is also specified in advance and known by the network. The fact that holding time and book-ahead time is known by the network allows the

network to more efficiently optimize resource usage. This is just one example of an AR request, we discuss the variations in Section III.

Advance reservation was initially proposed for non-optical networks, focusing on circuit-switches, packet-switched, and ATM. We briefly mention some of this work here. Initial work focused on traffic modeling and call admission for telecommunication systems (e.g. [12], [13]). Wolf et al. [14], [15] proposed advance reservation for quality-of-service of multimedia applications like video conferencing. Greenberg et al. [16], [17] focused on similar applications with some theoretical results concerning mixed immediate reservation (IR) and AR traffic. They assume that AR traffic has higher priority than IR and focus on admission control algorithms for the two types of traffic. Extensions to RSVP were proposed in [18]. A detailed discussion on path computation of advance reservation requests was presented in [19]. In this work, the authors focus on routing algorithms to handle both spatial and temporal aspects of AR.

Advance reservation for optical networks was first proposed by Zheng and Mouftah in [20], [11]. While some solution techniques may be adapted from the electronic domain to the optical domain, the advance reservation problem for optical networks presents new challenges, such as the wavelength continuity constraint, grooming, survivability, and others.

A. Organization

The paper is organized as follows. We begin by motivating the need for advance reservation in optical networks in Section II. We then discuss and classify the various types of advance reservations that have been proposed in the literature in Section III. We discuss network architectures to support advance reservation in Section IV. Next, we present our survey on problems and solution techniques proposed for advance reservation in Section V. Advance reservation for optical networks is a relatively new topic, so our survey will be comprehensive covering the first papers to the latest work. In Section VI we discuss the various advance reservation frameworks and architectures that have been implemented. In Section VII we discuss other related work on advance reservation that are not in the optical domain or not related to routing and wavelength assignment. Section VIII will discuss open problems and possible research directions for advance reservation. Finally, we conclude the paper in Section IX.

II. MOTIVATION

In this section we discuss the motivation for advance reservation over optical networks. Advance reservation has applications for both wide-area networks and Grid networks. We will discuss the applications specific to these types of networks in the following subsections. Some of these applications can be applied to both types of networks, but many advance reservation papers focus specifically on Grid networks. In general, advance reservation benefits the network because knowledge of future state information (due to declared arrival and holding times of data transmission) can be used to improve the admission control and planning/provisioning to increase network utilization and maximize profits. It also benefits the

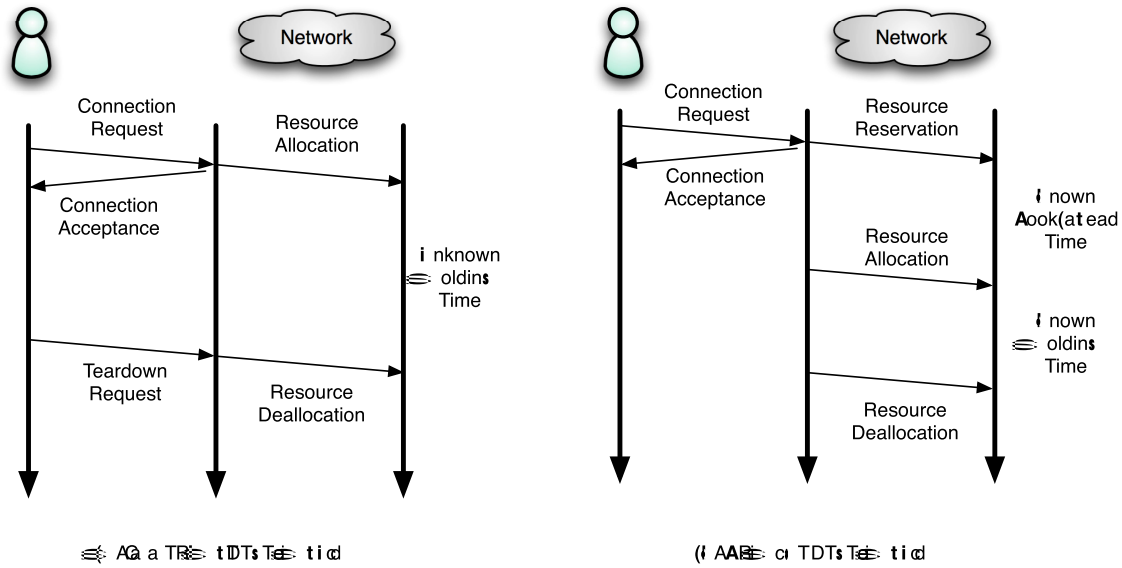


Fig. 2. The request and allocation of resources for immediate and advance reservation requests. In the figure, we assume that the requested resources are available. Before reservation/allocation the network must find an appropriate lightpath. For immediate reservation (a), the allocation is at the same time as the request arrival and the duration is unknown. For advance reservation (b) the allocation is some (known) time after the arrival and the duration is also known. Variations of this advance reservation model are discussed in Section III.

user because the network can provide better quality-of-service to requests that book-ahead.

A. Wide-area Networks

Here we are primarily concerned with network operators or ISPs that provide wavelength services to customers (e.g. other ISPs, large institutions). There are a number of applications where advance reservation is preferable to dynamic immediate reservation or static provisioning of lightpaths. For example, offsite backups or large data transfers can be scheduled overnight using advance reservation. These demands can specify a window or deadline to allow the network to choose the best start time. The knowledge of future network state and the new request's holding time allows the network to make better decisions compared to immediate reservation requests, especially for large demands which are difficult to allocate. Many real-time streaming applications that require large amounts of bandwidth can also benefit from advance reservation. IPTV, video conferencing, and video on demand are all examples of these applications. As a specific example, telepresence is currently being offered by Cisco [21] and Huawei [22] as an HD video conferencing solution over IP. These applications are well-suited for advance reservation since video conferences are typically scheduled for specific times in advance and require some guaranteed bandwidth and delay. By definition, since advance reservation demands book-ahead, they will have higher priority over other demands, allowing the network to be able to make better service guarantees compared to immediate reservation.

Advance reservation can also be used to request more VPN bandwidth during peak hours. For example, a VPN may use static requests for minimum connectivity, advance reservation for peak hour or scheduled demands, and dynamic immediate reservation for unexpected increases in bandwidth.

In a similar manner, advance reservation can be used for logical topology reconfiguration (for details about logical topology configuration, see [24], among others). While provisioning a network, a set of static demands may be used to setup initial lightpaths of a logical topology for some ISP. The traffic across the network fluctuates, therefore logical topology must either be over-provisioned, which wastes resources while traffic demand is low, or use dynamic IR traffic requests when IP layer traffic demands exceed the initial capacity. Using dynamic IR traffic demands may result in request blocking which can cause congestion for the IP layer since additional resources could not be reserved. Often, these traffic fluctuations are predictable. Fig. 3 shows traffic from New York to Washington over a backbone link. It is easy to see a pattern of traffic fluctuations. According to Cisco, peak Internet hours carry 20% more traffic than non-peak hours [25]. Advance reservation provides a good solution to this problem. With advance reservation we can reserve extra capacity only when it is needed according to the predictable pattern.

B. Grid Networks

A Grid network is a collection of geographically distributed resources, such as storage clusters, super computers, and scientific equipment, that are accessible to users over a network. Examples of e-Science Grids include the Large Hadron Collider Computing Grid Project [26], the Biomedical Informatics Research Network [27], and the George E. Brown Network for Earthquake Engineering and Simulation [28]. These networks typically deal with the transfer of large amounts of data in the terabytes, petabytes, and soon exabytes range. When the Grid resources are connected by application-configurable optical paths, the Grid can be considered a LambdaGrid [29]. These networks are an example of "service-oriented" networks in that they allow applications to directly request optical bandwidth




Fig. 3. Traffic on New York to Washington link of the Abilene backbone network from April 3, 2003 to April 10, 2003 [23].

resources. When we refer to Grid networks from now on, we will be referring to LambdaGrids.

There are a number of reasons that it is beneficial to provide advance reservation services for Grid applications. Since the traffic in a Grid is completely user driven, often times bandwidth requirements and request durations are known in advance due to requests being for specific tasks. Advance reservation requests allow applications to ensure network resources are available when certain computing resources are. Users may have access to certain Grid resources for specified times in the future. In order to access these resources, the user must be able to receive guarantees about network availability. This is known as resource co-allocation.

Also, many Grid applications involve delay-tolerant background or recurring tasks. For example, once a scientific instrument finishes an experiment, the data set usually must be transferred to other sites over the Grid. Instead of issuing these transfers as immediate reservation requests, the user can submit them as advance reservation requests that specify a deadline or window in which the transfer must take place. By providing advance reservation for such tasks, the network can achieve higher utilization while increasing the probability that the Grid applications will be able to successfully reserve the required network resources.

Collaboration is an important part of large scale scientific computing. Advance reservation can support real-time collaboration through real-time experimentation or high-definition video conferencing. It is easier to allocate these requests by booking-ahead instead of using immediate reservation.

There are a number of optical Grid networks that are beginning to, or already have, incorporated some form of advance reservation. We will discuss these in more detail later, but they include the U.S. Department of Energy's ESnet [30], the NSF funded EnLIGHTened project [31], the Japanese G-Lambda project [32], and the European Union's PHOSPHORUS project [33].

III. ADVANCE RESERVATION CLASSIFICATION

In this section we define advance reservation and consider the variations that have been presented in the literature. There are two defining characteristics of advance reservation requests. First, the holding time must be explicitly declared or must be able to be calculated based on other information. For example, a request may specify a file size, which can then

be used to determine the holding time. Second, the deadline, or the end of the data transfer, must be greater than then request arrival time plus the holding time. In other words, the transmission of data does not need to start immediately at the request arrival. This broad, informal, definition is able to classify a wide range of similar work as advance reservation, though different terminology has been used in the literature.

The two most common terms used for these types of demands are *advance reservation* and *scheduled demands*. Schedule demands, or scheduled traffic, is typically used when describing static traffic demands whereas advance reservation is typically used when describing dynamic traffic, particularly in Grid related papers. We will use the term advance reservation throughout the survey. Advance reservation can be classified into several types as denoted by [20]. Demands that specify a start time and duration are denoted STSD, demands that specify a start time but no duration are STUD, and demands that specify a duration but no start time are UTSD. Most research work assumes STSD advance reservation demands. STUD may be used when the user wants the network resources for as long as possible. UTSD may be used when the user requires service as soon as possible or with an undefined start time. We extend this classification in Fig. 4 and provide examples of each.

Before doing so, we define some terms. The *horizon* is the time range from the current time to the latest available time that the network allows resources to be reserved. The *booking-ahead time* is the time difference between the requested start time and the current time (the request arrival time). In the following subsections we assume we are given the network, $G = (V, E, W, H)$, where V is the set of switches, E is the set of links, W is the set of wavelengths available on each link, and H is the horizon. We will consider request tuples that describe each type of advance reservation. For traditional unicast *immediate* reservation, we can describe a request by a two-tuple, (s, d) , where $s, d \in V$ are the source and destination nodes, respectively.

A. STSD Requests

These advance reservation requests specify both a start time and a duration. The user may specify a fixed start time, meaning the request must start at the specified time, otherwise it is blocked. This request can be described as (s, d, α, τ) , where s and d are the source and destinations, $\alpha \in H$ specifies




Fig. 4. Extended advance reservation classification based on [20].

the start time, and τ is the duration. An example is given in Fig. 5(a). Assuming the current time is t_{now} , the figure shows that the request books-ahead some time in the future for a specific time and specifies its duration. Typical uses for this type of request are real-time streaming application. For example, setting up a high-definition video conference would require a specified start time and duration.

Fig. 5(b) shows another variation of STSD requests, STSD with flexible window. Instead of specifying a single start time, the user specifies a range of start times. This request can be defined by (s, d, L, R, τ) , where L and R are the initial and end times of the window. The request must be able to fit within the window, so we have $R - \tau \geq L$, and the request can start at anytime within this window. Flexible advance reservation requests can be used for large file transfers. The user may specify a window that allows the transfer to be scheduled anytime overnight. This added flexibility allows for efficient resource usage and lower blocking, as we will discuss later.

B. UTSD Requests

UTSD requests specify a duration and some deadline by which the request must be completed. The user does not explicitly state a start time. Deadline-driven requests can be described as (s, d, D, τ) , where $D \in H$ is the deadline and τ is the duration. There may or may not be incentives to minimize the delay between the request submission and the start of the data transmission. As with flexible advance reservation requests, the main motivation here is for large file transfers. Because a start time is not specified, it is possible to vary the bandwidth used by deadline-driven requests overtime, as long as the deadline is still met.

C. Variations

In this section we discuss some variations that have been proposed. These variations can be applied to any type of advance reservation request. The first we discuss is delay tolerance. Fig. 2(b) shows that after submitting an advance reservation request, the user gets an answer immediately. Alternatively, the user can specify a *delay tolerance* that allows the network to queue the request for some amount of time. This approach has two advantages. First, if the request would have been blocked, we can instead queue it in hopes of resources being freed before the delay tolerance ends. This is applicable if there are requests in the network that do not announce holding times (e.g. immediate reservation requests). Second, if enough requests specify a delay tolerance,

the network can perform batch optimizations, where multiple requests are scheduled at once instead of handling them individually. This should provide efficient solutions. Delay tolerance can be applied to any type of advance reservation request.

There have also been proposals for variable bandwidth advance reservation. This is also known as malleable or elastic reservations. In this case, the allocated bandwidth changes as a function of time. This can be taken to the extreme where it is allowed to send no data at all within some time frame. This is known as non-continuous advance reservation. For example, a request may specify a file size and a deadline and the network is free to assign different bandwidth at different times.

IV. NETWORK ARCHITECTURES AND IMPLEMENTATION

In this section we discuss network architectures and implementation issues to support advance reservation. We consider two broad classes of architectures. One is a centralized architecture where a single entity is responsible for handling incoming requests, scheduling, and configuring switching elements. The other option is a distributed approach where each node maintains some information and makes decisions independently when receiving a request.

In addition to deciding between centralized and distributed architectures, we must also take into consideration the length of the horizon, which determines how far we allow requests to book ahead. This impacts the amount of state information we must maintain. Another option to consider is whether or not the time-domain is slotted or continuous. If it is slotted, the duration of a timeslot is an important characteristic.

A. Centralized Architectures

Most work summarized in this paper consider centralized architectures. In this type of architecture, a centralized scheduler is responsible for call admission. The users (or applications) may interface with the scheduler through a web service API or extensions to the OIF User Network Interface (UNI) [34], for example. The scheduler authenticates the user to ensure they have proper credentials and permissions for the requested resources. The scheduler maintains global topology information and it uses this information to perform RWA for incoming requests. The scheduler is responsible for sending control messages to the network devices to reconfigure the switches (e.g. when a reserved request is about to begin). This can be accomplished with protocols like RSVP-TE. Similar mechanisms are used to tear-down requests. There is no need to maintain state information in the network switches for this architecture and no internal routing protocols (e.g. OSPF-TE) are required because the centralized scheduler handles all requests. This can greatly simplify the control plane. Another advantage of the centralized approach is that more complex algorithms can easily be incorporated and used. Synchronization is not required among switches in the network since the centralized scheduler sends out control messages when the switches must be reconfigured.

The downside of a centralized architecture are that handling link failures may be more difficult since nodes do not constantly send link-state updates. Centralized architectures are

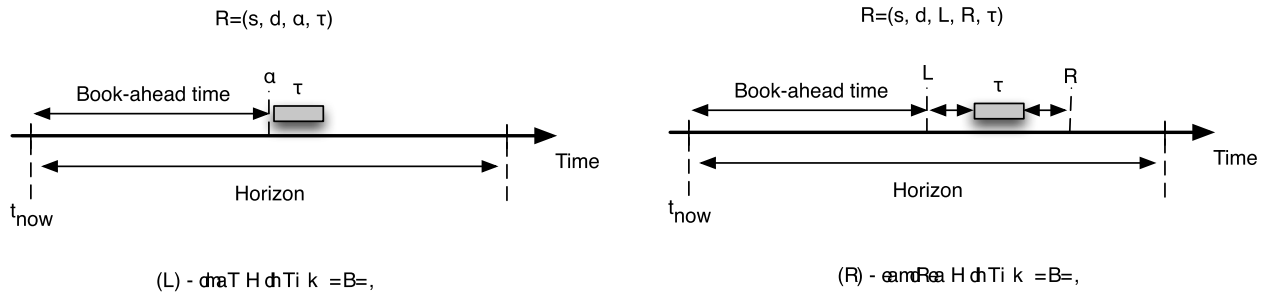


Fig. 5. STSD based advance reservation demands. Fixed window (a) specifies a single valid start time and duration while flexible window (b) specifies a time window within which the transfer must be completed.

typically considered impractical for WAN networks where the network must handle a large number of requests. A centralized approach is practical for LambdaGrids due to the relatively small number of resources and requests. There also must be replication in case the scheduler fails.

B. Distributed Architectures

The authors of [35], [36] provide some discussion about supporting advance reservation under a distributed architecture. In order to support a distributed architecture, each node must maintain some state information and must be able to perform path computation. Each node in the network could have an electronic controller that maintains state information. The controller must maintain state information about each wavelength-link incoming and outgoing from that node. In [36] this information is stored in the form of interval vectors. Each vector represents a gap (unused bandwidth) in the time domain (they assume the network is not time-slotted). In a time-slotted network, each node would have to maintain state information about each slot on each link.

The GMPLS signaling (RSVP-TE) and routing (OSPF-TE) protocols must be extended to support advance reservation demands. They need to incorporate time domain information from the reservation requests as well as link state updates. Using a modified OSPF-TE, each node in the network would know the global topology from link state updates. Using this information, the network could perform source routing for each advance reservation request. The impact of the additional temporal information on traditional RWA techniques such as fixed, fixed-alternate, and adaptive routing is discussed in [36]. Once the path computation is complete, the source can send RSVP-TE reservation (with time domain information) messages along the path. Each node updates its state information and sends out link state updates. The nodes are responsible for reconfiguring the optical switches when a reservation is about to start. Three phases of signaling, reservation, intermediate, and utilization, are proposed in [36].

More recently, [37] and [38] proposed detailed distributed routing algorithms to support advance reservation. Although the work is not directly applied to optical networks, the algorithms could be extended. A distributed distance-vector based algorithm for supporting advance reservation is proposed in [37], which discusses the state information and messages exchanged between nodes. The goal is to find the earliest possible start time for each request in order to minimize delay.

It is proven that in order to realize this, widest path routing in combination with path switching must be used. A novel loop-free distributed widest-path routing algorithm is proposed and shown to converge in finite time. Using the tables computed by this algorithm, a scheduling algorithm then finds the earliest start time for each request.

Alternatively, in [38] the authors propose modified link state routing algorithms for advance reservation routing. They assume that the nodes in the network use a modified OSPF type protocol. They propose modified link state data structures to incorporate the time dimension as well as update triggering policies for the link state updates. When a request arrives in the network, the source node uses this information to compute a path using a load-balancing technique. The source node then uses RSVP-TE to setup the path.

In the case of a distributed architecture, the nodes must be synchronized since each is responsible for configuring its own switches at the proper times. There may also be significant control plane traffic compared to a centralized approach.

C. Time Domain

An important topic in advance reservation is the management of the time domain. In [39], the authors classify two broad categories of resource management in the time domain. The first is a *reservation-based* approach which uses set of already accepted reservations for the admission control of an incoming reservation request. All provisioned requests that overlap the requested time interval of the current request are identified. In doing so, one can determine if enough resources are available to fulfill the current request. This method has low memory consumption as it only stores accepted requests which are needed for connection establishment. However, if one reservation request is handled after the other, up to $(i - 1)$ accepted requests have to be considered in the worst case when the i^{th} request is handled. This means that the time complexity to determine the available resources for n subsequent requests is $O(n^2)$. As a consequence, this approach is favorable if the number of requests is low. To cope with this complexity, a *timeslot-based* approach is introduced that maintains aggregated resource consumption information. Here, the time-domain is broken into a set of timeslots, which hold information about what resources are used or unused. The timeslot-based approach can further be classified as *static* or *dynamic*. The static timeslot approach breaks the time-domain into a fixed number of timeslots of constant length.

The amount of state information is independent of the number of requests and this approach is easy to implement. However, it is inefficient for networks with a small number of reservations. The dynamic timeslot approach allows the duration of a timeslot and the number of timeslots to change depending on the number of reservations in the network.

In addition to these two approaches, there is also the granularity of the time-domain to consider. *Infinitesimal granularity* allows the user to specify any starting/end-time while *non-infinitesimal granularity* forces requests to lie within some defined boundaries. In [39], the authors evaluate the performance impact of these options through analytical modeling and simulation.

The majority of works that we will discuss use the static timeslot-based approach, especially when considering a dynamic traffic model. We will refer to this case as a time-slotted network. On the contrary, some works, particularly work for the static traffic model, considers the reservation-based approach with infinitesimal granularity. We will refer to this approach as continuous-time.

In addition to determining how to manage the time-domain, the size of the horizon must also be specified. The length of the horizon also impacts the amount of state information and how far ahead requests are allowed to reserve bandwidth. In, [40], the authors examined both of these issues through analytical modeling of a simplified advance reservation model. They use a single link divided into a number of channels and define two types of advance reservation. One where the user specifies n starting timeslots (BA- n) and another where the user does not specify a starting time but instead accepts a range of possible start times (BA-all). The authors find that in the slotted time case, both types perform about the same, but for unslotted time case BA- n does not perform as well as BA-all. They also discuss that in their model, the required length of the horizon grows linearly with the average holding time. This work is for a simplified model and there are no studies on the traditional advance reservation requests discussed previously over wavelength-routed optical networks.

V. ADVANCE RESERVATION SURVEY

In this section we begin the survey on advance reservation. As discussed in the previous section, some authors use different terminology, but throughout this paper, the terminology introduced above will be used. We classify the work into two categories, those dealing with dynamic traffic demands and those dealing with static traffic demands. All of the work is summarized in Tables I-IX. We also discuss testbeds and frameworks as well as some work related to advance reservation scheduling, particularly in Grids. We defer discussion of network and implementation issues until Section IV.

Advance reservation for optical networks was first proposed by Zheng and Mouftah in [20], [11]. As mentioned earlier, they provide the initial classification of STSD, STUD, and UTSD requests. While they were the first to propose dynamic AR request for optical networks, Kuri et al. were the first to propose the static AR problem where the request set is given a priori [41], [23]. They focus on STSD AR requests and present heuristics and meta-heuristics to solve the static problem.

A. Dynamic Advance Reservation

We now begin our survey by discussing the work dealing with STSD fixed window requests. The studies in [20], [11] present simple heuristics for the STSD fixed window problem. They assume the network is under centralized control and the time-domain is broken into fixed timeslots. Each request requires one wavelength. They also assume no wavelength conversion. In [11], they use a fixed routing scheme where k -routes are precomputed. Each route is checked for a wavelength common to each link for each time interval in the duration of the request. In [20] an adaptive routing approach is used. This removes any links not available during the required fixed window from the network. After this step, the algorithm tries to find a path with the remaining links and assign a wavelength if there is a common wavelength available along the path.

Naiksatam et al. propose heuristics for STSD fixed window requests requiring multiple wavelengths [42]. They also assume a network under centralized control, fixed sized timeslots, and wavelength continuity constraint. In order to handle multiple wavelength requests, the heuristics proposed either concentrate all required wavelengths on a single path or spread them over multiple paths. k edge-disjoint paths are precomputed. For wavelength balancing, as requests arrive, the lightpaths are assigned on the first wavelength of the first path, first wavelength on the second path, and so on. Once all paths are examined, the algorithm checks the second wavelength on all paths. On the other hand, the wavelength concentrating algorithm tries all wavelengths on the first path, then all wavelengths on the second path, and so on. Both algorithms terminate once enough lightpaths have been allocated for the request. Results show that in networks where all links are requested uniformly, wavelength concentrating performs best, otherwise wavelength balancing should be used. In this work the authors also introduce an advance reservation traffic generator, the Flexible Optical Network Traffic Simulator (FONTS). Later in [43] the authors derive a simple analytical model for STSD fixed window requests. They model a single network link and assume each request requires a single timeslot (though they can use multiple wavelengths).

Wallace et al. apply lightpath migration to STSD fixed window requests [44], [45]. They assume a network under centralized control with the wavelength continuity constraint. The time-domain is not broken into discrete timeslots. The basic idea behind lightpath migration is to reassign resources to reserved lightpaths that have not yet begun transmission, in order to accommodate a newly arriving request. Two cost functions are evaluated. One that minimizes the number of existing paths that must be migrated and another one that minimizes the path length of the new request (with no restriction on how many existing requests will be migrated). To do this, they construct auxiliary graphs for each wavelength in the network and assign edge weights according to the required cost function. Given these auxiliary graphs, Dijkstra's algorithm is used to find a lightpath. The results show that there is no significant difference between the two cost functions for reducing blocking probability. The overall improvement compared to no migration is up to 23%.

We now discuss papers that propose RWA for STSD flexible window requests. Tanwir et al. consider RWA for STSD flexible window requests in networks with full wavelength conversion [46]. In this work it is assumed that the time is slotted into fixed length slots and that the network is under centralized control. In addition to traditional STSD flexible window requests, the authors also analyze the scenario with a non-blocking scheduler. In this case, instead of blocking a request that cannot be scheduled, the request can be moved outside its window until it can be scheduled. Two different routing strategies are proposed, both using k precomputed routes. The routes are computed by first selecting the shortest-path route, then checking wavelength availability on each link. If there is any link with no wavelengths available, it is removed and the shortest-path route is recomputed. This is done until a path is found or k links have been deleted. In the first strategy: Slide Window First (SWF), the algorithm tries all possible starting timeslots on one path and then moves to the next path trying all timeslots in order, and so on. The second strategy: Switch Path First (SPF), loops over the start time slots first. The algorithm tries the first start timeslot on path 1, then path 2, and so on up until path k . If the request cannot be accommodated, the next timeslot is checked. The algorithms are also modified to include load balancing, where the cost of the link is based on the number of wavelengths currently used. Given a path in the network, each link must be assigned a wavelength. They propose different wavelength assignment strategies to minimize fragmentation in wavelength usage on each link. The strategies are first-fit, min-leading-gap, min-trailing-gap, and best-fit. They also propose a network optimization technique where batches of requests (that have not begun data transmission) are re-scheduled periodically in order to find a better schedule. The re-optimization is simple in that it just tries to reassign each request sequentially according to earliest start time. However, the results show that it had little impact on blocking probability. The load-balanced version of their algorithms performed the best along with minimizing leading or trailing gaps for wavelength assignment. Moreover, the SWF algorithm performs slightly better than SPF since it favors shorter paths. In addition to the above, restoration techniques are also proposed, which will be discussed in more detail in the following subsection.

The authors in [47] investigate STSD flexible window requests with and without wavelength conversion using a continuous-time model. The state of each link is maintained by recording the times when the available bandwidth changes. For each incoming request, a start time list is computed for each wavelength/link. Next they compute a vector of start times that can be used to accommodate the request. The authors investigate two RWA algorithms. The extended Bellman-Ford (EBF) algorithm finds the lightpath that uses the shortest-path and the list sliding window (LSW) algorithm finds the lightpath with the earliest possible start time. These algorithms are compared to the algorithms proposed in [46], which were extended for the continuous-time model. The algorithms in [46] were not guaranteed to find a solution if one existed (i.e., the routing does not take wavelength availability into account until after a route is found), while the algorithms present in [47] make this guarantee. First-fit assignment is used when the wavelength

continuity constraint is also assumed (a layered wavelength graph is used for the routing algorithms) and min-leading-gap from [46] when wavelength conversion is assumed. A deferred wavelength assignment technique for networks with wavelength conversion is proposed. In this case, the actual wavelengths used on the lightpath are not selected until the request begins transmission. This reduces the complexity of the algorithms while not degrading performance. Results show that the EBF algorithm performed the best.

Shen et al. investigate both fixed and flexible STSD requests in a time-slotted network with no wavelength conversion. In their work [48], [49], [50] they propose RWA heuristics and also use a re-optimization technique. As for their RWA algorithms, k -shortest-paths are precomputed and a slotted first-fit wavelength assignment is used. For each possible start time (fixed requests only have one start time), and for each path the first wavelength available for the duration of the request (slotted first-fit) is added to a solution pool. Once all paths and start times are scanned, a lightpath is selected based on an objective function. The first objective function minimizes the path length and the second minimizes the load (load-balancing). If a request would be blocked, re-optimization is performed. Given the blocked request, all scheduled (but not yet transmitting) requests that overlap with this request in time are found. These requests are then ordered and RWA (with the load-balancing objective) is performed for each request one by one. If they can all be re-routed, then the new request is accepted, otherwise it is still blocked. For flexible window requests, this process is repeated for each possible start time. The set of overlapping requests are ordered by increasing start time, increasing minimum hop path, and increasing service durations. In addition to re-optimization on request arrival, the authors also propose periodic background re-optimization, which is performed before the start of each timeslot. The results show that re-optimization at blocking can improve performance by up to 50% (for a 7:3 ratio of fixed and flexible window requests) while periodic re-optimization has little impact (around 6% improvement).

In [51], the authors extend the work in [50]. In [50], re-optimization is done at blocking only. In [51] continuous re-optimization is proposed. To accomplish this, two independent algorithms are used that run in separate threads. One algorithm is used to schedule user requests when they arrive. This algorithm is based on the slotted first-fit algorithm in [50]. The other algorithm is a genetic algorithm that continuously tries to improve the requests that have already been scheduled. Both algorithms work on their own copy of the network state information. If the genetic algorithm finds a better solution at the end of a timeslot than the current solution the greedy algorithm has, the state information is copied over from the genetic algorithm to the greedy algorithm. The results show that this continuous optimization approach improves upon the performance of re-optimization at blocking.

Andrei et al. [52] consider deadline-driven requests for distributing data from multiple sources to a single destination. They consider a network with opaque OXCs (wavelength conversion is allowed) and sub-wavelength granularity requests, meaning grooming is also performed. A request specifies a destination (e.g. supercomputer), a set of files that must be

transferred with their source nodes, and a deadline by which they must all be transferred. For each file they must choose a route, wavelength assignment, start time, and must also groom the traffic. The problem is formulated as an MILP and present several objective functions. A heuristic is proposed for the problem that uses k -shortest paths for each file in the request. Given the paths, the heuristic finds the earliest start time possible. Different routing metrics based on hop count and congestion are evaluated as well as different wavelength assignment policies. These include random assignment, first fit, and integrated, where all wavelengths are examined as part of the RWA algorithm. The authors also propose to partition the files into pieces and schedule the pieces individually. The results show that the heuristic performs better than the MILP and that partitioning can provide a small performance improvement compared to the heuristic.

The authors in [20], [11] also present algorithms for UTSD and STUD in addition to STSD-fixed that we discussed earlier. For UTSD, the request specifies a holding time only. The authors assume there is a maximum ending time, which is equivalent to the deadline-driven model. They use the same algorithm for STSD-fixed, but now is run for every possible start time. The UTSD algorithm uses either fixed or adaptive routing. The lightpath with the earliest start time is selected. For STUD demands, the user specifies a start time but no duration. The authors assume that a request instead specifies a minimum duration and that there is an upper bound on the end time. The goal is to maximize the actual duration with the constraint that it must be at least as long as the minimum duration. They propose fixed routing algorithm in [11] examines all start times and durations on the precomputed routes. If one is found with a common wavelength it is stored. Once this is complete, the lightpath with the longest duration is selected. The adaptive routing algorithm [20] is similar except that it starts by removing all links with no available wavelengths and then computes k -alternate paths dynamically.

We summarize the papers discussed so far in Table I. In the network assumptions we denote wavelength continuity constraint as WCC and wavelength conversion as WC. We also differentiate between work that assumes time-slotted and continuous-time for the time-domain. Next we classify works with dynamic traffic into more specific topics, which are survivability, anycast and multicast, multi-domain, and quality of service.

1) Survivability: In this section we discuss work related to survivability of dynamic advance reservation demands. There are two approaches for providing survivability against link failures. One is known as protection where backup resources are provisioned along with primary resources. This increases resource usage but recovery time is very fast. The other approach is restoration where backup resources are found dynamically after a failure occurs.

As we discussed earlier, the authors in [46] also propose survivability for STSD flexible window requests on networks with full wavelength conversion. The authors propose to use a restoration technique. When a link fails, the requests that are currently active need to be restored. There will also be a set of future requests that have already been scheduled, but not yet active, using the failed link. The authors try to determine

how far into the future these requests should re-scheduled. The link will eventually be restored, therefore it may not be necessary to re-schedule all future requests. They define the re-routing interval as the amount of time after the failure for which they will re-schedule existing requests. In the paper, three re-routing intervals: a fixed interval, an adaptive interval based on the duration of past failures, and an unlimited interval where all future requests are re-scheduled are evaluated. The results show that the adaptive interval performs the best.

Cavdar et al. propose using delay tolerance with deadline-driven demands [53], [54], [55]. In this case, the deadline is the delay tolerance plus the duration of the request. When the user submits a request, the user specifies the duration and a delay tolerance. The request can start anytime between the arrival time and the arrival time plus the delay tolerance. The authors also provide survivability through shared path protection. A heuristic to find a backup path for each arriving request is proposed. As requests arrive, if they can be provisioned given the current state, the resources are setup and the connection begins. Otherwise, the requests are added to a queue to be processed later when the network's state changes, i.e., when a currently reserved request leaves the network. Requests can stay in the queue up until the customer's delay tolerance. To prioritize some requests over others, different queuing priorities are proposed. The priorities investigated are based on the arrival time, the holding time, and the customer's delay tolerance. The results show that at high loads, prioritizing by smallest holding time is the best whereas for smaller loads, prioritizing by delay tolerance performs the best. In [55] the authors also provide an analytical model of a single link network with multiple wavelengths using the user's delay tolerance.

2) Anycast and Multicast: The work up to now has dealt with unicast requests where a request specifies a single source and a single destination. There has also been work in AR for both anycast and multicast. In unicast, data is transferred between a specified source and destination. In anycast, a candidate set of destination nodes is given along with the source. Out of this set, a single node must be selected as the destination. In multicast communication [56], a single source transmits data to multiple destinations simultaneously. A destination set is given in the request and data must be transferred to all nodes in the set.

Anycast has been proposed in the context of Grid computing. Here, the focus is on scheduling both a node and the underlying network resources to perform some computation or storage task. The request specifies a set of candidate nodes and one of these nodes must be chosen. In addition to maintaining temporal information about all links, information about resource/node availability in the Grid must also be stored.

The authors of [57] consider deadline-driven anycast requests. The user submits a request specifying the number of CPUs, computation time, data transfer size, and a deadline. The network must then select a node, a lightpath, and a start time for this request so that it can finish by the deadline. They assume a grid network under a centralized scheduler where clusters of nodes have specific computational power. The time-domain is divided into discrete sized timeslots. The

TABLE I
SUMMARY OF AR RWA WITH DYNAMIC TRAFFIC.

Reference	AR type	Network Assumptions	Summary
[20]	STSD-fixed, UTSD, STUD	time-slotted, WCC	Adaptive routing heuristics. Remove links with no wavelengths available, compute routes for each start time. No performance evaluation. Brief discussion of co-existence of IR and AR.
[11]	STSD-fixed, UTSD, STUD	time-slotted, WCC	Static routing heuristics. Given precomputed paths, check all wavelengths and start times. No performance evaluation. Brief discussion of co-existence of IR and AR.
[42]	STSD-fixed	time-slotted, WCC	Users request multiple wavelengths. Heuristics use static routing and spread wavelengths across paths (balancing) or put as many wavelengths as possible on a single path (concentrating). Concentrating better when all links request uniformly.
[44], [45]	STSD-fixed	continuous-time, WCC	Allow previously reserved requests that have not begun transmission to be migrated to new lightpaths given arrival of new request. Create auxiliary graph with edge weights to minimize hops of new request or minimize number of migrated lightpaths. Both cost functions perform similarly, providing up to 23% improvement to no migration.
[46]	STSD-flexible	time-slotted, WC	Use adaptive routing by computing shortest path, removing any link with no available wavelengths, then re-computing path, k times. Two strategies. Try all starting times on one path before going to next path (SWF) and loop over paths before going to next start time (SPF). Propose four wavelength assignment techniques. SWF performs slightly better than SPF.
[47]	STSD-flexible	continuous-time, WCC and WC	Dynamic routing with goal of finding shortest path (EBF) or earliest start time (LSW). Compare to [46]. Find EBF performs the best. Propose technique called deferred wavelength assignment, reduces time complexity with no worse performance.
[48], [49], [50]	STSD-fixed/flexible	time-slotted, WCC	Static routing, check all start times, chooses first available wavelength. If multiple solutions found, use two cost functions, LB and MWL. If request would be blocked, do re-optimization. For all overlapping requests, order them and do RWA in with LB metric. Propose periodic re-optimization in background. Up to 50% improvement for re-optimization and 6% for background re-optimization.
[51]	STSD-fixed/flexible	time-slotted, WCC	Propose continuous optimization as an improvement to the work in [50]. A genetic algorithm runs continuously in a separate thread attempting to improve reservations in the network. At end of each time slot, if genetic algorithm is successful, the updated state information is copied to the normal RWA algorithm that handles requests as they arrive. Find improvement compared to re-optimization at blocking only.
[52]	Deadline-driven	continuous-time, WC	Proposed new data aggregation problem with grooming. Compare an MILP formulation with a heuristic (DARP). Heuristic uses k -shortest paths. Evaluate different link costs, wavelength assignment policies. Show DARP is better than MILP (which only optimizes one dimension of the problem).

authors turn the problem of finding a route and Grid node into multi-cost routing problem and propose two algorithms. One algorithm is used for immediate reservation requests and the other for advance reservation. The algorithm for immediate reservation is shown to be optimal for the given link metrics and runs in polynomial-time for the specific problem (it is NP-hard in general). The authors present an optimal algorithm for advance reservation, but due to its exponential complexity present a polynomial-time heuristic. The authors show that advance reservation achieves lower blocking than immediate reservation.

A similar problem was also investigated by [58] (it also investigates the static problem, discussed later). The authors propose anycast RWA for STSD-flexible anycast requests. Each request specifies a time window, a duration, the computing resources required, and a set of candidate destinations from which one node must be selected. They assume a fixed number of computing resource types that are grouped together in different resource groups. A request specifies a type of resource and a resource group, along with the necessary time

information for advance reservation. Then the algorithm tries to find the lightpath and select a node from the resource group. In addition to a primary lightpath, it also selects a backup lightpath for shared-path protection. They do not assume that the time-domain is divided into fixed timeslots. They assume there are no wavelength converters in the network. The algorithm works as follows: first it selects a number of possible start times for the demand. For each start time, it tries to find a path to each possible resource node with sufficient capacity by creating a layered wavelength graph. If a path is found, a backup path is also computed. Given all of the lightpaths found, the lowest cost solution is selected. In this case, the cost function takes into account network costs, computing resource costs, and time costs.

Regarding the multicast case, [59] investigate STSD-fixed multicast AR requests on a network with no wavelength conversion. The authors use a pre-existing multicast routing algorithm and first-fit wavelength assignment. They also use load-balancing by increasing the costs of links with used wavelengths. The routing is done dynamically to take advantage of

TABLE II
SUMMARY OF AR RWA WITH DYNAMIC TRAFFIC AND SURVIVABILITY.

Reference	AR type	Network Assumptions	Summary
[46]	STSD-flexible	time-slotted, WC	See heuristic summary in Table I. They propose restoration techniques. After a link failure, some number of requests need to be re-routed before the link is back up. Evaluate re-routing requests within a fixed time interval after failure, an adaptive interval, and an unlimited interval. Adaptive performs best.
[53], [54], [55]	Deadline-driven	No time-domain state, WC	Shared path protection. Customer specifies a delay tolerance, which is how long customer will wait for blocked/accepted response. If cannot be accommodated immediately, added to queue. Authors evaluate different queuing priorities.

load-balancing. In [60], the authors propose the provisioning of AR and On-demand (i.e., immediate reservation (IR)) requests in a dynamic optical circuit switched network. They consider an adaptive routing strategy (Delay constrained shortest path routing (DCSP)) and employ a dynamic wavelength assignment policy with minor variations to the policy of [61]. They apply these strategies to a mesh network and consider traffic grooming of the AR requests in order to achieve higher network utilization. In [62], the authors extend this work by using a multipath provisioning capability by using a Link Capacity Adjustment Scheme.

Andrei et al. propose RWA for STSD flexible window multicast AR requests [63], [61]. They assume all nodes are opaque (which allows wavelength conversion) with traffic grooming capabilities. A centralized scheduler is used and the time-domain is time-slotted. Requests specify an arrival time, file size, transmission rate, and end time. The authors propose a number of different heuristics. For all heuristics, the routing of a tree is based on a previously proposed Steiner tree heuristic. Two heuristics are based on pre-computed route trees where one heuristic generates a single pre-computed tree while the other generates a number of random trees. Another heuristic uses dynamic routing and examines all possible start times of the request. In addition, the authors propose heuristics to divide the tree into multiple subtrees where each destination can be reached independently (in space and time). The authors also consider the case where data can be buffered at intermediate nodes on the tree (using some spare storage capacity) in the event that some links are not available during particular time slots. The authors propose a heuristic that breaks the file into equal-size pieces. These heuristics are compared to using separate unicast requests to provision a multicast request. Finally, modifications are proposed to some of the heuristics to work in all-optical networks (i.e. no wavelength conversion).

3) *Multi-domain Advance Reservation*: He et al. consider advance reservation across multiple domains [64], [65]. The work focuses on STSD fixed and flexible window requests. They propose an architecture to support advance reservation by incorporating the time-domain into the information passed between domains. Multiple domain-level paths are explored in parallel by a photonic interdomain controller (PIN). These paths are based on abstract links topology summaries exchanged between domains. The authors propose a peer-to-peer model of topology information exchange between do-

main. For each of these paths, a photonic domain controller (PDC) finds a switch-level path through its own domain and produces a 2-D grid representing the wavelengths and the time-domain for the selected path. This is combined with the grids created by other domains and a final wavelength/start time is selected. The proposed architecture is called the unified flexible advance reservation model (FARM). Results show a performance improvement brought by allowing time-domain flexibility. Additionally, [65] discusses the addition of immediate reservation traffic to the advance reservation traffic. This aspect is discussed in the following section.

4) *Quality of Service*: The work referenced so far has considered all traffic to be of the same priority. Nonetheless, Different priority schemes can be incorporated into advance reservation. In addition most work considers only advance reservation traffic to be present in the network. Because advance reservation demands book-ahead, they will have higher priority than immediate reservation requests. It is therefore necessary to consider the impact of advance reservation on immediate reservation requests and it is likely necessary to have some resource broker or admission control mechanism to ensure immediate reservation requests can achieve reasonable blocking levels in the presence of advance reservation requests.

The mix of immediate reservation and advance reservation is discussed in [11], [65]. Both identify the problem of *preemption*, where advance reservation requests may interrupt immediate reservation requests that are currently in service. This is primarily a result of not knowing the holding time of the immediate reservation requests. Both works discuss introducing resource partitioning, where different demands use different wavelengths, as well as requiring immediate reservation requests to at least specify a minimum duration. Another option is to limit the number of advance reservation requests accepted by the network (as a form of admission control). [65] evaluates these strategies and analyze the sharing of wavelengths and partitioning of wavelengths between AR and IR demands, requiring IR demands to specify a minimum duration, and limiting the amount of admitted AR demands.

Escalona et al. propose RWA algorithms specifically to deal with the mix of AR and IR traffic in optical networks [35]. They show that since AR requests book ahead, they can preempt IR requests that are still active as previously discussed. They propose RWA algorithms to help minimize preemption of IR requests with both STSD fixed and flexible window

TABLE III
SUMMARY OF AR RWA WITH DYNAMIC TRAFFIC FOR ANYCAST AND MULTICAST TRAFFIC.

Reference	AR type	Network Assumptions	Summary
[57]	Deadline-driven	time-slotted, single wavelength	Anycast requests specify number of CPUs, computation time, data size, deadline. Network selects node, lightpath, start time. Multi-cost optimal routing algorithms for immediate and advance reservation and heuristic for AR (optimal AR algorithm has exponential runtime). Show that AR achieves lower blocking than IR.
[58]	STSD-flexible	continuous-time, WCC	Anycast requests specify time window, duration, computing resources, and candidate destinations. For each start time, find path to each possible destination with layered wavelength graph. Given all lightpaths found, choose lowest cost based on network costs, computing resource costs, and time costs.
[59]	STSD-fixed	WCC	Multicast AR. Use existing Steiner tree heuristic. For each request, assign link weights for load-balancing, run Steiner tree heuristic, do first-fit wavelength assignment.
[63], [61]	STSD-flexible	time-slotted, WC	Multicast AR with grooming. Propose heuristics for static routing and dynamic routing using existing Steiner tree heuristics. Extensions: breaking tree into independent subtrees, buffering data at intermediate nodes, breaking file into pieces that can use different trees, using unicast to reach all destinations, and modifications for heuristics to work in all-optical networks. Show unicast performs poorly, proposed extensions all provide small improvements over base algorithms.
[60], [62]	STSD-fixed	continuous time, WCC	AR and On-Demand (i.e. immediate reservation) in circuit switched optical networks. Delay constrained shortest path routing with adaptive wavelength assignment. Use of multiple paths for request provisioning.

TABLE IV
SUMMARY OF AR RWA WITH DYNAMIC TRAFFIC AND MULTI-DOMAIN.

Reference	AR type	Network Assumptions	Summary
[64], [65]	STSD-fixed/flexible	time-slotted, WCC	Propose multi-domain reservation architecture. AR-PIN module computes inter-domain paths while AR-PDC module computes intra-domain paths. Discuss how the modules interact, different RWA strategies. Evaluate impact of flexibility, routing strategies, and impact of AR vs. IR.

requests. This is done using inverse wavelength assignment, where AR requests begin using higher-index wavelengths while IR use lower-index wavelengths. This minimizes contention between the two types of traffic without resorting to fixed resource partitioning. They also propose to select wavelengths for IR requests that have the largest gap until the next AR start time to further help reduce the probability that the IR request will be preempted when the AR request begins transmission. The results show that these techniques can significantly reduce the preemption probability for immediate reservation requests. In addition to the proposed algorithms, the authors also discuss results obtained from a testbed.

Most previous work dealing with mixed AR and IR traffic (in electronic and optical networks) consider AR to be higher priority and allow them to preempt IR traffic. The authors of [66] discuss giving IR higher priority than AR. For example, customers paying for premium service may want the ability to submit requests at any time (IR) and still receive good QoS. They propose assigning priority levels to every request, giving IR the ability to preempt STSD-fixed AR requests. They do not provide performance evaluation. The same authors also propose prioritized STSD-fixed AR, where different requests

belong to different priority levels [67]. This last work is not based on optical networks.

As we will discuss later, the authors in [68] discuss QoS among advance reservation requests with static traffic demands. They define two priorities. One priority requires path protection while the other does not. The lower priority demands can also use backup resources assigned to the higher priority requests.

The work for dynamic advance reservation is summarized in Tables I-V. In the network assumptions we state whether the network is time-slotted or not (continuous time) and whether wavelength conversion (WC) is present or the wavelength continuity constraint (WCC) is enforced. We also state any problems beyond RWA the paper investigates as well as a brief summary of the solution technique.

B. Static Advance Reservation

Kuri et al. were the first to propose the static STSD-fixed AR problem [41], [23]. They assume a continuous-time network with the wavelength continuity constraint. For all algorithms, k-shortest-paths are precomputed for all source-destination pairs. The first algorithm is a branch and bound

TABLE V
SUMMARY OF AR RWA WITH QOS/PRIORITIZATION.

Reference	AR type	Network Assumptions	Summary
[65]	STSD-fixed/flexible	time-slotted, WCC	Study impact of AR on IR demands. Evaluate resource partitioning where AR and IR use different wavelengths. Evaluate forcing IR demands to specify a minimum duration. Evaluate admission control for AR to limit number of AR demands in network.
[35]	STSD-fixed/flexible	time-slotted, WCC	Consider AR as higher priority and propose heuristics to minimize IR preemption. Wavelength assignment starts at opposite indices for IR and AR (IFW). In addition, when provisioning IR they select a lightpath with the largest period between the current time and the start of the first AR request. They show IFW+LP reduces IR preemption by up to 30% while having no negative effects on blocking.
[66]	STSD-fixed	time-slotted	Propose adding a priority field to IR and AR demands. Use same heuristic to allocate both (duration of IR demands are estimated). Search k static paths for path with available bandwidth. If none found, try to preempt request with lower priority on one of the paths. No performance evaluation.
[69], [68]	STSD-fixed	continuous-time, WC	See details in Table VIII.

algorithm to find the optimal routing (given the precomputed routes). Because of the exponential runtime, they also propose a tabu search meta-heuristic to find a set of routes. Optimality is determined by different cost functions the authors propose. Given the routes, they use a graph coloring heuristic to assign wavelengths. These two approaches, branch and bound and tabu search, consider the entire request set at once. They also propose a sequential heuristic which first orders the requests by a weight function and then performs RWA one request at a time using k -shortest-paths with first-fit assignment.

The authors of [70] propose heuristics to improve upon those presented by Kuri in [23]. Their work is focused on continuous-time network and the wavelength continuity constraint. The proposed heuristic first creates groups of time-independent requests. Given the time-independent sets, a wavelength is assigned to the first set and they are all routed independently. The requests can use the same wavelength-links since they are independent in time. Then the heuristic tries to route requests in the remaining sets on the same wavelengths. If there are still requests left over, a new wavelength is selected and the process is repeated. They show their proposed algorithm is faster as the sequential algorithm previously described for [23] while having similar performance.

Saradhi et al. propose two heuristics for STSD-fixed AR [71], [72]. Again they assume continuous-time and the wavelength continuity constraint. One is based on finding independent (in the time-domain) sets of requests using circular-arc graph theory. These independent requests can then share resources during RWA. The other approach attempts to find time-independence by iteratively dividing the time-domain into fixed partitions called windows. Requests that fall completely within separate windows can share resources since they do not overlap in time. RWA is done with layered wavelength graphs using shortest-path routing. Demands that are found to be time-independent can share a wavelength.

Another work dealing with the static STSD-fixed AR problem is [73]. In this work, an additional heuristic, a meta-heuristic, and a lower bound are proposed. The same

constraints are considered here as in previous works. The main motivation for the tabu search meta-heuristic is to improve upon the tabu search proposed in [41]. k -shortest-paths are precomputed and the tabu search explores the solution space of assigning different paths to the requests as in [41]. The neighborhood generation differs in this approach. It is based on graph coloring. Given the set of routes and the resulting conflict graph used for graph coloring, they determine which routes (nodes in the conflict graph) determine the upper bound for the chromatic number. The neighborhood is created by changing these routes. This significantly reduces the complexity of the tabu search. They also propose sequential heuristics that creates sets of demands that are independent in time *or* space. The idea is similar to that proposed in [70]. Lastly, they provide a theoretical lower bound on the number of wavelengths required for a given network and request set.

The authors of [74], [75] propose a heuristic for the STSD-fixed AR problem and compare their results to the sequential heuristic from [23] and the heuristic in [70]. The requests are split into time independent groups similar to [70], but they propose new routing techniques that try to force time independent requests in a set to share as many links as possible. This allows time overlapping requests from other sets to find link-disjoint paths easily and hence decrease the number of wavelengths required. Results show that the new heuristic outperforms the previous heuristics.

Wallace et al. propose STSD-fixed AR with request service time (RST) [76]. The request specified by the user is the same as STSD, but the actual start time can be after the requested service time. The goal is to minimize tardiness, which is defined as the difference between the actual start time and requested service time. The authors propose an MILP formulation that uses precomputed paths. In addition to this, they present two heuristics. One heuristic is sequential heuristic that handles each request one by one. The paths are precomputed and for each request the heuristic checks all possible lightpaths and chooses the one that minimizes the tardiness. The other heuristic is a simulated annealing meta-

heuristic with an objective function that minimizes the average tardiness of all requests. The perturbation function chooses a request at random and tries to reroute it on an alternative lightpath.

Chen et al. study the STSD-fixed AR problem in [77], [78]. In their problem formulation, the goal is to maximize total revenue from admitted calls. They assume the wavelength continuity constraint and a time-slotted network. The problem is formulated as an ILP using precomputed paths. In order to speed up computation, a Lagrangian relaxation heuristic based on this ILP is presented along with three other simpler heuristics which allocate requests sequentially by their revenue, start time, or end time.

Wang et al. investigate the problem of static STSD-flexible AR [79] with continuous-time and the wavelength continuity constraint. In this work, an algorithm to minimize the amount of overlap by assigning fixed start times to all demands in the request set is proposed. Once this is complete, one of two heuristics are used to schedule these fixed time demands. For one heuristic, the time-domain is divided into time windows such that demands in separate windows do not overlap. The time windows are found using a maximum independent set algorithm. Separate virtual networks are used for each time window. Once this is complete, the algorithm does RWA for each demand sequentially by order of largest capacity requirement. This is done by creating a layered wavelength graph created from the window's virtual topology. If a demand straddles two windows, the algorithm creates an intersection of the virtual topologies. If demands cannot be routed, the algorithm attempts to rearrange these demands at the end. The authors also propose a matrix-based algorithm which transforms the demands into a set of traffic matrices that do not overlap each other in time. Both algorithms are compared to a sequential algorithm and a tabu search algorithm.

The authors of [80] consider STSD-flexible requests over a single time-slotted WDM link with W wavelengths. They assume the connections are periodic, meaning the bandwidth may be required everyday for a period of one year, for example. An example is nightly backups where a certain amount of bandwidth is required at the same time (or time-window) each day. The paper develops a model that relates properties of the requests and traffic parameters to how many wavelengths would be required. In addition to the modeling they propose a number of heuristics.

In [81], Saradhi et al. extend their previous work ([71], [72]) to handle STSD-flexible requests. They use a similar algorithm to their previously proposed time-window algorithm but now they try to adjust the start time of a demand (within its sliding window) to fit the demand within the created time windows.

Andrei et al. study the STSD-flexible AR problem in [82], [83] on a time-slotted network with the wavelength continuity constraint. A nonlinear integer program formulation is proposed that maximizes the number of accepted requests given a fixed number of wavelengths. The formulation solves both the time schedule and RWA problem jointly, using a set of precomputed candidate paths. The authors also provide an alternative objective function that maximizes the total bandwidth accepted. The authors then propose two heuristic approaches that solve the time assignment and RWA problem

jointly, instead of separately as in most previous work. Their heuristics are based on sorting the request set, then for each request generating k -shortest-paths on a layered wavelength graph. The earliest start time is selected. One heuristic is a modified version of this that uses information from a Lagrangian relaxation of their integer program formulation.

The authors in [84] propose a new model of non-continuous demands. They assume that for certain types of requests, like large file transfers, the user does not need continuous service. The network can break the request into multiple smaller requests and schedule them independently as long as the user gets the amount of time required by some deadline. They assume a time-slotted network and consider a static set of STSD-flexible requests. The problem is formulated as an ILP with precomputed routes and wavelength conversion. They also propose a heuristic (with wavelength continuity) to solve the problem. A set of shortest-paths is precomputed. The heuristic processes the requests sequentially, allocating a timeslot on the least-congested path for each timeslot during the requests window. If each of the paths during the current timeslot is above a certain congestion threshold, that timeslot is skipped. If a request is unable to find available timeslots to cover its duration, the algorithm will then try to use the congested links in a later phase. In both cases (ILP and heuristic), different segments of a request can use different routes. Segments cannot be transmitted in parallel, however, so segment i must use a later timeslot than segment $i - 1$. Results show that non-continuous demands have better performance than traditional STSD-flexible demands.

The authors of [85], [86] propose algorithms for provisioning both static STSD-fixed demands along with dynamic IR traffic on a continuous-time network with no wavelength conversion. They propose two main strategies. One is performing RWA on the STSD-fixed demands first, then with the remaining resources handling the dynamic demands sequentially. The other option is to handle both STSD-fixed and dynamic demands sequentially. That is, RWA is only performed for the static demands when a demand needs to begin transmission.

In [87], the authors consider the case of bifurcated and non-bifurcated routing when requests require multiple lightpaths. In [88] the authors consider re-routing of dynamic lightpath demands if an incoming dynamic request cannot be accommodated. In [89] the authors examine the impact of wavelength conversion in terms of cost and performance with the two types of demands. They also investigate shared path protection for the two types of demands in [90] using the backup-multiplexing strategy previously proposed in [91].

The same authors then extend their work for the provisioning static traffic, static STSD fixed window demands, and then handling dynamic immediate reservation demands arriving over time [92]. They define three stages of provisioning. They are network planning, which occurs over long time periods (static demands), network engineering, which occurs daily, weekly or monthly (static STSD demands), and lastly traffic engineering which handles other requests (dynamic immediate reservation traffic). The authors propose using standard MILP approaches for the network planning problem, then consider either network and traffic engineering jointly or separately. When done jointly, the same RWA algorithms used

TABLE VI
SUMMARY OF AR RWA WITH STATIC TRAFFIC

Reference	AR type	Network Assumptions	Summary
[41], [23]	STSD-fixed	continuous-time, WCC	Pre-compute k -shortest paths. Use branch and bound to find optimal solution given paths. Propose tabu-search with different cost functions. Propose sequential heuristic with low time complexity. Tabu-search performs similarly to branch and bound algorithm.
[70]	STSD-fixed	continuous-time, WCC	Heuristic to create groups of time-independent sets. Assign a wavelength to each set, then look for spatial reuse. Compared to [23]. Similar performance with lower runtimes.
[71], [72]	STSD-fixed	continuous-time, WCC	One heuristic to find time-independence by using circular arc graph theory (ISA) and one heuristic to find time-independence by breaking the time domain into smaller windows (TWA). Independent requests use same wavelengths. For most metrics, TWA is best.
[73]	STSD-fixed	continuous-time, WCC	Improve upon tabu-search from [23] by improving neighborhood generation. Propose two greedy algorithms based on finding independent sets. Derive lower bounds on number of wavelengths required. Algorithms perform better than those in [23] with lower runtimes.
[74], [75]	STSD-fixed	continuous-time, WCC	Propose heuristic to find time-independent groups, including new routing technique to further reduce wavelengths required. Compare heuristic to those proposed in [23], [70] and find it provides better results.
[76]	STSD-fixed	WCC	Use request service time (RST) as start time. Requests can start after this, but objective is to minimize difference between start time and RST. All approaches use k pre-computed paths. Propose MILP, sequential heuristic that scans all lightpaths and start times, and a simulated annealing heuristic with two annealing schedules. Simulated-annealing provides solutions with cost close to the MILP.
[77], [78]	STSD-fixed	time-slotted, WCC	Propose heuristics to maximize total revenue of admitted calls. Use k pre-computed paths. Formulate as ILP and present Lagrangian relaxation based heuristic. Propose simpler sequential heuristic with different strategies to order requests. Lagrangian relaxation performs best with results close to optimal from ILP.
[79]	STSD-flexible	continuous-time, WCC	Two-step heuristic. First assign fixed start times to all requests, problem becomes STSD-fixed. Then use heuristic to schedule set. One based on finding independent requests (WA), other is matrix based (TMA). WA performs better than TMA. Also compare WA to a tabu-search heuristic. Performance is similar, WA has lower runtimes.
[80]	STSD-flexible	time-slotted, single link	Develop model that relates properties of request and traffic parameters to number of wavelengths required. Evaluate number of heuristics for single-link. Show moderate amount of flexibility significantly increases performance.
[81]	STSD-flexible	continuous-time, WCC	Extensions to their TWA heuristic proposed in [71], [72]. New heuristics try to adjust start times of demand when creating time windows.
[82], [83]	STSD-flexible	time-slotted, WCC	Formulate as MILP that jointly finds start times and lightpaths for all requests. Uses pre-computed paths. Propose two heuristics that solve time and lightpath assignment jointly. One based on Lagrangian relaxation of MILP. Solving problem jointly better than assigning times first then solving resulting STSD-fixed problem.
[84]	STSD-flexible	time-slotted, WCC and WC	Propose non-continuous transmission. Break request into smaller requests that can be scheduled independently on different lightpaths. Formulate as ILP using pre-computed paths and wavelength conversion. Propose a sequential heuristic based on congestion (solves problem with wavelength continuity). Show that non-continuous can improve performance compared to continuous.

for STSD fixed window demands are used dynamic immediate reservation traffic. When done separately, offline scheduling of STSD fixed window demands is done using a random search algorithm.

1) *Anycast and Multicast*: As we discussed previously, the authors of [58] propose solutions to the static STSD-flexible anycast problem as well. In this work the authors present an ILP, with a time-slotted network, that finds the optimal solution for a request set with the objective of accepting the maximum number of requests. The ILP does shared-path protection for each request. The authors also propose a heuristic that attempts to spread the demands to reduce the time overlap.

While there are still demands to be scheduled, the heuristic alternates between choosing the earliest starting demand and the latest ending demand in the set. When choosing a demand with the earliest start time, the algorithm tries to schedule the demand as early as possible within its window, and similarly for the latest ending demand, the algorithm tries to schedule the demand as late in its window as possible. The actual RWA for each selected demand uses the heuristic proposed for the case of dynamic traffic (with a modified cost function).

The static multicast STSD-fixed AR problem was first investigated in [93] on a continuous-time network with the wavelength continuity constraint. The authors assume that

it is possible to renegotiate the start times of the static requests. This allows them to accommodate the entire set assuming there are a limited number of resources because they can always move requests later in the time-domain. The first proposed heuristic sorts the requests according to their demand size, which is the product of the destination set size and request duration. A multicast heuristic is then used to route each request sequentially, creating an auxiliary graph that removes any resources in use by other requests that overlap in time with the current request. If the RWA is successful (first-fit is used), the algorithm continues with the next request. Otherwise, the request is tagged and no resources are allocated. It then continues with the next request. In a final phase, the algorithm attempts to allocate all tagged requests by changing their start times. The second heuristic is based on aggregating time-disjoint multicast requests that share a source node into a larger multicast tree, then routing that tree. A similar final step where rearrangement is performed can also be used here.

Charbonneau et al. also investigate the static multicast advance reservation (MCAR) problem for all-optical wavelength-routed WDM networks [94]. They develop two efficient heuristics, independent set heuristic (ISH) and simulated annealing (SA), to solve the problem. They introduce two theoretical lower bounds on the number of wavelengths required. They also formally show the problem is NP-complete and formulate an integer linear program [95] solution. Through simulations, they show that the SA heuristic provides up to a 21% improvement over ISH on realistic networks and SA provides solutions 1.5-1.8x times the cost given their conservative lower bound on large networks.

2) *Survivability*: Kuri et al. were the first to propose survivability for static AR. They considered the case of STSD-fixed requests in [91]. The goal is to provision a set of requests with primary and arc-disjoint backup paths such that the total number of resources required is minimized. They assume wavelength conversion in the network and continuous-time. Channel reuse is applied for protection. A channel can be reused as long as the requests using it do not overlap in time. The authors also introduce the concept of backup-multiplexing which allows a channel to be used by multiple lightpaths for protection assuming they do not overlap in time and share a link on their primary paths. Backup-multiplexing is more efficient than just channel reuse. Two heuristics are proposed in the paper, one that takes advantage of only channel reuse and one that takes advantage of backup-multiplexing. Both problems are formulated as combinatorial optimization problems (with pre-computed paths) and use simulated annealing to find solutions.

Li et al. propose several ILP formulations for the STSD-fixed AR problem with both shared and dedicated path protection in wavelength convertible networks [96], [97]. They consider a single-link failure scenario with the goal to minimize network resource usage. In the first paper, ILP formulations are used with precomputed paths and STSD-fixed is compared with traditional static traffic, showing performance improvement for the former. In the second paper, joint ILPs are proposed for the problems that perform both routing and wavelength assignment for shared and dedicated path protec-

tion. In addition to the joint ILPs, the authors propose two separate ILPs for each problem that solve the routing and the wavelength problems independently. The same authors extend their work in [98], [99], [100] proposing a heuristic solution for shared-path protection with wavelength conversion. In this case, a matrix-based method is used where the demands are first sorted, according to one of several policies. Then a technique called iterative survivable routing is used to solve the problem. This approach iteratively tries to change the working and protection paths using state information stored in matrices until no better solutions can be found. This is compared to their previous ILP formulations. In [101], [102] the same authors propose a sequential heuristic for the shared-path protection problem using the same network assumptions and compare the results to their previously proposed joint ILPs. The heuristic first sorts the demands according to one of several policies and then uses a modified Dijkstra's algorithm to find the working and protection paths for each demand sequentially. The link weights are modified based on sharing and time constraints to encourage a more efficient resource usage. Finally, the authors then extend their work to the more general shared risk link group (SRLG) failure model in [103] using a modified version of their previously proposed iterative survivable routing algorithm and matrix-based method. They only examine the case of shared-path protection.

Jaekel et al. have also investigated the static STSD-fixed AR survivability problem on a time-continuous network with wavelength conversion. In [69], [68] they introduce priorities, or service levels, into their problem formulation. Two priorities for advance reservation demands are considered. High priority demands require both primary and backup lightpaths. Low priority demands require only a primary path (no protection) and they can also be routed using resources allocated to backup paths of high priority demands. This means that in the event of a failure, some low priority demands may be preempted. ILPs for both shared and dedicated path protection with the goal of minimizing congestion are formulated making use of precomputed routes. In addition to the ILPs, the authors propose a heuristic for both shared and dedicated path protection. The heuristic, which also uses precomputed routes, sorts the demands in order of start time and processes them sequentially. The route selection is based on congestion. Wavelength assignment is then done based on priority and type of protection required. The same authors also considered the STSD-flexible AR survivability problem [104] on wavelength convertible networks with continuous-time requests. ILPs for both shared and dedicated path protection using precomputed routes are proposed. The ILPs jointly schedule the demands within their windows and perform RWA. The authors also propose a two-step optimization process, where an ILP is formulated to schedule the demands optimally (minimal overlap) in time, then another ILP or heuristic could be used for RWA. In [105] the authors consider networks without wavelength converters and propose ILPs (with pre-computed routes) for both fixed and sliding demands with the wavelength continuity constraint and with and without shared-path protection. A two-step optimization technique is proposed as in the previous paper where they schedule the demands optimally then use a heuristic similar to that of [68] to perform RWA. Their work

TABLE VII
SUMMARY OF AR RWA WITH STATIC TRAFFIC FOR ANYCAST AND MULTICAST

Reference	AR type	Network Assumptions	Summary
[58]	STSD-flexible	time-slotted, WCC	Anycast problem. Present ILP to maximize accepted requests. Propose Time-Spread-Algorithm (TSA) to minimize overlap between demands, then use heuristic discussed for dynamic case to route the demands.
[93]	STSD-fixed	continuous-time, WCC, full splitting	Multicast problem. Propose one heuristic that processes requests sequentially using first-fit and a Steiner tree heuristic. If requests cannot be allocated, tries changing start times (re-negotiation). Propose second heuristic that aggregates time-disjoint requests into larger multicast trees.
[94] [95]	STSD-fixed	continuous-time, WCC, full splitting	Multicast problem. Propose two heuristics that processes requests based on the concept of independent set theory and simulated annealing. Introduce theoretical lower bounds on the number of wavelength required. Prove that the problem is NP complete and formulate an ILP solution.

with STSD-flexible demands is combined in [106] where they also propose an ILP for dedicated-path protection in networks with no wavelength conversion.

The authors in [107] also investigate static STSD-fixed AR survivability. They consider time-slotted networks with wavelength conversion and shared-path protection. An ILP (using precomputed routes) is proposed for provisioning primary and backup paths for a set of requests. The ILP allows the backup paths to be changed at the start of each timeslot. They call this re-optimization and results show that this can help reduce resource usage.

In [108] the dual-link survivability using shared-path protection for STSD-fixed AR demands is investigated on a time-slotted network with wavelength conversion. The authors assume the first failure is permanent and attempt to reduce the number of requests that would be unable to recover from a second failure. They formulate the problem as an ILP with precomputed paths.

Heuristics for the STSD-flexible AR survivability problem are presented in [109] on a continuous-time network without wavelength conversion. The authors propose providing shared-path protection for the STSD-flexible AR requests. First, the request set is divided into time disjoint windows, similar to [72]. Given the fixed start times, two RWA algorithms are proposed to find primary and backup paths for each request. One considers finding routes on a single graph with multiple wavelengths per link while the other uses a layered wavelength graph.

There is also work in multicast STSD-fixed AR survivability [110] on a time-continuous network both with and without wavelength conversion. The primary and backup trees are precomputed and an ILP is formulated to minimize cost.

3) *Grooming*: A heuristic for grooming of static STSD-flexible demands is proposed in [112] for the case of a continuous-time network with wavelength continuity. The authors consider two priorities of demands, high and low. They use the same techniques to minimize overlap and schedule demands within time windows as their work in [79]. They then perform RWA on demands starting with the high priority demands and in order of largest requested capacity. Grooming is incorporated by modifying link costs of a layered wavelength graph for each time window. They also attempt to rearrange demands if necessary, as in their previous work. This is

compared to a tabu search algorithm with precomputed paths that explores the solution space by randomly swapping paths.

In [111] an ILP for grooming of static STSD-fixed demands that also performs dedicated path protection is proposed on a continuous-time network with wavelength conversion. The main goal is to create a stable logical topology that is capable of handling all requests during all time intervals. Time disjoint demands can share resources. The proposed ILP uses pre-computed edge-disjoint paths to find a stable logical topology (with path protection), and grooms the requests in the request set onto the logical topology. The authors provide two objective functions and compare the solutions from the ILP to traditional holding time unaware static demands. In [113], the authors extend the framework of [111], to include shared path protection and also consider the case of having no wavelength conversion in the network.

The book chapter [114] also discusses the problem of grooming electrical demands onto lightpaths and lightpaths onto wavebands with static STSD-fixed requests. The authors consider scheduled lightpath demands that request one or more lightpaths and scheduled electric demands that request bandwidth at the sub-wavelength level. ILP formulations are provided for the electronic and optical grooming problems.

VI. ADVANCE RESERVATION FRAMEWORKS

We now discuss the various advance reservation frameworks¹ and architectures that have been implemented. Some projects propose supporting advance reservation across a single domain, whereas other projects focus on providing inter-domain support as well. There is also work on co-allocation of network and Grid resources (e.g., computing/storage) using advance reservation. Table X provides a comparison of the different frameworks discussed below.

The On-demand Secure Circuits and Advance Reservation System (OSCARS) is a project by the U.S. Department of Energy that supports dynamic end-to-end provisioning of network resources (layer 2/3 VCs) with support for advance reservation [116], [30]. OSCARS is used over the DOE's Energy Science Network (ESnet). OSCARS is implemented as a centralized service that provides a web-services based API

¹Some parts of this section was presented in [115]. We present it here for completeness.

TABLE VIII
SUMMARY OF SURVIVABLE AR RWA WITH STATIC TRAFFIC

Reference	AR type	Network Assumptions	Summary
[58]	STSD-flexible	continuous-time, WCC	Discussed in Table VII. Heuristics and ILP find paths for shared path protection after provisioning primary path.
[91]	STSD-fixed	continuous-time, WC	Arc-disjoint backup paths. Define backup-multiplexing. Propose two simulated annealing heuristics, one uses backup-multiplexing and one does not. Use pre-computed paths.
[96], [97]	STSD-fixed	continuous-time, WC	Shared and dedicated path protection. ILPs for solving routing and wavelength assignment jointly and separately.
[98], [99], [100]	STSD-fixed	continuous-time, WC	Shared path protection. Proposes an iterative survivable routing heuristic. Start by sorting according to some policy, then iteratively change working and protection paths based on state information to improve solution. Compared to previous ILP formulations in [96], [97].
[101], [102]	STSD-fixed	continuous-time, WC	Shared path protection. Sequential heuristic that sorts demands according to some policy. Use Dijkstra with modified link weights to find shared and protection paths. Compares to previously proposed joint ILPs [96], [97].
[103]	STSD-fixed	continuous-time, WC	Shared path protection with shared risk link groups (SRLGs). Modify the previously proposed iterative survivable routing technique [100].
[69], [68]	STSD-fixed	continuous-time, WC	Shared and dedicated path protection. Introduce priority levels, high priority requires protection, low does not. Low can use resources allocated to high priority backup paths and may get preempted if a failure occurs. Propose ILP formulations using pre-computed paths. Propose sequential heuristic that uses link weights based on congestion and priority.
[104]	STSD-flexible	continuous-time, WC	Shared and dedicated path protection. Using pre-computed paths, propose ILPs that perform joint time scheduling and RWA as well as ILPs that first schedule in time with minimum overlap then do RWA.
[105]	STSD-fixed/flexible	continuous-time, WCC	Shared path protection with and without flexibility. Propose two step optimization where overlap is minimized with an ILP, then a heuristic is used to solve STSD-fixed problem. Heuristic similar to [68].
[106]	STSD-fixed/flexible	continuous-time, WCC and WC	Combination of previous work with addition of dedicated path protection.
[107]	STSD-fixed	time-slotted, WC	Shared path protection. Propose ILP with pre-computed paths. Use re-optimization, where backup paths can be changed at beginning of each timeslot. Show this reduces resource usage.
[108]	STSD-fixed	time-slotted, WC	Shared path protection with dual link failure. Assume first failure is permanent, reduce number of requests that would be unable to recover from second. Propose ILP using pre-computed paths.
[109]	STSD-flexible	continuous-time, WCC	Propose heuristics for shared path protection. Find time-independent requests then use two RWA heuristics to find routes, using single graph and layered wavelength graph.
[110]	STSD-fixed	continuous-time, WCC and WC	Multicast problem. Propose ILP for shared protection that uses pre-computed trees. Compare to traditional static traffic.
[111]	STSD-fixed	continuous-time, WC	See details in Table IX.

to clients to make advance reservations. OSCARS consists of a web-based interface, an Authorization, Authentication, and Auditing (AAA) module, a bandwidth scheduling module, and a path setup module. Typically, a client will submit a request through the web interface. The client is authenticated and the client's request is then scheduled with the bandwidth scheduler. When the actual data transfer is about to begin, a path is setup using existing signaling protocols such as RSVP-TE. The basic architecture of OSCARS is shown in Fig. 6. OSCARS is designed for a single domain, but interoperability with other domains is also possible.

A. On-demand Secure Circuits and Advance Reservation System (OSCARS)

B. EnLIGHTened

EnLIGHTened (initially funded by the U.S. NSF) is a project that focuses on advance reservation of both Grid and network resources (lightpaths) [31]. The goal of the project is to allow Grid applications to request in-advance or on-demand compute, storage, instrumentation, and network resources. These resources have to be co-allocated and may be spread across multiple domains. The architecture includes a resource broker (ERB), resource monitoring, and network scheduling

TABLE IX
SUMMARY OF AR RWA WITH STATIC TRAFFIC WITH GROOMING

Reference	AR type	Network Assumptions	Summary
[112]	STSD-flexible	continuous-time, WCC	Use similar heuristics as discussed in [79]. Incorporate grooming by modifying link weights of a layered wavelength graph. Compared to a tabu search based approach.
[111], [113]	STSD-fixed	continuous-time, WC	Find stable logical topology that is capable of handling all requests. Present ILP with pre-computed paths that finds logical topology and grooms requests onto it. Compare to traditional static demands. Consider dedicated and shared path protection, with and without the availability of wavelength converters.
[114]	STSD-fixed	continuous-time, WC	Propose grooming for electronic demands over lightpaths and lightpaths over wavebands. Propose ILP formulations for both problems.




Fig. 6. Overview of OSCARS Architecture.

(in a Network Domain Manager). The architecture utilizes the Highly-Available Resource Co-allocator (HARC) system, which allows clients to reserve multiple distributed resources in one step. HARC consists of Acceptors that manage co-allocation and Resource Managers that are the interfaces used to make reservations. Custom acceptors and resource managers were implemented specifically for the project. The resource broker provides the primary interface to the client and the ERB uses HARC internally. The architecture of the ERB is shown in Fig. 7(a). The ERB accepts requests from the user, who is authenticated through the AAA module. From the user request, the broker then gets availability information from each of the local managers (i.e., the storage managers and network managers). With this information, it uses its own scheduler to find a resource set able to satisfy the request. Given a set of resources, it then uses HARC acceptors to co-allocate the required resources. The acceptors in turn use the HARC Resource Managers to schedule the final resources with the local managers. The local managers are then responsible for activating and deactivating resources at the proper times.

The Domain Network Manager (DNM) is responsible for controlling network resources, resource reservation and path computation, restoration, and dynamic teardown/setup of lightpaths for a single administrative domain. Each domain will have an DNM. An example of a DNM is shown in Fig. 7(b).

The Discovery and Monitoring System (DMS) is required to support the dynamic addition or removal of resources. It is also responsible for monitoring the performance and reliability of resources in real-time. There are three types of resource monitoring. At-reservation-time monitoring is used to verify lightpath establishment. During-reservation monitoring is used to assess QoS/SLA requirements are being met. Ongoing grid resource status monitoring is used to collect relevant

performance metrics from grid resources in real time. All the monitoring types use the perfSONAR framework. Details of this module can be found in [117].

Lastly, EnLIGHTened also emphasizes the importance of standardized interfaces. Interaction with all local managers is done through well defined interfaces, such as Grid Security Infrastructure (GSI) and Grid Network Infrastructure (GNI). The project as promotes collaboration with other groups to define a GNI as well as an application interface, Grid Application Interface (GAI).

C. G-Lambda

Another project with similar goals was the Japanese project, G-Lambda [32]. The main goal of the project was to define a standard web interface between a grid resource scheduler (similar to the ERB) and network resource management systems (similar to DNM). Network operators could then implement this interface to allow the grid resource coordinators to make advance reservations of network resources. The grid resource coordinators are implemented as middleware and coordinate both network and computation/storage resources at the request of grid applications.

An overview of the G-lambda architecture is shown in Fig. 8(a). A detailed view of the architecture can be seen in Fig. 8(b) showing the modules that will be discussed. As introduced above, the main focus on the project was to define the Grid Network Service-Web Service Interface (GNS-WSI) shown in the figure. The GNS-WSI allows a number of commands to be sent to the network resource manager (NRM). These include commands to reserve paths, modify previously reserved paths, release paths, and query path or status information. An example architecture of an NRM is shown in Fig. 8(c).

The Grid Resource Scheduler (GRS) provides a web service interface to Grid clients using the Web Service Resource Framework (WSRF). The GRS is implemented using the Globus Toolkit 4. The modules include a client API, the web service module, and a scheduling module that uses the CRMs and NRMs to co-allocate and reserve resources. The scheduling module consists of a co-allocator and planner. The planner takes the user requests and selects a candidate set of resources, which are then reserved simultaneously by the co-allocator. A two-phase commit protocol is used to reserve

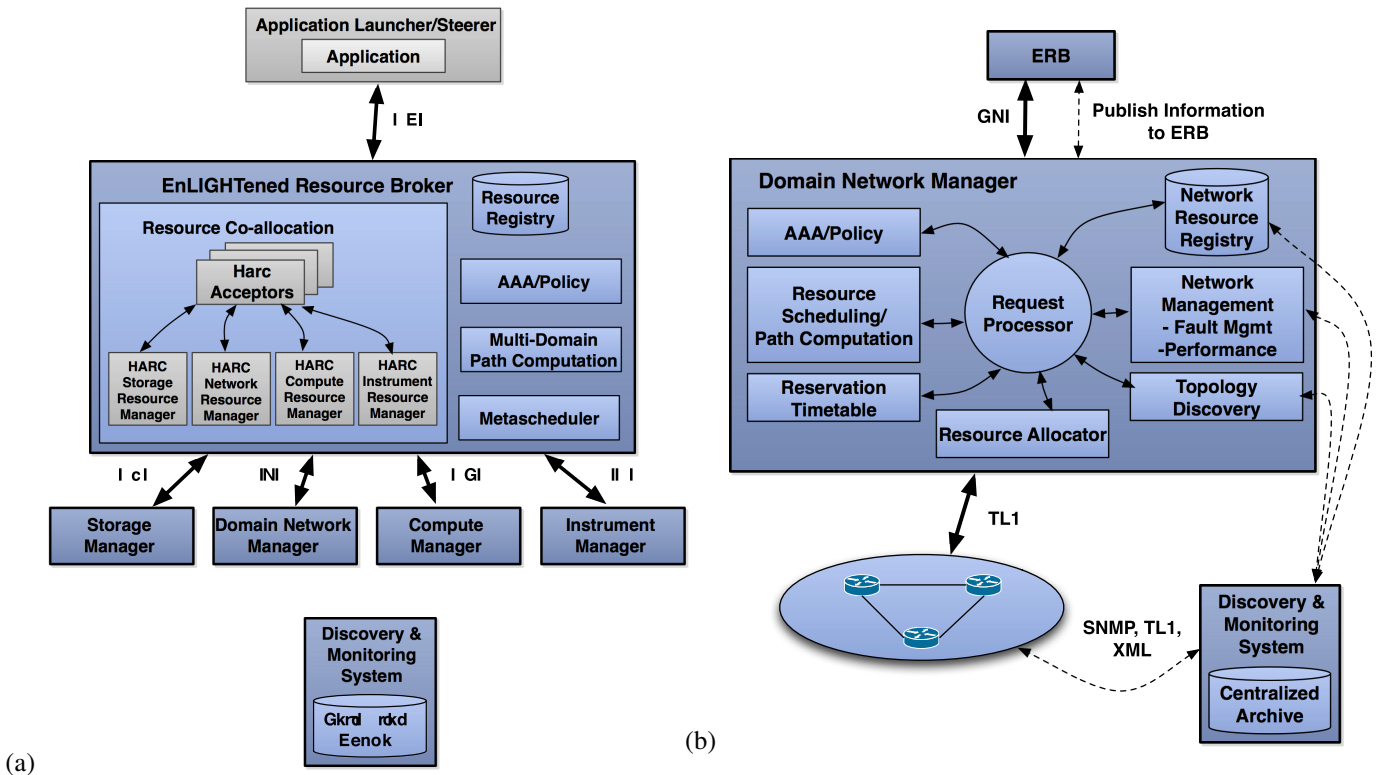


Fig. 7. Detailed view of the EnLIGHTened architecture. (a) EnLIGHTened resource broker (ERB) and (b) Domain network manager (DNM).

resources in parallel. A detailed description of the modules and protocols can be found in [118].

The CRMs are essentially wrappers around existing software like PluS and GridEngine. The project implemented two NRMs (one by KDDI and another by NTT). The NRMs are responsible for path virtualization between endpoints, local scheduling, and activation/de-activation of lightpaths. Paths are virtualized to hide implementation details. NRMs have a web service module (interface), a mediation module which virtualizes the GMPLS optical network and performs the scheduling based on requests from the GRS, a network control module for the GMPLS routers to manage state information. Details about NRMs can be found in [119].

EnLIGHTened and G-Lambda established a collaboration to prove network co-allocation across the two network domains. Software wrappers around components in the architectures were used so that the systems could interoperate. The technical details can be found in [120].

D. PHOSPHOROUS

The EU's PHOSPHOROUS project also incorporates the advance reservation of Grid and networking resources [33]. The goal of PHOSPHOROUS is to provide on-demand and in-advance end-to-end provisioning of network and grid resources across multiple domains and multiple vendors. The PHOSPHOROUS project comprises of two phases. In the first phase each independent domain is controlled by an existing Network Resource Provisioning System (NRPS), while in the second phase interoperability is added with other existing networks and Grid resources through standardized interfaces.

The NRPSs are similar to the NRM/DRM of the previous two projects.

PHOSPHOROUS defines an architecture composed of three planes. The first is the *service plane*, which consists of middleware extensions to allow applications to make advance reservations, co-allocate resources, provide AAA across multiple domains, etc. The next layer is the *network resource provisioning plane*. This is an adaptation layer between the service plane and existing NRPSs. Lastly, there is the *control plane*, which is used to control physical resources.

The general architecture is shown in Fig. 9. The project started with a centralized architecture with only one service plane (global broker). The service plane (network service plane, NSP) is responsible for finding end-to-end paths, managing AAA, etc. The individual NRPSs, shown underneath the NSP manage their own domains and intra-domain routing. Each NRSP has its own control plane, such as GMPLS. All interfaces are deployed as web services.

At present, PHOSPHOROUS architecture works with three existing NRPSs: ARGON [121] from Germany, DRAC from Nortel, and User Controlled Lightpaths (UCLP) from CANARIE in Canada. The NRSPs publish border end-points to the NSP, which can then do inter-domain routing while NRSPs perform intra-domain routing based on the selected end-points. Each NRPS provides the ability to make a reservation request, cancel a reservation, get the status of a request, and bind and activate requests. This functionality is provided by an NRPS adapter which abstracts each specific NRPS. A number of interfaces (NBI, EWI, SBI), shown in the figure, for communication between all the components are defined.

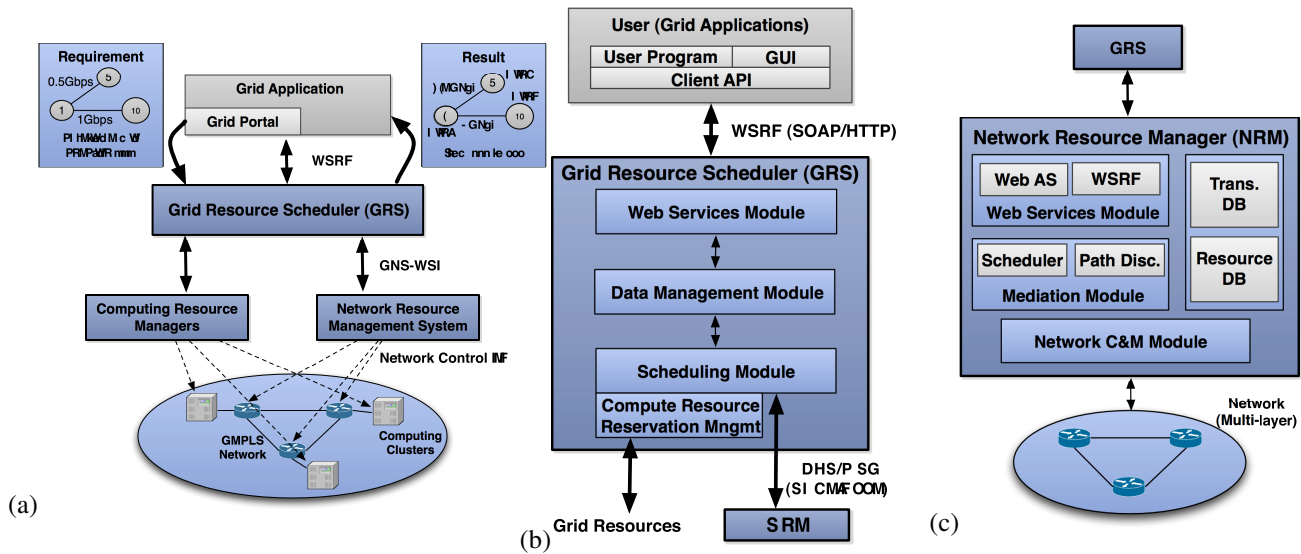


Fig. 8. (a) Overview of G-lambda architecture. Detailed view of the G-lambda architecture (b) Grid Resource Scheduler and (c) Network Resource Manager.

The general flow when a grid request is received is as follows. The grid application or grid middleware sends a request to the reservation web service of the NSP. Once validated by the AAA module and available resources are found, the NSP finds an inter-domain path. With the inter-domain path selected, each independent domain is notified and the lookup of an intra-domain path is initiated.

Harmony [122] is the network service provisioning system in PHOSPHORUS. Harmony assumes a group of independent NRPSs and provides an abstract service plane to do end-to-end co-allocation. It defines three types of architectures: a centralized approach, hierarchical, and distributed. In each case, a Harmony inter-domain broker sits above each of the NRPSs. In the case of a centralized system there is only one, whereas in the case of a distributed system, there is one for each NRPS. A Harmony Service Interface (HSI) is used to exchange abstract topology and other information between the brokers and the NRPSs.

As discussed previously, Harmony is responsible for inter-domain paths while each NRPS is responsible for intra-domain paths. The NRPSs provide border endpoints and abstract links to Harmony. A reservation specifies one or more services, where each service requires one or more connections. The requests can be fixed or malleable advance reservation requests. A simple iterative search over the inter-domain paths is used for malleable requests, trying different start times and different bandwidths. For each path, Harmony checks whether each individual domain can create an intra-domain path or not.

Grid Enabled GMPLS (G^2 MPLS) [123] is the new control plane defined in PHOSPHORUS which extends GMPLS to work with, both Grid and network resources. The goal is to incorporate *grid network services* (GNS) into GMPLS. GNS is a service that allows provisioning of network and grid resources in a single step. G^2 MPLS allows selection, co-allocation, and maintenance of grid and network resources. Some of the extensions include:

- Discovery and advertisement of Grid capabilities (e.g.

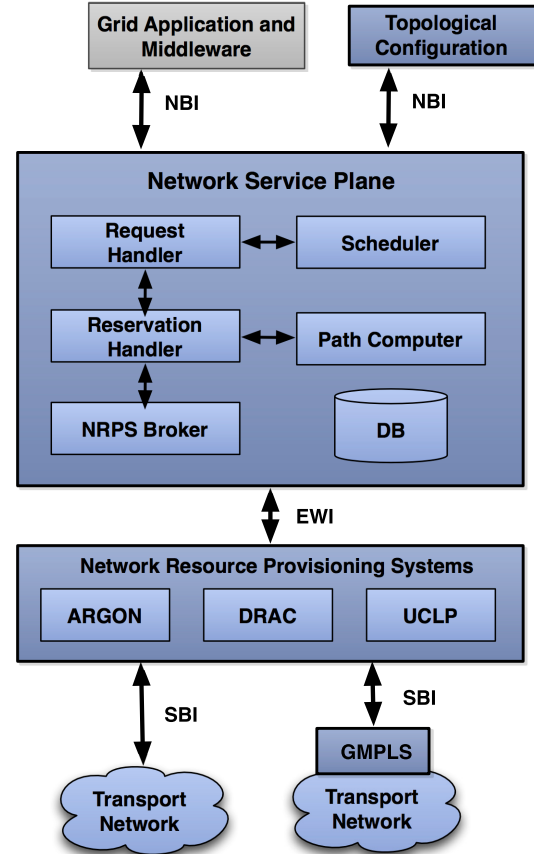


Fig. 9. Overview of PHOSPHORUS architecture.

number of CPUs) and network resources (e.g. bandwidth).

- Service setup, which involves coordination with Grid middleware, configuration of network connections, advance reservation, etc.
- Service monitoring.

G²MPLS provides two models of operation. One is an overlay model designed for legacy systems that only supports standard GMPLS. In this case, NRPSs still exist and can use G²MPLS in the same way they would use GMPLS. The other model is the integrated model, where G²MPLS provides an interface directly to the NSP. It can then be used to co-allocate grid and network resources for a given domain. The basic unit of work is a grid job.

VII. OTHER RELATED WORK ON ADVANCE RESERVATION

The focus of our survey is on routing and wavelength assignment algorithms for advance reservation over optical networks. In this section we discuss other work related to advance reservation that did not fall in the optical domain or did not perform routing and wavelength assignment. This may include advance reservation for IP/MPLS networks or grid scheduling advance reservation algorithms, for example. The references for these topics do not represent a complete survey. Interested readers should follow references within these papers for more details about the topics.

Burchard et al. have done work with advance reservation outside wavelength-routed WDM networks. We discussed earlier a distinction between STSD-fixed and STSD-flexible types of advance reservation. It has been shown that STSD-fixed advance reservation leads to resource fragmentation, which can lead to higher blocking. This was first mentioned in [124], where the concept of malleable reservations was also introduced. Other work by the authors related to advance reservation include failure recovery [125], [126], performance evaluation and algorithms [127], [128], and rerouting [129].

A number of path selection and computation algorithms have been proposed specifically for advance reservation over single channel networks. In [19], the authors propose algorithms for different AR demands with different goals such as earliest completion and maximum bandwidth. The authors in [130] propose a Bellman-Ford based algorithm to find the shortest hop path for STSD-flexible requests. Similarly, in [131] the authors propose another Bellman-Ford based algorithm that finds all time-slots during which a path is available with a specified bandwidth for a specified duration. In [132] several algorithms are presented with different goals from finding specific bandwidth between the start and end slots, to looking for the earliest available time with a specified bandwidth/duration by extending breadth-first search. The algorithm find the highest available bandwidth in a specified timeslot by extending Dijkstra's algorithm. Lastly, they find all available timeslots with a specified bandwidth/duration by extending Bellman-Ford's algorithm. Variable bandwidth path computation algorithms are discussed in [133]. They allow the bandwidth and/or path to change during each timeslot. Several algorithms are presented in [134] with the goal of either finding a path at a specified start time with a specified duration/bandwidth or finding the earliest such time. A number of time-based link metrics are proposed in [135].

The previous algorithms are based on global topology information. The authors in [37] propose a distributed advance reservation routing algorithm, as discussed previously. The algorithm is for STSD-flexible requests that specify a duration

as well as required bandwidth. The group also proposed a way to rank timeslot and path combinations when routing an advance reservation request [136] (assuming centralized routing). They also introduce the concept of path switching where a request may use a number of different paths over its duration.

The authors of [137] propose a re-routing algorithm based on load balancing. Re-routing, or re-optimization, has been discussed previously for RWA in [48].

Patel et al. have proposed modifications to the basic advance reservation concept. Their work does not consider wavelength assignment, only routing. In [138] they consider time-shift advance reservation. Here, delay elements are placed in the network that can buffer circuits between two links. This can be used to shift the circuit in the time domain. Instead of all links of a path having to use the same timeslots to transfer a circuit, different links can use different slots by shifting the circuit in time. They assume a time-slotted network and maintain state information about links and delay elements. They also use horizon scheduling, which only maintains state information about earliest available times. Requests specify their holding times and the network finds a start time based on the horizon schedule as well as buffer assignment that minimizes end-to-end latency. In [139] the same authors also propose variable bandwidth advance reservation. Again the network is time-slotted and no wavelength assignment is performed. The main contribution is to allow the bandwidth of the circuit change as a function of time, i.e. each timeslot can transmit data at different bandwidth. The requests specify a file size and the network finds a start timeslot and bandwidth schedule that minimizes file transfer completion time.

The authors of [140] explore the problem of logical topology design given a series of traffic matrices that change over time. The static AR RWA work we have discussed assumes that the set of advance reservation lightpath demands are given. This work will create a set of advance reservation lightpath demands from the time-dependent traffic matrices. The algorithm finds the virtual topology and flow routing over it. Routing and wavelength assignment is not performed (any of the previous work for static STSD-fixed demands could be used). Two MILP formulations are proposed that find a the virtual topology and flow routing on top of it. In one case, it is assumed that the virtual topology cannot change (this creates a static demand set) and in the other case it is assumed that the virtual topology can change over time (this creates a STSD-fixed demand set). A tabu search meta-heuristic is derived in [141] and additional variations to flow routing are proposed in [142], but now considering a static virtual topology.

The concept of delay tolerance was used in the following two papers [143], [144]. A batch mode scheduling technique was proposed in [143] where a customer specifies a waiting time (delay tolerance). The scheduler queues up requests and then schedules them in batch in order to find a better solution compared to scheduling them one at a time. The authors of [144] propose a notification interval. Batch scheduling is not performed in this work. The interval is used to queue requests that could not be provisioned upon arrival. The authors propose optimizing future reserved requests so that it may be possible to free up resources for queued requests.

TABLE X
COMPARISON OF ADVANCE RESERVATION FRAMEWORKS.

Framework	Region and Funding Body	Bandwidth Provision	Provisioning Layer	Network Resource Provisioning System	Grid Scheduling Capabilities
OSCARS	U.S.A. (U.S. Department of Energy)	Hybrid: packet/circuit	Layer 2 & 3	Integrated (MPLS-based)	No
EnLIGHTened	U.S.A. (U.S. NSF)	Circuit	Layer 1	Integrated (GMPLS)	Yes
G-Lambda	Japan (KDDI R&D, NTT, NICT and AIST)	Circuit	Layer 1 & 2	Integrated (GMPLS)	Yes
PHOSPHORUS	Europe (E.C. FP6)	Circuit	Layer 1 & 2	ARGON, DRAC and UCLP	Yes

Zhe et al. derived analytical models for blocking of advance reservation requests [145], [40]. They look at a simplified scenario with a single link consisting of a number of discrete channels. They show how flexibility can impact blocking as well as the relationship between the horizon size and blocking.

Lastly, we would like to discuss some work with scheduling bandwidth of advance reservation demands. These works consider a scenario where lightpaths have already been established and user demands can be divided into discrete size bandwidth blocks (e.g. size of a timeslot). The problem then becomes a scheduling problem of how to schedule the blocks in the wavelengths and timeslots of the already established lightpaths. For example, the authors in [146] consider sliding window demands mixed with immediate reservation traffic. They also propose categorizing advance reservation demands as preemptable and non-preemptable. To reduce IR blocking and to minimize fragmentation, some AR requests can be split up (preempted) and continued later so they are non-continuous in the time domain (similar to the idea proposed in [84]).

A similar scheduling problem is investigated in [147] where they divide each timeslot into smaller *bandwidth slots*. During each timeslot, each request must use some number of bandwidth slots. The paper proposes two types of request: streaming and elastic. A streaming request (e.g. real-time traffic) must use the same number of bandwidth slots for the duration of the request whereas an elastic request (e.g. file transfer) can use a variable number of bandwidth slots over time. In order to reduce blocking, the scheduler can take advantage of the elastic request's ability to use more or fewer bandwidth slots in any given timeslot.

In [148], the authors consider the scheduling problem with UTSD requests that specify a deadline. Similar to the work just discussed, lightpaths are already established and the scheduler must schedule timeslot sized bandwidth chunks to each request so that it completes by its deadline. The authors propose allowing requests with later deadlines to be pushed back to help accommodate requests with earlier deadlines. Initially, a request is scheduled using as much bandwidth as early as possible. If a request cannot be accommodated, the algorithm determines if another request with a later deadline can be pushed forward in the schedule so the current request does not have to be blocked.

Advance reservation is also a popular topic for job scheduling in Grid networks. Each job can have start/end times or deadlines similar to circuit requests we have seen for wavelength-routed networks. The job scheduler must determine how to assign these jobs to servers. An example of this work is [149]. The paper provides a number of references and small survey of work related to this area.

We have discussed circuit-switched networks so far. AR has also been proposed for OBS networks in [150]. They propose routing algorithms for UTSD and UTUD AR requests with the goal to minimize delay, or the difference between start time and request submission. The main focus of the paper is a multi-cost routing algorithm that also considers the temporal domain. Static STSD-fixed AR has also been proposed for the Light-trail architecture [151], [152]. They explore two solution techniques. One where light-trails are setup once and never changed and one where light-trails can setup and torn down over time. They propose ILP formulations and heuristics.

VIII. OPEN ISSUES

In this section we discuss several open areas for advance reservation in optical networks. Our intention is briefly point out promising areas of future work in this research area.

A. Admission Control and Quality of Service

In real-world networks, there will likely be a mix of immediate and advance reservation demands in the network. There are few works that discuss admission control and quality of service for optical networks in detail. While there are some works that start addressing this area, there is no work that makes any type of service guarantees for optical networks. We note that there has also been recent work for non-optical networks as well, e.g. [153].

Instead of assuming that AR is always higher priority compared to IR, it may also be interesting to investigate the case where IR demands are *urgent* and require higher priority. There are many interesting open problems in the area of admission control and QoS for advance reservation over optical networks.

B. Multi-domain Advance Reservation

Multi-domain, or inter-domain, setup of dynamic circuits is an important problem for wide area and grid-based networks. The authors of [154] provide a survey for this area. The work that has been presented needs to be extended for advance reservation demands to incorporate the time-domain information.

C. Analytical Modeling

We have reviewed a few papers that provide initial analytical results for simplified networks, e.g. a single link. There is no work that proposes a general model which incorporates time-domain information similar to [155].

D. Survivable Dynamic Advance Reservation

As we have discussed, there is significant work for survivable static advance reservation (see Table VIII). However, there are only a few works that discuss survivable routing with dynamic advance reservation. Both restoration and protection techniques need to be explored in more detail, as well as other communication paradigms such as anycast and multicast.

E. Grid and Network Layer Integration

The work we discuss in this paper deals with advance reservation of network resources. There is also significant work on advance reservation of Grid resources [156]. The co-allocation of network and grid resources is still an open problem. Some testbeds provide some form of this feature, but there is little theoretical work considering co-allocation of both types of resources. The work in [157], addresses the new evolving paradigms for application- driven networking within the optical layer in the context of Grid computing. The authors discuss open research issues of the optical network control plane and present the issues of interaction between the optical network control plane and applications.

IX. CONCLUSION

In this paper we have presented a comprehensive literature survey on advance reservation in optical networks. We provide a classification of the types of advance reservation that have been proposed in the literature. We discuss motivation for advance reservation in both WAN and Grid-based networks. We discuss architectural issues such as centralized and distributed scheduling and management of the time-domain. We then provide a survey of advance reservation for both static and dynamic traffic, we discuss testbeds and other networks supporting advance reservation, and we discuss other related work. Lastly, we provide some areas with open problems dealing with advance reservation.

REFERENCES

- [1] Gerd E. Keiser, "A review of WDM technology and applications," *Optical Fiber Technology*, vol. 5, no. 1, pp. 3 – 39, 1999.
- [2] Rajiv Ramaswami and Kumar N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Netw.*, vol. 3, no. 5, pp. 489–500, 1995.
- [3] I. Chlamtac, A. Ganz, and G. Karmi, "Lightpath communications: an approach to high bandwidth optical WAN's," *IEEE/ACM Trans. Netw.*, vol. 40, no. 7, pp. 1171–1182, Jul 1992.
- [4] H. Zang, J. P. Jue, and B. Mukherjee, "A review of routing and wavelength assignment approaches for wavelength-routed optical WDM networks," *SPIE Optical Networks Magazine*, vol. 1, 2000.
- [5] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 122–133, 2002.
- [6] Y. Lee and B. Mukherjee, "Traffic engineering in next-generation optical networks," *IEEE Commun. Surveys Tutorials*, vol. 6, no. 3, pp. 16 –33, 2004.
- [7] S. Azodolmolky, M. Klinkowski, E. Marin, D. Careglio, J. S. Pareta, and I. Tomkos, "A survey on physical layer impairments aware routing and wavelength assignment algorithms in optical networks," *Elsevier Computer Networks*, vol. 53, no. 7, pp. 926 – 944, 2009.
- [8] C.V. Saradhi and S. Subramaniam, "Physical layer impairment aware routing (PLIAR) in WDM optical networks: issues and challenges," *IEEE Commun. Surveys Tutorials*, vol. 11, no. 4, pp. 109 –130, 2009.
- [9] D. Banerjee and B. Mukherjee, "A practical approach for routing and wavelength assignment in large wavelength-routed optical networks," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 903 –908, jun 1996.
- [10] A.E. Ozdaglar and D.P. Bertsekas, "Routing and wavelength assignment in optical networks," *IEEE/ACM Trans. Netw.*, vol. 11, no. 2, pp. 259–272, 2003.
- [11] J. Zheng and H. T. Mouftah, "Routing and wavelength assignment for advance reservation in wavelength-routed WDM optical networks," in *Proc. IEEE International Conference on Communications (ICC)*, 2002, vol. 5, pp. 2722–2726.
- [12] James Roberts and Keqiang Liao, "Traffic models for telecommunication services with advance capacity reservation," *Computer Networks and ISDN Systems*, vol. 10, no. 3-4, pp. 221–229, 1985.
- [13] J.T. Virtamo, "A model of reservation systems," *IEEE Trans. Commun.*, vol. 40, no. 1, pp. 109 –118, jan 1992.
- [14] L. Wolf, L. Delgrossi, R. Steinmetz, and S. Schaller, "Issues of reserving resources in advance," in *International Workshop, NOSSDAV*, 1995.
- [15] L. Wolf and R. Steinmetz, "Concepts for resource reservation in advance," *Multimedia Tools and Applications*, Jan. 1997.
- [16] D. Wischik and A. Greenberg, "Admission control for booking ahead shared resources," in *Proc. IEEE INFOCOM*, Mar. 1998, vol. 2, pp. 873–882.
- [17] A. G. Greenberg, R. Srikant, and W. Whitt, "Resource sharing for book-ahead and instantaneous-request calls," *IEEE/ACM Trans. Netw.*, vol. 7, no. 1, pp. 10–22, 1999.
- [18] A. Schill, Sabine Kühn, and Frank Breiter, "Design and evaluation of an advance reservation protocol on top of rsvp," in *BC '98: Proc. IFIP TC6/WG6.2 Fourth International Conference on Broadband Communications*, 1998, pp. 23–40.
- [19] R. Guerin and A. Orda, "Networks with advance reservations: the routing perspective," in *Proc. IEEE INFOCOM*, Mar. 2000, vol. 1, pp. 118–127.
- [20] Jun Zheng and H. T. Mouftah, "Supporting advance reservations in wavelength-routed WDM networks," in *Proc. IEEE International Conference on Computer Communications and Networks (ICCCN)*, 2001, pp. 594 –597.
- [21] "Cisco telepresence," .
- [22] "Hauwei telepresence," .
- [23] J. Kuri, N. Puech, M. Gagnaire, E. Dotaro, and R. Douville, "Routing and wavelength assignment of scheduled lightpath demands," *IEEE J. Sel. Areas Commun.*, vol. 21, no. 8, pp. 1231–1240, Oct. 2003.
- [24] A. Gençata and B. Mukherjee, "Virtual-topology adaptation for WDM mesh networks under dynamic traffic," *IEEE/ACM Trans. Netw.*, vol. 11, no. 2, pp. 236–247, 2003.
- [25] "New Cisco study reveals peak Internet traffic increases due to social networking and broadband video usage," 2009.
- [26] "Worldwide lhc computing grid," .
- [27] "BIRN - biomedical informatics research network," .
- [28] "NEES George E. Brown Jr. network for earthquake engineering simulation," .
- [29] T. DeFanti, M. Brown, J. Leigh, O. Yu, E. He, J. Mambretti, D. Lil-lethun, and J. Weinberger, "Optical switching middleware for the OptiPuter," *IEICE Trans. Communications [Invited]*, 2003.
- [30] C.P. Guok, D.W. Robertson, E. Chaniotakis, M.R. Thompson, W. Johnston, and B. Tierney, "A user driven dynamic circuit network imple-mentation," in *IEEE GLOBECOM Workshop*, Dec. 2008, pp. 1 –5.

- [31] L. Battestilli, A. Hutanu, G. Karmous-Edwards, D.S. Katz, J. MacLaren, J. Mambretti, J.H. Moore, Seung-Jong Park, H.G. Perros, S. Sundar, S. Tanwir, S.R. Thorpe, and Yufeng Xin, "EnLIGHTened computing: An architecture for co-allocating network, compute, and other Grid resources for high-end applications," in *International Symposium on High Capacity Optical Networks and Enabling Technologies*, Nov. 2007, pp. 1–8.
- [32] Atsuko Takefusa, Michiaki Hayashi, Naohide Nagatsu, Hidemoto Nakada, Tomohiro Kudoh, Takahiro Miyamoto, Tomohiro Otani, Hideaki Tanaka, Masatoshi Suzuki, Yasunori Sameshima, Wataru Imajuku, Masahiko Jinno, Yoshihiro Takigawa, Shuichi Okamoto, Yoshio Tanaka, and Satoshi Sekiguchi, "G-lambda: Coordination of a Grid scheduler and lambda path service over GMPLS," *Future Generation Computer Systems*, vol. 22, no. 8, pp. 868–875, 2006.
- [33] S. Figuerola, N. Ciulli, M. de Leenheer, Y. Demchenko, W. Ziegler, and A. Binczewski, "PHOSPHORUS: single-step on-demand services across multi-domain networks for e-Science," in *SPIE*, 2007.
- [34] U.N. Interface, "1.0 Signaling Specification," in *Optical Networking Forum (OIF)*, 2001.
- [35] E. Escalona, S. Spadaro, J. Comellas, and G. Junyent, "Advance reservations for service-aware GMPLS-based optical networks," *Elsevier Computer Networks*, vol. 52, no. 10, pp. 1938–1950, 2008.
- [36] J. Zheng, B. Zhang, and H. T. Mouftah, "Toward automated provisioning of advance reservation service in next-generation optical Internet," *IEEE Commun. Mag.*, vol. 44, no. 12, pp. 68–74, Dec. 2006.
- [37] N. Fazlollahi and D. Starobinski, "Distributed advance network reservation with delay guarantees," in *IEEE International Symposium on Parallel Distributed Processing (IPDPS)*, 19–23 2010, pp. 1–12.
- [38] C. Xie, H. Alazemi, and N. Ghani, "Routing and scheduling in distributed advance reservation networks," in *Proc. IEEE GLOBECOM*, 2010.
- [39] C. Barz, U. Bornhauser, P. Martini, and M. Pilz, "Timeslot-based resource management in grid environments," in *Proc. IASTED International Conference on Parallel and Distributed Computing and Networks*, 2008, pp. 39–48.
- [40] Xiangfei Zhu and M. Veeraraghavan, "Analysis and design of book-ahead bandwidth-sharing mechanisms," *IEEE Trans. Commun.*, vol. 56, no. 12, pp. 2156–2165, Dec. 2008.
- [41] J. Kuri, N. Puech, M. Gagnaire, and E. Dotaro, "Routing foreseeable lightpath demands using a tabu search meta-heuristic," in *Proc. IEEE GLOBECOM*, Nov. 2002, vol. 3, pp. 2803–2807.
- [42] S. Figueira, N. Kaushik, and S. Naiksatam, "Advance reservation of lightpaths in optical-network based grids," in *IEEE Gridnets*, 2004.
- [43] S. Naiksatam, S. Figueira, S. A. Chiappari, and N. Bhatnagar, "Analyzing the advance reservation of lightpaths in lambda-grids," in *IEEE International Symposium on Cluster Computing and the Grid*, May 2005, vol. 2, pp. 985–992.
- [44] T. Daniel Wallace and Abdallah Shami, "Connection management algorithm for advance lightpath reservation in WDM networks," in *Proc. IEEE BroadNets*, Sept. 2007, pp. 837–844.
- [45] T. D. Wallace, A. Shami, and C. Assi, "Advance lightpath reservation for WDM networks with dynamic traffic," *IEEE/OSA J. Optical Commun. Netw.*, vol. 6, pp. 913–924, 2007.
- [46] S. Tanwir, L. Battestilli, H. Perros, and G. Karmous-Edwards, "Dynamic scheduling of network resources with advance reservations in optical grids," *Int. J. Netw. Manag.*, vol. 18, no. 2, pp. 79–105, 2008.
- [47] Eun Sung Jung, Yan Li, S. Ranka, and S. Sahni, "Performance evaluation of routing and wavelength assignment algorithms for optical networks," in *Proc. IEEE International Conference on Communications (ICC)*, Jul. 2008, pp. 62–67.
- [48] X. Yang, L. Shen, A. Todimala, B. Ramamurthy, and T. Lehman, "An efficient scheduling scheme for on-demand lightpath reservations in reconfigurable WDM optical networks," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2006, p. 3.
- [49] Lu Shen, A. Todimala, B. Ramamurthy, and Xi Yang, "Dynamic lightpath scheduling in next-generation WDM optical networks," in *Proc. IEEE INFOCOM*, April 2006, pp. 1–5.
- [50] Lu Shen, Xi Yang, A. Todimala, and B. Ramamurthy, "A two-phase approach for dynamic lightpath scheduling in WDM optical networks," in *Proc. IEEE International Conference on Communications (ICC)*, June 2007, pp. 2412–2417.
- [51] P. Angu and B. Ramamurthy, "Continuous and parallel optimization of dynamic bandwidth scheduling in WDM networks," in *Proc. IEEE GLOBECOM*, 2010.
- [52] D. Andrei, M. Tornatore, D. Ghosal, C.U. Martel, and B. Mukherjee, "On-demand provisioning of data-aggregation sessions over WDM optical networks," *IEEE/OSA J. Lightwave Technology*, vol. 27, no. 12, pp. 1846–1855, June 2009.
- [53] C. Cavdar, M. Tornatore, and F. Buzluca, "Availability-guaranteed connection provisioning with delay tolerance in optical WDM mesh networks," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2009, pp. 1–3.
- [54] C. Cavdar, F. Buzluca, M. Tornatore, and B. Mukherjee, "Dynamic scheduling of survivable connections with delay tolerance in WDM networks," in *IEEE INFOCOM Workshop*, Apr. 2009, pp. 1–6.
- [55] C. Cavdar, M. Tornatore, F. Buzluca, and B. Mukherjee, "Shared-path protection with delay tolerance (SDT) in optical WDM mesh networks," *IEEE/OSA J. Lightwave Technology*, vol. 28, no. 14, pp. 2068–2076, 2010.
- [56] H. Jingyi, S. G. Chan, and D.H.K. Tsang, "Multicasting in WDM networks," *IEEE Commun. Surveys Tutorials*, vol. 4, no. 1, pp. 2–20, 2002.
- [57] T. Stevens, M. De Leenheer, C. Develder, B. Dhoedt, K. Christodoulou, P. Kokkinos, and E. Varvarigos, "Multi-cost job routing and scheduling in Grid networks," *Future Generation Computer Systems*, vol. 25, no. 8, pp. 912–925, 2009.
- [58] Hong-Ha Nguyen, M. Gurusamy, and Luying Zhou, "Scheduling network and computing resources for sliding demands in optical grids," *IEEE/OSA J. Lightwave Technology*, vol. 27, no. 12, pp. 1827–1836, June 15, 2009.
- [59] F. Arshad, S.R. Ramay, S. Tanwir, L. Battestilli, and S.M.H. Zaidi, "Advance reservation and dynamic scheduling of point to multipoint lightpaths," in *International Symposium on High Capacity Optical Networks and Enabling Technologies*, Nov. 2008, pp. 69–74.
- [60] A. Munir, S. Tanwir, L. Battestilli, and S.M.H. Zaidi, "Holding time aware multicast requests provisioning algorithm for dynamic optical circuit switched (docs) networks," in *Proc. 6th International Conference on Frontiers of Information Technology*, 2009.
- [61] D. Andrei, M. Tornatore, C.U. Martel, D. Ghosal, and B. Mukherjee, "Provisioning Subwavelength Multicast Sessions With Flexible Scheduling Over WDM Networks," *Journal of Optical Communications and Networking*, vol. 2, no. 5, pp. 241–255, 2010.
- [62] A. Munir, S. Tanwir, L. Battestilli, and S.M.H. Zaidi, "Holding time aware dynamic bandwidth allocation algorithm for emerging bandwidth on demand multicast applications," in *6th International Symposium on High Capacity Optical Networks and Enabling Technologies, HONET*, 2009, pp. 16–21.
- [63] D. Andrei, M. Tornatore, C.U. Martel, and B. Mukherjee, "Flexible scheduling of multicast sessions with different granularities for large data distribution over WDM networks," in *Proc. IEEE GLOBECOM*, Dec. 2009, pp. 1–6.
- [64] E. He, X. Wang, V. Vishwanath, and J. Leigh, "AR-PIN/PDC: Flexible advance reservation of intradomain and interdomain lightpaths," in *Proc. IEEE GLOBECOM*, Dec. 2006, pp. 1–6.
- [65] E. He, Xi Wang, and J. Leigh, "A flexible advance reservation model for multi-domain WDM optical networks," in *Proc. IEEE BroadNets*, Oct. 2006, pp. 1–10.
- [66] Chongyang Xie and N. Ghani, "A prioritized routing algorithm for multi-domain optical networks supporting advance reservation," in *International Symposium on High Capacity Optical Networks and Enabling Technologies*, Nov. 2008, pp. 211–215.
- [67] Chongyang Xie and N. Ghani, "Admission control in networks with prioritized advance reservation," in *Proc. IEEE International Conference on Computer Communications and Networks (ICCCN)*, Aug. 2009, pp. 1–6.
- [68] A. Jaekel and Y. Chen, "Quality of service based resource allocation for scheduled lightpath demands," *Computer Communications*, vol. 30, no. 18, pp. 3550–3558, 2007.
- [69] A. Jaekel and Y. Chen, "Routing and wavelength assignment for prioritized demands under a scheduled traffic model," in *Proc. IEEE BroadNets*, Oct. 2006, pp. 1–7.
- [70] H. Ahn, T. Lee, M. Chung, and H. Choo, "RWA on scheduled lightpath demands in WDM optical transport networks with time disjoint paths," *Lecture notes in computer science*, 2005.
- [71] C. V. Saradhi, J. C. Wei, M. Shujing, and M. Gurusamy, "Circular arc graph based algorithms for routing scheduled lightpath demands in WDM optical networks," in *Proc. IEEE BroadNets*, Oct. 2005, vol. 1, pp. 320–322.
- [72] C.V. Saradhi and M. Gurusamy, "Graph theoretic approaches for routing and wavelength assignment of scheduled lightpath demands in WDM optical networks," in *Proc. IEEE BroadNets*, Oct. 2005, pp. 1291–1299, Vol. 2.
- [73] N. Skorin-Kapov, "Heuristic algorithms for the routing and wavelength assignment of scheduled lightpath demands in optical networks," *IEEE J. Sel. Areas Commun.*, vol. 24, no. 8, pp. 2–15, Aug. 2006.

- [74] S. Park, J. Yang, M. Kim, and Y.C. Bang, "RWA Algorithm for Scheduled Lightpath Demands in WDM Networks," *Parallel and Distributed Processing and Applications*, pp. 113–124, 2007.
- [75] S. Park, J.S. Yang, and Y.C. Bang, "On RWA Algorithms for Scheduled Lightpath Demands in Wavelength Division Multiplexing Networks," *IJCSNS*, vol. 7, no. 3, pp. 144, 2007.
- [76] T. D. Wallace, A. Shami, and C. Assi, "Scheduling advance reservation requests for wavelength division multiplexed networks with static traffic demands," *IET Communications*, vol. 2, no. 8, pp. 1023–1033, Sept. 2008.
- [77] A. Chen and S. S. W. Lee, "A near-optimal heuristic algorithm for advance lightpath reservation in WDM networks," in *European Conference on Optical Communication*, sept. 2008, pp. 1–2.
- [78] S. S. W. Lee, A. Chen, and M.C. Yuang, "A lagrangean relaxation based near-optimal algorithm for advance lightpath reservation in WDM networks," *Photonic Network Communications*, vol. 19, pp. 103–109, 2010.
- [79] B. Wang, T. Li, X. Luo, Y. Fan, and C. Xin, "On service provisioning under a scheduled traffic model in reconfigurable WDM optical networks," in *Proc. IEEE BroadNets*, Oct. 2005, vol. 1, pp. 13–22.
- [80] Wei Su, Galen Sasaki, Ching-Fong Su, and Ashok Balasubramanian, "Scheduling of periodic connections with flexibility," *Optical Switching and Networking*, vol. 3, no. 3-4, pp. 158–172, 2006.
- [81] C. V. Saradhi and M. Gurusamy, "Scheduling and routing of sliding scheduled lightpath demands in WDM optical networks," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2007, pp. 1–3.
- [82] D. Andrei, Hong-Hsu Yen, and M. Tornatore, "Integrated design for sliding scheduled traffic in WDM networks," in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2009, pp. 1–3.
- [83] D. Andrei, Hong-Hsu Yen, M. Tornatore, C.U. Martel, and B. Mukherjee, "Integrated provisioning of sliding scheduled services over WDM optical networks [Invited]," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 1, no. 2, pp. A94–A105, Jul. 2009.
- [84] Ying Chen, A. Jaekel, and A. Bari, "Resource allocation strategies for a non-continuous sliding window traffic model in WDM networks," in *Proc. ICST BroadNets*, Sept. 2009, pp. 1–7.
- [85] M. Koubaa, N. Puech, and M. Gagnaire, "Strategies for the routing and wavelength assignment of scheduled and random lightpath demands," in *Third European Conference, Universal Multiservice Networks*, 2004, pp. 91–103.
- [86] M. Koubaa, N. Puech, and M. Gagnaire, "Routing and wavelength assignment of scheduled and random lightpath demands," in *Proc. 1st IFIP Int. Conf. Wireless and Optical Communications Networks*, 2004, pp. 16–19.
- [87] M. Koubaa, N. Puech, and M. Gagnaire, "Routing and Wavelength Assignment for Scheduled and Random Lightpath Demands: Bifurcated Routing versus Non-Bifurcated Routing," in *IFIP First Optical Networks and Technologies Conference*, 2004, pp. 137–144.
- [88] M. Koubaa, N. Puech, and M. Gagnaire, "Lightpath rerouting for differentiated services in WDM all-optical networks," in *Design of Reliable Communication Networks, 2005.(DRCN 2005). Proceedings. 5th International Workshop on*, 2005, p. 8.
- [89] EA Doumith, M. Koubaa, N. Puech, and M. Gagnaire, "Gain and Cost brought in by Wavelength Conversion for the Routing and Wavelength Assignment of two Traffic Classes in WDM Networks," in *10th European Conference on Networks & Optical Communications*, 2005, pp. 147–154.
- [90] M. Koubaa, N. Puech, and M. Gagnaire, "Routing and spare capacity assignment for scheduled and random lightpath demands in all-optical networks," *Next Generation Internet Networks*, 2005, pp. 39–46, 2005.
- [91] J. Kuri, N. Puech, and M. Gagnaire, "Diverse routing of scheduled lightpath demands in an optical transport network," in *Proc. Design of Reliable Communication Networks*, Oct. 2003, pp. 69–76.
- [92] M. Gagnaire, M. Koubaa, and N. Puech, "Network dimensioning under scheduled and random lightpath demands in all-optical WDM networks," *IEEE J. Sel. Areas Commun.*, vol. 25, no. 9, pp. 58–67, Dec. 2007.
- [93] B. Wang, Y. Fan, and X. Luo, "Multicast service provisioning under a scheduled traffic model in WDM optical networks," in *IASTED International Conference on Communications and Computer Networks*, 2005.
- [94] N. Charbonneau and V.M. Vokkarane, "Multicast advance reservation RWA heuristics in wavelength-routed networks," in *Proc. IEEE GLOBECOM*, 2010.
- [95] N. Charbonneau and V.M. Vokkarane, "Static routing and wavelength assignment for multicast advance reservation in all-optical wavelength-routed WDM networks," in *UMass Dartmouth Tech Report, UMASSD-CIS-TR-2009006*, Available: <http://www.umassd.edu/engineering/cis/reports/report.cfm?r=31>, Dec. 2009.
- [96] T. Li, B. Wang, C. Xin, and X. Zhang, "On survivable service provisioning in WDM optical networks under a scheduled traffic model," in *Proc. IEEE GLOBECOM*, Nov. 2005.
- [97] T. Li and B. Wang, "On optimal survivability design in WDM optical networks under a scheduled traffic model," in *International Workshop on Design of Reliable Communication Networks*, Oct. 2005, p. 8.
- [98] Tianjian Li and Bin Wang, "Approximating optimal survivable routing in WDM optical networks under a scheduled traffic model," in *IEEE Sarnoff Symposium*, Mar. 2006, pp. 1–4.
- [99] Tianjian Li and Bin Wang, "Approximating optimal survivable scheduled service provisioning in WDM optical networks with iterative survivable routing," in *Proc. IEEE BroadNets*, Oct. 2006, pp. 1–10.
- [100] B. Wang and T. Li, "Survivable scheduled service provisioning in WDM optical networks with iterative routing," *Optical Switching and Networking*, vol. 7, no. 1, pp. 28–38, 2009.
- [101] Tianjian Li and Bin Wang, "On optimal survivability design under a scheduled traffic model in wavelength-routed optical mesh networks," in *Communication Networks and Services Research Conference*, May 2006, pp. 8–33.
- [102] Tianjian Li and Bin Wang, "Path-protection-based routing and wavelength assignment in wavelength-division multiplexing optical networks under a scheduled traffic model," *J. Opt. Netw.*, vol. 5, no. 7, pp. 575–588, 2006.
- [103] Tianjian Li and Bin Wang, "Approximating optimal survivable scheduled service provisioning in WDM optical networks with shared risk link groups," in *Proc. IEEE BroadNets*, Sept. 2007, pp. 601–610.
- [104] A. Jaekel, "Lightpath scheduling and allocation under a flexible scheduled traffic model," in *Proc. IEEE GLOBECOM*, Dec. 2006, pp. 1–5.
- [105] A. Jaekel and Y. Chen, "Demand allocation without wavelength conversion under a sliding scheduled traffic model," in *Proc. IEEE BroadNets*, Sept. 2007, pp. 495–503.
- [106] A. Jaekel and Y. Chen, "Resource provisioning for survivable WDM networks under a sliding scheduled traffic model," *Optical Switching and Networking*, vol. 6, no. 1, pp. 44–54, 2008.
- [107] Wenda Ni, M. Schlosser, Qingshan Li, Yili Guo, Hanyi Zhang, and Xiaoping Zheng, "Achieving optimal lightpath scheduling in survivable WDM mesh networks," in *Proc. Optical Fiber Communication Conference (OFC)*, Feb. 2008, pp. 1–3.
- [108] Qingshan Li, Wenda Ni, Yanhe Li, Yili Guo, Hanyi Zhang, and Xiaoping Zheng, "Improving the dual-failure restorability in scheduled WDM mesh networks," in *Communications and Photonics conference and Exhibition*, Nov. 2009, pp. 1–2.
- [109] C. Vijaya Saradhi, M. Gurusamy, and R. Piesiewicz, "Routing fault-tolerant sliding scheduled traffic in WDM optical mesh networks," in *Proc. IEEE BroadNets*, Sept. 2008, pp. 197–202.
- [110] Shuqiang Zhang and Chun-Kit Chan, "Multicast protection in WDM optical networks with scheduled traffic," in *European Conference on Optical Communication*, Sept. 2009, pp. 1–2.
- [111] A. Jaekel, Ying Chen, and A. Bari, "Survivable traffic grooming for scheduled demands," in *Proc. ICST BroadNets*, Sept. 2008, pp. 176–183.
- [112] B. Wang, T. Li, X. Luo, and Y. Fan, "Traffic grooming under a sliding scheduled traffic model in WDM optical networks," in *Proc. IEEE BroadNets*, 2004.
- [113] A. Jaekel, Y. Chen, and A. Bari, "Stable logical topologies for survivable traffic grooming of scheduled demands," *IEEE/OSA J. Optical Commun. Netw.*, vol. 2, pp. 782–792, 2010.
- [114] M. Gagnaire, J. Kuri, and E. Doumith, *Traffic Grooming for Optical Networks*, chapter Grooming of Scheduled Demands in Multi-Layer Optical Networks, Springer, 2008.
- [115] N. Charbonneau, C. Guok, I. Monga, and V. Vokkarane, "Advance reservation frameworks in hybrid IP-WDM networks," to appear in *IEEE Commun. Mag., Special Issue on Hybrid Networking: Evolution Towards Combined IP Services and Dynamic Circuit Network Capabilities*, May. 2011.
- [116] Chin Guok, D. Robertson, M. Thompson, J. Lee, B. Tierney, and W. Johnston, "Intra and interdomain circuit provisioning using the OSCARS reservation system," in *Proc. IEEE BroadNets*, Oct. 2006, pp. 1–8.
- [117] S. Tanwir, L. Battestilli, H. Perros, and G. Karmous-Edwards, "Monitoring and Discovery for EnLIGHTened Computing," in *High-Capacity Optical Networks and Enabling Technologies*, 2007.
- [118] A. Takefusa, H. Nakada, T. Kudoh, Y. Tanaka, and S. Sekiguchi, "GridARS: An advance reservation-based Grid co-allocation framework for distributed computing and network resources," in *International*

- Workshop Job Scheduling Strategies for Parallel Processing*, 2007, pp. 152–168.
- [119] M. Hayashi, H. Tanaka, and M. Suzuki, “Advance reservation-based network resource manager for optical networks,” in *Proc. Optical Fiber Communication Conference (OFC)*, 2008, pp. 1–3.
 - [120] Steven R. Thorpe, Lina Battestilli, Gigi Karmous-Edwards, Andrei Hutanu, Jon MacLaren, Joe Mambretti, John H. Moore, Kamaraju Syam Sundar, Yufeng Xin, Atsuko Takefusa, Michiaki Hayashi, Akira Hirano, Shuichi Okamoto, Tomohiro Kudoh, Takahiro Miyamoto, Yukio Tsukishima, Tomohiro Otani, Hidemoto Nakada, Hideaki Tanaka, Atsushi Taniguchi, Yasunori Sameshima, and Masahiko Jinno, “G-lambda and enLIGHTened: wrapped in middleware co-allocating compute and network resources across Japan and the US,” in *GridNets '07: Proc. first international conference on Networks for grid applications*, 2007.
 - [121] M. Pilz, C. Barz, U. Bornhauser, P. Martini, C. deWall, and A. Willner, “ARGON: Reservation in Grid-enabled networks,” in *DFN Forum on Communication Technologies*, 2008.
 - [122] A. Willner, C. Barz, J. A. G. Espin, J. F. Riera, S. Figuerola, and P. Martini, “Harmony-advance reservations in heterogeneous multi-domain environments,” *Lecture Notes in Computer Science*, 2009.
 - [123] G. Zervas, E. Escalona, R. Nejabati, D. Simeonidou, G. Carrozzo, N. Ciulli, B. Belter, A. Binczewski, M.S. Poznan, A. Tzanakaki, and G. Markidis, “Phosphorus grid-enabled GMPLS control plane (GMPLS): architectures, services, and interfaces,” *IEEE Commun. Mag.*, vol. 46, no. 6, pp. 128–137, Jun. 2008.
 - [124] L.-O. Burchard, H.-U. Heiss, and C.A.F. De Rose, “Performance issues of bandwidth reservations for Grid computing,” in *Symposium on Computer Architecture and High Performance Computing*, Nov. 2003, pp. 82–90.
 - [125] L.O. Burchard and M. Droste-Franke, “Fault tolerance in networks with an advance reservation service,” in *International Workshop on Quality of Service (IWQoS) 2003*, 2003, pp. 215–228.
 - [126] L.-O. Burchard, B. Linnert, and J. Schneider, “A distributed load-based failure recovery mechanism for advance reservation environments,” in *IEEE International Symposium on Cluster Computing and the Grid*, 9–12 2005, vol. 2, pp. 1071–1078 Vol. 2.
 - [127] L.-O. Burchard, “Networks with advance reservations: Applications, architecture, and performance,” *Journal of Network and Systems Management*, vol. 13, no. 4, pp. 429–449, 2005.
 - [128] L.-O. Burchard, “Analysis of data structures for admission control of advance reservation requests,” *IEEE Trans. Knowl. Data Eng.*, vol. 17, no. 3, pp. 413–424, march 2005.
 - [129] L.-O. Burchard, B. Linnert, and J. Schneider, “Rerouting strategies for networks with advance reservations,” in *e-Science and Grid Computing, 2005. First International Conference on*, 2005, p. 8.
 - [130] B. Wang and A. Deshmukh, “An all hops optimal algorithm for dynamic routing of sliding scheduled traffic demands,” *IEEE Commun. Lett.*, vol. 9, no. 10, pp. 936–938, Oct. 2005.
 - [131] N.S.V. Rao, Qishi Wu, Song Ding, S.M. Carter, W.R. Wing, A. Banerjee, D. Ghosal, and B. Mukherjee, “Control plane for advance bandwidth scheduling in ultra high-speed networks,” in *Proc. IEEE INFOCOM*, 23–29 2006, pp. 1–5.
 - [132] S. Sahni, N. Rao, S. Ranka, Yan Li, Eun-Sung Jung, and N. Kamath, “Bandwidth scheduling and path computation algorithms for connection-oriented networks,” in *Networking, 2007. ICN '07. Sixth International Conference on*, 22–28 2007, pp. 47–47.
 - [133] Yunyue Lin, Qishi Wu, N.S.V. Rao, and Mengxia Zhu, “On design of scheduling algorithms for advance bandwidth reservation in dedicated networks,” in *IEEE INFOCOM Workshops*, Apr. 2008, pp. 1–6.
 - [134] Eun-Sung Jung, Yan Li, S. Ranka, and S. Sahni, “An evaluation of in-advance bandwidth scheduling algorithms for connection-oriented networks,” in *Parallel Architectures, Algorithms, and Networks, 2008. I-SPAN 2008. International Symposium on*, 7–9 2008, pp. 133–138.
 - [135] C. Barz, M. Pilz, and A. Wichmann, “Temporal routing metrics for networks with advance reservations,” in *International Symposium on Cluster Computing and the Grid*, May 2008, pp. 710–715.
 - [136] R. Cohen, N. Fazlollahi, and D. Starobinski, “Path switching and grading algorithms for advance channel reservation architectures,” *IEEE/ACM Trans. Netw.*, vol. 17, no. 5, pp. 1684–1695, Oct. 2009.
 - [137] C. Xie, F. Xu, N. Ghani, E. Chaniotakis, C. Guok, and T. Lehman, “Load-Balancing for Advance Reservation Connection Rerouting,” *IEEE Commun. Lett.*, vol. 14, no. 6, pp. 1, 2010.
 - [138] A.N. Patel, Yi Zhu, and J.P. Jue, “Routing and horizon scheduling for time-shift advance reservation,” in *Proc. Optical Fiber Communication Conference (OFC)*, Mar. 2009, pp. 1–3.
 - [139] A. N. Patel, M. M. Hasan, Y. Zhu, and J. P. Jue, “Routing and scheduling for variable bandwidth advance reservation in elastic applications,” in *Proc. Optical Fiber Communication Conference (OFC)*, 2009.
 - [140] B. Garcia-Manrubia, R. Aparicio-Pardo, P. Pavon-Marino, N. Skorin-Kapov, and Garcia-Haro, “Milp formulations for scheduling lightpaths under periodic traffic,” in *International Conference on Transparent Optical Networks*, June 2009, pp. 1–4.
 - [141] N. Skorin-Kapov, P. Pavon-Marino, B. Garcia-Manrubia, and R. Aparicio-Pardo, “Scheduled virtual topology design under periodic traffic in transparent optical networks,” in *Proc. IEEE BroadNets*, 14–16 2009, pp. 1–8.
 - [142] P. Pavon-Marino, R. Aparicio-Pardo, B. Garcia-Manrubia, and N. Skorin-Kapov, “Virtual topology design and flow routing in optical networks under multihour traffic demand,” *Photonic Network Communications*, vol. 19, no. 1, pp. 42–54, 2010.
 - [143] Ch. Bouras and K. Stamos, “An adaptive admission control algorithm for bandwidth brokers,” in *International Symposium on Network Computing and Applications*, Sept. 2004, pp. 243–250.
 - [144] S. Schmidt and J. Kunegis, “Scalable bandwidth optimization in advance reservation networks,” in *IEEE International Conference on Networks*, Nov. 2007, pp. 95–100.
 - [145] Xiangfei Zhu, M.E. McGinley, Tao Li, and M. Veeraraghavan, “An analytical model for a book-ahead bandwidth scheduler,” in *Proc. IEEE GLOBECOM*, Nov. 2007, pp. 2280–2285.
 - [146] U. Farooq, S. Majumdar, and E.W. Parsons, “Dynamic scheduling of lightpaths in lambda Grids,” in *Proc. IEEE BroadNets*, Oct. 2005, vol. 2, pp. 1463–1472.
 - [147] S. Naiksatam and S. Figueira, “Elastic reservations for efficient bandwidth utilization in LambdaGrids,” *Future Gener. Comput. Syst.*, vol. 23, no. 1, pp. 1–22, 2007.
 - [148] H. Miyagi, M. Hayashitani, D. Ishii, Y. Arakawa, and N. Yamanaka, “Advanced wavelength reservation method based on deadline-aware scheduling for lambda grid networks,” *IEEE/OSA J. Lightwave Technology*, vol. 25, no. 10, pp. 2904–2910, Oct. 2007.
 - [149] C. Castillo, G.N. Rouskas, and K. Harfoush, “Efficient resource management using advance reservations for heterogeneous grids,” in *IEEE International Symposium on Parallel and Distributed Processing*, April 2008, pp. 1–12.
 - [150] E. Varvarigos, V. Sourlas, and K. Christodoulopoulos, “Routing and scheduling connections in networks that support advance reservations,” *Comput. Netw.*, vol. 52, no. 15, pp. 2988–3006, 2008.
 - [151] X. Luo and B. Wang, “Service provisioning under a scheduled traffic model using light-trails in WDM optical networks,” in *Proc. IEEE BroadNets*, Sept. 2007, pp. 385–393.
 - [152] X. Luo and B. Wang, “On service provisioning using light-trails in WDM optical networks with waveband switching,” *Photonic Network Communications*, pp. 1–9, 2010.
 - [153] I. Ahmad and J. Kamruzzaman, “Preemption-aware instantaneous request call routing for networks with book-ahead reservation,” *IEEE Trans. Multimedia*, vol. 9, no. 7, pp. 1456–1465, 2007.
 - [154] M. Chamania and A. Jukan, “A survey of inter-domain peering and provisioning solutions for the next generation optical networks,” *IEEE Commun. Surveys Tutorials*, vol. 11, no. 1, pp. 33–51, 2009.
 - [155] A. Birman, “Computing approximate blocking probabilities for a class of all-optical networks,” *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 852–857, Jun. 1996.
 - [156] I. Foster, C. Kesselman, C. Lee, B. Lindell, K. Nahrstedt, and A. Roy, “A distributed resource management architecture that supports advance reservations and co-allocation,” in *Proc. Workshop on Quality of Service*, 1999, pp. 27–36.
 - [157] A. Jukan and G. Edwards, “Optical control plane for the grid community,” *IEEE Commun. Surveys Tutorials*, vol. 9, no. 3, pp. 30–44, 2007.



Neal Charbonneau (S'08) received the B.S. degree and M.S. degree in Computer Science from the University of Massachusetts, Dartmouth in 2008 and 2010, respectively. He is currently working at The MITRE Corporation. His interests include computer networks and software design and development.



Vinod M. Vokkarane (S'02-M'04-SM'09) received the B.E. degree with Honors in Computer Science and Engineering from the University of Mysore, India in 1999, the M.S. degree in Computer Science from the University of Texas at Dallas in 2001, and the Ph.D. degree in Computer Science from the University of Texas at Dallas in 2004. Dr. Vinod Vokkarane is an Associate Professor of Computer and Information Science at the University of Massachusetts, Dartmouth. He is currently a Visiting Scientist at the Research Laboratory of Electronics (RLE) at Massachusetts Institute of Technology (MIT). His primary areas of research include design and analysis of architectures and protocols for

optical and wireless networks. Dr. Vokkarane is the co-author of a book, *Optical Burst Switched Networks*, Springer, 2005. He is a recipient of the Texas Telecommunication Engineering Consortium Fellowship 2002-03, and the University of Texas at Dallas Computer Science Dissertation of the Year Award 2003-04. He is a recipient of Best Paper Award at the IEEE GLOBECOM 2005 Symposium on Optical Systems and Networks. He currently serves on the Editorial Board of the *IEEE Communication Letters*. He has been on the technical program committees of several IEEE conferences including INFOCOM, ICC, GLOBECOM, ICCCN, HSN, and ANTS, and served as TPC Co-Chair for the Optical Networks and Systems symposia at ICCCN 2007 and 2010, GLOBECOM 2011, and ICC 2012 and INFOCOM High-Speed Networks (HSN) workshop 2011.