

Advance Reservation for Virtual Overlay Network Services

Feng Gu¹, Chongyang Xie¹, Min Peng², Çiçek Çavdar³, Samee Khan⁴, Nasir Ghani¹

¹University of New Mexico, ²Wuhan University, ³KTH Royal Institute of Technology, ⁴North Dakota State University

Abstract: Network advance reservation is an area of growing interest and focus today. However most proposed solutions in this field have only focused on the scheduling of point-to-point client connections. Now as applications and user communities continue to expand, there is a growing need to extend advance reservation support for multipoint services. To address this concern, a heuristic solution is presented for virtual overlay network scheduling. The scheme uses graph-theoretic schemes and is analyzed in detail using network simulation.

Keywords: Advance reservation, network scheduling, topology overlay, virtual network services

I. INTRODUCTION

Advance reservation (AR) services have seen increased focus and interest in recent years. In particular, the ability to schedule future connection requests is becoming an integral requirement for a range of applications in grid-computing, e-science, storage backup, special event broadcasting, etc. As a result, a wide range of AR scheduling algorithms have been studied in recent years [1]-[6], including optimization formulations and graph-theoretic heuristics. Others have also addressed broader topics such as AR survivability [7], re-routing schemes [8], and even distributed solutions [9].

Nevertheless, with expanding applications there is a growing need to extend AR capabilities to support multipoint service connectivity between dispersed users/sites. For example, ubiquitous *virtual private network* (VPN) and/or overlay network services have been widely used to connect organizations and communities of interest over common physical network infrastructures [10]-[15]. However related studies in VPN design have only looked at a-priori or “on-demand” provisioning models for *immediate reservation* (IR). As a result there is a critical need to address future scheduling of multipoint services as well, termed herein as *virtual overlay network scheduling* (VONS). To the best of our knowledge, there are no known studies that address this particular area, i.e., combining AR scheduling with VPN overlay provisioning. Hence this forms the key motivation for the work herein.

Overall, this paper is organized as follows. Section II first presents a brief overview of the existing work in AR provisioning and overlay services. Next, Section III presents a heuristic solution for the VONS problem using graph-theoretic schemes. Section IV then presents detailed performance analysis results for the scheme using network simulation. Finally conclusions and directions for future work are presented in Section V.

II. BACKGROUND

Researchers have studied many different AR schemes for bandwidth-provisioning networks. A broad and in-depth introduction to the area is presented in [1], including detailed discussions on various possible AR service models. Meanwhile [2] outlines scheduling schemes for users with variable starting time/fixed duration requests using an augmented Bellman-Ford approach. The goal here is to accommodate incoming requests at the earliest possible time. [3] extends this work to handle request variations with different path and bandwidth constraints, e.g., variable/fixed path, variable/fixed bandwidth. Detailed polynomial and pseudo-polynomial complexity scheduling algorithms are proposed here, but akin to [2], detailed simulation analysis studies are not done. Meanwhile [4] looks at two different AR schemes for scheduling fixed path/variable bandwidth and variable path/variable bandwidth AR requests. Both of these problems are shown to be NP-complete and related heuristics presented for each. Overall results show good performance gains over greedy algorithms here.

Others have also looked at AR scheduling for optical *dense wavelength division multiplexing* (DWDM) networks. For example, [5] studies *routing and wavelength assignment* (RWA) in the AR context and develops three lightpath scheduling algorithms. However performance analysis results are not presented here. Meanwhile [6] studies lightpath RWA scheduling in “all-optical” networks and compares several spatio-temporal optimizations against sequential RWA schemes (yielding good improvements in setup success rates). Finally, some have also looked at broader aspects of the AR problem, including protection [7], re-routing [8], and distributed implementations [9].

Meanwhile virtual overlay networking design has also seen much focus over the years, with researchers detailing a host of solutions based upon optimization and heuristic approaches. Notable examples include the *resilient overlay network* (RON) study [10] which develops a static overlay for best-effort Internet networks to improve resiliency. The *service overlay network* (SON) study [11] also addresses *quality of service* (QoS) support using queuing theory and optimization methods to implement static partitioning and oversubscription. Meanwhile “dynamic” topology overlays have also been considered. For example, [12] proposes a *virtual network* (VN) scheme for node/link selection with reconfiguration, whereas [13],[14] study multipoint Ethernet overlays in over next-generation SONET/SDH networks. Finally, numerous efforts have addressed *virtual topology* (VT) design for optical DWDM networks, see survey in [15]. Most of these schemes have pursued optimization-based formulations for a-priori virtual link sets.

In summary, existing AR studies have only looked a *point-to-point* connection scheduling. Meanwhile VPN and overlay network provisioning schemes have only treated immediate user demands. Indeed, no efforts have looked at future scheduling of multipoint “network overlay” type requests. Hence there is a strong need to investigate new schemes for *virtual overlay network scheduling* (VONS), i.e., combining the fields of AR and VPN/overlay provisioning. These issues are now investigated further.

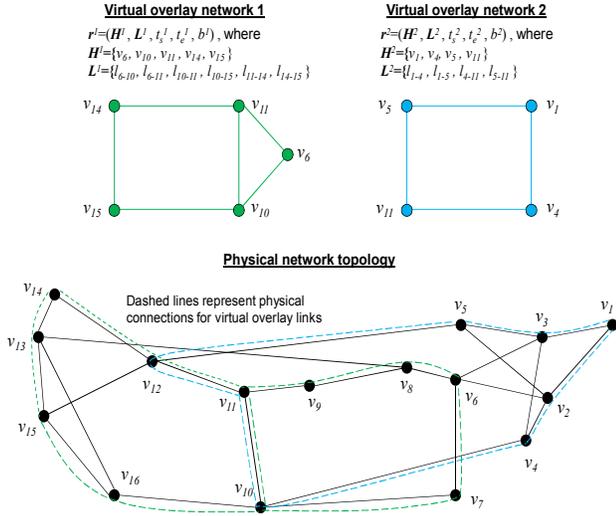


Figure 1: Virtual overlay network services

III. VONS HEURISTIC SOLUTION

A heuristic VONS solution is now presented for regular bandwidth-provisioning networks (with further considerations for optical DWDM settings left for future study). The proposed solution uses graph-theoretic algorithms and assumes that all overlay requests arrive in a random “on-demand” manner, i.e., reflective of real-world settings as opposed to idealized a-priori settings. However, before presenting this approach, the requisite notation is first introduced (all vector and set entities denoted in bold).

The overall physical network is modeled as a graph topology, $\mathbf{G}(\mathbf{V}, \mathbf{E})$, where $\mathbf{V} = \{v_1, v_2, \dots\}$ is the set of network router/switch nodes and $\mathbf{E} = \{e_{ij}\}$ the set of network links, i.e., e_{ij} is the link between v_i and v_j . Furthermore, without loss of generality, all links have fixed capacity, C , and connectivity is assumed to be bidirectional, i.e., two opposing uni-directional links between neighboring nodes. In addition, it is assumed that each link e_{ij} has an associated bandwidth-time function, i.e., $c_{ij}(t)$, which tracks the used capacity at future time instants. Now the n -th overlay request is denoted by the 5-tuple $\mathbf{r}^n = (\mathbf{S}^n, \mathbf{L}^n, t_s^n, t_e^n, b^n)$, where \mathbf{S}^n is the set of nodes/sites ($\mathbf{S}^n \in \mathbf{V}$), \mathbf{L}^n is the set of desired overlay links between node sites in \mathbf{S}^n , t_s^n is the start time, t_e^n is the stop time, and b^n is the requested bandwidth, $b^n \leq C$.

Hence, the client is basically requesting setup of a virtual network to interconnect its sites in \mathbf{S}^n via a set of “virtual” overlay links given in \mathbf{L}^n (i.e., implemented using physical network connections, see Figure 1). Note that this AR service model only assumes fixed request intervals, and future variants can relax this assumption.

The complete pseudocode for the VONS heuristic is shown in Figure 2 and assumes random “on-demand” request arrivals. The algorithm basically loops through all requested overlay links in the overlay topology and attempts to route each using a given *traffic engineering* (TE) strategy. Note that all of these computations are done over a temporary copy of the network graph, $\mathbf{G}'(\mathbf{V}, \mathbf{E})$, derived from $\mathbf{G}(\mathbf{V}, \mathbf{E})$ by removing *non-feasible* physical links, i.e., those without sufficient capacity in the request interval, $c_{ij}(t) < b^n$ in $[t_s^n, t_e^n]$. Furthermore, three different TE routing approaches are considered here for AR, i.e., *minimum hop count*, *maximum bottleneck capacity*, and *minimum distance*. These strategies are now detailed further:

```

/* Given incoming request  $\mathbf{r}^n = (\mathbf{S}^n, \mathbf{L}^n, t_s^n, t_e^n, b^n)$  */
Generate temporary graph copy  $\mathbf{G}'(\mathbf{V}, \mathbf{E}) = \mathbf{G}(\mathbf{V}, \mathbf{E})$ 
Remove non-feasible links in  $\mathbf{G}'(\mathbf{V}, \mathbf{E})$ , i.e.,  $c_{ij}(t) < b^n$  in  $[t_s^n, t_e^n]$ 

/* Loop and provision all overlay links in request */
for  $i = 1$  to  $|\mathbf{L}^n|$ 
  if (minimum hop count)
    - Assign unity weight to links in  $\mathbf{G}'(\mathbf{V}, \mathbf{E})$ 
    - Compute Dijkstra's shortest-path for  $i$ -th overlay link
  else if (minimum distance)
    - Assign dynamic weight to links in  $\mathbf{G}'(\mathbf{V}, \mathbf{E})$ , Eq. (1)
    - Compute Dijkstra's shortest-path for  $i$ -th overlay link
  else (maximum bottleneck)
    - Compute k-SP between end-points of  $i$ -th overlay link
    - Select path with maximum bottleneck bandwidth
  end

if (all overlay link connections in  $\mathbf{L}^n$  routed)
  Setup successful, copy  $\mathbf{G}'(\mathbf{V}, \mathbf{E}) \rightarrow \mathbf{G}(\mathbf{V}, \mathbf{E})$ 
else
  Reject request
  
```

Figure 2: VONS heuristic algorithm

Minimum hop count: This scheme assigns unity weights to all links and finds the shortest feasible route in $\mathbf{G}'(\mathbf{V}, \mathbf{E})$, i.e., via Dijkstra’s algorithm. This approach pursues pure resource minimization (as it does not account for link loads) and may cause saturation on certain network links.

Maximum bottleneck: This scheme ignores path lengths and instead chooses an end-to-end route with the maximum available capacity. Specifically, the *k-shortest paths* (k-SP) are computed between the source and destination nodes, and the one with the maximum residual capacity (on its bottleneck link during the request interval) is selected. This is deemed as a pure load-balancing approach and can give longer path lengths. Note that k-SP computation generally entails much higher complexity.

Minimum distance: This scheme achieves a balance between the above two strategies by assigning dynamic weights to all links. Namely, the cost of each link is set as inversely-proportional to the lowest residual capacity in the request interval, i.e., cost for link e_{ij} is given as:

$$\omega_{ij} = C / (\chi + \epsilon), \quad \text{Eq. 1,}$$

where

$$\chi = \min \{b_{ij}(t)\} \text{ in } [t_s^n, t_e^n] \quad \text{Eq. 2,}$$

and ϵ is a small value chosen to avoid division errors.

The overall algorithm in Figure 2 only returns success if all overlay link connections are successfully routed over the temporary graph. In this case the network resource graph is also updated, i.e., $G^*(V,E) \rightarrow G(V,E)$.

IV. PERFORMANCE ANALYSIS

The performance of the heuristic VONS solution is analyzed using specialized simulation models developed in *OPNET Modeler*TM. These tests consider two topologies, including the ubiquitous NSFNET with a moderate node degree of 3.12 (16 nodes, 25 links) and the denser *Deutsche Telekom* (DT) backbone with a higher node degree of 3.85 (27 nodes, 52 links). All network nodes are generic IP *multi-protocol label switching* (MPLS) routers with 10 Gbps links. Now the client overlay requests have random exponentially-distributed holding and inter-arrival times, with means μ and λ , respectively. In particular, a scaled mean holding time of $\mu=600$ seconds is used and the mean inter-arrival times are adjusted according to load. Carefully note that this mean holding time is just a relative figure and does not necessarily reflect real-world values.

Meanwhile all overlay topology sizes are uniformly varied from 4-6 nodes with corresponding (overlay link) bandwidth requests ranging uniformly from 200 Mb/s to 1.0 Gb/s in 200 Mbps increments, i.e., to model fractional Ethernet demands. Here the node/sites in the overlay topologies are randomly selected from the network nodes. Furthermore the virtual overlay links are selected using the random Inet topology generator [16] to achieve a node degree of about 2.5. All runs are averaged over 500,000 requests and a modified Erlang loading metric is used to account for connection group sizes, as in [14]:

$$\text{Modified Erlang load} = \sum_{n=4}^6 (n-1) \cdot \frac{\mu}{\lambda} \quad \text{Eq. (3),}$$

where the overlay topology sizes range from 4-6 nodes and the $1/\lambda$ represents the mean inter-arrival rate.

Initial tests are done to measure overlay request blocking rates for the various TE (overlay link) routing strategies, Figure 3. These results show that the minimum distance heuristic (Eq. 1) consistently gives the lowest blocking for both topologies. In particular this scheme yields less than half the blocking rate of the maximum bottleneck scheme at

low-medium input loads (as shown on log scale in Figure 3). By contrast the maximum bottleneck scheme only yields marginal (negligible) blocking reduction over the minimum hop count scheme, despite its sizeable increase in complexity, i.e., k -SP computations versus Dijkstra's shortest path.

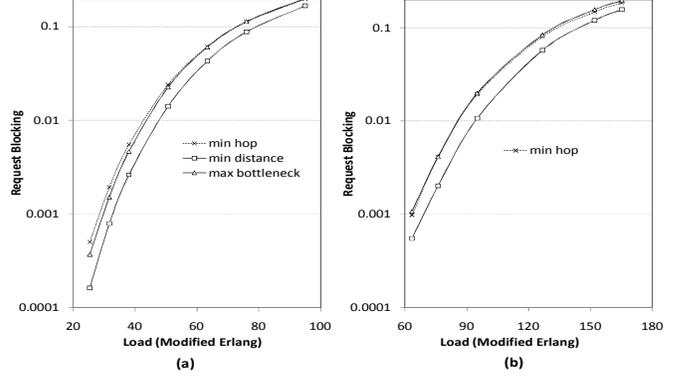


Figure 3: Overlay request blocking: a) NSFNET, b) DT

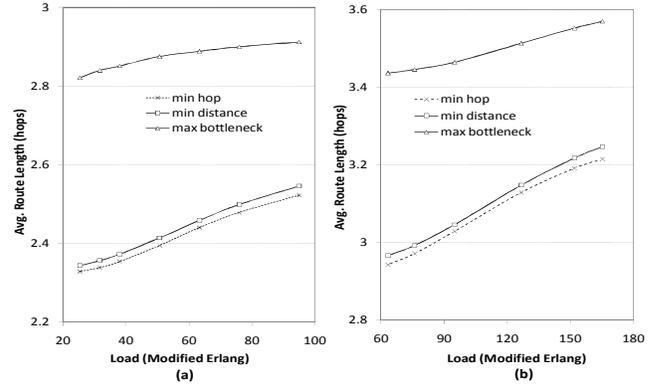


Figure 4: Average path lengths: a) NSFNET, b) DT

Next, the average path lengths of the computed routes for the overlay links (in the overlay topologies) are plotted in Figure 4. As expected, the minimum hop count routing scheme gives the lowest utilization of all. Nevertheless, the minimum distance heuristic provides very competitive performance here, with average path length values coming to within 2% of the minimum hop count scheme. Conversely, the hop count utilization for the maximum bottleneck scheme is notably higher, by about 15% for both topologies.

Carefully note that the VONS heuristic in Figure 2 tries to setup all overlay link (connections) in the order in which they are listed in the request, i.e., random. Hence modified versions of the heuristic were also tested to better “shuffle” the sequence in which the overlay links are attempted, e.g., longest-path first, etc. However, these renditions showed no noticeable reduction in blocking over the base random ordering scheme. Overall, these findings indicate that the minimum distance scheme gives the best performance tradeoff in terms of minimizing blocking and lowering resource consumption on overlay link routes.

V. CONCLUSIONS

There is a growing need to develop advance reservation scheduling algorithms for overlay network services. This paper addresses this concern and proposes a graph-theoretic heuristic solution to schedule random “on-demand” virtual overlay topology requests. The scheme incorporates a range of traffic engineering objectives, including resource minimization and load-balancing. Detailed simulation results show that dynamic link-load weighting gives the best performance/tradeoffs in terms of blocking reduction and resource utilization. Future efforts will look at developing more formalized optimization models for this problem and also incorporating survivability capabilities.

VI. ACKNOWLEDGEMENTS

This research has been supported by the US Department of Energy Office of Science and the National Science Foundation. The authors are very grateful to these sponsors for their generous support.

VII. REFERENCES

- [1] R. Guerin, A. Orda, “Networks with Advance Reservations: The Routing Perspective”, *IEEE INFOCOM 2000*, TelAviv, Mar 2000.
- [2] N. Rao, *et al*, “Control Plane for Advance Bandwidth Scheduling in Ultra High-Speed Networks”, *IEEE INFOCOM 2006 High Speed Networks Workshop*, Barcelona, Spain, April 2006.
- [3] Yunyue Lin, *et al*, “On Design of Scheduling Algorithms for Advance Bandwidth Reservation in Dedicated Networks”, *IEEE INFOCOM 2008 High Speed Networks Workshop*, Phoenix, AZ, April 2008, pp.1-6.
- [4] Y. Lin, Q. Wu, “Path Computation with Variable Bandwidth for Bulk Data Transfer in High-Performance Networks”, *IEEE INFOCOM 2009 High Speed Networks Workshop*, Rio De Janeiro, Brasil, April 2009.
- [5] J. Zheng, H. Mouftah, “Routing and Wavelength Assignment for Advance Reservation in Wavelength-routed WDM Optical Networks”, *IEEE ICC 2002*, New York, NY, May 2002.
- [6] J. Kuri, *et al*, “Routing and Wavelength Assignment of Scheduled Lightpath Demands”, *IEEE Journal on Selected Areas in Communications*, Vol. 21, No. 8, October 2003, pp.1231-1240.
- [7] T. Li, *et al*, “On Survivable Service Provisioning in WDM Optical Networks under a Scheduled Traffic Model”, *IEEE GLOBECOM 2005*, St. Louis, MO, December 2005.
- [8] C. Xie, *et al*, “Load-Balancing Connection Rerouting in Advance Reservation Networks”, *IEEE Communications Letters*, Vol. 14, No. 6, June 2010, pp. 578-580.
- [9] C. Xie, H. Alazemi, N. Ghani, “Routing and Scheduling in Distributed Advance Reservation Networks”, *IEEE GLOBECOM 2010*, Miami, FL, December 2010.
- [10] D. Andersen, *et al*, “Resilient Overlay Networks”, *ACM SOSP 2001*, October 2001, pp. 131-185.
- [11] Z. Duan, *et al*, “Service Overlay Networks: SLAs, QoS, and Bandwidth Provisioning”, *IEEE/ACM Transactions on Networking*, Vol. 11, No. 6, December 2003, pp. 870-883.
- [12] Y. Zhu, M. Ammar, “Algorithms for Assigning Substrate Network Resources to Virtual Network Components,” *IEEE INFOCOM 2006*, Barcelona, Spain, April 2006.
- [13] C. Xie, *et al*, “Multi-Point Ethernet Over Next-Generation SONET/SDH”, *IEEE ICC 2009*, Dresden, June 2009.
- [14] C. Xie, *et al*, “Traffic Engineering For Ethernet over SONET/SDH (EoS): Advances and Frontiers”, *IEEE Network*, Vol. 23, No. 3. May/June 2009, pp. 18-25.
- [15] R. Dutta, G. Rouskas, “A Survey of Virtual Topology Design Algorithms for Wavelength Routed Optical Networks,” *Optical Networks*, Vol. 1, No. 1, January 2000, pp. 73-89.
- [16] Inet topology generator, see <http://topology.eecs.umich.edu/inet/>.