

ERONs: Dynamic Lightpath Networking via Overlay Control of Static Optical Connections

(Invited Paper)

Gigi Karmous-Edwards, Arun Vishwanath, Douglas Reeves, Lina Battestilli, Priyanka Vegesna, George N. Rouskas
Department of Computer Science, North Carolina State University, Raleigh, NC, USA

Abstract—To bridge the gap between the current practice of setting up expensive, dedicated, lightpath connections (i.e., static topologies), and the distant future vision of inexpensive access to dynamically switched end-to-end lightpaths we propose a medium term solution in the form of *edge-reconfigurable optical networks (ERONs)*. An ERON is an overlay-control network created by installing readily available MEMS optical switches and implementing a GMPLS control plane at sites interconnected by static lightpaths. The switches and control software are deployed at the edge of the network and operated by the organization-user (i.e., outside the network provider’s control), hence the term “ERON.” By providing dynamic, automated control of end-to-end lightpaths, ERONs enable the sharing of expensive network resources among multiple users and applications that require sporadic access to these resources. We develop an algorithm for creating an ERON from an existing topology of static lightpaths, and we present simulation results that quantify the benefits of ERONs in terms of lightpaths compared to assembling a set of independent and dedicated circuits.

I. INTRODUCTION

Many existing and emerging classes of high-end applications involve complex, intensive computations on large data sets in a manner that requires coordination of resources residing at several geographically dispersed sites. Such applications arise in a wide range of domains, including e-Science and scientific discovery, distributed simulation and visualization, petascale data mining and analytics, education, intelligence gathering and analysis, and defense. Current research practice also tends towards collaborations among multiple groups and institutions that require high bandwidth connections to network-attached resources [7]. To enable these applications, lightpaths along end-to-end paths across multi-domain networks, must come up and go down based on user requirements and on short timescales (i.e., sub-seconds to seconds).

Ideally, the optical network would provide native support for establishing lightpaths on demand, making it possible to update the logical topology over time in response to traffic (i.e., application) demands. This model, which we refer to as *core-reconfigurable optical network (CRON)*, encompasses the vision of a dynamically reconfigurable optical core that has been contemplated for over a decade [4]. This vision underlies the DARPA CORONET project [5] which seeks to capitalize

on the maturing of optical switching technologies and the development of signaling protocol standards (e.g., GMPLS), so as to prototype commercially viable approaches for network providers to offer users the ability to set up and tear down optical connections dynamically.

Nevertheless, despite years of research and development, the vision of a dynamically reconfigurable optical core is far from realization. Today, setting up high bandwidth connections on demand is nearly impossible due to a fundamental lack of infrastructure capabilities to support the automatic and rapid establishment of end-to-end lightpaths. To establish such paths several network providers need to be coordinated (one regional optical network at each end, as well as one national-scale backbone provider) and multiple contracts negotiated. Due to this administrative burden and associated timescales for the establishment of lightpaths, it is typical for lightpath connections to be held in place for long periods of time (e.g., months or longer). The result is the creation of *static topologies* consisting of a collection of independent lightpaths, each dedicated to serving a specific pair of high-end users or devices. For instance, several national research and education networks (NRENs) provide point-to-point optical connections, often reaching across more than one administrative domains, to serve existing large-scale applications.

Since typical applications require *sporadic* rather than continuous use of these lightpaths, the utilization of static point-to-point connections can be extremely low, often below 1% [1]. This fact is illustrated in Figure 1 which plots the average utilization of an active 10 Gbps interface in use within the GLIF community; similar figures are available for several such interfaces [1]. As we can see, whereas the peak transmission rate often matches the interface rate of 10 Gbps, the average rate is significantly lower, less than 1 Gbps for the egress and just above 100 Mbps for the ingress. These average levels roughly correspond to utilization of 73 and 7.3 hours per month, respectively. Due to the substantial expenses involved in leasing and operating long-haul lightpaths, the only viable solution is to increase utilization by sharing these resources among multiple applications and users.

To bridge the gap between the current practice of setting up expensive, dedicated, lightpath connections (i.e., static topologies), and the distant future vision of inexpensive access

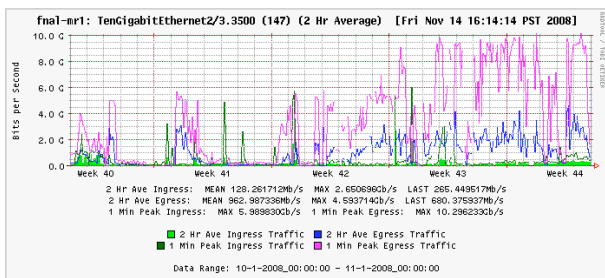


Fig. 1. Average and peak rates of a 10 Gbps interface

to dynamically switched end-to-end lightpaths we propose a medium term solution in the form of *edge-reconfigurable optical networks (ERONs)*. An ERON is an overlay-control network created by installing readily available MEMS optical switches and implementing a GMPLS control plane at sites interconnected by lightpaths. The switches and control software are deployed at the edge of the network and operated by the organization-user (i.e., outside the network provider’s control), hence the term “ERON.” By providing dynamic, automated control of end-to-end lightpaths, ERONs enable the sharing of expensive network resources among multiple users and applications that require occasional access to these resources. ERONs represent a practical and cost-effective solution that transforms the set of static lightpath connections owned or leased by a single organization into a flexible network topology that affords users the capability to reserve on demand, or in advance, lightpaths for any required duration.

The remainder of this paper is organized as follows. In Section II we explain how ERONs support dynamic lightpaths, and in Section III we develop an algorithm for creating an ERON from an existing topology of static lightpaths. We present simulation results to quantify the benefits of ERONs in Section IV, and we conclude the paper in Section V.

II. EDGE-RECONFIGURABLE OPTICAL NETS (ERONs)

An ERON, depicted in Figure 2, consists of:

- A collection of permanent lightpaths that connect users at geographically dispersed sites. The lightpaths are typically leased, provide static connections among the various sites, and define a logical interconnection topology that does not change over time.
- Additional equipment at each site (i.e., at the edge of the network), including a MEMS optical switch, that operates outside the network provider’s control.
- Control software, including GMPLS signaling protocols and a resource broker, that implements a control overlay among the various sites. The control software manages the MEMS switches to enable dynamic connections among users at diverse sites, allowing the sharing of the permanent lightpaths among multiple applications.

The ERON architecture shown in Figure 2 is inspired by the testbed of the EnLIGHTened Computing project [6], a nation-wide, GMPLS-enabled testbed demonstrating the viability of the proposed approach. The underlying philosophy




Fig. 2. A single-domain edge reconfigurable optical network (ERON)

is to enhance the simple dedicated transport service available from today’s optical network infrastructure by implementing dynamic connection functionality and associated intelligence at the edge devices. In essence, the addition of the reconfigurable edge optical switches transforms a sparse static optical network into a highly connected dynamic network.

An ERON relies on hardware and software to create and manage dynamic connections over the fixed logical topology. Each site is equipped with a fiber-based MEMS optical switch, such as the DiamondWave^(R) PXC product from Calient. Unlike an OEO switching fabric, a MEMS optical switching fabric provides optical transparency and scalability for future network growth. Data generating devices (including compute and storage servers, cameras, telescopes, etc.) attach to the optical switch directly, or through IP routers and/or Ethernet switches, via short reach optical interfaces. The MEMS device is capable of switching any input optical signal to any of the WDM output lightpaths, and vice versa on the other end, thus creating on-demand connections among arbitrary sets of users at the various sites. Note that it is possible to set up multi-hop paths even between sites that are *not* connected directly by permanent lightpaths. Returning to Figure 2, users at node A can be connected to users at node C by configuring the switch at node B to concatenate the lightpaths A-B and B-C to form a two-hop connection over which traffic from A to C may flow. Hence, it may not be necessary to lease expensive permanent connections between all pairs of users: by adding capacity sharing and reconfiguration capabilities, a smaller network may function as a network of much larger capacity.

The establishment and termination of the dynamic connections is performed by appropriate control software, e.g., GMPLS, at sub-second timescales. The ERON is also equipped with a resource broker, specialized software that coordinates access to the shared lightpath resources. Users contact the broker to establish or reserve lightpaths. The broker maintains an up-to-date timetable of existing reservations, and may accept, negotiate, or reject reservation requests based on resource availability, user priority, and other policy specifications.

We note that the ability to satisfy user requests is affected directly by both the amount of network resources (lightpaths) available to them and the properties of the logical topology implemented by these lightpaths. Therefore, the issue of logical topology design is integral to the design and implementation of ERONs, and is the subject of the next section.

III. ERON TOPOLOGY DESIGN

The problem of designing logical topologies of lightpaths given some information about the traffic demands among nodes of an optical network has been studied extensively in the literature; the interested reader is referred to [2], [3] and references thereof. The ERON topology design problem we consider in this paper differs from previous studies in several respects. First, ERONs are overlay networks, whereas earlier work dealt with the design of core networks. Second, most algorithms were developed for constructing a logical topology from scratch, whereas our approach is to build an ERON starting from an existing static topology. More importantly, the objective of most logical topology studies has been to minimize the maximum congestion, which is achieved by balancing the traffic load over the network links. On the other hand, our objective is to minimize blocking probability, which requires that the load be concentrated on a small number of high-capacity links. The more recent work in [8] is similar to ours in that it compares static lightpath topologies to dynamic optical networks. Although some of the conclusions of [8] are consistent with ours, our network and traffic models and design approach are more representative of realistic scenarios.

A. Current Practice: Static Topology of Lightpaths

Consider an organization with M users (i.e., high-end devices) distributed across $N < M$ geographic locations. In current practice, if two users at different locations wish to communicate, a permanent *end-to-end* (e2e) lightpath must be established between them. In the absence of switching capabilities, the e2e lightpath is *dedicated* to communication between these users only, and cannot be accessed by other users at the same locations. We assume that e2e lightpaths are unidirectional, hence, two e2e lightpaths, one in each direction, must be set up if bidirectional communication is required.

We make a distinction between e2e and *intra-domain* lightpaths. Since users of a national or global organization may reside at locations in different administrative domains, an e2e lightpath between two users may consist of a string of intra-domain lightpaths, each originating and terminating at the boundaries of a single administrative domain. We assume that the organization leases intra-domain lightpaths separately from each network provider, and connects (patches) them together at its own premises (i.e., at the edge, outside the operator's network) to form longer e2e lightpaths whenever necessary.

In the absence of reconfigurable elements inside the network, providing full connectivity among all M users would require $M(M - 1)$ permanent e2e lightpaths, each dedicated to a directed pair of users. Since leasing such a number of e2e lightpaths would be impractical and prohibitively expensive for

other than a small number M of users, the organization might lease a smaller number K of e2e lightpaths between a select set of user pairs; typically, $K \ll M(M - 1)$. The K pairs of users to be connected directly may be selected based on the amount or criticality of information exchanged, the priority of users, or some other relevant criteria. Let t_k denote the long-term traffic demand for the k -th user pair, expressed in Gbps. Let C denote the lightpath capacity, also expressed in Gbps. Then, the k -th pair is assigned $\lceil t_k/C \rceil$ e2e lightpaths.

Note that an e2e lightpath may consist of multiple intra-domain lightpaths, and all e2e lightpaths routed over a given intra-domain hop between some sites A and B require a dedicated intra-domain lightpath from A to B . Let L be the number of distinct intra-domain hops required by the e2e lightpaths and c_l be the number of intra-domain lightpaths set up over the l -th hop. These L hops are the links of the *static topology* defined among the K pairs of users. The capacities of the L links are such that $\sum_{l=1}^L c_l = \sum_{k=1}^K h_k t_k / C$, where h_k is the number of intra-domain lightpaths making up the e2e lightpath between the k -th user pair.

B. ERON Topology Algorithm

Dedicated lightpaths are expensive resources that may be economically justified only if their utilization remains at high levels throughout the lease period. Recent statistics (refer to Figure 1) indicate that the utilization of GLIF lightpaths is quite low, often as low as a few hours per month [1]. Alternatively, an organization may deploy an ERON to increase the utilization of its leased lightpaths. This increase is due to two factors: (1) higher connectivity, as all users at a given site have access to all incoming and outgoing lightpaths, and (2) sharing of the (previously dedicated) lightpaths among this larger number of user pairs.

An important question that arises is whether deploying an ERON would result in significant savings in the number of lightpaths to be leased to justify the (one-time) hardware and software expense. Note that, by allowing pairs of users to share a lightpath, an ERON introduces the possibility of blocking. To ensure that users receive a good quality of service, we impose an upper bound of 10^{-3} in *network-wide* blocking probability.

In order to quantify the benefits of ERON in terms of the number of lightpaths required to provide connectivity among the K pairs of users, we consider the following problem:

Problem 3.1: Given a static topology of L intra-domain links connecting K pairs of users, the long-term traffic demands $t_k, k = 1, \dots, K$, and the capacity C of a lightpath, design an ERON topology with the least number of lightpaths such that the network-wide blocking probability $BP \leq 10^{-3}$.

Logical topology design problems are typically NP-hard [2]. An additional challenge in this case is the difficulty of expressing the blocking probability in exact and closed-form manner. We also emphasize that the goal of this study is not to identify optimal solutions but rather to quantify the benefits that are practically achievable using the ERON model. Hence, to solve the above problem we use the greedy algorithm presented in Figure 3. The algorithm starts with the static topology

Algorithm for Removing Capacity from a Static Topology

Input: Static topology of L intra-domain links

Output: An ERON topology with fewer lightpaths

begin

1. $c_l \leftarrow$ capacity of link $l, l = 1, \dots, L$, in lightpaths
 2. Run simulation, obtain link utilizations $\{u_l\}$ and BP b
 3. **while** $b < 10^{-3}$ **do**
 4. **for** $l = 1$ **to** L **do** // Consider each link in isolation
 5. $\rho_l \leftarrow \text{Erl}^{-1}(10^{-4}, c_l)$ // Target link offered load
 6. $u'_l \leftarrow \frac{\rho_l}{c_l}$ // Corresponding target link utilization
 7. $\text{slack}_l \leftarrow \frac{u_l}{u'_l}$
 8. **endfor**
 9. $j \leftarrow$ link such that $\text{slack}_j = \min\{\text{slack}_l\}$
 10. $c_j \leftarrow c_j - 1$ // Reduce capacity of link j
 11. Run simulation, obtain link utilizations $\{u_l\}$ and BP b
 12. **endwhile**
 13. **return** the last topology for which $b < 10^{-3}$
- end**
-

Fig. 3. Algorithm for constructing an ERON topology by removing capacity from the links of the initial static topology

and selectively removes lightpaths as long as the blocking probability remains below 10^{-3} .

The algorithm first runs a simulation of the ERON network with a topology identical to the one of the static topology (Step 2 in Figure 3). From the simulation, we obtain the overall blocking probability in the ERON, as well as the utilization u_i of each link i in the network. The main operation of the algorithm consists of the **while** loop between Steps 3 and 12, which is executed as long as the blocking probability is below the threshold of 10^{-3} . Considering each link i of the network in isolation, Step 5 determines the target offered load ρ_i to the link so that the blocking probability on this link, computed using the Erlang-B formula with a number of servers equal to the capacity c_i of the link, be equal to 10^{-4} ; we have found that this target link probability is sufficient to guarantee the network-wide threshold for networks with a small diameter. The corresponding target link utilization u'_i is obtained by dividing the target offered load into the capacity c_i of the link (Step 6). The *slack* of a link is then defined as the ratio of its actual utilization u_i (obtained through simulation) and the target utilization u'_i , as shown in Step 7. Note that a low slack value implies excess capacity at a link relative to the actual (simulation) offered load. Therefore, in Steps 9 and 10, the algorithm reduces the capacity of the link with the minimum slack by one lightpath. The algorithm then runs the simulation again, and repeats the **while** loop to reduce capacity further, if the blocking probability remains below 10^{-3} .

IV. PERFORMANCE EVALUATION

A. Network Model

We consider the three-domain ERON network shown in Figure 4. The network consists of $N = 20$ ERON switches (nodes), each representing a site of a global organization.

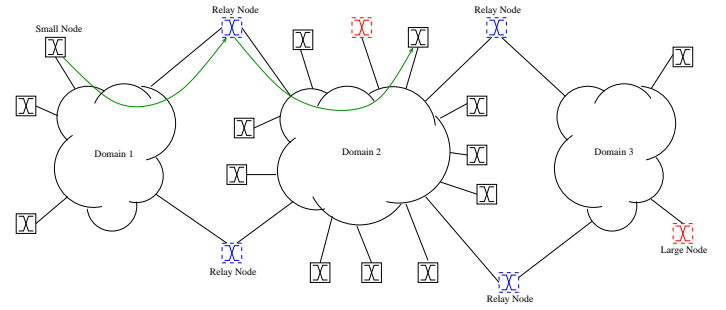


Fig. 4. Three-domain ERON network

Attached to the ERON nodes are M users that generate the traffic demands; note that the users are *not* shown in Figure 4. The users are distributed across the various ERON nodes based on the size of each node, as we describe shortly.

The ERON nodes belong to one of three domains based on their geographic locations. For instance, the three domains in Figure 4 might correspond roughly to the providers serving an organization's sites in Asia, US, and Europe, respectively. The middle domain is the largest and contains about 70% of the ERON nodes. The side domains are smaller and each contains about 15% of the nodes. Lightpaths terminate at domain boundaries. Thus, an e2e lightpath that connects a node in Domain 1 to a node in Domain 3 is made up of three intra-domain lightpaths. Hence, for the topology in Figure 4, an e2e lightpath may consist of one, two, or three intra-domain lightpaths, depending on the location of its endpoints.

ERON nodes are classified as either *small* or *large*. Large nodes, denoted by dotted lines in Figure 4, have more users connected to them than small nodes. Consequently, the total aggregate traffic demand from/to large nodes is higher than from/to small nodes. We assume that 6 of the ERON nodes are large and 14 are small. The three domains are interconnected by four of the large ERON nodes, referred to as the *relay* nodes. Just as any other large node, relay nodes have users connected to them that generate traffic. Unlike other large nodes, however, each relay node is attached to two adjacent domains, it is equipped with an OEO capability, and it terminates all intra-domain lightpaths that are part of an e2e lightpath passing through it. This capability is necessary so that the organization may concatenate lightpaths provided by independent service providers to form longer e2e lightpaths as needed, without intervention from the service providers.

B. Traffic Model

We assume that a lightpath is unidirectional and its capacity is $C = 10$ Gbps. Connections that generate traffic higher than 10 Gbps are assigned multiple lightpaths. Let D_{tot} denote the aggregate network traffic, in Gbps, generated by users attached to ERON nodes. We denote the aggregate traffic generated by large-to-large, large-to-small, and small-to-small connections as D_{ll} , D_{ls} and D_{ss} , respectively; clearly, $D_{tot} = D_{ll} + D_{ls} + D_{ss}$. Let K be the number of (unidirectional) connections that make up the static topology from which

the ERON topology is derived. We assign each of the $2K$ endpoints (users) of these connections probabilistically, such that users are twice as likely to be assigned to large node than to small node. The total amount of traffic D_{tot} is then divided into components D_{ll} , D_{ls} and D_{ss} in proportion to the number of connections (out of K) that fall into the large-to-large, large-to-small, and small-to-small classes, respectively. We also consider three traffic patterns: *uniform*, *distance-decreasing*, and *distance-increasing*. The individual traffic components are then distributed over the connections within each class according to one of these traffic patterns. The result is a list of K user connections and their specific (average) data rates.

We consider two scenarios. In the *low* traffic scenario, the average aggregate traffic D_{tot} ¹ over all connections does not exceed 300 Gbps, whereas in the *high* traffic scenario, the average aggregate traffic varies between 1 Tbps and 3 Tbps.

C. Simulation Design and Methodology

To carry out the evaluation study, we implemented our own event-driven simulator in C. Connection requests were generated according to a Poisson arrival process. Each traffic demand has a holding time (duration) and a traffic amount. Each traffic demand amount was in units of lightpaths, ranging from a low of 1 lightpath to a high of 4 lightpaths. Traffic amounts were generated according to a truncated power-law distribution, using an exponent of 1.5. The holding times of demands normally ranged from 10 seconds to one hour. These demands were probabilistically generated according to a truncated Pareto distribution, again with an exponent of 1.5.

We generated and simulated the routing of traffic that conforms to the average rates for each pair of users dictated by the traffic model above. This was accomplished by setting the value of the rate λ of the Poisson arrival process for each pair of users to a value such that the aggregated average matches the corresponding element in the traffic matrix.

Each simulation run lasted until a set of 500,000 traffic demands was generated and routed. For every data point shown on the graphs, 10 runs were executed, each with a different random seed. From these, the mean and the 95% confidence intervals were computed (shown on the graphs as error bars).

The routing algorithm simulated is a standard method of achieving good overall network utilization and low blocking. First, from among the available paths to the destination, only paths that had sufficient available (i.e., residual, not currently used) capacity were considered. From among such paths, the shortest (in terms of hop count) path was chosen. In the event of a tie, one path from among the shortest, highest available capacity ones was selected uniformly randomly.

D. Numerical Results

We present a set of experiments to quantify the benefits of ERONs compared to establishing a static topology of dedicated

¹We distinguish between *average* and *peak* traffic on a connection. The peak traffic may fully utilizing the capacity of a lightpath, whereas the average traffic can be quite low, especially if the specific connection is used sporadically, e.g., for a few hours per month. The ERON topology algorithm we described earlier is based on average traffic demands, not peak ones.




Fig. 5. Low traffic: effect of average aggregate traffic amount D_{tot} , $K = 100$ ($D_{tot} = 0.01$ Tbps equivalent to peak transmission an average of 7.3 hrs/mo)

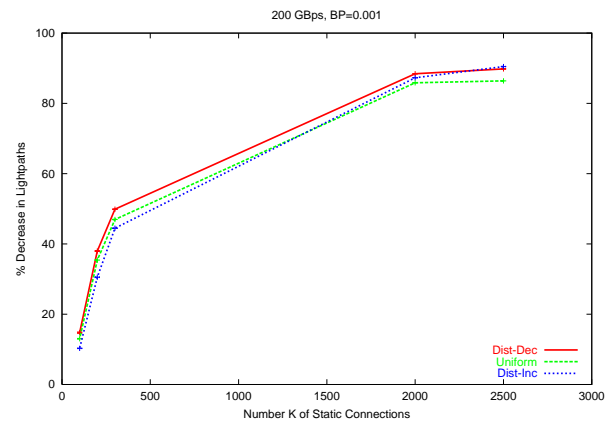


Fig. 6. Low traffic: effect of number K of connections, $D_{tot} = 200$ Gbps ($K = 2000$ equivalent to peak transmission an average of 7.3 hrs/mo)

connections among a given set of users. Figure 5 shows the achievable decrease in number of intra-domain lightpaths for the case of $K = 100$ connections. Each data point in the figure is an average over 10 problem instances generated according to one of the three traffic patterns we described earlier; 95% confidence intervals are also plotted. In this case, the static topology always consists of $K = 100$ e2e lightpaths (each made up of at least one intra-domain lightpath) between a set of users; this set is randomly generated for each problem instance and each user is probabilistically assigned to one of the ERON nodes, as we explained above. For a given instance, the ERON topology is constructed using the algorithm in Figure 3 to remove capacity from the static topology.

The figure plots the decrease in lightpaths against the average aggregate traffic D_{tot} generated by the K connections. To put the aggregate traffic values in perspective, note that the $K = 100$ static connections can generate a peak traffic equal to 1 Tbps (since the capacity of each e2e lightpath is 10 Gbps). Consequently, an average aggregate traffic $D_{tot} = 0.01$ Tbps corresponds to each connection transmitting at its peak rate for an average of 7.3 hours per month (and being idle the remaining of the time). At low utilization, the savings can be




Fig. 7. High traffic scenario: effect of average aggregate traffic amount, D_{tot} , $K = 1000$ ($D_{tot} = 1$ Tbps equivalent to peak transmission an average of 73 hrs/mo)




Fig. 8. High traffic scenario: effect of number K of connections, $D_{tot} = 2$ Tbps ($K = 2000$ equivalent to peak transmission an average of 73 hrs/mo)

substantial, more than 40% compared to the static topology. As utilization increases, the savings decrease accordingly, but even when $D_{tot} = 0.1$ Tbps (equivalent to peak transmission for an average of 73 hours per month or more than 3.5 hours per work day), the savings can exceed 20%. Also, the traffic pattern does affect the results, but not substantially.

For the results shown in Figure 6, we let the average aggregate traffic $D_{tot} = 200$ Gbps, and vary the number K of connections in the static topology. In this case, $K = 2000$ corresponds to each connection transmitting at peak rate for an average of 7.3 hours per month. Again, we observe that at utilization levels of 7.3 hours per month or lower, the achievable savings in lightpaths are quite significant, between 80-90% compared to the static topology. Even at higher utilization levels, e.g., $K = 200$ (or 73 hours per month), the savings are about 30% or more.

Figures 7 and 8 are similar to Figures 5 and 6, respectively, but present results for the high traffic scenario. In Figure 7 we let $K = 1000$ connections and vary the average aggregate amount D_{tot} of traffic from 1-3 Tbps. Note that $D_{tot} = 1$ Tbps corresponds to each connection transmitting at peak rate of

10 Gbps for an average of 73 hours per month, and remaining idle the rest of the time. At this level of utilization, the average savings in terms of lightpaths is approximately between 55-65%, depending on the traffic pattern, decreasing to between 35-45% as utilization increases to 219 hours per month (for $D_{tot} = 3$ Tbps). In Figure 8 we let $D_{tot} = 2$ Tbps, and vary the number K of connections from 1000 (corresponding to utilization of 146 hours per month) to 3000 (utilization of about 49 hours per month). Again, as utilization decreases, the average lightpath savings increase substantially, reaching 75% or more over the static topology for $K = 3000$.

By comparing Figures 7 and 8 (high traffic scenario) to Figures 5 and 6, respectively (low traffic scenario), we observe that under the high traffic scenario it is possible to achieve substantial more savings at the same level of utilization (equivalently, achieve the same amount of savings at higher utilization). Consider, for instance, the savings for $D_{tot} = 0.1$ Tbps in Figure 5 and for $D_{tot} = 1$ Tbps in Figure 7 and for $D_{tot} = 1$ Tbps; both correspond to a utilization level of 73 hours per month. The savings in Figure 5 are between 20-25%, whereas in Figure 7 are significantly higher, between 45-50%, depending on the traffic pattern. This result is due to the fact that in both cases the capacity of a lightpath is the same (i.e., 10 Gbps), but under the high traffic scenario, the links of the ERON topology have much higher capacity (number of lightpaths) than under the low traffic scenario. As a result, the blocking probability of the larger network can be substantially lower under the same offered load.

V. CONCLUDING REMARKS

We have proposed ERONs, overlay dynamic optical networks based on a collection of static lightpaths and MEMS optical switches with GMPLS control. An ERON enables the sharing of expensive lightpaths by providing reconfigurability at the edges of the optical network. We consider ERONs as a medium-term solution for enabling today's emerging high-end applications in the areas of science, defense, and enterprise. Simulation results demonstrate that, depending on average utilization, migration to an ERON network from a static topology may constitute a significant savings in the number of static lightpaths required to meet the needs of the traffic demands at an acceptable blocking probability.

REFERENCES

- [1] C. Guok, Sr. Network Engr., ESnet. Private correspondence, Dec. 2008.
- [2] R. Dutta and G. Rouskas. A survey of virtual topology design algorithms for wavelength routed optical networks. *Opt. Netw.*, 1(1):73-89, Jan. 2000.
- [3] E. Leonardi, et al. Algorithms for the logical topology design in WDM all-optical networks. *Optical Networks*, 1(1):35-46, January 2000.
- [4] A. A. Saleh and J. M. Simmons. Evolution toward the next-generation core optical network. *IEEE/OSA JLT*, 24(9):3303-3321, Sep 2006.
- [5] DARPA STO. The CORONET Project. <http://www.darpa.mil/STO/Solicitations/CORONET/index.htm>.
- [6] L. Battestilli et al. EnLIGHTened computing: An architecture for co-allocating network, compute, and other grid resources for high-end applications. In *Proceedings of IEEE Honet'07*, Dubai, UAE, Nov. 2007.
- [7] F. Travostino, J. Mambretti, G. Karmous-Edwards (Eds). *Grid Networks: Enabling Grids with Advanced Communication Technology*. Wiley, 2006.
- [8] A. Zapata and P. Bayvel. Dynamic vs. static wavelength-routed optical networks. *IEEE/OSA JLT*. (To appear).