Robust Network Design for Stochastic Traffic Demands

Matthew Johnston, *Student Member, IEEE*, Hyang-Won Lee, *Associate Member, IEEE*, and Eytan Modiano, *Fellow, IEEE*

***Abstract*—This paper addresses the problem of logical topology design for optical backbone networks subject to stochastic traffic demands. The network design problem is broken into three tasks: traffic routing, capacity allocation, and link placement. While the routing and capacity allocation subproblem can be formulated using convex optimization, it is prohibitive to add the link placement component to the nonlinear formulation since the link placement problem involves integer variables. To address this issue, we develop a linear formulation for the routing and capacity allocation subproblem by applying tools from robust optimization. We show that this linear formulation performs comparably to the optimal nonlinear formulation. Our formulation can then be used to solve the link-placement subproblem for stochastic traffic. We show that optimal logical topologies for deterministic traffic demands are not necessarily optimal for stochastic traffic demands. We develop algorithms for finding logical topologies optimized for stochastic traffic.**

***Index Terms*—Demand uncertainty, network design, robust optimization, routing and capacity assignment.**

I. Introduction

INTERNET service Pproviders (ISPs) have made optical wavelength division multiplexing (WDM) technology increasingly available to the network edge in order to meet the growing demand for high data rates. For instance, in fiber-to-the-home (FTTH), fiber connection from an ISP's central office reaches a home network, offering a very high data rate for the end users. However, the policy of merely boosting network capacity leads to a huge rise in operating and capital cost due to the inefficiencies of the current network architecture, where the Internet protocol (IP) traffic is transported through multiple layers of network elements.

Manuscript received March 6, 2013; revised August 5, 2013; accepted August 6, 2013. Date of publication August 7, 2013; date of current version September 10, 2013. This work was supported by NSF grants CNS-0626781, CNS-0830961, and CNS-1017800 and by DTRA grants HDTRA1-07-1-0004 and HDTRA-09-1-0050. The work of H.-W. Lee was supported by Basic Science Research Program through the National Research Foundation of Korea (NRF) funded by the Ministry of Science, ICT & Future Planning (2012R1A1A1012610). This work was presented in part at the IEEE GLOBECOM Conference, December 2011 [1]. *(Corresponding author: H.-W. Lee.)*

M. Johnston is with the Department of Electrical Engineering and Computer Science, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: mrj@mit.edu).

H.-W. Lee is with the Department of Internet and Multimedia Engineering, Konkuk University, Seoul 143–701, Korea (e-mail: leehw@konkuk.ac.kr).

E. Modiano is with the Laboratory for Information and Decision Systems, Massachusetts Institute of Technology, Cambridge, MA 02139 USA (e-mail: modiano@mit.edu).

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.

Digital Object Identifier 10.1109/JLT.2013.2277863

IP-over-WDM is considered one of the most promising architectures for next-generation networks that can address the aforementioned issues. In contrast to the conventional multilayer architecture, in IP-over-WDM, the IP layer is built directly on top of the WDM layer. This simple architecture offers greater flexibility, scalability, power efficiency as well as significant operating and capital cost savings [2]. An IP-over-WDM network consists of a logical (IP-layer) topology built on top of a physical (WDM-layer) topology. The physical topology is made up of a set of optical cross-connects (OXCs) and a set of fibers connecting the OXCs. The logical topology is made up of the interconnection of IP routers using optical lightpaths. Each logical link connecting a pair of IP routers corresponds to a lightpath possibly spanning multiple physical fibers. In this paper, we focus on the problem of designing a robust logical topology for IP-over-WDM networks under stochastic traffic demands.

The objective of logical topology design over an existing physical topology is to place logical links on the physical topology, assign capacity to these links, and route the IP traffic over the logical topology. The physical topology is known *a priori*, as is a matrix describing the long-term average traffic flow between any node-pair in the logical network. Capacity is assigned to the logical links such that the traffic routed on each link will not exceed its capacity. This problem has been studied extensively in the literature, with various design criteria such as minimizing the maximum congestion, total capacity, or the number of hops to route the demands [3]–[7]. All these works assume that the traffic demand between every node pair is known *a priori*.

However, the design problem is typically solved during the network configuration stage, where the exact traffic information is unavailable. For this reason, traffic demands must be modeled as random quantities. Unlike the deterministic case, network links must be provisioned with sufficient capacity to withstand possible demand fluctuations. For models with bounded traffic variability, the most common approach is to allocate sufficient link capacity to support the worst-case set of demands so that any demand fluctuations can be handled by the network [8], [9].

On the other hand, a demand is typically an aggregation of many independent traffic sources, and by the central limit theorem, it can be approximated using a Gaussian random variable. In this case, allocating capacity to support the worst case set of demands yields an unacceptable overprovisioning of capacity, since Gaussian demands have unbounded support. Therefore, a different capacity provisioning approach is required for unbounded demands. The work of Mitra and Wang [10] addresses this issue by formulating an optimization problem that incorporates capacity underutilization into the objective as a penalty, to detract from overprovisioning.

In this paper, we take a different approach by explicitly taking into account the probability of traffic exceeding the provisioned capacity on each logical link. A similar approach was also considered in [11], where the overflow probability is explicitly derived for Gaussian distributed traffic and used in the optimization problem for routing and capacity allocation. While this formulation leads to a convex optimization problem that is efficiently solvable, its nonlinearity prevents extensions to discrete optimization problems such as logical link placement. For this reason, Meesublak [12] formulates a linear program (LP) by changing probabilistic constraints to linear deterministic ones. This approach suffers from an overprovisioning of capacity, since it inherently ignores the effect of statistical multiplexing of many demands. There are other approaches dealing with traffic variability [13], [14]; however, they ignore the distribution of the random demands, which also results in over-provisioning of capacity.

In this paper, we formulate and analyze the problem of logical topology design under Gaussian-distributed demands. We develop a linear formulation for routing traffic and allocating capacity to support random demands by extending the robust optimization techniques of [15] to Gaussian random variables. Our formulation exploits the effect of traffic aggregation to reduce the amount of allocated link capacity required, while at the same time yielding the computational advantages of a linear approach. This study is an extension of our previous work [1], in which we first introduce robust optimization for solving the routing and capacity allocation problems. In this study, we include a new empirical study to highlight the importance of our robust optimization approach, and extend the formulation to consider the complete network design problem, by studying the logical link placement subproblem. Our contributions can be summarized as follows:

1) We extend the robust optimization techniques of [15] to the case of (unbounded) Gaussian random variables.
2) We develop a mixed integer linear program (MILP) formulation for robust routing and capacity allocation under Gaussian traffic demands, based on the extended robust optimization framework.
3) We develop a scalable algorithm that finds approximate solutions to the MILP.
4) We show that logical topologies optimized for deterministic traffic may not be optimal for stochastic traffic.
5) We formulate the joint routing, capacity allocation, and logical topology design problem as an MILP.
6) We develop algorithms for finding optimal logical topologies for stochastic traffic.

The rest of the paper is organized as follows. In Section II, we introduce the network model and motivate the development of a new approach to solving the routing and capacity allocation problems. In Section III, we introduce robust optimization as a technique to develop a new formulation for the network design problem. In Section IV, we discuss how our formulation can be used to solve the logical link placement subproblem in parallel to the capacity allocation and routing problems, and we conclude in Section V.

II. Network Model and Motivation

Let $\mathcal{D}$ be a set of demands over the network, where for $(s,d) \in \mathcal{D}$, λ^{sd} is the traffic demand from node s to node d. Each demand λ^{sd} represents a long-term average traffic flow between two nodes. Average traffic demands between node-pairs are often described using a Gaussian distribution [16]; therefore, we assume that all traffic demands are independent, and each λ^{sd} follows a Gaussian distribution with mean μ_{sd} and variance σ_{sd}^2.

For now, assume that the network is described by a graph $G = (V, E)$ where V is the set of logical nodes and E is the set of logical links connecting the logical nodes. In our formulation, we allow a traffic demand to be split among various paths in the network G. Let $0 \leq a_{ij}^{sd} \leq 1$ represent the fraction of the traffic flow from s to d traversing logical link $(i,j) \in E$. Let C_{ij} represent the capacity allocated to link[1] (i,j). In this paper, we study the problem of finding a routing (determined by $\{a_{ij}^{sd}\}$) and capacity allocation (determined by $\{C_{ij}\}$) that can support the demands in $\mathcal{D}$ while minimizing the maximum link capacity. The chosen objective function has the effect of splitting the traffic over multiple paths, and thus, balancing the load over the network.

Note that the routing and capacity allocation problem can be extended to include the link placement subproblem, by optimizing over the set E. Let b_{ij} be a binary variable which takes on the value 1 if and only if a directed link is placed from node i to node j. While this extension captures the complete network design problem, we will initially focus on the routing and capacity allocation problems. The network design problem will be added in Section IV. For now, we assume the design variable $\mathbf{b} = [b_{ij}, \forall i, j \in V]$ is fixed and thus the network is described by a graph $G = (V, E)$ with a set of edges $E = \{(i,j) | b_{ij} = 1\}$. Define the overflow probability of link (i,j) to be the probability that the traffic on link (i,j) exceeds its capacity, i.e., $\Pr[\sum_{(s,d)} a_{ij}^{sd} \lambda^{sd} > C_{ij}]$. The routing and capacity allocation problem can be formulated as follows:

$$\begin{aligned} \min_{C,a \geq 0} \quad & \max_{(i,j)\in E} C_{ij} \\ \text{subject to:} \quad & \Pr\left[\sum_{sd} a_{ij}^{sd} \lambda^{sd} > C_{ij}\right] \leq \epsilon, \ \forall (i,j) \in E \\ & FC(s,d,a^{sd}), \ \forall (s,d) \in \mathcal{D} \end{aligned} \tag{1}$$

The objective is to minimize the maximum link capacity. This objective results in minimizing the congestion in the network, or balancing the traffic load through the network. We note that the objective function can be linearized by replacing it with a new variable $C_{\max}$ and introducing new constraints $C_{\max} \geq C_{ij} \, \forall (i,j)$, but for brevity, we will use the nonlinear form throughout the paper. The function $FC(s,d,a^{sd})$ represents

[1] In the following, *link* refers to *logical link* unless otherwise specified.

a flow conservation constraint from s to d given as

$$\sum_{j:(i,j)\in E} a_{ij}^{sd} - \sum_{j:(j,i)\in E} a_{ji}^{sd} = \begin{cases} 1, & \text{if } s = i \\ -1, & \text{if } d = i \\ 0, & \text{o.w.} \end{cases} \quad \forall i \in V. \tag{2}$$

This constraint describes the routing of each demand.

The constraint in (1) requires sufficient capacity to be allocated to ensure the overflow probability at each link is no greater than a design parameter $\epsilon > 0$. Since the sum of Gaussian random variables is also Gaussian, the left hand side of this constraint can be rewritten as

$$\begin{aligned} &\mathbf{P}\left(\sum_{sd} a_{ij}^{sd}\lambda^{sd} > C_{ij}\right) \\ &= \mathbf{P}\left(\frac{\sum_{sd} a_{ij}^{sd}\lambda^{sd} - \sum_{sd} a_{ij}^{sd}\mu_{sd}}{\sqrt{\sum_{sd}(a_{ij}^{sd}\sigma_{sd})^2}} > \frac{C_{ij} - \sum_{sd} a_{ij}^{sd}\mu_{sd}}{\sqrt{\sum_{sd}(a_{ij}^{sd}\sigma_{sd})^2}}\right) \\ &= 1 - \Phi\left(\frac{C_{ij} - \sum_{sd} a_{ij}^{sd}\mu_{sd}}{\sqrt{\sum_{sd}(a_{ij}^{sd}\sigma_{sd})^2}}\right). \end{aligned}$$

Consequently, the routing and capacity allocation problem can be reformulated as

$$\begin{aligned} \min_{C,a\geq 0} \quad & \max_{(i,j)\in E} C_{ij} \\ \text{s.t.:} \quad & C_{ij} \geq \sum_{sd} a_{ij}^{sd}\mu_{sd} + \Phi^{-1}(1-\epsilon)\sqrt{\sum_{sd}(a_{ij}^{sd}\sigma_{sd})^2} \\ & FC(s,d,a^{sd}),\ \forall (s,d) \in \mathcal{D}, \end{aligned} \tag{3}$$

where $\Phi(\cdot)$ is the CDF of standard normal random variable. While this optimization problem is solvable, the nonlinearity of (3) prevents the addition of integer constraints, as required by the network design problem, since the resulting formulation will be a nonlinear integer program and effective solution methodologies are yet unknown for general, nonlinear, integer programs. Additionally, solving this convex optimization problem may be impractical for large networks.

To alleviate this problem, one can use the fact that the right hand side of the constraint in (3) is upper-bounded by

$$\sum_{sd} a_{ij}^{sd}\mu_{sd} + \Phi^{-1}(1-\epsilon)\sqrt{\left(\sum_{sd} a_{ij}^{sd}\sigma_{sd}\right)^2}.$$

In particular, consider the following linear constraint:

$$C_{ij} \geq \sum_{sd} a_{ij}^{sd}(\mu_{sd} + \Phi^{-1}(1-\epsilon)\sigma_{sd}). \tag{4}$$

This constraint, which allocates dedicated capacity for each flow sufficient to satisfy (1) without statistical multiplexing gains, is wasteful in terms of the maximum capacity on any link in the network. To understand the difference between constraints (3) and (4), we study the following example. Consider the network shown in Fig. 1. Traffic demands exist from s_1 to d_1 and from s_2 to d_2, and are independently and identically

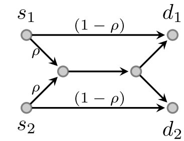

Fig. 1. Example network, with demands $\lambda^{s_1 d_1}, \lambda^{s_2 d_2} \sim \mathcal{N}(\mu,\sigma)$.

distributed according to $\mathcal{N}(\mu,\sigma^2)$. Due to symmetry, both demands have the same fraction of traffic (ρ) routed on the shared link. Clearly, the optimal solution allocates equal capacity on all links. Using the linear capacity constraint in (4), we can analytically show that the optimal routing is $\rho = \frac{1}{3}$. Additionally, this routing is optimal for any deterministic traffic. Alternatively, for the nonlinear constraint in (3), it can be shown that the optimal routing will send more traffic through the middle link (i.e., larger value of ρ), to take advantage of statistical multiplexing on that link. In particular, by letting $\kappa = \Phi^{-1}(1-\epsilon)$, it can be shown that

$$\rho = \frac{\mu + \kappa\sigma}{3\mu + (1+\sqrt{2})\kappa\sigma}. \tag{5}$$

Note that the nonlinear constraint in (3) exploits the statistical multiplexing gain from combining multiple traffic flows so that the resulting standard deviation is smaller than the sum of individual standard deviations. Consequently, the required capacity is $O(\sqrt{n})$ less than that with the linear constraint, which simply adds up individual variances, resulting in an overprovisioning of capacity, where n is the number of flows sharing a link. In the next section, we use robust optimization to develop a linear formulation exploiting the statistical multiplexing gain so that it can be applied to the network design problem in Section IV without significant capacity overprovisioning.

III. Robust Optimization Approach

Robust optimization is a method for solving optimization problems with parameter uncertainty. A novel robust optimization approach was presented in [15] based on the idea that it is unlikely for all of the random parameters to simultaneously take values far above their means. Therefore, resources are allocated to protect against scenarios where only a fraction of the parameters exceed their mean, and the probability of constraint violation can still be kept small. The technique in [15], however, applies only to bounded uncertainty, so cannot be directly applied to our model. In this section, we extend the robust optimization approach of [15] to Gaussian random variables by considering a truncated Gaussian random variable with a large truncation point.

A. *Robust Formulation for Truncated Gaussian Demands*

Initially, we assume that demands are drawn from a truncated Gaussian distribution, i.e., the traffic λ_B^{sd} between demand pair (s,d) satisfies $\lambda_B^{sd} \sim \mathcal{N}(z|z \leq \mu_{sd} + k\sigma_{sd}; \mu_{sd}, \sigma_{sd}^2)$ for some constant k, and has a PDF

$$f_{\lambda_B^{sd}}(z;\mu_{sd},\sigma_{sd}^2) = \frac{\frac{1}{\sigma_{sd}}\phi\left(\frac{z-\mu_{sd}}{\sigma_{sd}}\right)}{\Phi(k)}, \quad z < \mu_{sd} + k\sigma_{sd} \tag{6}$$

where $\phi(\cdot)$ and $\Phi(\cdot)$ are the PDF and CDF of a standard normal random variable, respectively. We now apply the approach of [15].

For each link (i,j), a parameter Γ_{ij} is introduced taking a value between 0 and n_{ij}, where n_{ij} is the number of demands traversing link (i,j). Capacity is assigned to protect against any scenario where Γ_{ij} of the demands exceed their means. The other demands are allocated a capacity equal to their mean. Using this idea, we formulate a new capacity constraint to replace those in (3) and (4). Let $S_{ij} \subseteq \mathcal{D}$ be a subset of demands of size $\lfloor \Gamma_{ij} \rfloor$, and let $t_{ij} \in \mathcal{D} \setminus S_{ij}$ be another demand. On each link (i,j), capacity is allocated to support the mean $a_{ij}^{sd}\mu_{sd}$ for each demand, as well as the $\lfloor \Gamma_{ij} \rfloor$ largest values of $a_{ij}^{sd}k\sigma_{sd}$ to guarantee support for the $\lfloor \Gamma_{ij} \rfloor$ demands exceeding their mean. In the case where Γ_{ij} is not an integer, the fraction $(\Gamma_{ij} - \lfloor \Gamma_{ij} \rfloor)$ of $a_{ij}^{t_{ij}}k\sigma_{t_{ij}}$ is also supported. The constraint is presented below:

$$C_{ij} \geq \sum_{sd\in\mathcal{D}} \mu_{sd}a_{ij}^{sd} + \max_{S_{ij}\cup\{t_{ij}\}:S_{ij}\subseteq\mathcal{D},|S_{ij}|=\lfloor\Gamma_{ij}\rfloor,t_{ij}\in\mathcal{D}\setminus S_{ij}} \left\{ \sum_{(s,d)\in S_{ij}} k\sigma_{sd}a_{ij}^{sd} + (\Gamma_{ij} - \lfloor\Gamma_{ij}\rfloor)k\sigma_{t_{ij}}a_{ij}^{t_{ij}} \right\}. \tag{7}$$

The first term in the right hand side is the mean of the demands traversing link (i,j), and the second term is the sum of Γ_{ij} largest possible fluctuations from the mean for those demands. This formulation exploits the statistical multiplexing effect discussed above by assuming that the demands are unlikely to simultaneously require all of the provisioned link capacity. Note that the solution is feasible for any realization of random demands such that no more than Γ_{ij} demands exceed their mean. Hence, the robustness of the formulation is parameterized by the value of Γ_{ij}. For example, when $\Gamma_{ij} = n_{ij}$, the solution yields the most conservative allocation, and all the demands are supported with probability 1. On the other hand, for small Γ_{ij}, a small additional capacity is allocated at the expense of higher overflow probability.

The link overflow probability can be bounded as a function of Γ_{ij}, and the value of Γ_{ij} that satisfies the probabilistic constraint in (1) can be characterized. In [15], an upper bound on that probability is derived for general symmetric, bounded random variables. We modify this bound for truncated Gaussians to relate ϵ to Γ_{ij} for each link. This bound is shown in the following theorem.

Theorem 1: Let λ_B^{sd} be the traffic from source s to destination d. Further, let λ_B^{sd} be a continuous random variable with density $\mathcal{N}(\lambda_B^{sd} | \lambda_B^{sd} \leq \mu_{sd} + k\sigma_{sd}; \mu_{sd}, \sigma_{sd}^2)$. Let $0 \leq a_{ij}^{sd} \leq 1$ and $C_{ij} \geq 0$ satisfy (7). Let n_{ij} be the number of demands routed over link (i,j). Then, the probability that link (i,j) overflows its allocated capacity is bounded by

$$\Pr\left[\sum_{sd} \lambda_B^{sd}a_{ij}^{sd} > C_{ij}\right] \leq \exp\left(-\frac{\Gamma_{ij}^2 k^2}{2n_{ij}}\right). \tag{8}$$

Proof: See Appendix A. ■

Therefore, given a desired probability ϵ of overflow on link (i,j), the value of the parameter Γ_{ij} required on that link can be computed by

$$\Gamma_{ij} = \frac{\sqrt{-2n_{ij}\log\epsilon}}{k}. \tag{9}$$

Hence, with the value of Γ_{ij} in (9), the capacity constraint (7) satisfies the overflow probability requirement in (1). Note that capacity is assigned to a link as a linear function of Γ_{ij}, and grows with $\sqrt{n_{ij}}$. This is similar to the optimal convex capacity constraint in (3). Therefore, our formulation results in routings in which traffic is multiplexed to reduce capacity.

B. Robust Formulation for Gaussian Demands

So far, we have assumed that demands obey a truncated Gaussian distribution, and thus, the constraint (7) is valid only for truncated Gaussian demands. In this section, we extend the robust optimization formulation and associated probability bounds on constraint violation developed in the previous section to unbounded Gaussian demands.

Recall that the demands λ_B^{sd} are upper bounded by $\mu_{sd} + k\sigma_{sd}$. In the capacity constraint in (7), Γ_{ij} corresponds to the number of demands that require additional capacity, and a capacity of $k\sigma_{sd}a_{ij}^{sd}$ must be allocated to account for the worst case realization for each of those demands. It is easy to see that if $\{X_k, k \geq 0\}$ is an independent sequence of continuous random variables, with X_k distributed according to a truncated gaussian described by (6), then the sequence of X_k converges in distribution to a Gaussian random variable with mean μ_{sd} and variance σ_{sd}^2.

In equation (9), k and Γ_{ij} are inversely proportional, implying that as k increases, fewer demands are allocated extra capacity, but more capacity is allocated to those demands. Since Theorem 1 holds for all finite values of k, Γ_{ij} eventually satisfies $0 \leq \Gamma_{ij} < 1$ as k becomes large so that S_{ij} in (7) is the empty set and the maximization is taken over one s–d pair representing t_{ij}. Consequently, equation (7) can be rewritten as

$$\begin{aligned} C_{ij} &\geq \sum_{(s,d)\in\mathcal{D}} a_{ij}^{sd}\mu_{sd} + \Gamma_{ij}k \max_{(s,d)\in\mathcal{D}}\{\sigma_{sd}a_{ij}^{sd}\} \\ &= \sum_{(s,d)\in\mathcal{D}} a_{ij}^{sd}\mu_{sd} + \sqrt{-2n_{ij}\log\epsilon} \max_{(s,d)\in\mathcal{D}}\{\sigma_{sd}a_{ij}^{sd}\} \end{aligned} \tag{10}$$

where we used the relationship (9). Using this constraint, the routing and capacity allocation problem can be formulated as follows:

$$\begin{aligned} \min_{C,a\geq 0} \quad & \max_{(i,j)\in E} C_{ij} \\ \text{s.t.:} \quad & C_{ij} \geq \sum_{(k,l)\in\mathcal{D}} a_{ij}^{kl}\mu_{kl} + \alpha\sqrt{n_{ij}}\sigma_{sd}a_{ij}^{sd}, \forall(s,d), \forall(i,j) \\ & FC(s,d,a^{sd})\, \forall(s,d) \in \mathcal{D} \end{aligned} \tag{11}$$

where $\alpha = \sqrt{-2\log\epsilon}$. Note that n_{ij} is a variable determined by a_{ij}^{sd}'s, and therefore, (11) is a nonlinear formulation. While this can be reformulated as a MILP that exactly computes the

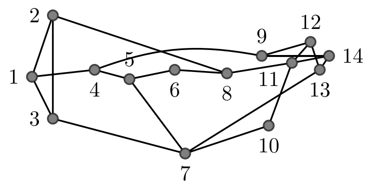

Fig. 2. 14-node NSFNET backbone network (1991).

exact values of n_{ij} and a_{ij}^{sd} in the problem (11), it quickly becomes intractable as the network size grows since it introduces additional integer variables (see Appendix B for more details). Therefore, we investigate a heuristic approach to solving the problem in (11) that scales better for large networks. In particular, our heuristic is an iterative algorithm that estimates the values n_{ij}, computes the routing variables a_{ij}^{sd} based on the estimated n_{ij}'s, and updates the values n_{ij} based on the computed routing variables a_{ij}^{sd}. The details are described as follows.

1) Every n_{ij} is initialized to 1.
2) The optimization problem in (11) is solved, and the number of demands traversing each link is counted.
3) If more than n_{ij} demands are routed on link (i, j), then n_{ij} is increased by 1.
4) The algorithm repeats until each link (i, j) has n_{ij} higher than the number of demands routed on it.
5) Link capacities are computed after termination by plugging the obtained solution a_{ij}^{sd}'s into (3).

This iterative scheme is called the increasing cost algorithm (ICA), and is guaranteed to converge in at most N^2 iterations, where N is the number of links. The value of n_{ij} can increase during each iteration, but never decrease. Note that since some values of n_{ij} can become large and cannot be lowered, ICA may find a suboptimal solution to the problem in (11) since it can end up searching a restricted space that does not contain an optimal solution. However, each iteration only requires solving an LP, and thus, the entire algorithm can be run in polynomial time.

C. Empirical Study

We compare the performance of the optimal nonlinear formulation in (1) (denoted by Optimal), linear formulation with constraint (4) (denoted by LP), and ICA. The maximum capacity $C_{\max}$ of linear formulation and ICA is recomputed using (3) after the routing has been computed. We use the NSFNET in Fig. 2 where demands are distributed as $\mathcal{N}(100,\ 35^2)$ and exist between randomly chosen node pairs. Fig. 3 shows that the optimal nonlinear formulation only provides a savings of 5% over the LP approach. Additionally, the LP approach outperforms the ICA algorithm by 5%.

An explanation for the good performance of the LP approach compared to the performance of the ICA algorithm is that the NSFNET does not provide enough opportunities for traffic from different sources to share link resources. To test this claim, we consider another network shown in Fig. 4. Note that this network allows for ample sharing of resources, allowing for a clear comparison of the different formulations. We run the same simulation, with we randomly generated, normally distributed

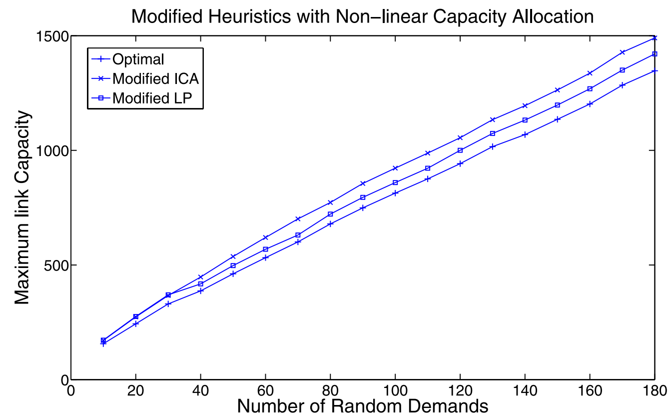

Fig. 3. Comparison of LP and ICA to the optimal approach for NSFNET. Each algorithm is given the same set of random demands, each normally distributed with $\mu = 100$ and $\sigma = 35$. Each measurement is an average of six simulations.

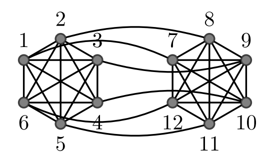

Fig. 4. Example network providing ample opportunities for splitting traffic over multiple paths. Links in this network are bidirectional.

demands with $\mu = 100$, $\sigma = 35$, and compare the results of our different approaches.

Fig. 5 compares the performance of three different approaches. Fig. 5(a) shows that the maximum link capacity for the optimal solution offers a 15% savings over that required by the linear approach. Our ICA approach allocates the same capacity as in the optimal approach. Fig. 5(b) plots the running time of each algorithm. Timing results are highly variable, and depend on the solver and computing hardware available. The purpose of Fig. 5(b) is to illustrate that the linear approach quickly finds a solution, whereas the running time of the nonlinear approach dramatically increases as the number of demands increases. Note that our algorithm finds a solution much faster than the nonlinear approach. This shows that even though convex optimization problems can be efficiently solved in general, it is desirable to develop a formulation with better scaling properties for problems over large networks.

IV. Robust Logical Topology Design

In this section, we study the logical topology design problem and develop algorithms for optimal network topologies for stochastic traffic. The goal here is to find a routing, capacity allocation, and logical network topology that minimizes the maximum link capacity. The routing and capacity allocation problem studied in the previous section provides clear insights on the network design. Namely, the optimal topology should allow for ample sharing of link resources among traffic demands since each demand is better to be split in order to minimize the maximum link capacity. This is true for both cases of deterministic and stochastic traffic demands, and leads to an interesting

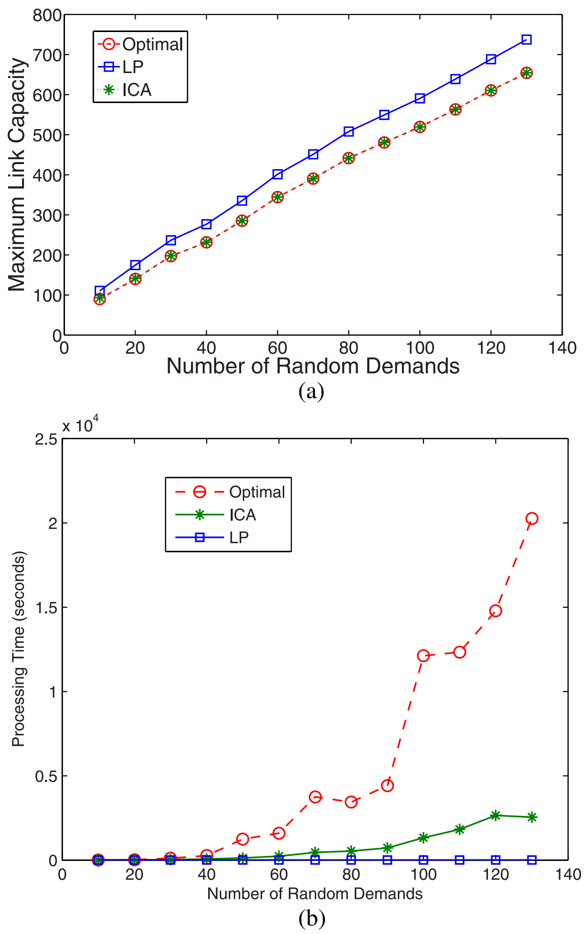

Fig. 5. Maximum link capacity and running time of each approach for the topology in Fig. 4. (a) Maximum link capacity (Average of ten simulations). (b) Running time (Average of six simulations, using LOQO for NLP and CPLEX for LP).

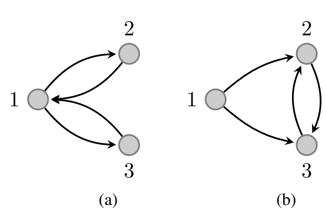

Fig. 6. Possible designs for a three-node network with four links, when there are two i.i.d demands from nodes 1 to 2 and 1 to 3. (a) Optimal deterministic topology. (b) Optimal stochastic topology.

algorithm for finding optimal topologies for stochastic traffic, as will be discussed later.

A. Motivating Example

We start with the three-node network example shown in Fig. 6. Suppose we have i.i.d. traffic ($\mathcal{N}(10,\sigma)$) from node 1 to node 2 and node 1 to node 3. The goal is to design a network and route the traffic over such a network such that max link capacity is minimized, and the link placement is restricted to a limit of four links, allowing no parallel links. If traffic is deterministic ($\sigma = 0$), the optimal link placement only requires a link from 1 to 2 and 1 to 3, while the other two links can be placed arbitrarily as in Fig. 6(a). With this link placement, each demand is routed on the one hop path to its destination, with a maximum link capacity of $C_{\max} = 10$.

Now consider the stochastic traffic case, by letting $\sigma = 1$. The optimal topology is shown in Fig. 6(b). On this topology, half of each demand can be sent on each link, and a lower $C_{\max}$ is achievable than on the network in Fig. 6(a). To be precise, the solution to (1) is $C_{\max} = 10 + \frac{\Phi^{-1}(1-\epsilon)}{\sqrt{2}}$. If the first topology is used to route the stochastic traffic instead, it would be sent directly to the destination, and a larger capacity of $C_{\max} = 10 + \Phi^{-1}(1-\epsilon)$ would be required. This example shows that network designs optimized for deterministic traffic are not necessarily optimal for stochastic traffic, and therefore, it is desirable to develop algorithms for optimal network design for stochastic traffic.

In the following, we discuss two approaches to finding optimal topologies for stochastic traffic. We note that the topology design problem can also be formulated as an MILP by adding the link-place component to the MILP for routing and capacity allocation problem in Appendix B. However, the formulation is intractable even for a network of moderate size and thus it is only presented in Appendix C for brevity.

B. Enumerate-and-Pick Approach

We have shown that an optimal topology for deterministic traffic may not be optimal for stochastic traffic. However, the converse was true in the example in Fig. 6. That is, the topology in Fig. 6(b) is also optimal for deterministic traffic ($\sigma = 0$). This example suggests that an optimal topology for stochastic traffic may also be optimal for deterministic traffic. While this claim will not be rigorously proved, there is strong intuitive evidence to suggest its validity. Recall from Section II that to minimize $C_{\max}$, it is better to split the traffic. The amount of traffic split onto each link varies with σ, but the number of demands routed over a link remains constant whether σ is positive or zero. Hence, a topology that is optimal for traffic demands with positive variance is likely to be optimal for deterministic traffic as well. A weaker version of this claim is proved through the following theorem.

Theorem 2: Given a set of nodes and a traffic matrix where the demand between each node pair is normally distributed with mean μ_{sd} and standard deviation σ, $\exists \sigma_0 > 0$ such that for any σ satisfying $0 \le \sigma \le \sigma_0$, the optimal topology designed for that σ is also optimal for deterministic traffic.

Proof: See Appendix D. ■

The aforementioned theorem can be generalized to the case where the variance on each link is different with a similar "small" σ assumption necessary on the vector σ. While this theorem does not hold for general large values of σ, it gives us a useful insight on optimal network design. In particular, given that the theorem is valid for general values of σ, the set $\mathcal{T}_\sigma$ of optimal topologies for stochastic traffic is contained in the set $\mathcal{T}$ of optimal topologies for deterministic traffic. Hence, once $\mathcal{T}$ is identified, an optimal topology for stochastic traffic can be found by computing $C_{\max}$ over each topology in $\mathcal{T}$ and choosing the one with minimum $C_{\max}$. In summary, the following procedure finds an optimal topology for stochastic traffic (provided that Theorem 2 holds for general $\sigma \geq 0$).

1) Identify $\mathcal{T}$.
2) Compute $C_{\max}$ over each $T \in \mathcal{T}$ for stochastic traffic.
3) Pick the one with minimum $C_{\max}$.

For step 2, the convex optimization problem in (3) can be solved over each topology in $\mathcal{T}$.

Step 1 can be completed as follows. Initially, the optimal deterministic topology is found for the set of nodes and demands provided, using existing methods such as the MILP in [5] (It is equivalent to the special case of the MILP in Appendix C with the variables x_{ij}^m, y_{ij}^m, and f_{ij}^{sd} and their related constraints removed). Once this topology is found, represented by variables $b_{ij}^{(1)}$, the following constraint is added to the MILP formulation.

$$\sum_{\{(i,j)|b_{ij}^{(1)}=0\}} b_{ij} + \sum_{\{(i,j)|b_{ij}^{(1)}=1\}} (1 - b_{ij}) \geq 1 \qquad (12)$$

This constraint enforces any feasible solution of the current optimization problem (b_{ij}) to differ from the solution already obtained ($b_{ij}^{(1)}$) by at least one link. If the solution to the modified MILP has the same objective function as the solution to the original problem, then the new topology is also an optimal topology for deterministic traffic. This is repeated by adding an additional constraint with every topology found, until the objective value of the modified program is larger than the optimal solution, implying that all the optimal deterministic topologies have been found.

Alternatively, since the formulation in (1) calculates the optimal routing and capacity assignment for any given topology, we can exhaustively search through every possible topology for that with the smallest max link capacity. Every possible permutation of links that meet the design requirements must be enumerated, and the convex routing formulation in (1) (it is a linear program for deterministic traffic) is applied to each. This is obviously inefficient, but returns the optimal topologies.

For example, consider an i.i.d. set of 11 demands over a five node network, where each demand is normally distributed with mean 100 and variance σ^2. Numbering the nodes from 1 to 5, assume traffic demands exists between the following node pairs: $\mathcal{D} = \{(1,2), (1,3), (2,1), (2,3), (2,4), (2,5), (3,2), (3,4), (3,5), (4,1), (4,5)\}$. The design constraint is that no more than ten links may be used, with no parallel links. As discussed previously, every topology meeting this requirement is enumerated, and the convex optimization in (1) is solved over each topology, assuming $\sigma = 0$. Table I summarizes all the topologies that are found to be optimal for $\sigma = 0$. The same technique is used for the traffic when $\sigma = 10$ and $\sigma = 30$. In both stochastic cases, there is only one optimal topology, namely topology 4 in Table I. The difference in the values of $C_{\max}$ for these topologies grows with the demand variance. For $\sigma = 30$, the optimal topology has $C_{\max}$ 5% lower than the others, and for $\sigma = 100$ this grows to 10%. This example again suggests that the stochastically optimal topology set $\mathcal{T}_\sigma$ is a subset of the deterministically optimal topology set $\mathcal{T}$.

TABLE I
OPTIMAL TOPOLOGIES FOR DETERMINISTIC TRAFFIC FOR THE DEMANDS IN $\mathcal{D}$: THE FIRST COLUMN IS THE TOPOLOGY NUMBER

T	Links
1	(1,2),(1,3),(2,1),(2,4),(2,5),(3,2),(3,4),(3,5),(4,1),(4,3)
2	(1,2),(1,3),(2,1),(2,4),(2,5),(3,2),(3,4),(4,1),(4,3),(4,5)
3	(1,3),(1,4),(2,1),(2,3),(2,5),(3,2),(3,4),(3,5),(4,1),(4,2)
4	(1,2),(1,3),(2,1),(2,3),(2,4),(3,2),(3,4),(3,5),(4,1),(4,5)
5	(1,2),(1,3),(2,1),(2,4),(2,5),(3,1),(3,2),(3,4),(4,3),(4,5)
6	(1,2),(1,4),(2,1),(2,3),(2,5),(3,2),(3,4),(3,5),(4,1),(4,3)
7	(1,2),(1,3),(2,3),(2,4),(2,5),(3,1),(3,2),(3,4),(4,1),(4,5)
8	(1,2),(1,3),(2,1),(2,3),(2,4),(2,5),(3,2),(3,4),(4,1),(4,5)

C. Simulated Annealing

The aforementioned approach is limited by the execution time of the MILP, which can be prohibitively high for large networks. Additionally, networks can have thousands of optimal topologies for deterministic traffic. As a more scalable approach, we consider simulated annealing (SA). Briefly, SA is a random search heuristic which can be used to find near optimal solutions to optimization problems. The algorithm begins with an arbitrary feasible solution, and a cost computed with respect to an objective function. Then, a random perturbation is applied to the solution, and the cost is reevaluated. The new solution is probabilistically accepted based on the relationship between the two costs. A positive probability of moving to a worse solution avoids the problem of being trapped in a local minima.

Simulated annealing (SA) can be used in this context to compute the link placements. We assume there is a constraint on the maximum number of links. From any feasible network topology, the optimization problem in (1) is solved to find $C_{\max}$. Then, a random perturbation is applied to the topology. Specifically, a link (i,j) in the topology and a link (k,l) not in the topology are randomly chosen uniformly from all the links in their respective sets, and link (k,l) is added to the topology while link (i,j) is removed.

For the new topology, the problem in (1) is solved, where the value of the objective function at the solution is $C'_{\max}$. If $C'_{\max} < C_{\max}$, the new topology allows for a smaller maximum link capacity and is accepted unconditionally. If the new topology is worse, it is accepted with probability q, where

$$q = \exp\left(\frac{C_{\max} - C'_{\max}}{T}\right). \qquad (13)$$

The probability of acceptance is a function of the difference between solutions so that topologies that are much worse are less likely to be accepted than topologies which are only slightly worse. The parameter T is referred to as the *temperature*, borrowing terminology from the physics literature. Initially, the temperature of the system is large, such that worse topologies are still likely to be accepted. As the algorithm progresses, temperature is lowered slowly, so less “bad” topologies will be accepted. Specifically, SA iterates for a fixed number of perturbations to simulate reaching a steady state, and then, temperature is modified according to $T' = \beta T$ for some $0 \leq \beta \leq 1$. The value of β should be close to 1, so that the temperature is

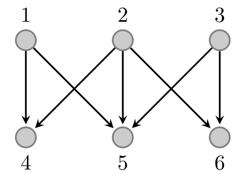

Fig. 7. Illustration of demands throughout a six-node cluster. A directed edge represents a demand in that direction. All demands are i.i.d.

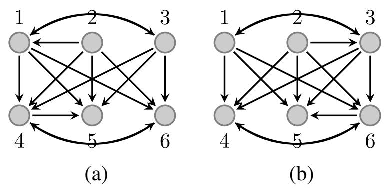

Fig. 8. The two optimal topologies with 15 links for the demand pattern in 7. (a) Topology 1. (b) Topology 2.

reduced slowly. The algorithm starts with a high value of T, and it is stopped when the probability of escape from some topology is small enough that new topologies are no longer considered.

The SA algorithm should not depend on the initial topology used to begin the search, since SA has measures to protect itself from local minima. Therefore, M links are randomly chosen initially, where M corresponds to the design requirement on the maximum number of links. If the initial links cannot support the demands, the topology is labeled as infeasible. As long as the topology is infeasible, a new topology is accepted with probability 1.

D. Empirical Study

In the previous sections, we have shown that while the optimal stochastic topology is also an optimal deterministic topology, some deterministic topologies perform poorly under stochastic traffic. We would like to quantize this difference.

Consider a six node network with the i.i.d demands between the set of node pairs shown in Fig. 7. Our goal will be to place up to 15 links between these six nodes, such that the traffic can be optimally routed. Traffic demands are normally distributed with mean 100 and variance σ^2. The optimal topology is found by first locating the set of optimal topologies for deterministic traffic, and then, pruning that set for the optimal stochastic topology for $\sigma = 35$. The two optimal stochastic topologies result in $C_{\max} =$ 123.89. These two topologies were chosen from a set of 264 deterministically optimal topologies. Note that each MILP for deterministic design takes between 2 and 3 s to solve, whereas each convex optimization takes between 1.5 and 2 s. The two optimal topologies are shown in Fig. 8.

By finding the optimal topologies rather than an arbitrary deterministically optimal topology, we save on max link capacity. The stochastically optimal topologies save 8% of the max link capacity that would be necessary for the worst case choice of deterministic topology. On average, the stochastic topologies save 5% of the max link capacity. We can expect to see larger numbers for bigger networks although this is difficult to verify due to the difficulty of solving large MILP's.

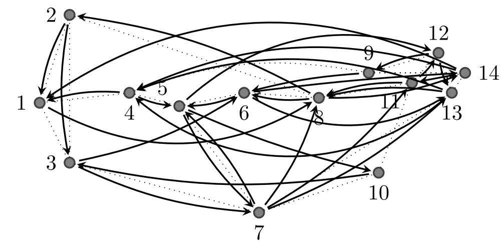

Fig. 9. Resulting 32-link network (solid) for i.i.d $\mathcal{N}(100, 35^2)$ demands following NSFNET (dotted). The network is designed for $\epsilon = 0.01$, using Simulated Annealing with $\beta = 0.95$ for the link placement, and ICA for the capacity allocation and routing.

For larger networks, the approach of finding all optimal topologies for deterministic traffic, then searching through those topologies for that which is optimal for stochastic traffic is computationally intensive. In order to design networks for these instances, the SA approach is used to calculate the link placements, and the routing and capacity allocation problems can be solved using the ICA heuristic developed in Section III-B.

Consider the NSFNET in Fig. 2. Assume each bidirectional link corresponds to a demand normally distributed with mean 100 and standard deviation 35. The goal of this simulation is to place 32 links on a 14 node network to best route those demands. This can be thought of as redesigning the NSFNET backbone with fewer directed links. The resulting network is shown in Fig. 9, and requires a maximum link capacity of 456.162. Note that this network may not be optimal, due to the suboptimalities of the SA approach and the ICA heuristic.

V. Conclusion and Future Work

Modern backbone networks must be designed to be robust to random traffic fluctuations and errors in traffic estimation. Previous attempts to solve the routing and capacity allocation problems have either resulted in formulations which are too conservative in terms of allocated capacity, or formulations which are nonlinear. By extending results from robust optimization, we have developed a formulation to allocate capacity in a less conservative fashion, while computing near optimal routes for traffic flows. The resulting routing for stochastic traffic takes advantage of multiplexing opportunities in the network, such that shared capacity is efficiently utilized. We provide an iterative scheme to solve our robust optimization problem, and show that our formulation finds near-optimal routings and capacity allocations, while remaining computationally tractable.

While our formulation allows an extension to the link-placement problem, the resulting MILP formulation is still computationally intractable. Our analysis shows that networks supporting random traffic should be designed in a manner to allow link capacity to be shared by many traffic flows. Therefore, it would be interesting to develop algorithms and heuristics to solve the link placement subproblem exploiting these strategies.

APPENDIX A
PROOF OF THEOREM 1

We need the following lemma to prove the theorem.

Lemma 1: Let X be a continuous random variable with PDF $\mathcal{N}(x|a \le x \le b; \mu, \sigma^2)$. Let $Y = cX + d$ for $c > 0$. Then, the random variable Y is distributed with density $\mathcal{N}(x|ac + d \le x \le bc + d; c\mu + d, (c\sigma)^2)$.

Proof: X has PDF $f_X(x; \mu, \sigma^2, a, b) = \frac{\frac{1}{\sigma}\phi\left(\frac{x-\mu}{\sigma}\right)}{\Phi\left(\frac{b-\mu}{\sigma}\right)-\Phi\left(\frac{a-\mu}{\sigma}\right)}$, where $\phi()$ and $\Phi()$ are the PDF and CDF of a standard normal random variable, respectively. Since Y is a linear function of X, we can write the density of Y as $f_Y(y) = \frac{1}{|c|} f_X(\frac{y-d}{c}; \mu, \sigma^2, a, b)$. By plugging in the definition of $f_X()$,

$$f_Y(y) = \frac{\frac{1}{c\sigma}\phi\left(\frac{y-d-c\mu}{c\sigma}\right)}{\Phi\left(\frac{b-\mu}{\sigma}\right) - \Phi\left(\frac{a-\mu}{\sigma}\right)} \tag{14}$$

Define $\mu' \triangleq c\mu + d$ and $\sigma' \triangleq c\sigma$

$$f_Y(y) = \frac{\frac{1}{\sigma'}\phi\left(\frac{y-\mu'}{\sigma'}\right)}{\Phi\left(\frac{bc+d-\mu'}{\sigma'}\right) - \Phi\left(\frac{ac+d-\mu'}{\sigma'}\right)} \tag{15}$$

$$= \mathcal{N}(x|ac + d \le x \le bc + d; c\mu + d, (c\sigma)^2) \tag{16}$$

□

Proof of Theorem 1: Define the following two variables

$$z_B^{sd} \triangleq \frac{\lambda_B^{sd} - \mu_{sd}}{\sigma_{sd}} \quad \text{and } \eta^{sd} \triangleq \frac{1}{k} z_B^{sd}. \tag{17}$$

By Lemma 1, z_B^{sd} has a PDF given by $\mathcal{N}(z|z \le k; 0, 1)$, and η^{sd} has a PDF given by $\mathcal{N}(\eta|\eta \le 1; 0, 1/k^2)$. It follows from (17) that

$$\mathbf{P}\left(\sum_{sd} a_{ij}^{sd}\lambda_B^{sd} > C_{ij}\right) = \mathbf{P}\left(\sum_{sd} a_{ij}^{sd}\mu_{sd} + \sum_{sd} a_{ij}^{sd}\sigma_{sd} z_B^{sd} > C_{ij}\right)$$
$$= \mathbf{P}\left(\sum_{sd} a_{ij}^{sd}\mu_{sd} + \sum_{sd} a_{ij}^{sd}\sigma_{sd} k\eta^{sd} > C_{ij}\right)$$

Since C_{ij} satisfies (7),

$$\mathbf{P}\left(\sum_{sd} a_{ij}^{sd}\lambda_B^{sd} > C_{ij}\right)$$
$$\le \mathbf{P}\left(\sum_{sd} a_{ij}^{sd} k\sigma_{sd}\eta^{sd} > \sum_{(s,d)\in S_{ij}} a_{ij}^{sd} k\sigma_{sd} + (\Gamma_{ij} - \lfloor\Gamma_{ij}\rfloor) a_{ij}^{t_{ij}} k\sigma_{t_{ij}}\right)$$
$$= \mathbf{P}\left(\sum_{sd} a_{ij}^{sd}\sigma_{sd}\eta^{sd} > \sum_{(s,d)\in S_{ij}} a_{ij}^{sd}\sigma_{sd} + (\Gamma_{ij} - \lfloor\Gamma_{ij}\rfloor) a_{ij}^{t_{ij}} \sigma_{t_{ij}}\right)$$
$$= \mathbf{P}\left(\sum_{(s,d)\notin S_{ij}} a_{ij}^{sd}\sigma_{sd}\eta^{sd} > \sum_{(s,d)\in S_{ij}} a_{ij}^{sd}\sigma_{sd}(1 - \eta^{sd}) + (\Gamma_{ij} - \lfloor\Gamma_{ij}\rfloor) a_{ij}^{t_{ij}} \sigma_{t_{ij}}\right)$$

Let $r = \arg\min_{r'\in S_{ij}\bigcup t_{ij}} a_{ij}^{r'}\sigma_{r'}$. Since $\eta^{sd} \le 1$, the aforementioned probability can be bounded as

$$\le \mathbf{P}\left(\sum_{(s,d)\notin S_{ij}} a_{ij}^{sd}\sigma_{sd}\eta^{sd} > a_{ij}^{r}\sigma_r\left(\sum_{(s,d)\in S_{ij}} (1 - \eta^{sd}) + \left(\Gamma_{ij} - \lfloor\Gamma_{ij}\rfloor\right)\right)\right)$$
$$= \mathbf{P}\left(\sum_{(s,d)\notin S_{ij}} a_{ij}^{sd}\sigma_{sd}\eta^{sd} + \sum_{(s,d)\in S_{ij}} a_{ij}^{r}\sigma_r\eta^{sd} > a_{ij}^{r}\sigma_r\Gamma_{ij}\right)$$
$$= \mathbf{P}\left(\sum_{(s,d)} \gamma_{ij}^{sd}\eta^{sd} > \Gamma_{ij}\right) \tag{18}$$

where

$$\gamma_{ij}^{sd} = \begin{cases} 1, & \text{if } (s,d) \in S_{ij} \bigcup t_{ij} \\ \frac{a_{ij}^{sd}\sigma_{sd}}{a_{ij}^{r}\sigma_r}, & \text{if } (s,d) \notin S_{ij} \bigcup t_{ij}. \end{cases} \tag{19}$$

Since r is chosen to be an element of $S_{ij} \bigcup t_{ij}$, then $a_{ij}^{r}\sigma_r \ge a_{ij}^{sd}\sigma_{sd} \quad \forall(s,d) \notin S_{ij}$. Therefore, γ_{ij}^{sd} defined in (19) is always less than or equal to 1. An additional Lemma is needed to complete the proof.

Lemma 2:

$$\mathbf{E}[e^{t\eta^{sd}}] = M^{sd}(t) = \exp\left(\frac{t^2}{2k^2}\right)\left(\frac{\Phi(k - t/k)}{\Phi(k)}\right). \tag{20}$$

Proof: Let $Y \sim \mathcal{N}(0, \frac{1}{k^2})$,

$$M^{sd}(t) = \mathbf{E}[e^{tY}|Y \le 1]$$
$$= \int_{-\infty}^{1} \frac{e^{ty} f_Y(y) dy}{\Phi(k)}$$
$$= \frac{1}{\Phi(k)} \int_{-\infty}^{1} e^{ty} \frac{k}{\sqrt{2\pi}} e^{-\frac{y^2k^2}{2}} dy$$
$$= \frac{k}{\Phi(k)\sqrt{2\pi}} \int_{-\infty}^{1} \exp\left(-\frac{k^2}{2}\left(y^2 - \frac{2ty}{k^2}\right)\right) dy$$
$$= \frac{k}{\Phi(k)\sqrt{2\pi}} \int_{-\infty}^{1} \exp\left(-\frac{k^2}{2}\left(y^2 - \frac{2ty}{k^2} + \frac{t^2}{k^4}\right) + \frac{t^2}{2k^2}\right) dy$$
$$= \frac{k}{\Phi(k)\sqrt{2\pi}} \int_{-\infty}^{1} \exp\left(\frac{t^2}{2k^2}\right) \exp\left(-\frac{k^2}{2}\left(y - \frac{t}{k^2}\right)^2\right) dy$$
$$= \frac{e^{\frac{t^2}{2k^2}}}{\Phi(k)} \int_{-\infty}^{1} \frac{k}{\sqrt{2\pi}} e^{-\frac{k^2}{2}(y-\frac{t}{k^2})^2} dy.$$

This can be reduced in terms of the standard normal CDF.

$$= \frac{e^{\frac{t^2}{2k^2}}}{\Phi(k)} \left(\Phi\left(\frac{1 - \frac{t}{k^2}}{\frac{1}{k}}\right) - \Phi\left(\frac{-\infty - \frac{t}{k^2}}{\frac{1}{k}}\right)\right)$$
$$= e^{\frac{t^2}{2k^2}} \left(\frac{\Phi(k - \frac{t}{k})}{\Phi(k)}\right).$$

□

This lemma can be used to complete the proof. Starting with the bound in (18),

$$\begin{aligned}
&\mathbf{P}\Big(\sum_{(s,d)} \gamma_{ij}^{sd}\eta^{sd} > \Gamma_{ij}\Big) \\
&= \mathbf{P}\Big(t\sum_{(s,d)} \gamma_{ij}^{sd}\eta^{sd} > t\Gamma_{ij}\Big) \qquad \forall t>0 \\
&= \mathbf{P}\Big(\exp\Big(t\sum_{(s,d)} \gamma_{ij}^{sd}\eta^{sd}\Big) > \exp\left(t\Gamma_{ij}\right)\Big) \\
&\le \frac{\mathbf{E}[e^{t\sum_{(s,d)}\gamma_{ij}^{sd}\eta^{sd}}]}{e^{t\Gamma_{ij}}} \qquad \text{by Markov's Inequality.} \\
&= \prod_{sd} \frac{\mathbf{E}[e^{t\eta^{sd}\gamma_{ij}^{sd}}]}{e^{t\Gamma_{ij}}} \qquad \text{by independence.} \\
&= \prod_{sd} \frac{M^{sd}(t\gamma_{ij}^{sd})}{e^{t\Gamma_{ij}}} \\
&= \prod_{sd} \frac{e^{\frac{(t\gamma_{ij}^{sd})^2}{2k^2}}\left(\frac{\Phi(k-\frac{t\gamma_{ij}^{sd}}{k})}{\Phi(k)}\right)}{e^{t\Gamma_{ij}}} \qquad \text{by Lemma 2} \\
&\le \prod_{sd} \frac{e^{\frac{(t\gamma_{ij}^{sd})^2}{2k^2}}}{e^{t\Gamma_{ij}}}
\end{aligned}$$

Since $\Phi(k) \ge \Phi(k - \frac{t\gamma_{ij}^{sd}}{k}) \quad \forall t \ge 0, \gamma_{ij}^{sd} \ge 0$.

$$\begin{aligned}
&= \frac{\exp\left(\frac{t^2}{2k^2}\sum_{sd}(\gamma_{ij}^{sd})^2\right)}{e^{t\Gamma_{ij}}} \\
&\le \exp\left(\frac{n_{ij}t^2}{2k^2} - t\Gamma_{ij}\right) \qquad \text{since } 0 \le \gamma_{ij}^{sd} \le 1.
\end{aligned}$$

The above bound is minimized when its derivative with respective to t is 0, i.e.,

$$\begin{aligned}
\frac{d}{dt}\left(\frac{n_{ij}t^2}{2k^2} - t\Gamma_{ij}\right) &= \frac{n_{ij}t}{k^2} - \Gamma_{ij} = 0 \\
\Rightarrow \quad t &= \frac{\Gamma_{ij}k^2}{n_{ij}}.
\end{aligned}$$

Plugging this t into the bound yields

$$\begin{aligned}
\mathbf{P}\Big(\sum_{(s,d)} \gamma_{ij}^{sd}\eta^{sd} > \Gamma_{ij}\Big) &\le \exp\left(\frac{n_{ij}}{2k^2}\frac{\Gamma_{ij}^2k^4}{n_{ij}^2} - \frac{\Gamma_{ij}k^2}{n_{ij}}\Gamma_{ij}\right) \\
&= \exp\left(-\frac{\Gamma_{ij}^2k^2}{2n_{ij}}\right)
\end{aligned}$$

□

APPENDIX B

MILP FORMULATION FOR GAUSSIAN DEMANDS

Let f_{ij}^{sd} be a binary variable satisfying $f_{ij}^{sd} \ge a_{ij}^{sd}$. Thus, f_{ij}^{sd} equals 1 if $a_{ij}^{sd} > 0$, and 0 otherwise. Let $x_{ij}^m = 1$ if there are m demands traversing link (i,j), and 0 otherwise. These variables satisfy

$$\sum_{sd} f_{ij}^{sd} = \sum_{m=0}^{D} m x_{ij}^m \qquad \forall(i,j) \tag{21}$$

where D is the total number of demands in the network. By combining this constraint with one forcing only one of $\{x_{ij}^0, x_{ij}^1, \ldots, x_{ij}^D\}$ to be equal to one for each link (i,j), the variables x_{ij}^m specify the number of flows on each link. Therefore, the capacity constraint of (11) can be rewritten as

$$C_{ij} \ge \sum_{sd\in\mathcal{D}} \mu_{sd} a_{ij}^{sd} + z_{ij}\alpha \sum_{m=0}^{D} x_{ij}^m \sqrt{m} \qquad \forall(i,j). \tag{22}$$

Equation (22) is nonlinear, since z_{ij} and x_{ij}^m are both optimization variables. However, (22) can be linearized by introducing a variable y_{ij}^m to represent the product $z_{ij}x_{ij}^m$ with the following constraints.

$$y_{ij}^m \ge z_{ij} + M(x_{ij}^m - 1) \qquad \forall(i,j), m \tag{23}$$

$$y_{ij}^m \le M x_{ij}^m \qquad \forall(i,j), m \tag{24}$$

$$y_{ij}^m \ge 0 \qquad \forall(i,j), m. \tag{25}$$

In the aforementioned equations, M is a large number such that $M > \max_{sd}\sigma_{sd}$. When $x_{ij}^m = 0$, then $x_{ij}^m z_{ij} = 0$, and constraints (24) and (25) force y_{ij}^m to be 0. On the other hand, if $x_{ij}^m = 1$, constraint (23) will force $y_{ij}^m \ge z_{ij}$, which at the optimal solution will be satisfied with equality. Therefore, the complete MILP to solve the routing problem is formulated as

$$\begin{aligned}
&\text{Min } C_{\max} \\
&\text{Subject To: } C_{ij} \ge \sum_{sd\in\mathcal{D}} \mu_{sd}a_{ij}^{sd} + \alpha\sum_{m=0}^{D}\sqrt{m}\,y_{ij}^m, \forall(i,j) \\
&\sum_j a_{ij}^{sd} - \sum_j a_{ji}^{sd} = \begin{cases} 1, & \text{if } s=i \\ -1, & \text{if } d=i, \forall(s,d), i \\ 0, & \text{otherwise} \end{cases} \\
&C_{ij} \le C_{max} \quad \forall(i,j) \\
&f_{ij}^{sd} \ge a_{ij}^{sd} \quad \forall(i,j),(s,d) \\
&\sum_{sd} f_{ij}^{sd} = \sum_{m=0}^{D} m x_{ij}^m \quad \forall(i,j) \\
&\sum_{m=0}^{D} x_{ij}^m = 1 \quad \forall(i,j) \\
&y_{ij}^m \ge \sigma_{sd}a_{ij}^{sd} + M(x_{ij}^m - 1) \quad \forall(s,d),(i,j),m \\
&y_{ij}^m \le M x_{ij}^m \quad \forall(i,j),m \\
&y_{ij}^m \ge 0 \quad \forall(i,j),m \\
&a_{ij}^{sd} \ge 0 \quad \forall i,j,s,d \\
&f_{ij}^{sd}, x_{ij}^m \in \{0,1\} \quad \forall(s,d),(i,j),m.
\end{aligned}$$

Appendix C
MILP for Robust Topology Design

Let $b_{ij} = 1$ if a link is placed between nodes i and j, and $b_{ij} = 0$ otherwise. We add the constraint $b_{ij} \geq a_{ij}^{sd}$ requiring that link (i,j) cannot be used for routing unless there is a link between nodes i and j. Let Δ^- and Δ^+ be the out-degree and in-degree of each node, respectively. The node degree constraints are written as $\sum_i b_{ij} = \Delta_i$ and $\sum_j b_{ij} = \Delta_o$. Note that these two constraints can be replaced with a constraint limiting the maximum number of links, or potentially other design restrictions. The complete MILP formulation is presented below. Note that only the last three constraints have been added to the routing and capacity allocation MILP in Appendix B.

Minimize $\quad C_{\max}$

$$\text{Subject To:} \quad C_{ij} \geq \sum_{sd\in\mathcal{D}} \mu_{sd} a_{ij}^{sd} + \alpha \sum_{m=0}^{D} \sqrt{m} y_{ij}^m \quad \forall (i,j)$$
$$FC(s,d,a^{sd}) \quad \forall (s,d) \in \mathcal{D}$$
$$C_{ij} \leq C \quad \forall (i,j)$$
$$f_{ij}^{sd} \geq a^{sd} \quad \forall (i,j),(s,d)$$
$$\sum_{sd} f_{ij}^{sd} = \sum_{m=0}^{D} mx \quad \forall (i,j)$$
$$\sum_{m=0}^{D} x_{ij}^m = 1 \quad \forall (i,j)$$
$$y_{ij}^m \geq \sigma_{sd} a_{ij}^{sd} + M(x_{ij}^m - 1) \quad \forall (s,d),(i,j),m$$
$$y_{ij}^m \leq M x^m \quad \forall (i,j),m$$
$$y_{ij}^m \geq 0 \quad \forall (i,j),m$$
$$a_{ij}^{sd} \geq 0 \quad \forall i,j,s,d$$
$$f_{ij}^{sd}, x_{ij}^m, b_{ij} \in \{0,1\} \quad \forall (s,d),(i,j),m$$
$$b_{ij} \geq a_{ij}^{sd}, \forall (i,j) \quad (s,d)$$
$$\sum_i b_{ij} = \Delta_i \quad \forall j$$
$$\sum_j b_{ij} = \Delta_o \quad \forall i$$

Appendix D
Proof of Theorem 2

Let T_0 and T_σ be the optimal topologies for deterministic traffic and Gaussian traffic with variance σ^2, respectively. Further, let $C_{\max}^{T,\sigma}$ be the optimal maximum link capacity in routing traffic with demand variance σ^2 over topology T. For example, $C_{\max}^{T_0,0}$ is the solution for routing deterministic traffic over the topology that is optimal for deterministic traffic. Throughout this proof, we will use $\kappa = \Phi^{-1}(1-\epsilon)$, given a design parameter ϵ. We will prove our claim by contradiction.

Suppose T_σ is worse than T_0 in terms of routing deterministic traffic. That is,

$$C_{\max}^{T_0,0} < C_{\max}^{T_\sigma,0}. \tag{26}$$

The difference between these two solutions Δ is defined as

$$\Delta = C_{\max}^{T_\sigma,0} - C_{\max}^{T_0,0} > 0. \tag{27}$$

Let a_{ij}^{sd} and b_{ij}^{sd} be the optimal routing for traffic with variance σ^2 and deterministic traffic on topology T_0 respectively. Let (i,j) be a link achieving max capacity when stochastic traffic is optimally routed over T_0, and (k,l) be a link achieving max capacity when routing b is used for stochastic traffic over T_0. Then, by the capacity constraint in (3),

$$C_{\max}^{T_0,\sigma} = \sum_{sd} a_{ij}^{sd}\mu_{sd} + \kappa\sigma\sqrt{\sum_{sd}(a_{ij}^{sd})^2} \tag{28}$$
$$\leq \sum_{sd} b_{kl}^{sd}\mu_{sd} + \kappa\sigma\sqrt{\sum_{sd}(b_{kl}^{sd})^2} \tag{29}$$
$$\leq C_{\max}^{T_0,0} + \kappa\sigma\sqrt{\sum_{sd}(b_{kl}^{sd})^2} \tag{30}$$
$$\leq C_{\max}^{T_0,0} + \kappa\sigma\sqrt{D} \tag{31}$$

where D is a constant representing the maximum value of $\sum_{sd}(b_{kl}^{sd})^2$ for any feasible routing.

Consider the following optimization problem for routing deterministic traffic over any fixed topology T.

Minimize: $C_{\max}$

$$\text{Subject To: } C_{\max} \geq \sum_{sd} a_{ij}^{sd}\mu_{sd} \qquad \forall (i,j) \in \mathcal{L}$$
$$\sum_j a_{ij}^{sd} - \sum_j a_{ji}^{sd} = \begin{cases} 1, & \text{if } s = i \\ -1, & \text{if } d = i, \forall (s,d), i \\ 0, & \text{otherwise} \end{cases}$$
$$a_{ij}^{sd} \geq 0, \forall (s,d),(i,j). \tag{32}$$

The Lagrangian of problem (32) is

$$L(\nu,\theta,C_{\max},a) = C_{\max} + \sum_{ij}\nu_{ij}\left(\sum_{sd}\mu_{sd}a_{ij}^{sd} - C_{\max}\right) + \sum_{s,d,i}\theta_i^{sd}(\beta_i^{sd} - h_i^{sd}(\mathbf{a^{sd}})) \tag{33}$$

where $\nu_{ij} \geq 0$ and θ_i^{sd} are dual variables, $h_i^{sd}(\mathbf{a^{sd}}) = \sum_j a_{ij}^{sd} - \sum_j a_{ji}^{sd}$ and

$$\beta_i^{sd} = \begin{cases} 1, & \text{if } s = i \\ -1, & \text{if } d = i \\ 0, & \text{otherwise.} \end{cases} \tag{34}$$

First, we minimize over the primary variables.

$$g(\nu,\theta) = \inf_{C_{\max}, a_{ij}^{sd} \geq 0} \left(C_{\max} + \sum_{ij} \nu_{ij} \left(\sum_{sd} \mu_{sd} a_{ij}^{sd} - C_{\max} \right) + \sum_{s,d,i} \theta_i^{sd} \left[\beta_i^{sd} - h_i^{sd}(\mathbf{a^{sd}}) \right] \right). \quad (35)$$

Now, the dual problem can be written as

$$\begin{aligned} d^*(0) = \text{Maximize:} \quad & g(\nu,\theta) \\ \text{Subject To:} \quad & \sum_{(i,j)} \nu_{ij} = 1 \\ & \theta_i^{sd} - \theta_j^{sd} \leq \nu_{ij} \quad \forall (s,d),(i,j) \\ & \nu_{ij} \geq 0 \quad \forall (i,j) \\ & \theta_i^{sd} \quad \text{free} \quad \forall i,(s,d). \end{aligned}$$

By the min–max theorem and by strong duality for linear programs, there is zero duality gap between the dual and primal problems.

$$\begin{aligned} C_{\max}^*(0) = d^*(0) &= \min_{C_{\max}, a_{ij}^{sd}} \max_{\nu,\theta} L(\nu,\theta,C_{\max},a) \\ &= \min_{C_{\max}, a_{ij}^{sd}} \Bigg(C_{\max} + \sum_{ij} \nu_{ij}^*(0) \left(\sum_{sd} \mu_{sd} a_{ij}^{sd} - C_{\max} \right) \\ &\quad + \sum_{s,d,i} \theta_i^{*sd}(0) [\beta_i^{sd} - h_i^{sd}(\mathbf{a^{sd}})] \Bigg). \end{aligned}$$

Consider the point $(C_{\max}^*(\sigma), a^*(\sigma))$, which is the optimal routing and max link capacity under traffic with variance σ^2 over the fixed topology T. $(C_{\max}^*(\sigma), a^*(\sigma))$ is obviously feasible and since it is potentially sub-optimal for the deterministic problem, it satisfies

$$C_{\max}^*(0) \leq C_{\max}^*(\sigma) + \sum_{ij} \nu_{ij}^*(0) \left(\sum_{sd} \mu_{sd} a_{ij}^{*sd}(\sigma) - C_{\max}^*(\sigma) \right) + \sum_{s,d,i} \theta_i^{*sd}(0) [\beta_i^{sd} - h_i^{sd}(\mathbf{a^{*sd}}(\sigma))]. \quad (36)$$

Since the aforementioned routing is feasible, $\beta_i^{sd} = h_i^{sd}(\mathbf{a^{*sd}}(\sigma))$ for all $i,(s,d)$. Furthermore, for each link (i,j), the point $(C_{\max}^*(\sigma), a^*(\sigma))$ satisfies

$$C_{\max}^*(\sigma) \geq \sum_{sd} \mu_{sd} a_{ij}^{*sd}(\sigma) + \kappa\sigma \sqrt{\sum_{sd} (a_{ij}^{*sd}(\sigma))^2} \quad (37)$$

By combining (37) with (36), we get the following inequality.

$$C_{\max}^*(0) \leq C_{\max}^*(\sigma) - \kappa\sigma \sum_{ij} \nu_{ij}^*(0) \sqrt{\sum_{sd} (a_{ij}^{*sd}(\sigma))^2} \quad (38)$$

$$C_{\max}^*(\sigma) \geq C_{\max}^*(0) + \kappa\sigma \sum_{ij} \nu_{ij}^*(0) \sqrt{\sum_{sd} (a_{ij}^{*sd}(\sigma))^2}. \quad (39)$$

The aforemenntioned result holds for all topologies, so we consider the specific topology T_σ. Let $X = \sum_{ij} \nu_{ij}^*(0) \sqrt{\sum_{sd} (a_{ij}^{*sd}(\sigma))^2}$, then equation (39) becomes

$$C_{\max}^{T_\sigma,\sigma} \geq C_{\max}^{T_\sigma,0} + \kappa\sigma X = C_{\max}^{T_0,0} + \Delta + \kappa\sigma X \quad (40)$$

where equation (40) results from the assumption in (27). Assume $\Delta + \kappa\sigma X \geq \kappa\sigma\sqrt{D}$. If $X \geq \sqrt{D}$, this assumption is valid regardless of the value of σ. On the other hand, if $X < \sqrt{D}$, then the assumption holds as long as

$$\sigma < \frac{\Delta}{\kappa(\sqrt{D} - X)} = \sigma_0. \quad (41)$$

Note that $\sigma_0 > 0$ since $0 < X < \sqrt{D}$. Therefore, assuming $\sigma < \sigma_0$ implies

$$\begin{aligned} C_{\max}^{T_\sigma,\sigma} &> C_{\max}^{T_0,0} + \kappa\sigma\sqrt{D} \\ &\geq C_{\max}^{T_0,\sigma} \end{aligned} \quad (42)$$

Equation (42) follows from equation (31), and is a contradiction, since T_σ is optimal for stochastic traffic. Therefore, the assumption made in (26) is false and $C_{\max}^{T_0,0} \geq C_{\max}^{T_\sigma,0}$. However, since T_0 is optimal for deterministic traffic,

$$C_{\max}^{T_0,0} = C_{\max}^{T_\sigma,0}. \quad (43)$$

REFERENCES

[1] M. Johnston, H.-W. Lee, and E. Modiano, "Robust network design for stochastic traffic demands," in *Proc. IEEE Global Telecommun. Conf. (GLOBECOM 2011)*, 2011, pp. 1–5.

[2] C. Whitepaper. (2007). Cost comparison between IP-over-DWDM and other core transport architectures. [Online]. Available: http://www.cisco.com

[3] R. Krishnaswamy and K. Sivarajan, "Design of logical topologies: a linear formulation for wavelength-routed optical networks with no wavelength changers," *IEEE/ACM Trans. Network.*, vol. 9, no. 2, pp. 186–198, Apr. 2001.

[4] E. Leonardi, M. Mellia, and M. Marsan, "Algorithms for the logical topology design in WDM all-optical networks," *Opt. Netw. Mag.*, vol. 1, no. 1, pp. 35–46, 2000.

[5] R. Ramaswami and K. N. Sivarajan, "Design of logical topologies for wavelength-routed optical networks," *IEEE J. Sel. Areas Commun.*, vol. 14, no. 5, pp. 840–851, Jun. 1996.

[6] A. Arora, S. Subramaniam, and H.-A. Choi, "Logical topology design for linear and ring optical networks," *IEEE J. Sel. Areas Commun.*, vol. 20, no. 1, pp. 62–74, Jan. 2002.

[7] B. Wu, K. Yeung, and P.-H. Ho, "Virtual topology design for minimizing network diameter and average hop count in WDM networks," *IEEE/OSA J. Opt. Commun. Network.*, vol. 2, no. 12, pp. 1077–1086, Dec. 2010.

[8] W. Ben-Ameur and H. Kerivin, "Routing of uncertain traffic demands," *Optim. Eng.*, vol. 6, no. 3, pp. 283–313, 2005.

[9] O. Karasan, M. Pinar, and H. Yaman, "Robust DWDM routing and provisioning under polyhedral demand uncertainty," presented at the INOC, Libson, Portugal, 2005

[10] D. Mitra and Q. Wang, "Stochastic traffic engineering for demand uncertainty and risk-aware network revenue management," *IEEE/ACM Trans. Network.*, vol. 13, no. 2, pp. 221–233, Apr. 2005.

[11] I. Widjaja, I. Saniee, A. Elwalid, and D. Mitra, "Online traffic engineering with design-based routing," presented at the ITC Spec. Workshop, Wuerzburg, Gerrmany, 2002.

[12] K. Meesublak, "Network design under demand uncertainty," in *Proc. Asia-Pacif. Adv. Netw. Meet.*, 2008.

[13] M. Kodialam, T. Lakshman, and S. Sengupta, "Efficient and robust routing of highly variable traffic," presented at the HotNets III, San Diego, CA, 2004.

[14] R. Zhang-Shen and N. McKeown, "Designing a predictable internet backbone network," presented at the HotNets III, San Diego, CA, USA, 2004.

[15] D. Bertsimas and M. Sim, “The price of robustness,” *Oper. Res.*, pp. 35–53, 2004.

[16] J. Cao, D. Davis, S. Wiel, and B. Yu, “Time-varying network tomography: Router link data,” *J. Amer. Statist. Assoc.*, vol. 95, no. 452, 2000.

Matthew Johnston (S’10) received the B.S. degree in electrical engineering and computer science from the University of California, Berkeley, CA, USA in 2008 and the S.M. degree in electrical engineering from the Massachusetts Institute of Technology (MIT), Cambridge, MA, USA, in 2010. He is currently working toward the Ph.D. degree at MIT with the Laboratory for Information and Decision Systems. His research interests include network control, wireless resource allocation, and network survivability.

Hyang-Won Lee (S’02–AM’08) received the B.S., M.S., and Ph.D. degrees in all electrical engineering and computer science from the Korea Advanced Institute of Science and Technology (KAIST), Daejeon, Korea, in 2001, 2003, and 2007, respectively. He was a Postdoctoral Research Associate at the Massachusetts Institute of Technology during 2007–2011, and a Research Assistant Professor at KAIST between 2011 and 2012. He is currently an Assistant Professor in the Department of Internet and Multimedia Engineering, Konkuk University, Seoul, Korea. His research interests include congestion control, stochastic network control, robust optimization and reliable network design.

Eytan Modiano (S’90–M’93–SM’00–F’12) received the B.S. degree in electrical engineering and computer science from the University of Connecticut, Storrs, CT, USA, in 1986 and the M.S. and Ph.D. degrees, both in electrical engineering, from the University of Maryland, College Park, MD, USA, in 1989 and 1992, respectively. He was a Naval Research Laboratory Fellow between 1987 and 1992 and a National Research Council Postdoctoral Fellow during 1992–1993. Between 1993 and 1999, he was with MIT Lincoln Laboratory where he was a Project Leader for Massachusetts Institute of Technology (MIT) Lincoln Laboratory’s Next Generation Internet (NGI) project. Since 1999, he has been on the faculty at MIT, where he is a Professor in the Department of Aeronautics and Astronautics and the Laboratory for Information and Decision Systems (LIDS). His research interests include communication networks and protocols with emphasis on satellite, wireless, and optical networks.

He is currently an Associate Editor for the IEEE/ACM TRANSACTIONS ON NETWORKING. He had served as an Associate Editor for the IEEE TRANSACTIONS ON INFORMATION THEORY, and as a Guest Editor for the IEEE JSAC SPECIAL ISSUE ON WDM NETWORK ARCHITECTURES; the *Computer Networks Journal special issue on Broadband Internet Access*; the *Journal of Communications and Networks special issue on Wireless Ad-Hoc Networks*; and for the IEEE JOURNAL OF LIGHTWAVE TECHNOLOGY special issue on Optical Networks. He was the Technical Program Cochair for the IEEE Wiopt 2006, the IEEE Infocom 2007, and the ACM MobiHoc 2007. He is a Fellow of the IEEE and Associate Fellow of the AIAA.