

Cooperative Data-Optical InterNetworking: Distributed Multi-Layer Optimization

Anwar Elwalid Debasis Mitra Qiong Wang
 Bell Laboratories, Lucent Technologies, Murray Hill NJ 07974
 {anwar mitra,chiwang}@lucent.com

Abstract - We present a novel approach for joint optical network provisioning and IP traffic engineering, in which the IP and optical networks collaboratively optimize a combined objective of network performance and lightpath provisioning cost. We develop a framework for distributed multi-layer optimization. Our framework is built upon the IP-over-Optical (IPO) overlay model, where each network domain has a limited view of the other. Our formulation allows the two domains to communicate and coordinate their decisions through minimal information exchange. Our solution is based on a novel application of Generalized Bender’s Decomposition, which divides a difficult global optimization problem into tractable sub-problems, each solved by a different domain. The procedure is iterative and converges to the global optimum. We present case studies to demonstrate the efficiency and applicability of our approach in various networking scenarios. Our work builds a foundation for “multi-layer” grooming, which extends traditional grooming in the optical domain to data networks. The data networks are active participants in the grooming process with intelligent homing of data traffic to optical gateways.

I. INTRODUCTION

We present a novel approach for joint optical network provisioning and IP traffic engineering, in which the IP and optical domains collaboratively optimize a combined objective of network performance and the cost of provisioning capacity. The context for this work is the rapid transitioning of the Internet transport infrastructure towards a model of high-speed router networks that are directly interconnected by reconfigurable optical core networks. The work in this paper is premised on just such a model. This IP over Optical (IPO) architecture, when coupled with the emerging ASON/GMPLS (Automatically Switched Optical Network/Generalized Multi-Protocol Label Switching) optical control plane, offers network operators opportunities for dynamic multi-layer optimization that will give significant savings in capital and operating expenses [3], [4], [8], [12].

This work was supported in part by the Office of Naval Research (ONR) under grant N00014-1-0256.

A key conceptual contribution of the present work is the notion of “multi-layer grooming”, which embraces both the optical and IP layers, whereas conventional grooming functions are confined to the former [7], [18], [22]. The traditional goal of grooming is the minimization of stranded, i.e., unutilized, bandwidth in optical lightpaths. This has also been viewed as optimal packing of the wavelengths. This goal is recognized in the present work. However, the goal here is broader and to achieve this goal we make the data networks active participants in the grooming process. Specifically, intelligent homing of data traffic to optical gateways is an integral mechanism for achieving our overall objectives. However, the homing gains must be weighed against other performance factors in the data networks, such as load balancing. Thus our overall objective stated above unavoidably combines performance in the data as well as the optical networks. The objective reflects the value from carrying data traffic from source to destination, as well as the cost of provisioned wavelengths in the optical core.

However great are the potential benefits of converged data-optical networks, the concept is only interesting if the implementation is scalable and distributed. This paper develops a framework for distributed implementation that converges to the global optimum. The implementation is premised on cooperation between the data and optical networks. More specifically, in each iteration these networks perform local optimizations, which is followed by an exchange of the computational results. A key feature is that the information exchanged is kept to a minimum. Yet the iterations converge to the global optimum, i.e., the solution is as good as if all the networks were administered as a single entity.

Our framework is enabled by GMPLS, which facilitates the convergence of data and optical networks and supports different levels of cooperation and information exchange. At the opposite ends of the integration spectrum are the Peer and Overlay models [3], [17]. In the Peer model, optical and IP nodes act as peers such that a single routing protocol instance runs over both the IP and optical domains. Hence, the optical network elements become IP addressable entities. The advantage of the Peer model is that the

entire network can be managed and traffic engineered as if it is a single network; its drawback is that routing and resource information need to be globally advertised.

In the Overlay model, IP/MPLS routers do not participate in the routing protocol instance that runs among the optical nodes; in particular, the routers are unaware of the topology of the optical domain. The optical network primarily offers high bandwidth connectivity in the form of lightpaths according to a client-server model. A standard GMPLS user-network interface (UNI) based on RSVP-TE ([21]) has been developed at the IETF. The GMPLS UNI enables signaling and information exchange between the IP and the optical domains.

Our framework is built upon the Overlay model. This choice is motivated by the observation that the optical and data networks are typically operated as distinct organizations, either with separate owners or as separate divisions within the same corporate entity. Such an organizational structure is likely to be even more pervasive in the future due to increasing disaggregation of vertically integrated service providers and the emergence of new models, such as the carrier's carrier [2]. The Overlay model of this paper reflects the implications of organizational separation of the optical and individual data networks. We omit the discussion on the Peer model due to space limitations, and observe in passing that the solution procedure presented here is an attractive candidate for numerically solving the Peer model.

When building this framework, we address several important issues that lead to the following features. First, the purpose of the framework is to enable the data and optical networks to cooperatively optimize a given objective, which is broadly defined to be the surplus of the utility from carrying end-to-end traffic demand over the cost of optical lightpath provisioning. A nonlinear utility function is used to measure the value of carrying data traffic from source to destination. We do not confine ourselves to any specific form of the utility function, except for requiring that it is concave and monotonically increasing in the amount of traffic carried. We capture several scenarios in the problem formulation, including random traffic demands with known distributions and price-demand relationships. Consequently, our framework allows many different forms of nonlinear utility functions.

Another notable feature of the framework is the significant separation of scales in the data and optical networks. The bandwidth of links in the data network are of considerably lower capacity. A typical data network rate is T1, DS3 or OC3, while in the optical core it is OC48 or OC 192. On the other hand, the number of routers is typically orders of magnitude greater than the number

of optical cross-connects (OXC). The indivisibility of the wavelength as a unit of provisionable bandwidth in the optical transport network imposes integrality constraints on the decision variables in the optimization problem. This feature adds significantly to the difficulty of solving the optimization problem.

Our framework implements a division of tasks that allows each network to focus on its own domain. Decision-making in the data networks aims to make the most efficient use of resources in the optical network by routing and admission control. Decision-making in the optical network is concerned with providing necessary resources to transport traffic at minimum lightpath provisioning cost. Nevertheless, the division of the tasks does not imply a complete separation of decision-making. The framework facilitates the communication and coordination between different networks to maximize the global objective function. For instance, the optical network has the implied task of inducing the data network to home traffic in such a manner that facilitates efficient packing of wavelengths.

Finally, service providers incur an implied "cost" associated with information transfer between the data and optical networks. In the case when the two networks are organizationally separate, the cost may be in strategic terms, i.e., loss of competitive advantage, as for example, with the transfer of infrastructure capacity information. Another cost to the service provider is the operating expense for collecting and transferring detailed information on the network. For this reason, while information exchange in our scheme is sufficient for achieving global optimality, communications between the networks are kept at a minimum.

The mathematical foundation of our framework is Generalized Bender's Decomposition. Bender's method is well-known for mixed integer linear programming. It has also been proposed previously for the solution of nonlinear programming problems with continuous variables. The Generalized Bender's Decomposition presented in this paper combines a nonlinear objective function with integer variables. Also of note is the mapping of the decomposition into a cooperative and iterative internetworking solution procedure. In each iteration, each data network passes the net and marginal values of the proposed provisioned lightpath bandwidths to the optical network. The optical network passes information on the proposed provisioned lightpath bandwidth to the data networks. The procedure is proven to converge to the global optimum in a finite number of iterations. Moreover, each iteration refines an upper and a lower bound on the global solution. The procedure terminates when the two bounds coincide.

The treatment in this paper is restricted to the case of

a single pattern of end-to-end traffic demands. Nevertheless, we recognize that in reality network traffic patterns are constantly changing. Our framework can be extended to accommodate the latter, starting from decomposing the problem according to scale. For relatively small scale changes, the corresponding network response is only at the data networks. That is, the routing and admission control at the data networks are affected, while the reprovisioning of wavelengths to optical pipes is not necessitated. However, sufficiently large changes in traffic patterns will justify the need for undertaking the relatively heavy load of re-calculating the optical provisioning process and with it, of course, the data network traffic engineering solution as well. An important element of this strategy is the design of automatic thresholds that trigger the latter on the basis of on-line measurements. Such procedures are outside the scope of this paper. However, as the results from the case study in Section V-B shows, the insights from this study are useful and the tools these generate will be essential ingredients of such procedures.

Interaction between different network domains has been discussed in other contexts (e.g., [9], [16]). The main distinction of our work is that we focus on cooperative rather than selfish behavior. We model and analyze cross-layer network cooperation instead of interactions of multiple data networks within the same layer. The techniques developed here are quite likely to be effective for a broad range of applications.

In Section II we describe the data and optical networks under consideration. In Section III we formulate the optimization model. In Section IV we discuss a novel application of Generalized Bender's Decomposition. We present sample results from case studies in Section V and conclude in Section VI.

II. INTERNETWORKING MODEL

The network under consideration is composed of an optical core connected to multiple data subnetworks. The latter are indexed by $j = 1, \dots, J$. Let \mathcal{V}_j be the set of nodes in data network j and \mathcal{V}_o be the set of all optical nodes. Data networks interconnect with the optical core at a set of gateway nodes \mathcal{V}_g where

$$\mathcal{V}_g \equiv \cup_{j=1}^J (\mathcal{V}_j \cap \mathcal{V}_o).$$

The link set of the network is \mathcal{L} , i.e.,

$$\mathcal{L} = \mathcal{L}_1 \cup \dots \cup \mathcal{L}_J \cup \mathcal{L}_o,$$

where $\mathcal{L}_j (j = 1, 2, \dots, J)$ is the set of data links in subnet j and \mathcal{L}_o is the set of optical links. $\mathcal{L}_j (j = 1, 2, \dots, J)$ are pairwise mutually exclusive, as are \mathcal{L}_o and \mathcal{L}_j for all j .

Traffic has its sources and destinations in nodes of the data networks, and it is carried by the composite network. A traffic-carrying route is typically composed of a sequence of data, optical and data network links, in that order. Usually there are many routes from source to destination; however, from policy and technical restrictions only a subset of these routes may be eligible to carry traffic. The collection of eligible routes between a (source,destination) node pair is defined to be the *admissible route set* for the pair.

A. Data Network Model

Let

$$\mathcal{S}_j = \{\sigma : \sigma = (v_1, v_2), v_1, v_2 \in \mathcal{V}_j\}$$

be the set of all (source,destination) node pairs in subnet $j (j = 1, \dots, J)$. Note that we exclude cases where sources and destinations are in different data networks. The symbol for a node pair σ does not indicate the data network to which the node pair belongs, but the context should make it clear. For each $\sigma \in \mathcal{S}_j$, $\mathcal{R}_j(\sigma)$ is the admissible route set. Each route $r \in \mathcal{R}_j(\sigma)$ contains a subset of data network links $l \in \mathcal{L}_j$, and possibly an optical segment. The data networks have no knowledge of the optical network's internal structure. Therefore, the entire optical segment is treated as a single optical pipe that connects the (ingress, egress) gateway nodes pairs. Define

$$\mathcal{G} = \{\varsigma : \varsigma = (\varsigma_1, \varsigma_2), \varsigma_1, \varsigma_2 \in \mathcal{V}_o\}$$

as the set all gateway node pairs. We say $\varsigma \in r$ if route r enters and exits the optical core at node pair ς . We assume for any route r , $|r \cap \mathcal{G}| \leq 1$, i.e., a route enters the optical core at most once.

Let y_σ be the total bandwidth provisioned to carry traffic between the node pair σ and x_r be the bandwidth provisioned on route $r \in \mathcal{R}_j(\sigma)$. Then the following constraints apply for end-to-end (E2E) routing:

$$y_\sigma = \sum_{r \in \mathcal{R}_j(\sigma)} x_r, \quad \sigma \in \mathcal{S}_j, \quad j = 1, 2, \dots, J. \quad (1)$$

Furthermore, denote the capacity of data link $l \in \mathcal{L}_j$ by q_l and bandwidth between the gateway pair ς provisioned to data subnet j by $w_{\varsigma,j}$, then

$$\begin{aligned} \sum_{r \in \mathcal{R}_j: l \in r} x_r &\leq q_l, \quad l \in \mathcal{L}_j, \quad j = 1, 2, \dots, J. \\ \sum_{r \in \mathcal{R}_j: \varsigma \in r} x_{r,j} &\leq w_{\varsigma,j}, \quad \varsigma \in \mathcal{G}, \quad j = 1, 2, \dots, J. \end{aligned} \quad (2)$$

In this paper, q_l are given parameters. The size of optical pipes, $w_{\varsigma,j}$, are controlled by the optical core, as explained in the following subsection.

We end this subsection by noting that in the following discussions it is convenient to express the two systems of inequalities in (2) in matrix form:

$$\mathcal{A}_j \vec{x}_j \leq \vec{q}_j, \quad \mathcal{B}_j \vec{x}_j \leq \vec{w}_j, \quad (3)$$

where $\vec{x}_j \equiv (x_r, r \in \mathcal{R}_j)$, $\vec{q}_j \equiv (q_l, l \in \mathcal{L}_j)$, and $\vec{w}_j \equiv (w_{\varsigma,j}, \varsigma \in \mathcal{G})$.

B. Optical Network Model

The optical core decides on $w_{\varsigma,j}$, the amount of bandwidth between gateway node pair ς to be allocated to the data network j . Assume that the core treats all traffic between gateway pairs ς uniformly, i.e., obliviously of their source-destination node pairs in the data subnets. Then for the optical core,

$$W_{\varsigma} = \sum_{j=1}^J w_{\varsigma,j}$$

is the end-to-end demand that is identified with an optical pipe connecting the node pair ς . It is realized by bandwidth provisioned in possibly multiple paths, where paths (abbreviated from lightpaths) in the optical network are analogous to routes in the data networks. We define $\mathcal{P}(\varsigma)$ to be the set of candidate paths between node pair ς and let $\chi_p (p \in \mathcal{P}(\varsigma))$ be the bandwidth provisioned on path p . It follows that

$$\sum_{p \in \mathcal{P}(\varsigma)} \chi_p = W_{\varsigma}. \quad (4)$$

Denote the number of wavelengths deployed on link l by z_l , and b as bandwidth per wavelength. We omit the straightforward generalization that allows b to depend on l . Thus z_l is a non-negative integer, and bz_l is the aggregate bandwidth available on the optical link l . The sum of bandwidths provisioned on all paths that use link l cannot exceed the link capacity, i.e.,

$$\sum_{p:l \in p} \chi_p \leq bz_l, \quad l \in \mathcal{L}_o. \quad (5)$$

Let c_l be the cost of deploying each wavelength on link l . Thus, if z_l wavelengths are deployed, the deployment cost for link l is $c_l z_l$.

In the above formulation, while the bandwidth availability implied by the number of wavelengths deployed on each link is a key feature, the identities of the wavelengths are not tracked from link to link. This is done deliberately. First, the methodology introduced in Section II-C, on "routing constraints in the optical network" is sufficiently general to allow the latter feature to be taken into account if necessary. Second, this issue is related to wavelength conversion on which a great deal has been written

and much is already known. Introducing this topic on an already over-extended paper would place an unreasonable burden on the reader.

C. Routing Constraints in the Optical Network

Equations (4) and (5) are conditions that have to be satisfied by the gateway-to-gateway ($G2G$) routing. In general, the routing may be subject to other constraints implied by the diverse capabilities of the optical nodes to groom (i.e., unpack/pack between lower-rate and higher-rate data streams) and switch traffic. These capabilities depend on the presence of sophisticated electronics and come at considerable cost. Nodes at the gateways to the optical network are more likely to have these capabilities than internal nodes. We will consider several configurations in a unified framework. At one extreme, every optical node has these capabilities. In this case, traffic can be arbitrarily split at all nodes and routed on different paths between a pair of optical nodes. At the other extreme, no node has these capabilities and the gateway has to select one path from the candidates in the admissible set to carry all traffic to the destination. An interesting intermediate case is where the gateway nodes have the capabilities to split and switch traffic, while the internal nodes do not.

To reflect these instances of constraints in a unified formulation, let $\Omega(\vec{W})$ be the set of all wavelength configurations, $\{z_l, l \in \mathcal{L}_o\}$, each of which makes \vec{W} realizable. The set $\Omega(\vec{W})$ is the collection of all non-negative vectors that satisfy three conditions. These conditions include (4) and (5); additionally, another set of defining conditions of $\Omega(\vec{W})$ arises in case the use of a path excludes the use of others. This exclusion property can be given as a condition on the paths' indicator functions (a path's indicator function takes value 1 if the path is used and 0 otherwise).

$$\sum_{p \in P'} 1(\chi_p > 0) \leq 1, \quad (6)$$

where P' can be any collection of paths in which the use of one path excludes the use of any other path.

For example, in the first of the two aforementioned extreme examples, $\Omega(\vec{W})$ is defined as the collection of \vec{z} that are feasible for both (4) and (5), i.e., there is no need for any exclusion condition (6). In contrast, in the second extreme case, $\Omega(\vec{W})$ is restricted by an exclusion condition,

$$\sum_{p \in \mathcal{P}(\varsigma)} 1(\chi_p > 0) = 1, \quad (7)$$

which indicates that only one path (from the admissible route set) can be used between a gateway node pair. Now consider the example in which only gateway nodes have




Fig. 1. An Example of Routing Constraints

grooming and switching capabilities, so that traffic may be split at the ingress gateway, but not at any internal node. Then the definition of $\Omega(\vec{w})$ is determined by (4), (5), and the exclusion condition which applies on subsets of all admissible routes that have a common initial link. For example, consider the gateway pair (A, C) in Figure 1. In the figure, A and C are gateway nodes and E is an internal optical network node. Suppose $\mathcal{P}(A, C)$ has three paths $A - E - F - C$, $A - E - D - C$, $A - B - C$. Because the internal node E has no grooming capability, the usages of routes $A - E - F - C$ and $A - E - D - C$ are mutually exclusive. Therefore, the following exclusion condition applies

$$\sum_{p \in \mathcal{P}(A, C|E)} 1(\chi_p > 0) \leq 1, \quad (8)$$

where $\mathcal{P}(A, C|E) = \{A - E - F - C, A - E - D - C\}$

In conclusion, the concept of $\Omega(\vec{W})$ is versatile and allows quite general routing constraints, including mutual exclusion, to be modeled. However, it should be noted that additional exclusion conditions in $\Omega(\vec{W})$ require the introduction of binary variables that makes the subsequent optimization problem more burdensome.

Table I gives the key variables that have been introduced in this section.

TABLE I
KEY ROUTING VARIABLES

	Traffic Demand	Route/Path Sets	Provisioned Capacity on Rte./Path
E2E	y_σ	$\mathcal{R}_j(\sigma)$	x_r
G2G	W_ζ	$\mathcal{P}(\zeta)$	χ_p

III. OPTIMIZATION PROBLEM FORMULATIONS

In this section, we formulate various optimization problems, which have in common the objective of maximizing the surplus of total utility over the deployment cost in the optical transport network. We give various choices for the

utility function, which in all cases reflects the value of carrying traffic. The deployment cost is proportional to the number of wavelengths deployed on the optical links.

A. Utility Functions, Traffic Demand Characterizations

The total utility is the sum of the utility for each data network, where the utility $U_j(\vec{y}_j)$ is a function of provisioned bandwidth $\vec{y}_j = (y_\sigma, \sigma \in \mathcal{S}_j)$, $(j = 1, 2, \dots, J)$ (see (1)). Our model accommodates various traffic demand characterizations, as well as utility functions.

1. Suppose that a “traffic matrix” specifies demand between source-destination pairs, i.e.,

$$\vec{D}_j = (D_\sigma, \sigma \in \mathcal{S}_j), \quad (9)$$

where D_σ is the deterministic traffic demand between the node pair $\sigma \in \mathcal{S}_j$. In this model the carried traffic for node pair σ is the minimum of provisioned bandwidth, y_σ , and the traffic demand D_σ . Since it is waste of resource to provision bandwidth beyond the demand, we require

$$y_\sigma \leq D_\sigma, \quad \sigma \in \mathcal{S}_j,$$

which implies that the carried traffic equals the provisioned bandwidth. The utility for the j^{th} data network is

$$U_j(\vec{y}_j) = \sum_{\sigma \in \mathcal{S}_j} \pi_\sigma y_\sigma, \quad (10)$$

which is the weighted sum of carried traffic between source-destination pairs. We may interpret π_σ as the revenue per unit of demand carried between node pair σ , in which case (10) gives the total revenue.

2. We may also adopt the concurrent flow problem formulation, induce fairness among source-destination pairs, and let

$$U_j(\vec{y}_j) = \min_{\sigma \in \mathcal{S}_j} \frac{y_\sigma}{D_\sigma}. \quad (11)$$

3. Demands D_σ may also be given as a set of random variables characterized by their distribution functions [20], i.e.,

$$F_\sigma(d) = \Pr(D_\sigma \leq d). \quad (12)$$

In this case, we define the utility function to be

$$U_j(\vec{y}_j) = \sum_{\sigma \in \mathcal{S}_j} \pi_\sigma \left[\int_0^{y_\sigma} \theta dF_\sigma(\theta) + y_\sigma \bar{F}_\sigma(y_\sigma) \right], \quad (13)$$

where $\bar{F}_\sigma(d) = 1 - F_\sigma(d)$, the bracketed term is the expected carried traffic for node pair σ and $U_j(\vec{y}_j)$ is the expected total revenue. See [14] for details and properties.

4. Price-demand relationship may also be incorporated into the model. Let the price $p_\sigma(y_\sigma)$ be a decreasing function of the carried traffic y_σ and define the utility to be

the revenue, which is the product of price and demand, summed over all node pairs, i.e.,

$$U_j(\vec{y}_j) = \sum_{\sigma \in \mathcal{S}_j} y_\sigma p_\sigma(y_\sigma). \quad (14)$$

It is required that $p_\sigma(y_\sigma)$ is such that $y_\sigma p_\sigma(y_\sigma)$ is a monotonically increasing, concave function of y_σ . An important example of price-demand relation that has been extensively used is the constant elasticity demand function,

$$p_\sigma = A_\sigma y_\sigma^{-1/\epsilon_\sigma}, \quad (15)$$

where $\epsilon_\sigma > 1$ is the constant price elasticity of demand.

5. With a little variation, we can also use the formulation to address the problem of minimizing aggregate delays subject to the condition that all demand is carried (assuming the problem is feasible). In this case, y_σ , which is given exogenously, is demand that must be carried. This condition is expressed as

$$\sum_{r \in \mathcal{R}_j(\sigma)} x_r = y_\sigma,$$

If each link is approximated by a $M/M/1$ queue, then the aggregate delay in data networking is

$$\sum_{l \in \mathcal{L}_j} \frac{1}{q_l - \sum_{r: l \in r} x_r}.$$

It is desirable to have a smaller value of the above quantity. Correspondingly,

$$U_j(\vec{x}_j) = \sum_{l \in \mathcal{L}_j} \frac{1}{\sum_{r: l \in r} x_r - q_l},$$

which is a monotonically increasing and concave function of x_r . We also require

$$\sum_{r: l \in r} x_r < q_l \quad \forall l \in \mathcal{L}_j.$$

B. The Optimization Problem

The global optimization problem can be formulated as

$$\max_{x_r, y_\sigma, z_l, w_{\varsigma, j}} \left\{ \sum_{j=1}^J U_j(\vec{y}_j) - \sum_{l \in \mathcal{L}_o} c_l z_l \right\} \quad (16)$$

subject to

$$\begin{aligned} \sum_{r \in \mathcal{R}_j(\sigma)} x_r &= y_\sigma \quad \sigma \in \mathcal{S}_j, \quad j = 1, 2, \dots, J, \\ \sum_{r \in \mathcal{R}_j: l \in r} x_r &\leq q_l \quad l \in \mathcal{L}_j, \quad j = 1, 2, \dots, J, \\ \sum_{r \in \mathcal{R}_j: \varsigma \in r} x_r &\leq w_{\varsigma, j} \quad \varsigma \in \mathcal{G}, \quad j = 1, 2, \dots, J, \\ \{z_l, l \in \mathcal{L}_o\} &\in \Omega(\vec{W}), \quad \vec{z} \text{ integral}, \end{aligned} \quad (17)$$

where $\Omega(\vec{W})$ is defined in Section II-C. The utility function $U_j(\cdot)$ in (16) is monotonically increasing and concave. The necessary condition for $\vec{z} \in \Omega(\vec{W})$ is the existence of $\chi_p \geq 0$ ($p \in \mathcal{P}(\varsigma)$) such that

$$\begin{aligned} \sum_{p \in \mathcal{P}(\varsigma)} \chi_p &= W_\varsigma = \sum_{j=1}^J w_{\varsigma, j} \quad \varsigma \in \mathcal{G}, \\ \sum_{p: l \in p} \chi_p &= b z_l \quad l \in \mathcal{L}_o. \end{aligned} \quad (18)$$

Additional conditions reflecting optical routing constraints may also need to be satisfied, as discussed in Section II-C. Note the nonlinear objective function and integer variables in the formulation of the optimization problem.

C. Introduction to Generalized Bender's Decomposition

In the following, we follow Schrijver [19] and Geoffrion [6] to introduce Generalized Bender's Decomposition. Schrijver's treatment is for the mixed integer linear programming problem and Geoffrion's is for nonlinear programming problems with continuous variables. However, our problem involves both integer variables and a nonlinear objective function. The development here is a synthesis of the two approaches.

Consider the following canonical optimization problem,

$$\max_{\vec{y}, \vec{z}} \{U(\vec{y}) - \vec{c}\vec{z} | M\vec{y} \leq \vec{z}, \vec{z} \in Z\} \quad (19)$$

where $U(\vec{y})$ is a concave and monotonically increasing function of \vec{y} , M is a constant matrix, and Z is a finite set of integral vectors to which \vec{z} is restricted. The problem can be decomposed as follows:

$$\max_{\vec{z}} \{G(\vec{z}) - \vec{c}\vec{z}, \vec{z} \in Z\}, \quad (20)$$

where

$$G(\vec{z}) \equiv \max_{\vec{y}} \{U(\vec{y}) | M\vec{y} \leq \vec{z}\}. \quad (21)$$

We shall refer to the problems in (20) and (21) as master and slave, respectively. In general $G(\vec{z})$ is only given implicitly, so that (20) cannot be solved directly. Generalized Bender's Decomposition is an approach to deal with this problem. In this approach a sequence of slave problems is solved for different values of \vec{z} . The solutions to these problems are used to construct *Bender's cuts* that define approximations to $G(\vec{z})$.

Let \vec{z}_k ($k = 1, 2, \dots, K$) be a set of given values in Z . Suppose \vec{y}_k maximizes (21) for $\vec{z} = \vec{z}_k$. Then (\vec{y}_k, \vec{z}_k) $k = 1, 2, \dots, K$ is a feasible solution to (19). Therefore,

$$\max_{k=1, \dots, K} \{U(\vec{y}_k) - \vec{c}\vec{z}_k\} \quad (22)$$

gives a lower bound to the solution of the original problem, and moreover the bound is non-decreasing in K .

Furthermore, by the duality theorem of convex programming,

$$G(\vec{z}) \leq U(\vec{y}_k) + \vec{\lambda}_k(\vec{z} - M\vec{y}_k), \quad \forall \vec{z} \in Z, \quad (23)$$

where $\vec{\lambda}_k$ denote the vector of Lagrange multipliers associated with constraints $M\vec{y}_k \leq \vec{z}$. Let

$$\Gamma_k = U(\vec{y}_k) - \vec{\lambda}_k M\vec{y}_k. \quad (24)$$

A *Bender's cut* is defined as

$$G(\vec{z}) \leq \Gamma_k + \vec{\lambda}_k \vec{z}, \quad k = 1, 2, \dots, K. \quad (25)$$

An upper bound to the solution of the original problem (19) is obtained by solving the following surrogate problem

$$\max_{\gamma, \vec{z}} \{\gamma - \vec{c}\vec{z} \mid \gamma \leq \Gamma_k, \quad k = 1, 2, \dots, K, \quad \vec{z} \in Z\}. \quad (26)$$

The upper bound decreases as K increases since increasing the number of Bender's cuts reduces the feasible region of the solution.

Generalized Bender's Decomposition is an iterative process. In each iteration, a new instance of the slave problem is solved. This solution is used to construct a new Bender's cut that augments the set of previous cuts. Each expansion of the set of cuts defines a refinement to the approximation of $G(\vec{z})$ for which the corresponding master problem is next solved. The process continues until the decreasing upper bound and the increasing lower bound coincide, which then defines the optimal solution.

Geoffrion has shown in [6] Theorem 2.4 that if Z is in a finite discrete set and $U(\vec{y})$ is concave and defined on a convex, compact set, then the procedure is guaranteed to terminate in a finite number of iterations.

D. A Simple Example

To illustrate how Generalized Bender's Decomposition applies to data-optical internetworking, consider the simple case of a single optical link and no data links. The optical link connects a pair of nodes. Then the problem in (19) is reduced to

$$\max_{y, z} \{U(y) - cz \mid y \leq bz, \quad z \text{ integral}\} \quad (27)$$

and the master and slave problems in (20) and (21) are

$$\max_z \{G(z) - cz\} \quad (28)$$

and

$$G(z) = \max_y \{U(y) \mid y \leq bz\}, \quad (29)$$

respectively. Both $G(z)$ and cz are illustrated in Figure 2. Note that $G(z)$ is explicitly shown in the figure, while in reality it is defined implicitly. The problem in (27) is to find z that maximizes the vertical distance between $G(z)$ and cz .

We assume that there exists a sufficiently large value y_{\max} such that

$$U(y) \leq \bar{U} \text{ for } y \geq 0 \text{ and } U(y) = \bar{U} \text{ if } y \geq y_{\max},$$

i.e., once the provisioned capacity is sufficiently large, adding new capacity will not improve the utility.

The procedure is as follows. First, let z_1 be some value greater than y_{\max}/b . Solving (29) with $z = z_1$ gives a solution that

$$y_1 = z_1, \quad U(y_1) = \bar{U}, \quad \text{and } \lambda_1 = \frac{\partial U(y)}{\partial y} = 0.$$

From (24) and (25), the first Bender's cut is

$$G(z) \leq \Gamma_1 = \bar{U},$$

which is the (dashed) horizontal line in Figure 2. Next the surrogate problem,

$$\max_{\gamma, z} \{\gamma - cz \mid \gamma \leq \Gamma_1, \quad z \text{ integral}\},$$

is solved to give the solution $z = 0$. Solving (29) with this value gives

$$y_2 = 0, \quad U(y_2) = 0, \quad \text{and } \lambda_2 > 0,$$

and the second cut

$$G(z) \leq \lambda_2 bz,$$

which is also shown in the figure. Solving the surrogate problem again with both cuts enforced

$$\max_{\gamma, z} \{\gamma - cz \mid \gamma \leq \Gamma_1, \quad \gamma \leq \lambda_2 bz, \quad z \text{ integral}\},$$

gives z_3 , shown in Figure 2, which generates a third cut, $\Gamma_3 + \lambda_3 bz$.

Notice in the figure that as the number of cuts increases, the approximation to $G(\vec{z})$ becomes increasingly refined. As previously stated, this process continues until the upper and lower bounds converge.

IV. DISTRIBUTED INTERNETWORKING PROCEDURE

We now explain how we obtain a distributed internetworking solution procedure based on Generalized Bender's Decomposition. Each data network j ($j = 1, 2, \dots, J$) makes admission and routing decisions based on information from the optical network on capacities provisioned on




Fig. 2. Illustration of Generalized Bender's Decomposition for a Simple Example

optical pipes. A parsimonious representation of the result is transferred to the optical network, which uses it to make decisions on the provisioning of wavelengths on optical links and bandwidth on optical pipes. In this procedure, the master problem in the decomposition to the optimization problem in (16) and (17) is

$$\max_{\vec{z}, \vec{w}} \sum_{j=1}^J \{G_j(\vec{w}_j) + \vec{\lambda}_j \vec{w}_j - \vec{c} \vec{z} \mid \vec{z} \in \Omega(\vec{W}), \vec{z} \text{ integral}\}. \quad (30)$$

For each data network $j = 1, 2, \dots, J$, there is a corresponding slave problem

$$G_j(\vec{w}_j) \equiv \{\max_{\vec{x}_j, \vec{y}_j} \{U_j(\vec{y}_j) \mid \mathcal{A}_j \vec{x}_j \leq \vec{q}_j, \mathcal{B}_j \vec{x}_j \leq \vec{w}_j, y_\sigma = \sum_{r \in \mathcal{R}_j(\sigma)} x_r\}\}, \quad (31)$$

where \mathcal{A}_j and \mathcal{B}_j are defined in (3). Decision-making in the optical and data networks is associated with solving the master and slave problems, respectively. The procedure is iterative with information exchange (to be described below) at each iteration between each data network and the optical network. Figure 3 shows a schematic of this procedure.

A. Data Network Optimizations

The slave problem (31) is solved by the j^{th} data network for fixed value of \vec{w}_j . This optimization maximizes the utility, $U_j(\vec{y}_j)$, from carrying traffic, with respect to admission control (\vec{y}_j) and routing (\vec{x}_j), based on the fixed capacity of data network links (\vec{q}_j) and the capacity of the optical pipes (\vec{w}_j). The latter information is transferred from the optical network.

Recall that $U_j(\vec{y}_j)$ is a concave increasing function of \vec{y}_j , so (31) is a concave maximization problem and can be

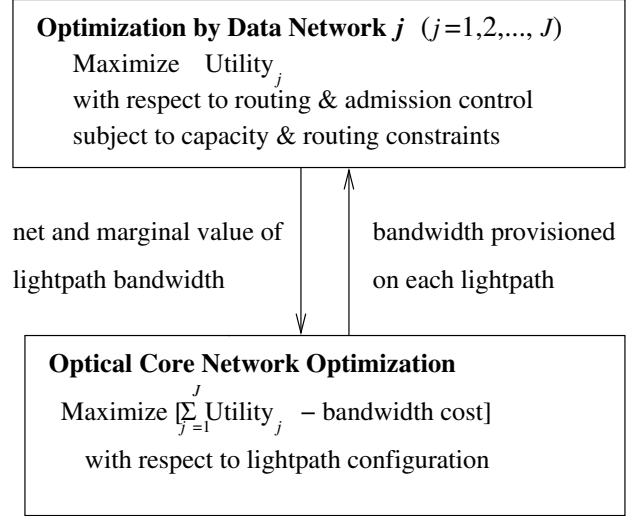


Fig. 3. Schematic of the Distributed Internetworking Procedure

transformed into

$$\min_{\vec{\lambda}_j} \{\max_{\vec{x}_j, \vec{y}_j} \{U_j(\vec{y}_j) + \vec{\lambda}_j(\vec{w}_j - \mathcal{B}_j \vec{x}_j) \mid \mathcal{A}_j \vec{x}_j \leq \vec{q}_j, y_\sigma = \sum_{r \in \mathcal{R}_j(\sigma)} x_r\}\}, \quad (32)$$

where $\vec{\lambda}_j$ is the Lagrange multipliers associated with the capacity constraints on the optical pipes, $\mathcal{B}_j \vec{x}_j \leq \vec{w}_j$. By the KKT optimality condition,

$$U_j^* = U_j^* + \vec{\lambda}_j(\vec{w}_j - \mathcal{B}_j \vec{x}_j) = (U_j^* - \vec{\lambda}_j \mathcal{B}_j \vec{x}_j) + \vec{\lambda}_j \vec{w}_j \quad (33)$$

where U_j^* is the optimal solution. We interpret $\lambda_{j,\varsigma}$ as the marginal value of the optical pipe between gateway node pair ς , $\vec{\lambda}_j \vec{w}_j$ as the total value of optical pipes, and

$$\Gamma_j \equiv U_j^* - \vec{\lambda}_j \mathcal{B}_j \vec{x}_j$$

as the “net value” of the data network j , defined as the surplus of the utility of the network over the “shadow cost” of using the optical network pipes.

B. Optical Network Optimizations

The master problem (30) is solved by the optical network; it incorporates information provided by all the data networks. The information from the j^{th} network is on the net (Γ_j) and the marginal ($\vec{\lambda}_j$) values of the optical capacity to this data network. These values represent minimal and necessary feedback from the data networks, based on which the optical network optimizes its internal design problem for provisioning, without knowing details of any data network's topology, configurations, capacities and traffic demands.

The optical network solves the master problem in the form of

$$\begin{aligned}
 & \max_{\gamma, \vec{w}_j, \vec{z}} (\gamma - \vec{c}\vec{z}) \\
 \text{s. t. } & \gamma \leq \sum_{j=1}^J (\Gamma_{k,j} + \vec{\lambda}_{k,j} \vec{w}_j) \quad k = 1, 2, \dots, K. \\
 & \sum_{j=1}^J w_{j,s} = W_s, \quad \vec{z} \in \Omega(\vec{W}), \quad \vec{z} \text{ integral.} \quad (34)
 \end{aligned}$$

C. Solution Procedure with Information Exchange

To summarize, the internetworking procedure has the following steps.

1. Let K denote the iteration index. The optical network starts from an initial provisioning solution, i.e., wavelengths on each link and bandwidth allocation to each path in conformance with exclusion conditions, from which the capacity of optical pipes are determined. Sizes of optical pipes, \vec{w} , are communicated to the data networks.
2. Data network j ($j = 1, 2, \dots, J$) makes admission and routing decisions to optimize its utility function for the given size of optical pipes. The values of $\Gamma_{K,j}$ and $\vec{\lambda}_{K,j}$ that are obtained from this optimization are transferred to the optical network. These values generate a new cut that augments the set of cuts derived from all the previous iterations. The augmented set is used by the optical network for solving the master problem in the next iteration.
3. The optical network obtains a new provisioning solution by solving the problem in (34). The values \vec{w}_j that are obtained from the optimization are transferred to data network j ($j = 1, 2, \dots, J$).
4. Increment $K = K + 1$. Repeat steps 2 and 3 until the termination condition for Generalized Bender's Decomposition stated in subsection (III-C) is satisfied.

V. CASE STUDIES

In this section, we present numerical case studies based on the network topology shown in Figure 4. Circles denote data network nodes, squares denote gateway nodes, and ovals denote switching nodes inside the optical network. Bandwidth on each data link is fixed. The links that connect data and optical nodes have 20 units of capacity and other links that connect data nodes have 10 units. Wavelength deployment cost is assumed to be 5, 10 or 15, depending on the location of the optical link. Each wavelength carries 40 units of capacity.

A. Case 1: Convergence

In the first example, we consider one data network connected to the optical core. The utility for the service




Fig. 4. Network Diagram

provider, $U_\sigma(y_\sigma)$, is defined to be the revenue $y_\sigma p_\sigma$, where y_σ and p_σ are the volume and the unit price for carried demand for each node pair σ , respectively. In general, p_σ is a decreasing function of y_σ . Here we assume that the relationship is characterized by the well-known function,

$$y_\sigma = A_\sigma p_\sigma^{-\epsilon}, \quad (35)$$

where the parameter ϵ is the constant price elasticity of demand and is set at 1.5. It reflects the rate of traffic demand change with respect to price change. A_σ are scalars that parameterize the potential demand volume, and their values are randomly generated in the range between 10 and 70. The total utility is the sum of revenue over all node pairs,

$$\sum_{\sigma} U_\sigma(y_\sigma) = \sum_{\sigma} y_\sigma p_\sigma = \sum_{\sigma} A_\sigma^{1/\epsilon} y_\sigma^{1-1/\epsilon}. \quad (36)$$

Figure 5 shows the performance of the procedure described in Section IV for this example. At each iteration we obtain an upper bound on the optimal solution and a feasible solution. The figure shows rapid convergence of the two values. In the following we normalize by dividing the difference between the upper bound and the feasible solution by the former. This normalized difference is 62% initially, and drops to less than 5% after the first ten iterations. The observed convergence indicates that the iteration converges to the global optimal solution.

Figure 6 shows the configuration of the optical network in the final solution. One wavelength (with capacity 40) is installed on all optical links except the link between nodes D and T. As can be seen in the figure, these capacities are used to configure optical lightpath between nine out of the ten gateway pairs. The bottom figure shows the size of each lightpath that is communicated to the IP network. The top figure shows internal routing of these paths that is kept within the optical domain. The lightpaths are shown as dashed lines in the figure, and the accompanying numbers

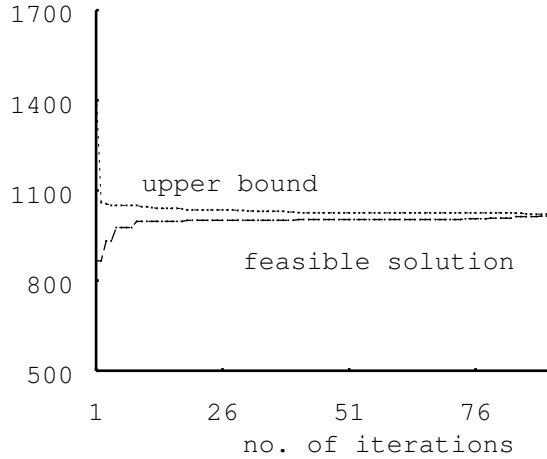


Fig. 5. Convergence of the Algorithm

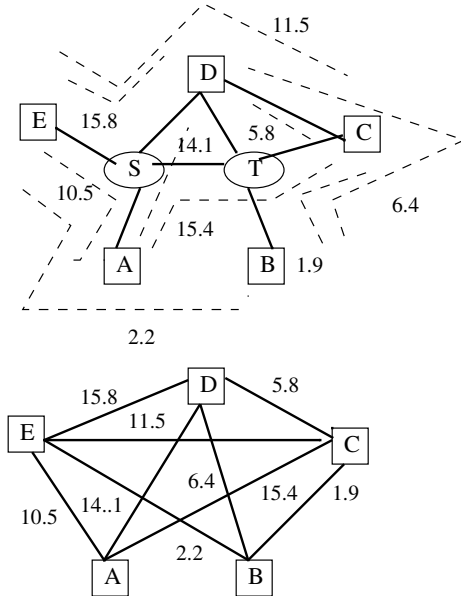


Fig. 6. Optical Network Configuration. The bottom figure shows the lightpath size that is communicated to the IP network; the top figure shows the internal routing of lightpaths.

indicate the capacity of the lightpaths. Note that bandwidth provisioning is not unique. It is possible to provision bandwidth to the path connecting gateway pair (C,E) on both routes C-T-S-E and C-D-S-E, instead of provisioning all bandwidth on the latter as in the displayed solution. The new provisioning arrangement reduces the burden on the heavily loaded link S-D and spreads the load on links S-T and T-C, which have the same capacity as link S-D but carry less traffic. However, doing so will require traffic splitting at node S, which may complicate the management of the optical network. In our scheme, the optical network deals with this tradeoff without burdening the data network.

We use the above procedure to conduct experiments to identify changes in optical network configuration as wavelength capacity increases and cost of unit bandwidth decreases. We start from the base case and scale capacity per wavelength by a factor of k_1 and cost per wavelength by a factor k_2 . We let $k_2/k_1 < 1$ to reflect economy of scale, i.e., the cost per unit of bandwidth decreases as bandwidth per wavelength increases. Capacity on each lightpath for different values of k_1, k_2 is shown in Table II. In the base case ($k_1 = k_2 = 1$), many lightpaths are configured in the optical core. As k_1, k_2 increase, the core bandwidth has to be deployed in increasingly large bundles. As a result, only a small number of paths are provisioned, reflecting a higher degree of capacity concentration in the core.

G2G-pair	$k_1 = 1$ $k_2 = 1$	$k_1 = 1.5$ $k_2 = 1.2$	$k_1 = 2$ $k_2 = 1.5$	$k_1 = 3$ $k_2 = 2$
(A,C)	15.4	25.6	30	50
(A,D)	14.1	20.3	50	42
(A,E)	10.5	14.1	0	0
(B,C)	1.9	5.0	0	0
(B,D)	6.4	9.1	0	0
(B,E)	2.2	0	0	0
(C,D)	5.8	14.2	15.7	0
(C,E)	11.5	0	0	0
(D,E)	15.8	0	0	0

TABLE II
CHANGE OF OPTICAL NETWORK CONFIGURATION WITH
WAVELENGTH GRANULARITY

B. Case 2: Random Demands and Shadow Costs

In the solution described earlier, shadow costs of optical pipes, $\bar{\lambda}$, are critical quantities that indicate marginal values of the optical bandwidth to the data networks. The passing of these values by a data network to the optical core coordinates separate optimizations performed by the two networks. Note that the concept of shadow costs has been widely adopted in distributed network management and utility maximization that involve optimal scheduling, routing, and admission decisions [10], [11]. The approach developed below presents a novel use of shadow costs for capacity deployment and inter-network coordinations.

In this case study we demonstrate that the shadow cost not only provides the basis for optimizing wavelength deployment and bandwidth provisioning in the optical network, but can also determine when re-optimizations should take place. We first solve a base case that optimizes the optical network configuration for some given

demands. We keep this configuration fixed while imposing various changes of the traffic demand to the data network. The question is whether the optical network needs to be re-configured to accommodate these changes. We show that shadow costs provide an excellent indicator to answer this question. Specifically, if changes in demand cause large changes in shadow costs, then re-optimization is in order. Otherwise, the changed demand can be adequately accommodated by the existing optical network configuration.

We consider one data network and use the same topology as in the previous example. However, we switch to different demand characterizations and utility function from the preceding case in Section V-A. We let the demand volume between a node pair σ be random and characterized by the following *truncated Gaussian distribution*

$$F_\sigma(d) = \Pr\{D_\sigma \leq d\} = \int_{-\infty}^d \frac{e^{-(x-\mu_\sigma)^2/2s_\sigma^2}}{\sqrt{2\pi}s_\sigma H_\sigma} dx, \quad (37)$$

for $d \geq 0$. The normalizing constant,

$$H_\sigma = \frac{Erfc(-\tau_\sigma)}{2}, \quad \tau_\sigma = \frac{\mu_\sigma}{\sqrt{2}s_\sigma}, \quad (38)$$

and $Erfc() = 1 - Erf()$, where μ_σ and s_σ are the mean and standard deviation of the untruncated normal distribution. Under certain conditions, e.g., when the ratio of μ_σ to s_σ is sufficiently large, these values are also good approximations to the mean and standard deviation of the above truncated distribution. The utility to be maximized is the expected revenue, given by (also see (13)),

$$\begin{aligned} U(\vec{y}) &= \sum_{\sigma} \pi_{\sigma} E[\min(y_{\sigma}, D_{\sigma})] \\ &= \sum_{\sigma} \pi_{\sigma} \left[\int_0^{y_{\sigma}} \theta dF_{\sigma}(\theta) + y_{\sigma} \bar{F}_{\sigma}(y_{\sigma}) \right], \end{aligned}$$

where π_{σ} is the revenue per unit of traffic carried between node pair σ . We generate values of μ_{σ} , s_{σ} , and π_{σ} by letting

$$\mu_{\sigma} = \bar{\mu} + \delta_{\sigma}^{\mu}, \quad s_{\sigma} = \bar{s} + \delta_{\sigma}^s, \quad \pi_{\sigma} = \bar{\pi} + \delta_{\sigma}^{\pi}. \quad (39)$$

where $\bar{\mu}$, \bar{s} , $\bar{\pi}$ are given constants, and δ_{σ}^{μ} , δ_{σ}^s , δ_{σ}^{π} are values generated randomly to reflect differences in demand distribution and unit revenue between different node pairs.

In the base case,

$$\bar{\mu} = 5, \quad \bar{s} = 2, \quad \text{and} \quad \bar{\pi} = 8, \quad (40)$$

and δ_{σ}^{μ} , δ_{σ}^s , and δ_{σ}^{π} are assumed to be uniformly distributed over $[-0.5, 0.5]$, $[-0.5, 0.5]$, and $[0, 4]$, respectively. The solution to the problem from our procedure is shown in Table III. The table gives both the optimal number of

link	# of wavelength	G2G-pair	bandwidth
A-S	4	(A,B)	0
B-T	1	(A,C)	60
C-D	1	(A,D)	60
C-T	3	(A,E)	40
D-S	2	(B,C)	6
D-T	0	(B,D)	4
E-S	3	(B,E)	10
S-T	3	(C,D)	36
		(C,E)	50
		(D,E)	20

TABLE III
OPTIMAL CONFIGURATION OF THE OPTICAL CORE

wavelengths deployed on each link as well as the optimal amount of bandwidth provisioned on each optical pipe.

Next we keep the configuration unchanged but allow demands to deviate from the base case in the following three different ways.

1. We increase the mean μ_{σ} by 20% to 100% at 20% increment to simulate a systematic increase of demands between all node pairs.
2. We increase s_{σ} by the same percentages so that the average volume of traffic demand between each node pair does not change but the demand uncertainty increases.
3. We keep both $\bar{\mu}$ and \bar{s} fixed but enlarge the support for δ_{σ}^{μ} by the same percentages as above. In this case, the aggregate demand in the data network remains unchanged, but the differences in mean demand between node pairs increase.

For each scenario, we solve the data network optimization problem (31) in Section IV (with the configuration of the optical network fixed as before). Figure 7 shows the resulting shadow costs λ_{ζ} , averaged over all optical gateway node pairs ζ .

Recall that shadow costs reflect marginal values of bandwidth of the optical pipes to the data network. The figure shows that the shadow costs are quite sensitive to the first type of demand change: when the demand increases across all node pairs of the data network, then there is a clear need for more optical bandwidth, which is reflected in the increases of the shadow costs. Note the rate of the shadow cost increase is approximately 7, which is quite close to π_{σ} , the unit revenue per unit of carried demand (by (39) and (40), π_{σ} is in the neighbourhood of 8). This indicates that an additional unit of optical bandwidth added can be used to carry close to one unit of traffic and to earn a unit of revenue. In the second case, the demand does not

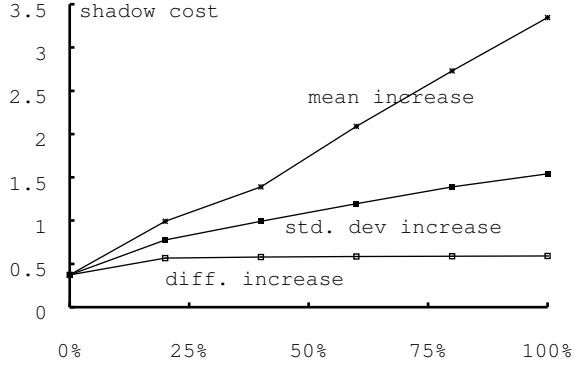


Fig. 7. Impact of Demand Change on Shadow Costs

increase on average, but the data network still needs more bandwidth to handle increased demand uncertainty. Using bandwidth to back up uncertain demands is less profitable than using it to carry new demand, which explains why the increase of shadow costs is smaller than in the first case. In the third case where the demand increases between some node pairs and decreases between others, the data network can absorb the fluctuation by adjusting routing and admission control in its own domain without requiring more changes in the provisioning of the optical core. Consequently, shadow costs are much less sensitive to this type of demand change, as is shown in the figure.

Intuitively, one would expect if the increases in shadow costs are small, then there is little need to reconfigure the optical core. This intuition is verified in Table IV in which we show the total number of wavelengths the comes from re-optimizing the problem with new demands. The first row corresponds to the base case where the number of deployed wavelength is 17. The other rows show the optimal wavelength deployment for different demand changes. For instance, a 40% increase in mean demand requires increasing the number of wavelength to 21. In the case of increasing standard deviation, the threshold is 60%. Furthermore, the number of wavelengths need to be increased to 23 when mean demand is increased beyond 80% while no such increase is needed in the other cases. Comparing Figure 7 with Table IV, we see that there is a strong correlation between the increase of shadow costs and the need to add new wavelengths. Therefore, it is conceivable that in a dynamic environment where the demand is constantly changing, the optical core should keep communicating with the data network on these shadow costs, and restart the aforementioned optimization process once it detects a significant change.

% increase	mean	std. dev.	difference
0%	17	17	17
20%	17	17	17
40%	21	17	17
60%	21	21	17
80%	23	21	17
100%	23	21	17

TABLE IV
OPTIMAL NUMBER OF WAVELENGTHS FOR DEMAND CHANGES

VI. DISCUSSIONS AND CONCLUSIONS

In this paper, we have presented a framework for efficient multi-layer IP traffic engineering and optical network configuration in the IP-over-Optical Overlay model. Our framework accommodates various traffic demand formulations and utility functions. We have shown that the distributed implementation of the overlay model achieves global optimum. Our approach is derived from the Generalized Bender’s Decomposition, where the sub-problems (slave and master) correspond to separate decision-making by the data and optical domains. The information exchange between the two domains is kept at a minimum and may be mapped into standard communications through the UNI.

Our work introduces the concept of “multi-layer” grooming, which broadens the traditional grooming in the optical domain to data networks, where now the latter are active participants in the grooming process with intelligent homing of data traffic to optical gateways.

While the treatment in this paper is restricted to the case of a single pattern of end-to-end traffic demands, our framework can be extended to accommodate a dynamic environment with changing demands. As we demonstrated in the numerical case studies, shadow costs of optical paths provide an excellent indicator as to when re-optimization of the optical network is required. This result points to an interesting direction for future work.

REFERENCES

- [1] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja and X. Xiao, Overview and Principles of Internet Traffic Engineering, *IETF RFC-3272*, May 2002
- [2] D. Awduche, Integrated Traffic Engineering in IP over Optical Networks: Fundamental Concepts, *SPIE ITCOM-2002*, Boston, Mass, July 31, 2002
- [3] B. Rajagopalan, D. Pendarakis, D. Saha, R. S. Ramamoorthy and K. Bala, IP over optical networks: Architectural aspects, *IEEE Communications Magazine*, vol. 38, no. 9, September 2000 pp. 94-102

- [4] L. Cheng, J. Ellson, A. Jukan, P. Lamy, E. Varma, Network Engineering - Control of Dynamic Link Topology in User Networks, *Bell Labs Technical Journal*, 8(1), pp. 207-218, 2003.
- [5] A. G. Chiu and E. H. Modiano, Traffic grooming algorithms for reducing electronic multiplexing costs in WDM ring networks, *IEEE Journal of Lightwave Technology*, vol. 18, no. 2, pp. 2-12, Jan. 2000.
- [6] A. M. Geoffrion, Generalized Benders Decomposition, *Journal of Optimization Theory and Applications*, No.10, vol. , pp. 237-260, 1972.
- [7] O. Gerstel, R. Ramaswami and G. Sasaki, Cost-effective traffic grooming in WDM rings, *IEEE/ACM Transaction on Networking*, vol. 8, no. 5, pp. 618-630, Oct 2000.
- [8] ITU-T Rec. 8080/Y.1304 *Architecture of the Automatically Switched Optical Network*, Nov. 2001, Amendment 1 Mar. 2003, Amendment 2 Feb. 2005.
- [9] Y. Liu, H. Zhang, W. Gong, D. Towsley, On the Interaction Between Overlay Routing and Underlay Routing, *IEEE Infocom 2005*.
- [10] F. P. Kelly, A. K. Maullo and D. K. H. Tang, Rate Control in Communication Networks: Shadow Prices, Proportional Fairness and Stability, *Journal of the Operational Research Society*, vol. 49, pp. 237-252, 1998.
- [11] S. H. Low, A Duality Model of TCP and Queue Management Algorithm, *IEEE/ACM Transaction on Networking*, vol. 11, No. 4, pp. 525-536, August 2003.
- [12] Generalized Multi-Protocol Label Switching (GMPLS) Architecture, *IETF RFC-3945*, Mannie eds., October 2004
- [13] R. Mahajan, D. Wetherall and T. Anderson, Towards Coordinated Interdomain Traffic Engineering, *HotNets-III*, Nov. 2004
- [14] D. Mitra and Q. Wang, Stochastic Traffic Engineering for Demand Uncertainty and Risk-Aware Network Revenue Management, *IEEE/ACM Transactions on Networking*, vol. 13, No. 2, pp. 221-233, April 2005.
- [15] M. Pioro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*, Morgan Kaufmann, San Francisco, 2004.
- [16] L. Qiu, Y. R. Yang, Y. Zhang and S. Shenker, On Selfish Routing in Internet-Like Environments, *Proceedings of the ACM SIGCOMM*, 2003.
- [17] B. Rajagopalan, J. Luciani and D. Awduche, IP over Optical Networks: A Framework, *IETF RFC-3717*, March 2004
- [18] R. Ramaswami and K. N. Sivarajan, *Optical Networks: A Practical Perspective*, Morgan Kaufmann Publishers, San Francisco, 1998.
- [19] A. Schrijver, *Theory of Linear and Integer Programming*, John Wiley and Sons, 1998
- [20] S. Sen, R. D. Doverspike and S. Cosares, Network planning with random demand, *Telecommunications Systems*, vol. 3, pp. 11-30, 1994.
- [21] G. Swallow *et al*, GMPLS User-Network Interface (UNI): RSVP-TE) Support for the Overlay Model, *draft-ietf-ccamp-gmpls-overlay-05.txt*, (work in progress) Oct. 2004.
- [22] I. Widjaja, I. Saniee, L. Qian, A. Elwalid, J. Ellson and L. Cheng, A new approach for automatic grooming of Sonet circuits to optical express links, *IEEE ICC 2003*.