Multi-Layer High-Speed Network Design in Mobile Backhaul Using Robust Optimization

Ashish Mathew, Tamal Das, Prasad Gokhale, and Ashwin Gumaste

***Abstract*—The multi-layer network design problem and the mobile backhaul problem are both interesting from the perspective of choosing the correct technology and protocol as well as choosing the appropriate node architecture to meet a wide variety of overlay traffic demands. Network operators encounter two variants of the multi-layer/backhaul problem: 1) For a given set of uncertain traffic demands, which set of technologies would minimize the network cost? 2) Would these technology choices be invariant to the changing traffic demands? This problem of technology choice can essentially be abstracted to a switching and grooming problem with the added complexity of unknown traffic demands, which at best may be approximated to some known statistical parameters. To solve this complex switching and grooming problem, our goal is to make use of the theory of robust optimization with the assumption of known boundary conditions on traffic. We present a comprehensive optimization model that considers technology choices in terms of protocols, physical layer parameters, link boundary conditions, and transmission layer constraints. Validated by simulations, our model shows the stand-off conditions between various technologies and how a network operator must take proactive steps to be able to meet requirements of the next-generation networks and services. Our main result showcases network design using two technology alternatives [1) Multi-Protocol Label Switching (MPLS) + Optical Transport Network (OTN) + Reconfigurable Optical Add–Drop Multiplexer (ROADM) and 2) Carrier Ethernet (CE) + OTN + ROADM] and the effect of robustness on these choices. A heuristic is used for comparative purposes as well as to exhaustively model the dynamic case of brownfield networks.**

***Index Terms*—Carrier Ethernet; Multi-layer optimization.**

I. Introduction

A common problem experienced globally by service providers is that revenues are not proportional with bandwidth usage [1]. While this aspect of revenue stagnation is valid for most wired services, when it comes to carrying mobile backhaul traffic as well as enterprise traffic, there is a significant deviation from this trend. Mobile backhaul and enterprise traffic represent revenue-bearing services that are important for the provider due to the rapid explosion of mobile devices and mobile applications [2] as well as cloud computing and data center usage by enterprises. Applications, particularly video-centric ones, are becoming an important aspect of traffic in metropolitan and core networks. The underlying technologies that govern metropolitan and core network traffic are important, especially as we migrate from 2G to 3G and 4G services across mobile networks and enhance the usage of cloud and data centers in enterprise environments.

Traditionally, Synchronous Optical Networking (SONET)/Synchronous Digital Hierarchy (SDH) circuits are used to connect base stations (BSs) to mobile switching centers (MSCs) in the mobile backhaul. Similarly, enterprises were connected using T1/E1 circuits that were aggregated to SONET/SDH networks. SONET/SDH was popular, especially when traffic was primarily voice-oriented, as there was a direct functional relationship between the fronthaul (radio network), enterprise network, and backhaul/core traffic type, both being primarily circuit centric. However, with the change of revenues from voice to data in the cellular fronthaul and the direct migration towards 3G and 4G, there is an immediate need to be able to better provision, plan, and exploit the backhaul network. Enterprise networks are also moving towards data center and cloud-based systems implying the requirement of packet-based fat pipes that interconnect enterprises to the rest of the cloud. Specifically, given the packetized requirement of data traffic and the unpredictability of users, there is a requirement to be able to provision the metro/core network using packet technologies that can lead to dynamic bandwidth provisioning in order to meet an uncertain (and unknown *a priori*) set of demands.

The problem of backbone/core design gets complicated when we consider the range of enabling technologies as plausible candidates in the backbone/core. A wide range of technologies are available such as Multi-Protocol Label Switching (MPLS), G.709 Optical Transport Network (OTN), wavelength-division multiplexing (WDM), and Carrier Ethernet (CE), all of which support the overlay Internet Protocol (IP) network. The design of the backbone/metro/core is critical in terms of decision-making as to

Manuscript received June 20, 2014; revised February 3, 2015; accepted February 16, 2015; published April 1, 2015 (Doc. ID 214142).

Ashish Mathew is with the Department of Computer Science, Stanford University, USA.

Tamal Das is with the Institute of Computer and Communication Networks, TU Braunschweig, 38106 Braunschweig, Niedersachsen, Germany.

Prasad Gokhale and Ashwin Gumaste (e-mail: ashwing@ieee.org) are with the Department of Computer Science and Engineering, Indian Institute of Technology, Bombay, Maharashtra 400076, India.

http://dx.doi.org/10.1364/JOCN.7.000352

which technology should be used and at which layer the switching should be performed.

In a practical scenario, two other considerations are important: 1) defining the inventory for such a network by making the right technology choices and 2) guaranteeing to meet traffic churn. The aspect of churn is of significance. Given that the number of users and their bandwidth requirement at a particular location/cell-site is unknown/uncertain, is it possible to perform network planning, and hence optimization? We solve this problem of uncertain traffic churn by using the theory of robust optimization that we apply to the specific case of multi-layer network design. In particular, our solution takes into consideration multiple technology approaches that a provider would deploy and consider cost and performance trade-offs associated with these approaches.

The premise of this paper is to aid service providers in network planning to decide which technology or set of technologies should be deployed to meet the service-centric requirements of the mobile backhaul. In addition, we assume that the service requirements are not known *a priori*. We use the tool of robust optimization [3] to solve the uncertainty problem experienced in planning networks for both mobile backhaul and enterprise traffic. Our formulation is independent of the exact number and absolute bandwidth consumption of the demands (at a particular BS, or from an enterprise) ahead of time.

Specifically, we assume knowledge of well-known network statistics to be available: for a particular cell or in a particular enterprise, there are, say, a maximum of n users active at a given time, and the maximum bound on n is known. Now, let us assume that k ($<n$) users will be accessing data at the full line rate provided by the network (we term this rate as the peak rate). Naturally, this means that the remaining $n-k$ users would be accessing bandwidth at some lower line rate. For the sake of convenience and preserving generality as well as correctness, say, the $n-k$ users access bandwidth at some average rate that is known (we term this rate simply as average rate). Given the burstiness of the traffic and the maximum line rate, it is possible to compute both the peak and average access rates ahead of time. Now, given the value of n, k, the peak rate, and the average rate, we seek to design the backbone/metro/core network choosing among the multiple technology options in the most cost-effective manner [from both a capital expenditure (CapEx) and operational expenditure (OpEx) minimization perspective].

Our optimization approach is as follows. Our goal in this paper is to reduce both CapEx and OpEx in the network design problem. Specifically, we seek to reduce and better utilize the ports in telecommunication equipment by

1) choosing the appropriate technology to reduce the total cost of ownership, and
2) intelligently aggregating traffic, leading to better port utilization.

These two strategies lead to CapEx reduction. In parallel, we attempt to route traffic to reduce overall OpEx; the fundamental principle being to keep the data in the lower layers, as far as possible. Specifically, we investigate the conditions under which it makes sense to keep the data in the optical, optical + Ethernet, and optical + Ethernet + MPLS layers, so as to reduce both OpEx and CapEx leading to an overall reduction in the total cost of ownership (as illustrated in Fig. 1).

The main contributions of the paper are as follows: 1) use of robust optimization to solve uncertain traffic demands in a backhaul network, 2) use of statistical multiplexing properties of different technologies in being able to select an appropriate technology solution, 3) use of statistical methods for computing uncertainty of demands in a backhaul network (proposed use of peak and average usage by a user), and 4) enumeration of the technology model through a state machine that enables the optimization model to choose the correct technology while preserving the protocol working.

The rest of this paper is organized as follows. Section II discusses the related work in the area of mobile backhaul and multi-layer network design. Section III elaborates the problem statement by describing a set of approaches that are instructive towards the robust optimization solution, specifying our optimization strategies, cost considerations, and the corresponding state machine. Section IV showcases our optimization model deploying robust optimization to solve uncertainty in the enterprise traffic/backhaul network, while computing the best technology choices for multi-layer design. Section V proposes fast heuristic algorithms for network design for the static and dynamic cases of traffic. Section VI describes the simulation model that validates our study through rigorous numerical evaluation, while Section VII concludes our paper.

II. Related Work

In this section, we summarize the related literature in the domain of mobile backhaul and multi-layer network design.

The problem of capacity planning has been considered in almost every aspect of telecommunication networks and at every layer in the protocol stack. ATM terminal capacity

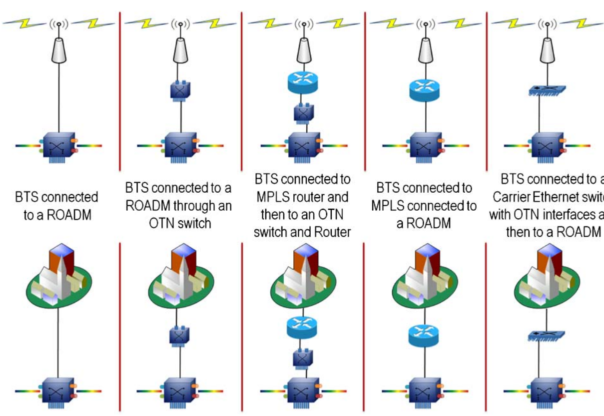

Fig. 1. Mobile backhaul options (top) and multi-layer technology options (bottom). BTS, base transceiver station.

problems were considered in [4] and the references therein. WDM network optimization was considered in [5] and continues to be a topic of debate with new technologies and added constraints due to gridless contentionless reconfigurable optical add–drop multiplexers (ROADMs). MPLS optimization has been considered in [6] and the references therein, as well as models based on that work. The theory of robust optimization to networks and flows was well described in [7]. One of the earlier works on applying robust optimization to networks is described in [8]. In [9], a study of applying robust optimization to optical networks is presented. Reference [9] discusses a given set of lightpaths and their behavior under uncertainty and applies a robust optimization technique to the specific domain of optical networks.

With regard to the work on multi-layer optimization, the work in [10] is important as it considers the WDM, SONET/SDH/OTN, and IP/MPLS layers. On a broader level, [10] focuses on network optimization through subsystem design. It however does not consider traffic statistics, which we do. An interesting work on using link metrics is reported in [11], which presents a mixed integer linear program (MILP) based on concave link metrics. Another static traffic multi-layer model is presented in [12], which also does not consider traffic or service statistics. The role of OTN as a distinct layer is considered by [13], whereby an optimization model for static traffic is presented. The authors therein examine the layer correlation effect and provide an important perspective on how multi-layer networks can be modeled. In fact, the impact of OTN on provider networks is immense especially for packet-oriented traffic requirements which can be modeled through only statistical means. In this regard our work is critical in mapping statistical traffic estimation on multi-layer networks with due cognizance of all the underlying protocols (CE, MPLS, OTN, and WDM).

In [14], the authors quantify CapEx savings of homing architectures using optical network equipment and show a linear relationship between equipment and traffic demand. This forms an interesting guideline to our work towards extending to multi-technology options with the added robustness constraints.

Our work builds upon the basic premise in [15], and our mathematical formulation is based on the one therein. In particular, we have developed the model around a solid application case of core multi-layer networking and mobile backhaul. In doing so, we have considered more pragmatic technology options, such as the introduction of CE, a practical choice in backhaul networks. Our work also includes the additional constraint of minimizing OpEx, reducing energy consumption, which is known to out-price the CapEx over a period of 5–7 years (lifetime of equipment). To solve the optimization problem, we also propose a heuristic of $O(Q(E + V \log V))$ complexity, making the solution tractable. The work in [15] considers the specific case of MPLS and to some extent OTN; [15] does not consider CE and ROADM at all which are intrinsic to any provider solution. Further, [15] does not propose any solvable heuristic that measures well with the robust optimization approach.

Our primary inspiration is that [15] is perhaps the first work that uses robust optimization in the theory of network design for core networks. In this paper, we use the same principles in a more appropriate use case of mobile backhaul (or even enterprises), where robustness is more valuable, due to the unpredictable nature of traffic.

III. PROBLEM DESCRIPTION

Let us assume a network graph and statistical information about traffic requests between each source–sink pair in the network. For each request, we assume the peak-rate and average-rate bandwidth requirements, but not the exact knowledge of which request would be at an average or peak value. The components that are considered for CapEx computation are illustrated in Fig. 2. The costs of these components are averaged across four vendors over a 2-year period. Our goal is to minimize sum of CapEx on the equipment and recurring OpEx in maintaining the ports (energy costs, incidentals to keep the ports active) by making decisions at the network planning stage (for a greenfield network) or for incremental upgrades (for a brownfield network).

We next discuss the decisions related to routing, higher-layer protocols, and cost considerations that a network operator has to contend with. At the end of this section, we also present a state machine corresponding to our system that enables the optimization model to take into consideration the protocol interaction and switching granularity.

A. Criteria for Network Design

We now list the variety of choices that network operators encounter in the process of backbone/core network design, along with associated criteria and trade-offs.

1) Criteria for Routing: We need not necessarily provision a request along the shortest path, as discussed next.

Routing Criterion 1. Network paths through which a large set of requests transit are preferred over shorter paths, so that routing/regeneration equipment need not be deployed extensively. For example, consider the situation

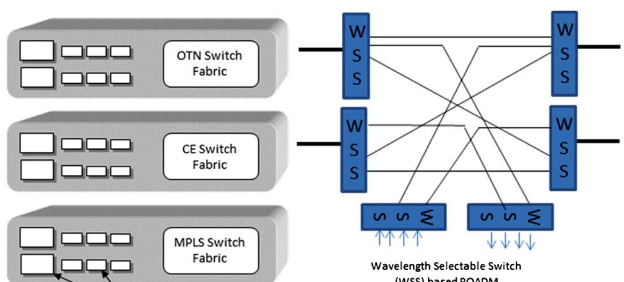

Equipment	Price (in 1000 USD)
10G ROADM IO	10
10G MPLS IO	25
10G OTN IO	18
10G CE IO	12
MPLS Switch fabric port-to-port	18
OTN switch fabric ODU switching	12
CE switch fabric	8
1 Gbps CE	4
ODU3 (40 Gbps)	27

Fig. 2. Network components contributing to CapEx [16].

described in Fig. 3. Assume that there exists some traffic along the paths *A–B–C–F* and *A–B–C–G*, and a new request must travel from *A* to *D*. The distance from *A* to *D* (via the shortest path *A–E–D*) being too large to send the packets via all-optical (OOO) nodes due to optical SNR (OSNR) limitations, there must be minimal acceptable signal regeneration at intermediate node *E*. Given existing traffic through *C*, using the shorter path (*A–E–D*) would require deploying extra equipment at *E*. However, if the new request follows the path *A–B–C–D*, i.e., its path overlaps with existing paths, it can use the equipment already deployed for these paths (at *C*).

Routing Criterion 2. Despite the advantage of multiplexed traffic, longer paths are, however, discouraged by attenuation constraints and loss in resilience to equipment failure, unless unavoidable. Such paths would then have optical–electrical–optical (OEO) regenerators to rejuvenate the signal.

2) Criteria for Propagation Across a Node: Once a path is selected to provision a request, we must choose between an OOO port-to-port connection or an OEO conversion at each intermediate node along the path.

Preference towards OOO. Existence of OEO capability at a node does not imply that all requests that pass through will use such OEO functionality.

Necessity for OEO. OEO is, however, necessary for 1) regeneration of a distorted and/or severely attenuated signal, 2) rerouting a new request due to unavailability of a single contiguous wavelength, and 3) demultiplexing a multiplexed signal into constituent requests that need to be forwarded along different paths.

3) Criteria for Deployment of Higher-Layer Protocols: Due to limited physical layer resources (e.g., signal power/OSNR limitations, limited number of wavelengths) and design considerations (maximizing traffic supported by the network), there is a need to deploy higher layer protocols along most paths. This requires deploying extra equipment in conjunction with a ROADM, the primary transport element in the optical layer. We consider the following practical issues when deploying higher-layer equipment.

CapEx Increase. We assume that with electronic regeneration and protocol conversion, the capital cost substantially increases. This is due to expensive chipsets needed for protocol conversion and signal regeneration.

Optimization Model. In the optimization model, the following technology choices are used to provide higher-layer functions:

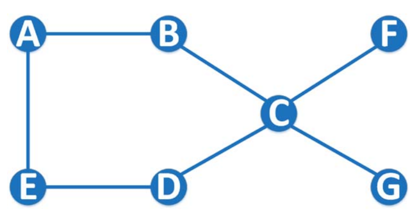

Fig. 3. Illustrative example for maximizing path overlap.

- OTN (ITU G.709) is used to increase the range of a signal by adding forward error correction (FEC) to a bit stream. OTN is also used for inter-domain networking, especially when several providers want to create an end-to-end optical circuit. OTN-based tandem channel monitoring (TCM) plays a key role in facilitating a user/provider to monitor the signal across multiple domains.
- MPLS is used for statistical multiplexing of multiple packet flows as label-switched paths (LSPs), thereby resulting in conservation of bandwidth.
- The combination of OTN and MPLS is popular when the distance is large and there is a need to obtain gain through statistical multiplexing.
- OTN is used with CE to obtain packet–optical integration, thereby resulting in the advantages of Layer 2 packet switching coalesced with long-range communication (through OTN-based FEC enhancement). We specifically neglect the impact of OTN–ODU switching (G.709 OTN optical data unit), focusing only on the FEC-based benefits of OTN, as the former provides a redundant switching layer and hence the switching aspect is only relegated to CE technology.

4) Criteria for Use of Deployed Higher-Layer Protocols: Once higher-layer equipment is deployed at a node, we must decide how to distribute traffic as it transits between the node's ports. In this regard, we assume the following.

Low Capacity. Low capacity ports, such as at 1- and 2.5-Gbps line rates, have lower CapEx and OpEx since the cost per bit at these line rates is even today much less than the cost-per-bit at higher line rates due to the voluminous use of lower-speed optics.

Allocation. Allocation of ports must ensure that the traffic does not exceed the port capacity and also include a sufficient margin to provide robustness against traffic variations.

Margin. A sizable margin is required to provide robustness against traffic variations and is defined by a parameter k that denotes the maximum number of requests from a common source that will be at a peak value.

The constraints considered in our formulation include 1) signal attenuation that limits the geographic range of transmission, 2) port capacity in terms of bandwidth, 3) switching constraints that restrict the routing of requests at nodes that do not have suitable (corresponding protocol-compliant) equipment, and 4) wavelength constraints that limit the number of flows through a fiber.

B. Cost Considerations

We make the following practical assumptions about the parameters related to the technology and traffic. Each fiber link can support up to 176 wavelengths at line rates of 1, 2.5, and 10 Gbps across the *C* and *L* bands. While we currently do not include 40/100 Gbps as an option, due to both cost of multiplexing and an increase in cost per bit, there is a provision to include 40/100 Gbps with almost no extra modification in the solution. Also for the case of mobile backhaul, aggregating BS traffic into 40/100 Gbps pipes makes no sense as the bandwidth juxtaposed by the BS would be much lower implying that an all-optical solution (of directly mapping BS traffic onto a wavelength) is not feasible. For the mobile backhaul case, the traffic emanating from a base station can vary from <2 Mbps (conventional T1/E1) to a few multiples of 1 Gbps (in support of aggregated multiple-input–multiple-output-based technologies).

In our solution for cost computation, the CE switches have 2 × 10G ports and 10 × 1G ports. MPLS label-switched routers have stackable line cards, each with 4 × 10G ports and 8 × 1G ports. OTN switches have 4 × ODU2 ports and 8 × ODU0 ports. The ROADMs can have a degree of connectivity of at most 8, supporting a full mesh configuration and supporting colorless, contentionless, and directionless provisioning at 50-GHz spacing.

In our model, CapEx is computed as the sum of the costs of three types of network components: 1) the router (MPLS/CE switch fabric), 2) the transponders (OTN), and 3) the wavelength selective switch (WSS) with amplifiers as well as peripheral optics that make up the ROADM. OpEx is calculated as the cost of total power consumed by the deployed equipment. Other energy costs such as amplifier, electronic-level processing, and so on are not considered as they cancel out for the purpose of comparison. In fact, the reason these cancel out is because these operational costs do not change with network load. For example, an amplifier is required irrespective of the number of channels, and its operating power remains almost the same irrespective of the number of channels to be amplified. We neglect service provider salaries, marketing, right-of-way, and bandwidth cost as shown in [17].

By solving the design problem for reasonably sized networks (80–100 nodes) in the metro and the regional area using robust optimization, it is evident (from Section IV) that the optical + CE family of solutions, i.e., ROADMs + CE + OTN solution, is most preferred in terms of reducing CapEx and OpEx. It is further evident from the solution of the robust optimization problem that the OpEx is the lowest when we attempt to keep the data in the optical or the CE layers. Hence, by minimizing the CapEx through a ROADM + CE solution, we are also able to minimize the OpEx.

C. Modeling the Network Design Problem as a State Machine

In this subsection, a state machine (see Fig. 4) that models the behavior of a request from a BS to another BS or from a BS to an MSC across a multi-layer multi-technology network or within any two enterprise end points in a high-speed network is presented. The states correspond to the various protocol levels and technology choices that a connection encounters while traversing from its source to destination, while the transitions between the states represent the movement of data between nodes complying with various technology choices. The state machine is instructive to the working of the robust optimization program.

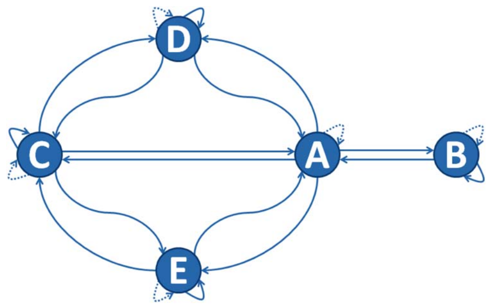

Fig. 4. State machine of the various transitions.

We define so-called base flow as the bit stream that emanates from a mobile tower (antenna) towards the backhaul network or from a customer edge to a core interconnection point in a multi-domain network. The five states that exhaustively define our system are defined in Table I. The arcs of the state machine exhaustively cover all the possible protocol changes/transitions that we model and retransmissions undertaken by a request as it passes through a node. The five states would lead to a five-tuple model, whose solution is obtained through the robust optimization process (described in Section IV). Hence, the state diagram is critical towards building the robust optimization model.

Each state has two self-loops. The dashed-line loop represents an OOO at a node (indicating that the signal passes through a ROADM, while a solid-line loop represents an OEO (for 3R retransmission and/or regeneration of the request provisioned on a wavelength for MPLS/CE flows). Note that as per our state machine, at most one transition can occur in a single hop. The set of states reachable in one hop, i.e., the protocol changes that can occur after a single OEO, depends on the current state. In order to keep the optimization problem tractable, we have made the simplifying assumption that the state of a request depends only on the port from which it originates. The state machine provides the bigger picture of the system and is instructive to our optimization model in Section IV. It is a precursor to the robust optimization formulation. Specifically, the state machine gives an exhaustive, yet feasible list of options that the formulation can achieve. The state machine is, hence, a guiding path for correctly framing the constraints.

IV. Robust Optimization

In this section, we first provide an overview of our robust optimization formulation, and then define the same.

The route of every request is divided into multiple sections. A section is a contiguous subset of the route that lies between two adjacent OEO junctions (along the route). Hence, signal passes as OOO through all those nodes that are within a section and OEO conversions occur at the endpoints of a section. A section is uniquely defined by a three-tuple (λ, a, b) where λ is the wavelength used, and a and b

TABLE I
DESCRIPTION OF STATES FROM THE STATE DIAGRAM

State	Description	Technology Choices
A	Base byte stream state, i.e., customer edge protocol mapped directly onto Ethernet frames	WSS-based ROADMs that support an all-optical switch fabric
B	Base flow encapsulated in MPLS labels	MPLS + ROADMs supporting statistical multiplexing of LSPs onto a wavelength and further supporting optical switching within the ROADM
C	Base flow encapsulated in Ethernet and then mapped onto OTN	OTN + ROADM supporting ODU/OTU-k multiplexing at the OTN layer and wavelength switching at the optical layer
D	State C type flow earlier encapsulated in MPLS (MPLS over OTN)	OTN + MPLS + ROADM supporting packet–optical integration by multiplexing LSPs mapped into ODU containers that are further mapped into wavelengths and supporting switching at the MPLS, OTN, and wavelength layers
E	State C type flow earlier encapsulated in Carrier Ethernet (CE over OTN)	OTN + CE + ROADM providing a fully multiplexed packet-based system with additional ODU switching, Ethernet (packet) switching, and wavelength-level switching (using the ROADM)

are the first and last link in the section, respectively. a and b overlap for a one-link section. Note that each three-tuple is unique and will result in an optimal solution. A state (as defined in Subsection III.C) is defined for a particular section. Requests are assigned to a section depending on its state. This implies that if a section is in state B, D, or E, then multiple requests can be assigned to the section. Each section is made of links, while each link comprises of an edge between two nodes. A link cannot be simultaneously used in more than one section on the same wavelength. Thus, a section exists irrespective of a feasible free path.

The technology choices that can be installed at any node in the mobile backhaul network are summarized in Table I. Table II summarizes the variables and parameters used in our formulation of the robust optimization MILP.

A. Objective Function

Our objective is to minimize the sum of the following costs:

TABLE II
LIST OF NOTATIONS

Parameter	Definition
V	Total number of nodes in the network
L	Total number of links in the network
$\bar{Q}$	Maximum number of requests in the network
$src(q)$	Source node of request q
$snk(q)$	Sink node of request q
$org(l)$	Source node of directed link l
$dst(l)$	Destination node of directed link l
$x^q_{\max}$	Maximum value of traffic for request q
x^q_{avg}	Average value of traffic for request q
δ_q	$x^q_{\max} - x^q_{\text{avg}}$
$AoL(l)$	Attenuation on link l
P	Maximum number of ports on any technology option
D	Five-tuple whose ith component is the transmission power of a port when using technology option i
$[a, b]$	$\bar{Q}$-component vector, whose ath component is −1, bth component is +1, and other components are 0
Variable	**Explanation**
$A^{\lambda,q}_{a,b}$	Boolean to determine whether request q is assigned to section (λ, a, b)
$P^{\lambda,T,p}_{a,b}$	Boolean to determine whether section (λ, a, b) uses port p of technology T
$SU^{\lambda}_{a,b}$	Boolean to determine whether section (λ, a, b) is realizable
$U^{k,v}_p$	Boolean to determine whether port p of technology k is used at node v
$LU^{\lambda,l}_{a,b}$	Boolean to determine whether link l is used to build section (λ, a, b)
$\mu^{\lambda,q}_{a,b}$	Boolean to determine whether deviation of q from average is maximum in section (λ, a, b)
T_v	Five-tuple Boolean vector variables representing all technology options that can be installed at node v

Switch CapEx

$$\sum_{\forall v}\sum_{i=1}^{5}\begin{pmatrix}\text{Cost of deploying}\\ \text{technology option } i\end{pmatrix}\times (T_v)_i. \tag{1}$$

Port CapEx (for MPLS, CE, and OTN)

$$\sum_{\forall v}\sum_{\forall T}\sum_{\forall \text{ ports in } T}\begin{pmatrix}\text{Cost of transponder/switch}\\ \text{for port } p \text{ of technology } T\end{pmatrix}\times U_p^{T,v}. \tag{2}$$

WSS CapEx (for ROADMs)

$$\sum_{\forall v}\sum_{\forall \text{ sections}(\lambda,a,b) \text{where } v \text{ is source/sink}}(\text{Cost of WSS})\times SU^{\lambda}_{a,b}. \tag{3}$$

Port OpEx

$$\sum_{\forall v}\sum_{\forall T}\sum_{\forall \text{ ports in } T}\begin{pmatrix}\text{Cost of operating}\\ \text{port } p \text{ of technology } T\end{pmatrix}\times U_p^{T,v}. \tag{4}$$

B. Constraints

We now list the various constraints of our MILP.

Strictly one of the five technology options must be installed at each node. Therefore

$$\sum_{k=1,\ldots,5}(T_v)_k = 1. \tag{5}$$

A particular technology option must be installed at node v if at least one of its ports is used for that technology:

$$P\times (T_v)_k \le \sum_{\forall \text{ port } p \text{ of tech. option } k}\ \sum_{\forall \text{ sections where } v \text{ is source}} P^{\lambda,T,p}_{a,b}. \tag{6}$$

A request can only be assigned to a physically realizable section. Conversely, a section that has been assigned to a request must be physically realizable:

$$\bar{Q}\times SU^{\lambda}_{a,b} - \sum_{\forall \text{ request } q} A^{\lambda,q}_{a,b} \ge 0. \tag{7}$$

At most one port can be assigned to a realizable section. Conversely, a section that has been assigned to a port must be physically realizable

$$SU^{\lambda}_{a,b} - \sum_{\forall \text{ tech. option } T}\ \sum_{\forall \text{ port } p \text{ of tech. option } T} P^{\lambda,T,p}_{a,b} > 0. \tag{8}$$

A port is in use if and only if it is assigned to some section:

$$U_p^{k,v} = \sum_{\forall \lambda}\sum_{\forall \text{ link } b}\ \sum_{\forall \text{ link } a \text{ where } v=org(a)} P^{\lambda,T,p}_{a,b}. \tag{9}$$

Each request q must be composed of sections that can form a valid path from $src(q)$ to $snk(q)$:

$$[src(q), snk(q)] = \sum_{\forall \lambda}\sum_{\forall \text{ link } l} A^{\lambda,q}_{org(l),dst(l)}\times [org(l), dst(l)]. \tag{10}$$

A section from a to b on wavelength λ can be used only if we can build a path from a to b using wavelength λ:

$$SU^{\lambda}_{a,b}\times [org(a), dst(b)] = \sum_{\forall \text{ link } l} LU^{\lambda,l}_{a,b}\times [org(l), dst(l)]. \tag{11}$$

If a section is in use, then its links are blocked on the concerned wavelength:

$$\begin{aligned} LU^{\lambda,a}_{a,b} + SU^{\lambda}_{a,b} &\ge 1,\\ LU^{\lambda,b}_{a,b} + SU^{\lambda}_{a,b} &\ge 1. \end{aligned} \tag{12}$$

For two overlapping sections, the same wavelength should not be used (similar to wavelength continuity [18]):

$$\sum_{\forall \text{ link } b}\sum_{\forall \text{ link } a} LU^{\lambda,l}_{a,b} \le 1. \tag{13}$$

1) Attenuation Constraints: These constraints are on a particular section. Therefore

$$\left\{\begin{matrix} D_i - \sum_{\forall \text{ link } l} AoL(l)\times LU^{\lambda,l}_{a,b}\\ -\text{Receiver Threshold}\end{matrix}\right\} + (1-(SU^{\lambda}_{a,b})_i)\times 1000 \ge 0, \tag{14}$$

where receiver threshold denotes the threshold that indicates the minimum received power required to obtain the requisite OSNR. The first term of Constraint (14) will be negative if the section is loss limited when it operates with a technology option i. If it is loss limited, then the corresponding state decision variable is forced to be 0. If all the state decision variables are forced to be 0, then by Eq. (5), the section cannot be used.

2) Capacity Constraints for Multiplexing Technologies: Similar to [15], we impose one capacity constraint per section:

$$\begin{aligned} &\sum_{\forall \text{ multiplexing tech. } T}\ \sum_{\forall \text{ port } p \text{ in } T} P^{\lambda,T,p}_{a,b}\times\begin{pmatrix}\text{Max. capacity}\\ \text{of port } p\end{pmatrix}\\ &\quad\ge \sum_{\forall q} A^{\lambda,q}_{a,b}\times x^{q}_{\text{avg}} + (\Delta^{\lambda}_{a,b})_{\max}, \end{aligned} \tag{15}$$

where $(\Delta^{\lambda}_{a,b})_{\max}$ is a scalar, representing the margin against failure (which in our case is congestion within the section). The margin should be greater than what is required in the worst case scenario as elaborated in the subsequent text. It must be noted that $(\Delta^{\lambda}_{a,b})_{\max}$ implies the slackness between an exact formulation, where the traffic conditions and hence constraints are known ahead of time, and the robust optimization formulation, whereby the constraints are expressed only as a worst-case measure. The role of $(\Delta^{\lambda}_{a,b})_{\max}$ is crucial to computing the degree of over-provisioning that we must indulge in a network to take into consideration the uncertainty that arises due to robustness.

For each request q, we denote the maximum and average bandwidth granularity as $x^q_{\max}$ and x^q_{avg}, respectively. We group the requests into subsets based on their source node, and order them in decreasing order of maximum deviation from their respective average values, i.e., $x^q_{\max} - x^q_{\text{avg}}$. This results in an ordered list of requests at each source node. The worst-case scenario for this model is when the first K requests from each of these lists are at their maximum value and all the other requests are at their average value.

The question we ask is, whether we can determine the worst-case margin for a section ahead of time? This is, however, not possible due to the uncertainty in traffic. The margin depends on which requests were assigned to a particular section. Based on the previous discussion, we describe intuitively a procedure for obtaining the worst-case margin $(\Delta^{\lambda}_{a,b})_{\max}$ for the section on wavelength λ from node a to node b:

- Consider the set of all requests assigned to this section.
- Partition this set based on the source of each request. Requests with the same source will be in the same partition.
- From each partition, choose the top K requests (based on maximum deviation from average).
- The margin for this section is the sum of the deviations of all requests chosen this way (across all partitions).

This procedure can be captured in the form of the following linear program:

$$(\Delta^{\lambda}_{a,b})_{\max} \geq (A^{\lambda,q}_{a,b} \times \mu^{\lambda,q}_{a,b} \times \delta_q) \quad \forall q, \tag{16}$$

subject to

$$\forall \text{ node } i, \quad \sum_{\forall \text{ request } q \text{ originating from node } i} \mu^{\lambda,q}_{a,b} \leq K. \tag{17}$$

3) Capacity Constraints for Nonmultiplexing Technologies: Capacity constraints are required for each section, since we do not know the technology used for a particular section:

$$\sum_{\substack{\forall \text{ nonmultiplexing} \\ \text{tech. } T}} \sum_{\substack{\forall \text{ port} \\ p \text{ in } T}} P^{\lambda,T,p}_{a,b} \times \begin{pmatrix} \text{Max. capacity} \\ \text{of port } p \end{pmatrix} \geq \sum_{\forall q} A^{\lambda,q}_{a,b} \times x^q_{\max} + (\Delta^{\lambda}_{a,b})_{\max}. \tag{18}$$

Only one request should be assigned to a section using a nonmultiplexing technology. For a section that uses a multiplexing technology, no such bounds are imposed apart from the capacity constraints:

$$A^{\lambda,q}_{a,b} \leq 1 \times \sum_{\forall \text{ port } p \text{ in } T=2} P^{\lambda,2,p}_{a,b} + Q \times \sum_{\forall \text{ port } p \text{ in } T\neq 2} P^{\lambda,T,p}_{a,b}. \tag{19}$$

The equipment at the customer edge or at the cell tower should be simple. It is desired that there will be just one wavelength leaving the customer edge/BS tower that will be shared by all requests originating from the customer/tower and one wavelength entering the customer/tower shared by all requests destined to the customer/tower.

In this model, a customer edge or a backhaul tower at a node v is treated as a separate node v^t at a distance of 0 from v and connected only to v by a bidirectional link:

$$\text{for the link } l_{\text{out}} = (v^t, v), \quad \sum_{\forall a} \sum_{\forall b} SU^{\forall\lambda}_{l_{\text{out}},b} = 1, \tag{20}$$

$$\text{for the link } l_{\text{in}} = (v, v^t), \quad \sum_{\forall a} \sum_{\forall b} SU^{\forall\lambda}_{a,l_{\text{in}}} = 1. \tag{21}$$

C. Achieving Multi-Technology Oriented Networks

The result of the robust optimization is a network that deploys the minimum cost resources (equipment) while maintaining the ability to provision traffic with the given degree of robustness. The working of the optimization technique is seen in the results section and the proposed heuristic in the next section actually mimics the optimization loosely. The optimization attempts to achieve a cost-wise lower bound by keeping the data in the lower layers as much as possible without compromising on the advantages obtained through statistical multiplexing. In that sense, whenever there is a possibility of achieving statistical multiplexing, the optimization does so by inculcating higher-layer technologies.

V. Heuristic Approach

Our robust optimization technique (in Section IV) is computationally tractable for networks of up to 200 nodes, typical in the metro area, using multicore servers and CPLEX solver [19]. Thus, there is a need for a fast network planning heuristic, i.e., to assign equipment and provision traffic efficiently, especially for large networks. This heuristic must also be tractable in polynomial time to be able to be adapted to service provider networks.

To this end, we propose a compact heuristic that solves the static traffic assignment problem, which we then extend to solve the dynamic assignment problem.

The uniqueness of our heuristic is that it proposes a recourse function between the optical and the CE/CE + OTN layers. We consider CE + OTN + optical because this solution works out to be cheaper than using MPLS as seen in Section VI. This recourse function facilitates a toggle between the two layers attempting to reach a pricing minimum.

We begin by considering an all-optical solution, to which we assign CE/CE + OTN interfaces in a stepwise manner, such that any addition of equipment will

- result in the overall reduction of equipment cost, and
- be followed by a recourse function, whereby existing connections are moved to best (i.e., maximally) utilize the added equipment.

It must be noted that when all the assignments followed by recourse have been accomplished, the algorithm tapers to a pricing minimum. It is at that point that we take the assignment matrix as our solution.

This heuristic also has a dynamic component to it (Subsection V.B), which is a simple extension of this heuristic. It allows traffic to be added and provisioned, thereby replicating a realistic backhaul scenario.

A. Static Case

Let $G(V,E)$ denote the network graph and q_{ij} the static traffic value of the flow from node i to j ($i,j \in V$). Start with request q_{ij}:$q_{ij} \geq q_{i'j'}$, $\forall\, i',j' \in V$ and allocate it to wavelength λ. Repeat until either of the following cases arise.

1) Case 1: All the Wavelengths Are Allocated: In this case, we have at least one request that was not allocated wavelength. Two actions can now be performed:

- free existing allocations using CE, and
- free existing allocations using OTN + CE.

Without differentiating between these two actions, we perform the actions in following three steps.

Step 1. Select the heaviest flow with available residual capacity (on its wavelength), and add those unassigned flows that are graphically incident on this heavy flow to maximally utilize the CE/CE + OTN equipment. The heaviest flow is chosen to maximize statistical multiplexing. While choosing graphically incident unassigned flows, selecting the lightest flows would enable us to pack more graphically incident flows into the same port, thereby conserving wavelengths. Graphically incident flows are defined as shown in Fig. 5 and Algorithm 1. Note that for single-edge flows, the concept of graphically incident flows is not applicable.

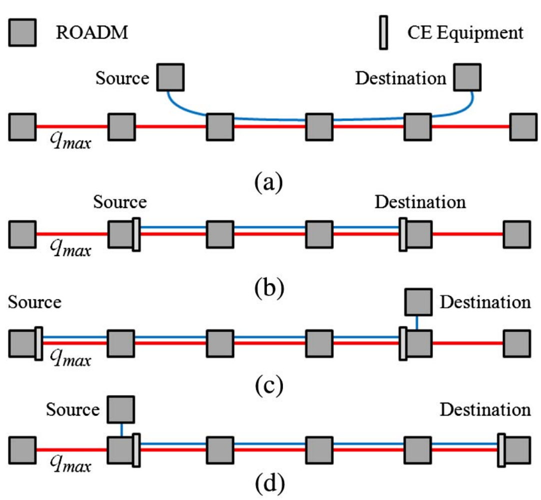

Fig. 5. Illustrative example of graphical incidence as defined. (a) Pair of requests is not graphically incident. (b)–(d) Pairs of requests satisfy the conditions and are graphically incident.

Algorithm 1 How to Identify Graphically Incident Flows?

$path(q)$:= path of request q
$subpath_{src}(q)$:= first two edges in $path(q)$
$subpath_{dst}(q)$:= last two edges in $path(q)$
$Q \leftarrow$ set of unallocated requests
$q_{max} \leftarrow$ largest flow request in Q
for all $q \in Q$ **do**
  if $\{path(q) \subseteq path(q_{max})\} \vee$
  $\{src(q) = src(q_{max}) \wedge subpath_{src}(q) \subseteq path(q_{max})\} \vee$
  $\{snk(q) = snk(q_{max}) \wedge subpath_{dst}(q) \subseteq path(q_{max})\}$ **then**
    allocate CE to q or to the entire subpath as appropriate
    $Q \leftarrow Q \backslash q$
  end if
end for

Step 2. Now that we have provisioned CE equipment in the network, we attempt to reroute preprovisioned requests, so as to maximally utilize the CE infrastructure.

- Denote by a and b the pair of nodes with CE equipment that were installed in the paths of two heavy requests.
- Consider the set of all requests that first pass through a and then through b. Transfer the heaviest request in this set to the CE layer.
- Multiplex the largest possible number of light requests from the same set to this heavy request.
- Repeat the previous two procedures for all such CE node pairs.

Step 3. Repeat Steps 1 and 2 until all connection requests are provisioned or robustness is achieved, i.e., $q = q_{peak}, \forall\, q$.

With this design, we will have a fully provisioned optical layer. However, this may not always result in the lowest-cost solution. Hence, we continue the algorithm as follows. At the culmination of Step 3, let us denote the cost of the network CapEx by C_3. Let us further denote the cost of the interface at node v that uses Ethernet ports, CE, and CE + OTN by $C_{Opt}(v)$, $C_{\mathrm{CE}}(v)$, and $C_{\mathrm{CE+OTN}}(v)$, respectively. Then

$$C_3 = \sum_{v} C_{Opt}(v) + C_{\mathrm{CE}}(v) + C_{\mathrm{CE+OTN}}(v). \tag{22}$$

The individual costs in the above equation will depend on the number of standard capacity transponders of each technology provisioned at a node (for example, at a node we may require 3×1 and 4×2.5 Gbps transponders). Now, can we reduce the number of customer-facing Ethernet interfaces used and increase the number of CE/CE + OTN interfaces such that C_3 decreases? For this, we continue the iteration of Steps 1 and 2 and terminate the algorithm when the network cost either saturates to a minimum or begins to increase. Finally, we will add mandatory OEO equipment to handle power/OSNR issues based on threshold margins.

2) Case 2: All the Requests Are Satisfied: In this case, the number of requests is such that we could successfully assign wavelengths to each of the requests. The network then contains only ROADMs. We do not require CE equipment to free up more wavelengths. Hence, we proceed with Step 3 in which we add CE equipment only as a measure to reduce the cost of the network.

3) Computational Complexity: The complexity of Step 1 is $O(Q)$ as it involves two linear scans (in the simplest case) through all the requests in the network. In Step 2, we search for the shortest path for each request that has not yet been provisioned. In an implementation using Dijkstra's algorithm, the complexity our heuristic is $O(Q(E + V \log V))$.

B. Dynamic Case

An extension of the previous heuristic can be used to provision requests as they arrive dynamically. In the dynamic case, there is no possibility to install new CE equipment after a request arrives. This reduces the problem to a dynamic routing and wavelength assignment problem [20] with the following enhancements:

- We can use wavelengths with CE flows, provided they have multiplexing capacity that can map new requests into a provisioned wavelength.
- Requests can move from one wavelength to another at those nodes, where CE equipment are deployed, using CE equipment as electronic wavelength translators.

1) Procedure: Start with the original topology and consider the first wavelength λ_1. If λ_1 is blocked on a link, then delete that link. If there is a CE flow from node a to node b that can accommodate the new request, then add a link from a to every node along the path of the flow that has CE equipment with free capacity. Repeat to create similar graphs for each wavelength. Call the set of such graphs G_λ.

Each of these graphs represents the free links for their respective wavelengths. As a request may change wavelengths after an OEO translation, we must search for a route in all of these graphs. This is accomplished by toggling between graphs. The toggling can be done only at the nodes with CE equipment, where a change of wavelength is possible. To enable such toggling, we merge the graphs as follows. Any node with CE equipment that has free capacity will have only one realization in the combined (superimposed) graph and all other nodes will have multiple realizations, one for each wavelength. If a path exists in this graph, we provision the request, else we drop it. We state the previously mentioned procedure in Algorithm 2.

Algorithm 2 Heuristic for Dynamic Case

v_{OEO} := number of nodes allowing OEO
v_{OOO} := number of nodes not allowing OEO
W := total number of wavelengths
$N := v_{OEO} + W \times v_{OOO}$
$Comp_G$:= zero matrix of size $N \times N$
for all $i = 1, 2, \ldots, N$ **do**
  for all $j = 1, 2, \ldots, N$ **do**
    if $i \le v_{OEO}$ and $j \le v_{OEO}$ **then**
      if $\exists \lambda$: $G_\lambda(i, j) = 1$ **then**
        $Comp_G(i, j) = 1$
      end if
    else if $i \le v_{OEO}$ and $j \le v_{OEO}$ **then**
      $\lambda = (j - v_{OEO}) \bmod W$
      $Comp_G(i, j) = G_\lambda(i, \frac{j - v_{OEO}}{W})$
    else if $i > v_{OEO}$ and $j \le v_{OEO}$ **then**
      $\lambda = (j - v_{OEO}) \bmod W$
      $Comp_G(i, j) = G_\lambda(\frac{i - v_{OEO}}{W}, j)$
    else if $(i - j) \bmod W = 0$ **then**
      $Comp_G(i, j) = G_\lambda(\frac{i - v_{OEO}}{W}, \frac{j - v_{OEO}}{W})$
    end if
  end for
end for

2) Example: Consider the network topology shown in Fig. 6, with the traffic to be mapped as shown in Table III, and assume we have three wavelengths (red, blue, and green). Assume CE ports of nonstandard capacity such as 5 and 20 Gbps. After initial allocation of as many requests as possible using only optical equipment, the free link paths on each wavelength are indicated by bold lines (Fig. 7). The provisioned and unprovisioned requests, after the optical allocation, are listed in Table III.

Take the lightest request, i.e., Request 1 and transfer it to a CE-enabled path. Subsequently, we deploy CE equipment at A and G. Of all the provisioned requests, only Request 2 is graphically incident on Request 1. Hence, we install CE equipment at D. Requests 1 and 2 share the blue wavelength on DG, thereby freeing the red wavelength on DG. The two requests have a total granularity of 3. Hence, we use a CE port with 5-Gbps capacity. We proceed with Step 2. No new requests could be provisioned in Step 2.

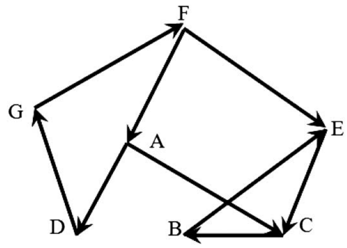

Fig. 6. Initial topology.

TABLE III
List of Requests

Request	Source	Destination	Granularity	Path	Color
1	A	G	1	ADG	Blue
2	D	G	2	DG	Red
3	A	B	2.5	ACB	Green
4	G	E	3	GFE	Blue
5	C	E	4	CBE	Red
6	F	D	4.5	GFAD	Green
7	E	B	5	ECB	Blue
8	F	C	6.5	FEC	Red
9	F	E	7	FE	Green
10	G	E	7.5	–	–
11	D	C	8	–	–

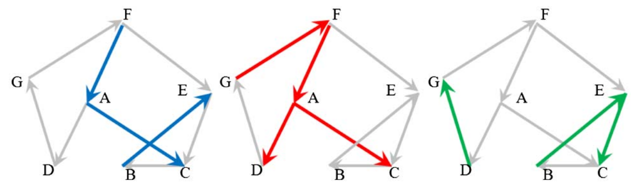

Fig. 7. Free link paths after optical allocation.

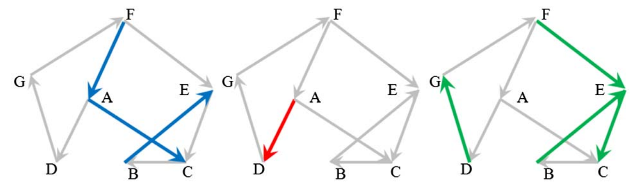

Fig. 9. Final state of the assignment.

Using the free wavelength, we can provision Request 11 along the path *D–G–F–A–C* on the red wavelength. This ends the first iteration and the network state is as shown in Fig. 8.

In the second iteration, we choose Request 3 on *A–C–B* to transfer to the CE layer. However, this action will not result in any immediate freeing up of wavelengths. Hence, Request 3 is provisioned at the optical layer. We then move Request 4 to the CE layer. Request 9 can share a channel with Request 4, thus freeing the green wavelength on link *FE*. For this, we use a CE port with capacity 10 Gbps at *F* and of capacity 5 Gbps at *G*. Now, we notice CE equipment at *G* and *E*, a path (*G–F–E* using the blue wavelength used by Request 4), and an unprovisioned request that can make use of this free path (Request 10). Hence, we provision Request 10 along *G–F–E* on the color blue. The total granularity of this flow is 17.5. Hence, we would need a 20-Gbps port at *G* and *F*. The final state of the network is then as shown in Fig. 9.

VI. Simulation Model and Numerical Results

We now present a simulation study to validate our proposed robust optimization model and heuristics. The simulation is organized in two parts: 1) individually compare the different technology options from a robust optimization perspective and 2) holistically compare network-wide technology alternatives. Our main result showcases network design using the two technology alternatives (MPLS + OTN + ROADM and CE + OTN + ROADM) and the effect of robustness on these choices.

A. Simulation Model

We simulate a large service provider network deploying mobile backhaul (3G/4G including long-term evolution services) as well as enterprise points of presence. The simulation models a greenfield network, where the provider augments an existing broadband network with a new roll-out of an independent 3G/4G mobile network. Our model assumes a 2000-site network spread over a 600-km region such that the farthest sites are 600-km apart. The network is modeled with a core mesh of number of nodes (16–64) in dense interconnection topology; each node in the core supports metro rings (each ring of 8–14 nodes); each node in a metro ring further supports an access network of 64–128 BS sites. The access network is also modeled as a ring (although a star overlay is the preferred option), to enable protection switching through ring-wrap-around. While sites can be either mobile towers or enterprise terminations, we assume the former as it represents stricter QoS requirements and finer granularity.

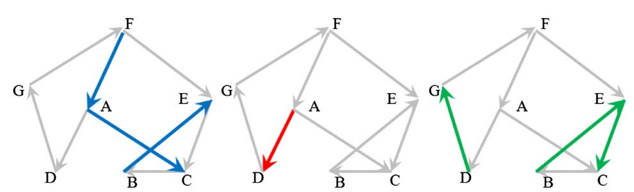

Fig. 8. Network state at the end of the first iteration.

Fiber with loss of 0.25 dB/km is assumed and the OSNR cut-offs for 1, 10, and 40 Gbps are assumed proportional to a BER of 1×10^{-13} with FEC. The region can be assumed to be circular with a diameter of 600 km. The longest optical path in the network is 1100 km. We are able to achieve all-optical paths up to 1000 km using OTN-based FEC, and beyond 1000 km using electronic regeneration. At each site, we can have a number of base stations generating traffic between 100 Mbps and 1 Gbps per base station. The distance between nodes is 1–8 km (with an average of 3.5 km). There are 50 MSCs. All the traffic from the BSs is brought to the MSC and then sent to the respective Internet gateways. There are 10 Internet gateways, five of which are active peer transit points to other providers and the remaining five are used for purpose of protection as well as internal routing. Table IV gives us the average and peak values of the traffic generated at a base station.

Table IV has been considered after discussions with two diverse service providers, one in the United States, the other in India [17,21], and the numbers are thereafter adjusted for future growth. At a node (site), the default architecture supports a ROADM option. If less than four channels are to be added/dropped, then the ROADM is replaced by an optical add–drop multiplexer (OADM; without WSS). The Layer 2-2.5 equipment can be MPLS or CE. OTN can be used in conjunction with MPLS or just standalone deploying ODU switching. The ROADM is assumed to support wavelengths at 25-GHz spacing resulting in

TABLE IV
Peak-to-Average Bandwidth Ratios

Provisioned Traffic at a Base Station	Peak Value	Average Value
100 Mbps	100 Mbps	40 Mbps
2 Mbps	4 Mbps	1 Mbps
1 Gbps	1 Gbps	500 Mbps
2.5 Gbps	1 Gbps	500 Mbps
10 Gbps	10 Gbps	1 Gbps

about 400 channels across the C and L bands. The CE and MPLS equipment has line cards that support 10 and 1 Gbps and a backplane that can scale up to a 720 Gbps nonblocking cross connect.

We simulated randomly arriving traffic demands. For a particular site, we assume that the value of N is known and corresponds to the number of traffic connections for that site. We assume k of those connections to be at their peak value (denoted by $x^q_{\max}$). We further assume $x^q_{\max}$ and x^q_{avg} values as shown in Table IV. Table IV represents the combined demands of all the active mobile users within the base station at a given time. At a site, we support multiple base stations. The average call request holding time (at a base station) was 65 s, while the maximum request holding time was 30 min. Requests could be 1) voice calls, 2) video files, 3) text downloads, and 4) interactive applications. Latency for each request had an upper bound based on the physical distance between the source–destination pair (computed randomly with an additional consideration of probability 0.75 that the destination would be outside the provider's network). If the request could not be provisioned within the latency bound, it was dropped. We assumed handoffs to occur as a random process during a session between adjacent cells. Due to the randomness of the handoffs, we did not model the mobility of the user.

Network load is computed as a percentile over the bandwidth supported with the maximum number of connections. This means that a load of 1 signifies the maximum bandwidth supported in the network (also implying the largest number of connections/bandwidth in the network). In essence, the number of mobile users attached to a base station is proportional to load.

The robust optimization was performed using CPLEX, and values obtained were fed to a MATLAB-based discrete event simulator that would compute complete network-wide parameters (blocking, cost). The simulation assumed a 3-year traffic pattern. The heuristic was simulated in MATLAB for the same traffic profile used in the optimization study. Equipment costs are listed in Fig. 2.

B. Preliminary Results

Shown in Fig. 10 is the CapEx (with equipment cost as shown in Fig. 2) as a function of network load for a value of $K = 0.5 \times N$. Shown in Fig. 11 is the amount of CapEx as a function of network load for $K = 0.3 \times N$, and shown in Fig. 12 is the CapEx cost for $K = 0.25 \times N$. Put together, Figs. 10–12 enable us to measure the impact of the choice of K. The take-away is that the choice of K is intrinsic to capacity planning, not just in the backhaul, but also in conventional transport networks. Hence, the use of robust optimization is well justified. There is no direct method to compute K, and only traffic statistics leading to forecasting of future growth can compute the value of K. The following conclusions are drawn from Figs. 10–12.

- The ROADM-only-based solution is the most expensive irrespective of the value of K except at very heavy loads.

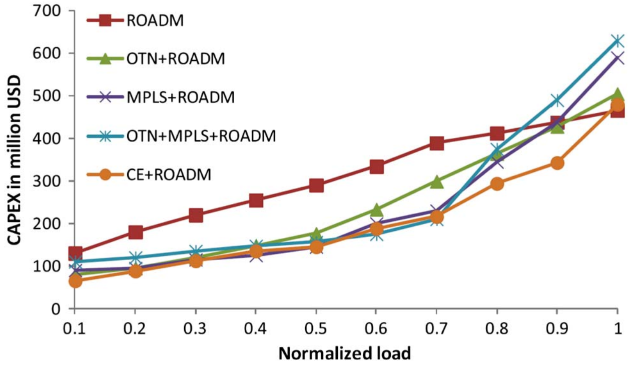

Fig. 10. CapEx as a function of load for $K = 0.5 \times N$.

At heavy loads, it makes sense to allocate wavelengths to sites, as the granularity mismatch between requests and wavelength capacity is minimum. Note that despite the high CapEx of the ROADM-based solution, the OpEx for this solution is the lowest.

- The CE solution (CE + ROADM) is on average the least expensive across all load configurations for all values of K.
- The OTN + MPLS + ROADM solution is the most expensive amongst all solutions, primarily due to the cost involved in packet processing, expensive MPLS routing gear, as well as OTN transponders.
- In certain cases of traffic, it makes sense to be all-optical beyond a load of 0.7. These cases are when the demands are such that the peak-to-average values are close. In such a situation, it makes more sense to have an all-optical core. Traffic uncertainty is detrimental to an all-optical backhaul network, and some degree of packet processing is then mandatory.

1) Effect of K on CapEx: There is a proportional relationship between CapEx and the value of K as seen in Figs. 10–12. This is of particular importance from the standards-based proposals pertaining to packet–optical integration, giving more emphasis on packet technologies as mandatory for lowering CapEx for moderate values of K (realistic).

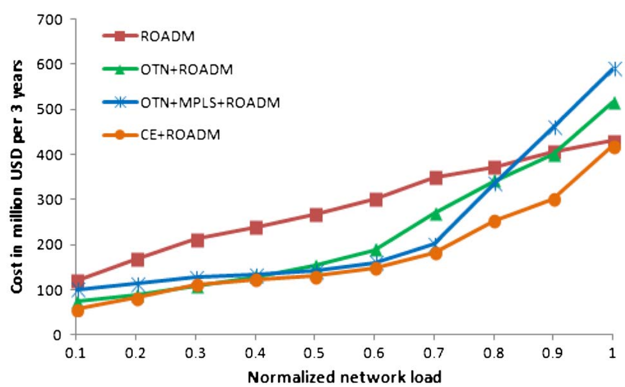

Fig. 11. CapEx as a function of load for $K = 0.3 \times N$.

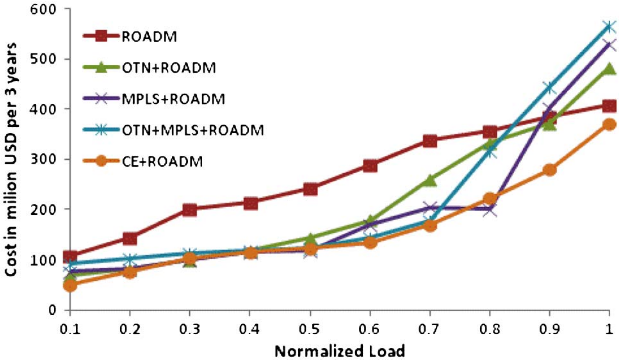

Fig. 12. CapEx as a function of load for $K = 0.25 \times N$.

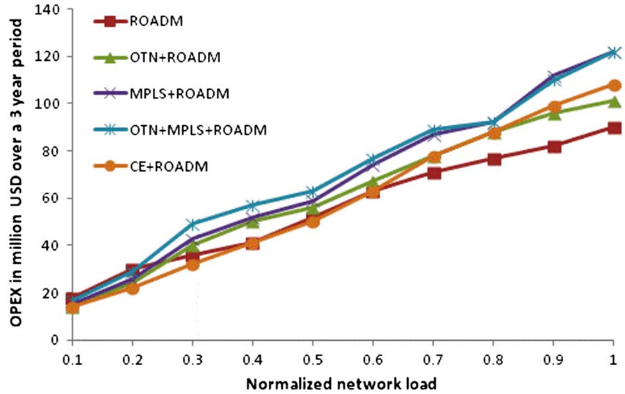

Fig. 14. OpEx cost for $K = 0.25 \times N$.

2) OpEx Results: While computing CapEx, we also compute the OpEx over a 3-year period. The OpEx was calculated based on power consumption. Except at lower loads, the OpEx is the lowest for the ROADM-based all-optical solution. OpEx viewgraphs are shown in Figs. 13 and 14 for values of $K = 0.5$ and $K = 0.25$, respectively. The energy consumption of the optical equipment is much less than that of the electronic equipment (for a 10-Gbps circuit, the energy consumption using ROADMs is 19 W, while for MPLS it is 170 W, and using CE it is 80 W [22,23]). The reliability and availability imply that the port down-time for optical is less than electronic equipment. We model the reliability values based on commercially available numbers [2], and the optical solution (ROADMs) leads to a lower OpEx value. This is not shown in Figs. 13 and 14 as the impact of reliability is negligible from a numerical value, but still is worth mentioning for sake of completeness. It is also worth noting that packet technologies, both MPLS and CE, consume significant OpEx, although CE is slightly lower (on average 5%–12%, varying with K). The reason why CE is lower is because of better statistical multiplexing obtained resulting in a lower number of ports implying lower OpEx.

Shown in Fig. 15 is a viewgraph of the blocking ratio (rate) of a connection request as a function of load. Figure 15 can be viewed as follows: the circuit technologies compete closely in a separate group, while the packet technologies compete in another group in the viewgraph area. The finely granular CE performs the best in terms of requests being met, while the ROADM-based wavelength switching saturates most rapidly, resulting in maximum blocking. The performance of OTN is similar to ROADMs, just that with ODU switching, there is more support for finer granularity (with subwavelength switching using OEO-based ODU switching). This results in slightly less blocking as compared to ROADMs. Amongst the packet technologies, MPLS LSPs have higher blocking than Ethernet-only frames due to higher packet processing complexity resulting in packet-drops.

Shown in Fig. 16 is a viewgraph of the amount of K that can be handled as a function of load. The y-axis represents the K to N ratio. The viewgraph shows how much K a particular technology can handle, while enabling a blocking of <0.00001 per request for a particular load. CE and ROADM solutions perform better than the other technologies. The CE-based solution is the best for low-to-medium loads, while the ROADM solution is the best for medium-to-heavy loads. The MPLS + ROADM solution starts off well, with similar performance to CE; however, it recedes significantly and consistently performs worse than all the technologies at medium-to-heavy loads. When OTN is added to the

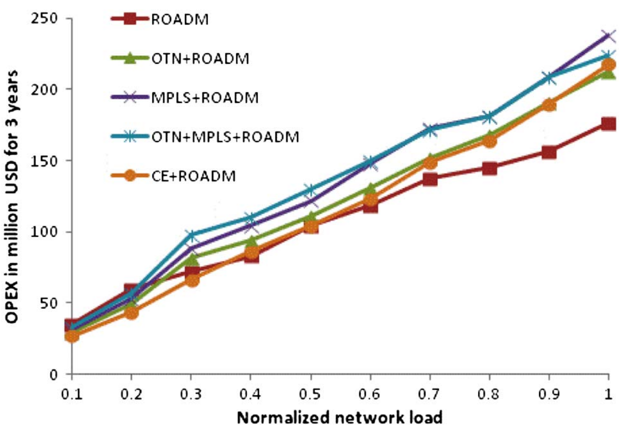

Fig. 13. OpEx cost for $K = 0.5 \times N$.

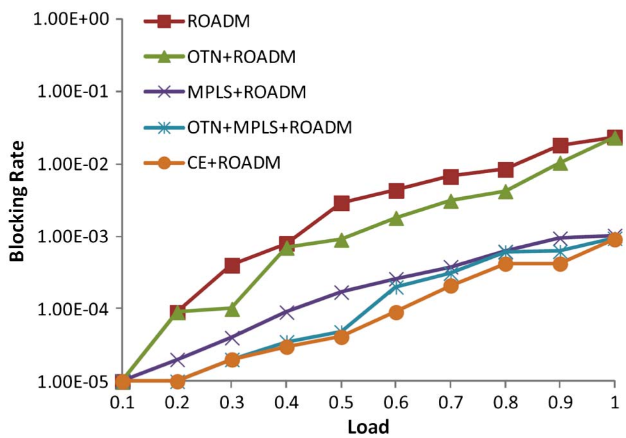

Fig. 15. Blocking rate as a function of load.

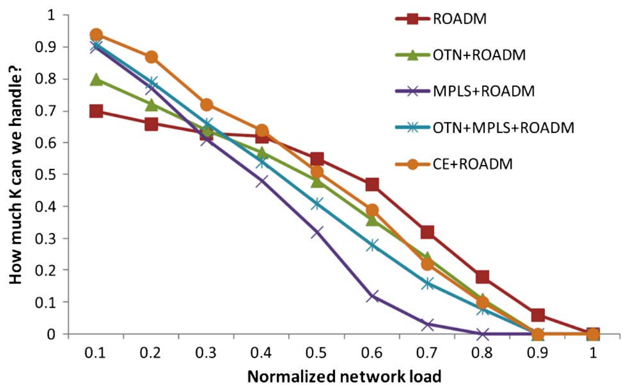

Fig. 16. How much K can we handle?

MPLS + ROADM solution (leading to MPLS + OTN + ROADM) the performance improves by about 10%, but is still worse than all other solutions.

C. Main Results

Shown in Figs. 17 and 18 are the technology compositions for the optimal mobile backhaul/core network using robust optimization.

We essentially compare the MPLS + OTN + ROADM (in Fig. 17) solution with the CE + OTN + ROADM solution (in Fig. 18). The idea is to understand what percentage of which technology does the optimal solution (CapEx + OpEx) use (at a blocking rate of 0.0001 indicating an acceptable service level agreement). For example, at a load of 0.5, we may use 35% ROADMs, 23% OTN equipment, and 42% MPLS equipment in the network. In the case of MPLS + OTN + ROADM, the use of MPLS steadily decreases, while the use of ROADMs steadily increases, crossing at about a load of 0.58. The OTN equipment in Fig. 17 is consistently used between 22% and 28%. It is also interesting to note that about 18.3% of all paths exceed 800-km reach, implying good utilization of OTN for reach enhancement.

Shown in Fig. 18 is the performance of the CE + OTN + ROADM-based solution. CE provides better statistical multiplexing, and hence its requirement in the optimal

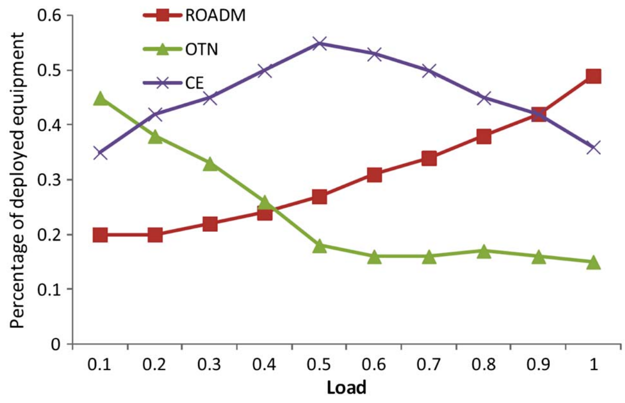

Fig. 18. CE, ROADM, and OTN solution.

solution is higher than MPLS. CE equipment usage increases with load, peaking at about 55% usage (starting from an initial usage of 35%) at a load of 0.5. In this solution, the ROADM usage steadily increases, from 20% to 49%. OTN is popular in this solution at the beginning, with good mapping of Ethernet frames onto OTN tunnels, but then, its usage recedes at medium to higher loads, when OTN acts primarily as an encapsulation technology for reach enhancement, rather than as a switching solution. Note that in this (CE + OTN + ROADM) solution, the ROADM usage increases proportionally with load; this is an important result implying that we can steadily add wavelengths (ports) to a ROADM in a deterministic manner.

From Figs. 17 and 18, we can conclude that packet technologies perform well under uncertainty, when there is significant room for so-called play; i.e., the peak-to-average bandwidth is large. Although we keep the peak-to-average bandwidth as the one shown in Table IV, we must note that with packet technologies, there is more flexibility in the choice of K. We also conclude that among all the solution costs shown in Figs. 17 and 18, the cost in Fig. 18 is less, implying that the CE + OTN + ROADM solution is lower in cost compared to the MPLS + OTN + ROADM solution.

Figure 19 compares the cost of the MPLS + OTN + ROADM network with the CE + OTN + ROADM network. This is an important result for greenfield networks, where

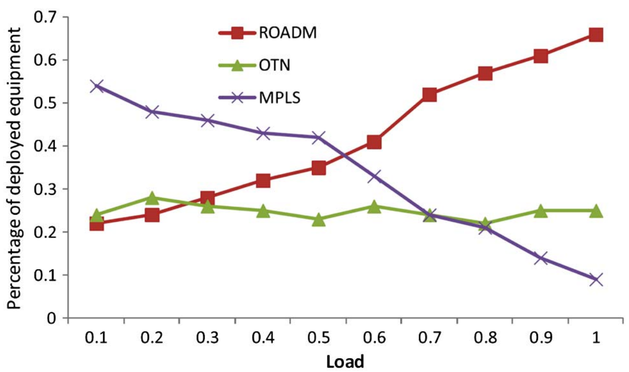

Fig. 17. MPLS, ROADM, and OTN solution.

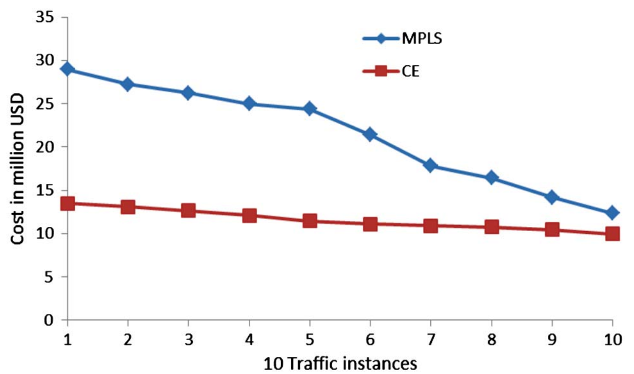

Fig. 19. Cost comparison of the two solutions at small traffic loads and for a 60-node network.

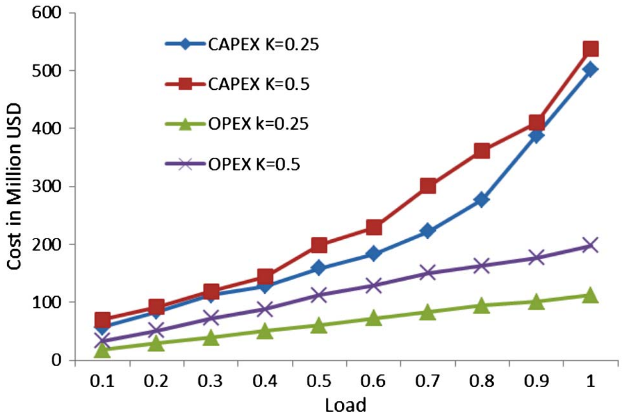

Fig. 20. Performance of the heuristic algorithm.

there is a critical choice to be made in terms of technology selection. The x-axis in Fig. 19 represent medium-to-full load conditions for a smaller 60-node network.

Shown in Fig. 20 is the performance of the heuristic algorithm for $K = 0.25$ and $K = 0.5$ for the 2000-node network. To plot these results, load in terms of the actual traffic matrix was the input to the discrete event simulator (DES) in MATLAB. In the DES, we assume that the circuit arrival rate is Poisson and the holding time is exponentially distributed. We further assume three traffic classes. These three traffic classes are 1) voice circuits (average holding time 3 min), 2) video download (average size 7.5 MB), and 3) interactive services (two-way circuits of average holding time 12 min). The heuristic implements only the CE + OTN + ROADM scheme, since from the optimization, we learned that this choice would lead to lower CapEx.

Shown in Fig. 21 is the comparison between the robust optimization solution and the heuristic solution for both CapEx and OpEx. Shown on the y-axis is the difference between the optimal and heuristic performance. In the CapEx results, for $K = 0.25$, the heuristic is on average 26% worse than the optimal, while for $K = 0.5$ the heuristic is 20% worse than the optimal. The OpEx performance of

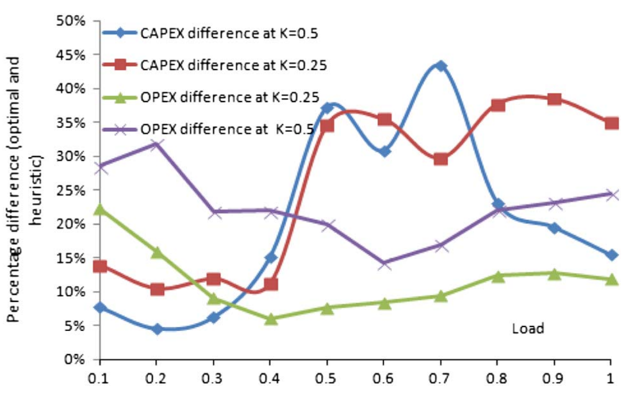

Fig. 21. Comparison between the dynamic heuristic and the optimal.

the heuristic is slightly better (as compared to its performance for CapEx). For $K = 0.25$, the heuristic is an average 12% worse than the optimal, while for $K = 0.5$, the heuristic is 23% worse than the optimal. It must be noted that solving the robust optimization problem for network sizes greater than 200 nodes takes several hours on a quad-core CPU, while the heuristic generates results within a few seconds on the same machine. It must also be noted that as the number of nodes increases, the time required for computation increases exponentially for the robust optimization solution.

VII. Conclusion

We have presented a solution to the multi-layer design problem by considering multiple technology choices and also applied it to the specific case of mobile backhaul design. To account for the unpredictable nature of traffic, we assumed knowledge of only the average value and the maximum value of each flow. We formulated a mixed integer linear program to find the optimal backhaul network that was robust to traffic variations. The robustness and cost of the network increased, when we allowed a larger number of requests to attain their peak values. Using this model, we plotted the optimal CapEx and OpEx when deploying a network of a single technology as we increase the load. We also plotted the percentage composition that each technology would have in the optimal configuration for a multi-layer network. The analysis suggests that packet–optical integrated technologies are more cost-efficient for networks with higher loads.

A problem of interest would be to consider a brownfield deployment that is completely orthogonal to a desired scenario (a poorly planned network) and examine how our robust optimization multi-layer multi-technology model can evolve to a plausible network architecture.

References

[1] S. Elby, "Bandwidth flexibility and high availability," presented at the Service Provider Summit, OFC/NFOEC, San Diego, Mar. 2009.

[2] IEEE 802.3 Ethernet Working Group, "IEEE industry connections Ethernet bandwidth assessment," IEEE 802.3 BWA Ad Hoc Report, 2012.

[3] D. Bertsimas, D. B. Brown, and C. Caramanis, "Theory and applications of robust optimization," *SIAM Rev.*, vol. 53, no. 3, pp. 464–501, 2011.

[4] T. G. Robertazzi, *Planning Telecommunication Networks*. Wiley–IEEE, 1998.

[5] Z. Zhang and A. S. Acampora, "A heuristic wavelength assignment algorithm for multihop WDM networks with wavelength routing and wavelength re-use," *IEEE/ACM Trans. Netw.*, vol. 3, no. 3, pp. 281–288, June 1995.

[6] K. Kar, M. Kodialam, and T. Lakshman, "Minimum interference routing of bandwidth guaranteed tunnels with MPLS traffic engineering applications," *IEEE J. Sel. Areas Commun.*, vol. 18, no. 12, pp. 2566–2579, Dec. 2000.

[7] S. Gast and A. Keimer, "Robust optimization in network flows," Tech. Rep., University of Erlangen, Nürnberg, 2009.

[8] D. Bertsimas and M. Sim, "Robust discrete optimization and network flows," *Math. Program.*, vol. 98, no. 1–3, pp. 49–71, 2003.

[9] R. Aparicio-Pardo, P. Pavon-Marino, and B. Mukherjee, "Robust upgrade in optical networks under traffic uncertainty," in *16th Int. Conf. on Optical Networks Design and Modeling (ONDM)*, 2012.

[10] D. A. Schupke, "Multilayer and multidomain resilience in optical networks," *Proc. IEEE*, vol. 100, no. 5, pp. 1140–1148, May 2012.

[11] A. Riedl and D. A. Schupke, "Routing optimization in IP networks utilizing additive and concave link metrics," *IEEE/ACM Trans. Netw.*, vol. 15, no. 5, pp. 1136–1148, Oct. 2007.

[12] M. Scheffel, R. G. Prinz, C. G. Gruber, A. Autenrieth, and D. A. Schupke, "NGL01-6: Optimal routing and grooming for multilayer networks with transponders and muxponders," in *IEEE Global Telecommunications Conf. (GLOBECOM)*, 2006.

[13] I. Katib and D. Medhi, "A study on layer correlation effects through a multilayer network optimization problem," in *Teletraffic Congr.*, 2011, pp. 31–38.

[14] E. Palkopoulou, D. A. Schupke, and T. Bauschert, "Quantifying CapEx savings of homing architectures enabled by future optical network equipment," *Telecommun. Syst.*, vol. 52, no. 2, pp. 1123–1129, Feb. 2013.

[15] P. Belotti, K. Kompella, and L. Noronha, "A comparison of OTN and MPLS networks under traffic uncertainty," submitted to *IEEE/ACM Trans. Netw.*

[16] "ROADMs and the future of metro optical networks. Heavy reading," May 2005 [Online]. Available: http://kt.uniza.sk/ktam/download/Studijne_materialy/Projektovanie_sieti/transport_network02.pdf

[17] A. Gumaste, P. Gokhale, T. Das, M. Purohit, and P. Agrawal, "Using global content balancing to solve the broadband penetration problem in the developing world: Case Study–India," *IEEE Commun. Mag.*, vol. 50, no. 5, pp. 74–81, May 2012.

[18] R. Ramaswami and K. N. Sivarajan, "Routing and wavelength assignment in all-optical networks," *IEEE/ACM Trans. Netw.*, vol. 3, no. 5, pp. 489–500, Oct. 1995.

[19] IBM Corporation, "IBM ILOG CPLEX V12.1: User manual for CPLEX," p. 157, 2009.

[20] O. Gerstel and S. Kutten, "Dynamic wavelength allocation in all-optical ring networks," in *IEEE Int. Conf. on Communication*, 1997.

[21] A. Gumaste and S. Akhtar, "Evolution of packet-optical integration in backbone and metropolitan high-speed networks: A standards perspective," *IEEE Commun. Mag.*, vol. 51, no. 11, pp. 105–111, Nov. 2013.

[22] A. Gumaste, S. Mehta, I. Arora, P. Goyal, S. Rana, and N. Ghani, "Omnipresent Ethernet technology choices for future end-to-end networking," *J. Lightwave Technol.*, vol. 28, no. 8, pp. 1261–1277, Apr. 2010.

[23] J. Baliga, R. Ayre, K. Hinton, W. V. Sorin, and R. S. Tucker, "Energy consumption in optical IP networks," *J. Lightwave Technol.*, vol. 27, no. 13, pp. 2391–2403, July 2009.