Towards a Robust and Green Internet Backbone Network

Xuezhou Ma
Department of Computer Science
North Carolina State University
Email: xuezhou_ma@ncsu.edu

Khaled Harfoush
Department of Computer Science
North Carolina State University
Email: harfoush@cs.ncsu.edu

***Abstract*—Networking infrastructure consumes a sizable fraction of the electricity supply. A network design model aimed at maximizing energy savings by *aggregating* traffic demand at a small set of resources, to put under-utilized resources to sleep, is offset by legacy models aimed at maximizing the network throughput by *spreading* the load across network resources. Traffic fluctuations and sudden spikes further complicate the problem. In this paper, we show that the problem is NP-hard, and propose a heuristic which targets optimal power usage for common traffic demand, while accommodating traffic fluctuations. The efficacy of our design is tested on two published backbone networks: NLR and NSFNET. Simulation results reveal that the proposed heuristic closely matches the results of the optimal algorithm in both energy savings and network utilization, while enjoying polynomial time complexity.**

***Index Terms*—WDM optical networks, Internet backbone, virtual topology design, energy efficiency.**

I. INTRODUCTION

The explosion of the Internet in reach and capacity is driven by a huge boost of the networking technology. The use of *wavelength division multiplexing (WDM)* exploits the full potential of optics and it is the primary choice to meet the growing demand [1]. WDM divides fiber bandwidth into tens of *wavelengths* each able of carrying traffic in the order of Gb/s. A virtual channel, *lightpath*, can be established between nodes by using one wavelength on each link along the path. Once established, it delivers information transparently such that the signal cuts through intermediate nodes without electronic switching. The set of all lightpaths then forms a *virtual graph* to deliver data packets. Determining an appropriate virtual graph and the associated traffic routing algorithm is referred to as WDM *virtual topology design*.

Previous design models mainly aim at maximizing the network *throughput* by *spreading* out the load evenly among network fabrics. That is reasonable because the throughput is limited by the bottleneck elements, whether at routers or links. Recently, however, *energy* concerns are highlighted as a sizable fraction of total electricity supply in U.S. was used by network equipment – nearly 10% in 2009 [2], and a significant part was devoured by Internet backbone infrastructure. A green Internet with as little as 1% lower power can save more than ten billion dollars per year. Most recent design models thus aim at minimizing energy consumption by *aggregating* traffic along fewer routes while allowing devices on other routes to sleep [3].

Hence, there is an obvious trade-off. On one side, backbone networks are typically over provisioned with bandwidth redundancy. A large part of electricity bill and heat dissipation costs are wasted by under-utilized resources with balanced traffic load. On the other side, Internet traffic is highly fluctuant, containing spikes that ramp up quickly on any links and/or nodes [4]. Concentrating the load on few active routes to save energy may cause the network to become vulnerable to sudden spikes, resulting in severe congestion. A virtual topology design that can handle both network *robustness* and energy conservation is desirable for backbone networks. Unfortunately, existing design models explore only one factor and not the other.

In this paper, we investigate network power and congestion models, and introduce a *linear programming (LP)* formulation to minimize energy consumption subject to traffic congestion constraints. The optimal solution is shown to be NP-hard. To make the solution feasible for real-sized networks, we propose a heuristic by decomposing the problem into two more tractable subproblems: 1) bounded congestion level for traffic fluctuations; and 2) optimal power usage for common traffic demand. The efficacy of our design is tested on two published backbone networks: NLR and NSFNET. The simulation results reveal that the proposed heuristic leads to energy savings and resource utilization closely matching the optimal solution but only with polynomial time complexity.

The rest of this paper is organized as follows. In Section II, we survey related research work. In Sections III and IV, we discuss the network power and congestion models. In Sections V and VI, we show the key idea and propose our design model. In Section VII, we present the results on two published ISP networks. We finally conclude in Section VIII.

II. RELATED WORK

For years, WDM virtual topology design has been formulated as an optimization problem maximizing network throughput given the input traffic. Approximate solutions were computed using linear programming and heuristic algorithms [1], [5], [6]. With the advent of the concept of green Internet [3], research efforts now mainly target network power models and LP formulations that minimize energy

978-3-901882-48-7 ©2012 IFIP

footprint [2], [7]–[9]. Unfortunately, none of them can handle both factors.

Moreover, existing design models consider only common demand, i.e., optimize over a single traffic matrix, and ignore the inherently dynamic nature of Internet traffic. Traffic dynamics were extensively studied and mitigated by optimizing traffic routing through *Traffic Engineering (TE)* techniques. For example, in [10], multiple representative traffic matrices were considered, and an optimal routing was found to minimize the maximum link utilization over the representative matrices. In [4], the authors design a heuristic to cope both traffic variations and common demand. While there is a large body of research in the TE field, the proposed solutions do not apply directly to the WDM virtual topology design due to the lack of consideration to the optical layer.

III. Network Architecture and Power Consumption Model

We target a backbone network running IP-over-WDM. As shown in Figure 1, a typical backbone structure consists of nodes interconnected by WDM fiber links. Each node is equipped with both *optical cross-connect (OXC)* and *core router*. At WDM-layer, the OXCs switch lightpaths transparently from input links to output links, and at IP-layer, core routers route data packets (when converted into electronic domain) over the lightpaths. Associated with a core router are several *access routers* that aggregate low-rate flows from local areas. Other devices essential to a WDM system are as follows: 1) A pair of *transponders* (labeled as T/R) are needed at the endpoints of each lightpath for data transmission and *EO/OE* conversions. 2) A pair of optical *multiplexer* and *demultiplexer* (labeled as *De/Mux*) are deployed at fiber ends to multiplex/demultiplex different wavelengths. 3) *Erbium-doped fiber amplifiers (EDFAs)* are placed along a fiber at certain distance intervals performing amplification.

In IP-over-WDM networks, virtual channels are configured in two different ways: *bypass* and *non-bypass*. In the bypass scheme, a lightpath directly connects the source and the destination (e.g., the solid line connecting nodes 1 and 3 in Figure 1). Information is delivered end-to-end without in-transit processing. In the non-bypass scheme, traffic undergoes OEO conversions and IP routing at every intermediate nodes (e.g., the dashed line connecting nodes 1 and 3 in Figure 1).

An IP-over-WDM network is a complex engineering structure containing different network devices. We next identify the devices that consume power most.

- Core routers are major contributors to total power usage. Energy is consumed for packet level processing such as memory access, scheduling, and table lookups. Modern technology clusters several components together to form one multi-chassis router whose power level increases in discrete steps depending on the number of ports (i.e., line cards) activated. For example, each working *OC-48* port in *Cisco 12008 GSR* consumes 70 watts [8]. We thus approximate router power usage, e^R, as two terms: a fixed term caused by the base system (i.e., chassis plus processor plus switching fabric), and a traffic-dependent term proportional to the amount of traffic passing through.
- Another primary power contributor is the transponder. According to the product data of Alcatel-Lucent WaveStar OLS [11], a pair of transponders of 10 Gbps lightpath consume $e^L = 146$ watts when working and $e^L = 0$ if the lightpath carries no traffic (i.e., standby mode).

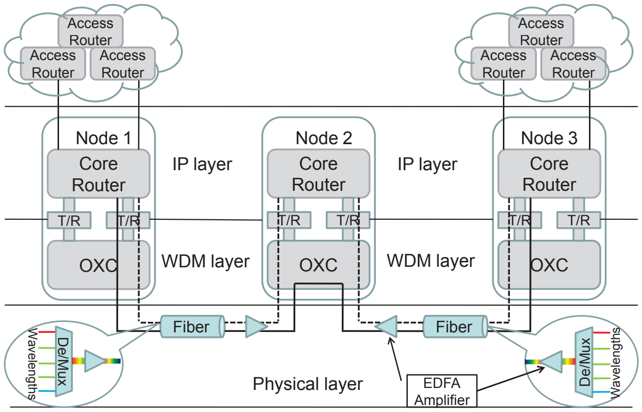

Fig. 1. Architecture of an IP-over-WDM backbone network.

Other devices consume minor power. For example, [7] estimates that one 8-watt EDFA is needed for every 80 km of fiber reach. Each MEMS-based OXC consumes power in the order of 0.45 watt per connection [9], which is negligible compared to core routers. Access routers are also not considered because they are out of our scope as we focus on backbone infrastructure.

In summary, the overall network power usage, E, is expressed as:

$$E = \sum_{i \text{ is a node}} e_i^R + \sum_{(i,j) \text{ is a channel}} e_{(i,j)}^L \tag{1}$$

where e_i^R is the router power usage at node i and $e_{(i,j)}^L$ is the power consumed by the channel between nodes i and j.

IV. Network Robustness to Traffic Spikes

In general, the Internet intra-domain traffic is predictable. It is not difficult for ISPs to estimate the average traffic volume with reasonable accuracy by considering customer subscription, daily peak hours, *etc.* However, estimating traffic fluctuations is difficult. To this end, representative traffic patterns are extracted based on history data and observed trends. These patterns serve as possible traffic spikes within next time window with granularity as fine as hour-to-hour [10]. If a design model does not incorporate this information, it may cause congestion at network devices, creating bottlenecks that limit a network throughput. We refer to the ability to handle traffic spikes as network *robustness*.

To investigate the robustness, we use the following model. Given traffic matrix T and a virtual topology design f, the utilization of a lightpath is the percentage of wavelength capacity that is used by traffic crossing the lightpath; the

TABLE I
NOTATIONS

Notation	Definition
$G = (V, E)$	physical topology with nodes V and links E
w	number of wavelengths per fiber
c^L	capacity of a wavelength
c^R	electronic switching capacity of a node
$u = \max\{u^R, u^L\}$	maximum (link and node) utilization
E	total network power consumption
$T = \{t^{sd}\}$	traffic matrix for common demand
$\mathbf{X} = \{X\}$	traffic matrices for possible traffic spikes
$\Omega = \{\omega_{ij}\}$	number of parallel lightpaths between nodes i, j
$\Pi = \{\pi^{ij}_{mn}\}$	$\pi^{ij}_{mn} = 1$ if lightpath (i, j) employs link (m, n) $\pi^{ij}_{mn} = 0$ otherwise.
$\Lambda = \{\lambda^{sd}_{ij}\}$	fraction of traffic between s and d that traverses lightpath (i, j)
$\widehat{\Lambda} = \{\widehat{\lambda}^{sd}_{ij}\}$	fraction of traffic between s and d that traverses lightpath (i, j) as part of a non-shortest path
$\widehat{\Omega} = \{\widehat{\omega}_{ij}\}$	number of working lightpaths of ω_{ij}

utilization of a router is the percentage of router capacity that is used by traffic crossing the router. The *maximum lightpath utilization (MLU)*, denoted by $u^L(f, T)$, is the maximum utilization of all lightpaths. The *maximum router utilization (MRU)*, denoted by $u^R(f, T)$, is the maximum utilization of all routers. MLU and MRU were widely used in previous network designs [1], [4], [6].

An optimal design f, which is most robust to traffic T, is the one that minimizes the maximum (lightpath and router) utilization $u(f, T)$. The resulting optimal utilization, $u^*(T)$, is given by

$$u(f, T) = \max\{u^L(f, T),\ u^R(f, T)\} \tag{2}$$

$$u^*(T) = \text{Minimize}_f\ u(f, T) \tag{3}$$

To compare different designs, the *performance ratio* of an arbitrary f is defined as:

$$p(f, T) = \frac{u(f, T)}{u^*(T)} \geq 1 \tag{4}$$

where $p(f, T)$ measures how far f is from being optimal.

Now, to account for different traffic patterns, we extend the defined metric to a set of traffic matrices $\mathbf{X}$. In particular, a design is based on an optimization minimizing $u(f, X)$ for all $X \in \mathbf{X}$, formally:

$$u(f, \mathbf{X}) = \max\{u(f, X) \mid \forall X \in \mathbf{X}\} \tag{5}$$

$$u^*(\mathbf{X}) = \text{Minimize}_f\ \ u(f, \mathbf{X}) \tag{6}$$

The performance ratio of an arbitrary f on $\mathbf{X}$ is:

$$p(f, \mathbf{X}) = \frac{u(f, \mathbf{X})}{u^*(\mathbf{X})} \tag{7}$$

Lower $p(f, \mathbf{X})$ translates to a more robust design with regard to the whole set $\mathbf{X}$.

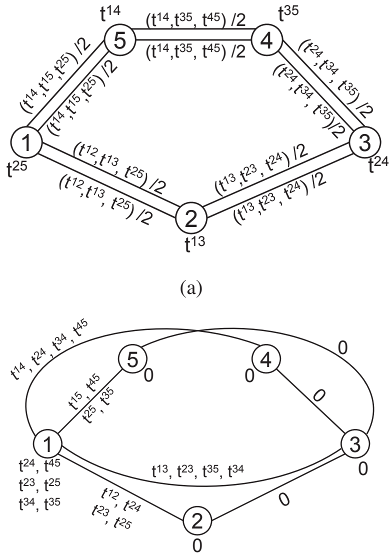

(a)

(b)

Fig. 2. Two virtual topology designs on a five-node two-wavelength network. (a) minimizing the maximum utilization; (b) minimizing the power usage. Assume $t^{sd} = 5$ Gbps for all node pairs, $w = 20$ Gbps and each router port has 15 Gbps capacity. Traffic flows that traverse each lightpath/node are exhibited in the figures.

V. PROBLEM STATEMENT

A. Terminology

Refer to Table I, network physical topology is represented by graph, G, with the set of nodes, N, and the set of links, L. It is assumed that each link is bi-directed composing of a single fiber supporting w wavelengths. c^L is the wavelength capacity, and c^R is the capacity of a routing node. Traffic matrix $T = \{t^{sd}\}$ denotes the predicted common demand between node pairs. $\mathbf{X}$ is a set of matrices with each element, $X = \{x^{sd}\}$, representing one possible traffic spike scenario. Ω describes which lightpaths (in terms of end nodes) to be established and Π describes how they are routed over physical links. The routing of traffic over lightpaths is then described by Λ. A virtual topology design f depends on Ω, Π, and Λ.

B. Problem Formulation

While both network robustness and energy conservation are desirable, the optimization aiming at individual objectives can lead to distinctly different designs. Refer to Figure 2 for an example. Figure 2(a) uses purely *non-bypass* scheme with no lightpath bypassing any intermediate node. Flows are balanced across the network. The resulting lightpath utilization $u^L = \frac{3 \cdot t^{sd}}{2 \cdot 20} = 37\%$, and router utilization $u^R = \frac{t^{sd}}{15} = 33\%$. The maximum utilization is thus $u = \max\{u^R, u^L\} = 37\%$. There are totally ten lightpaths and five router ports in use in Figure 2(a). On the other hand, Figure 2(b) combines *non-bypass* and *bypass* schemes. Flows are aggregated on fewer channels with $u^L = \frac{4 \cdot t^{sd}}{20} = 100\%$. Node 1 processes 30 Gbps in-transit tafffic while all other nodes are in idle. There are totally four lightpaths and two router ports activated in this case. The maximum utilization of Figure 2(b) is much worse than Figure 2(a), while its power consumption is 60% less.

A simple way to handle two objectives is to combine them into a single objective function, that is to optimize some function of both objectives. However, there is a significant trade-off between the size of the considered traffic set and energy optimality. In one extreme case, as the set expands to a complete space containing all possible demands, the resulting design is robust to arbitrary traffic spike but likely to produce poor energy savings on common demand. In our survey on several ASs, there are two primary policies towards a realistic design: (1) Profit-driven ISPs are not likely to compromise the guarantee of service for power reduction. Energy savings are pursued on top of the congestion. (2) Common demand lasts for most of the operation time while traffic spikes have much shorter duration. It is best to focus energy saving on normal operation.

We thus formulate the problem by separating the energy optimization for common traffic demand T and the bounded congestion level for traffic spikes $\mathbf{X}$:

$$\boxed{\begin{array}{ll}\textbf{Minimize}_f \text{ on } T: & E = \sum_i e_i^R + \sum_{(i,j)} e_{(i,j)}^L \\ \textbf{Subject to:} & (1)\ f \text{ is a virtual topology design} \\ & (2)\ u(f, \mathbf{X}) \leq 100\% \end{array}} \tag{8}$$

Formulation (8) is *reducible* to another NP-hard problem – minimizing the maximum lightpath utilization [6] – because testing $u^L \leq 100\%$ for all fs has the same rank as finding the minimum u^L. Solving this problem is numerically intractable for real-sized networks. This inherent complexity leads us to decompose the complete design into two subproblems: virtual graph layout (VGL) and traffic routing (TR) [1].

VI. A TWO-PHASE HEURISTIC

A. Key Idea

It should be noted that energy savings and robustness are mostly conflicting objectives when solving the TR subproblem while they are mostly in agreement when solving the VGL subproblem. In particular, VGL relies on *end-to-end* lightpaths to reduce (power and bandwidth) usage at routers; VGL relies on *hop-by-hop* lightpaths to improve (power and bandwidth) efficiency at link channels through electronic *traffic grooming*. So the optimality of the two factors is uniform when it comes to VGL. Meanwhile, we observe that a distinctly different TR is deployed as we target one factor instead of the other. For example, Figure 2(a) will concentrate the flows at one of the two parallel lightpaths on each link to save energy. Figure 2(b) will balance the load among links/nodes to reduce the maximum utilization. The idea is summarized in Figure 3.

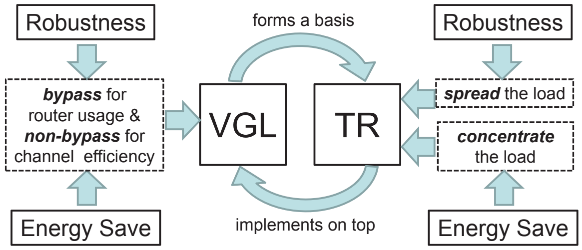

Fig. 3. Design strategies of two factors on solving each subproblem.

B. Phase I – Virtual Graphs with Bounded Congestion

In phase I, we find out the virtual graphs with bounded worst-case MLU/MRU against traffic spikes given physical network (i.e., G, c^L, c^R, w) and traffic matrices, $\mathbf{X}$. The MLU/MRU values are computed by assuming the *shortest-path-first (SPF)* traffic routing. The SPF algorithm is implemented in major ISP networks [12].

Recalling Equation (7), a virtual topology design f is said to have *penalty envelope* ϕ if the performance ratio of f on $\mathbf{X}$ is no more than ϕ ($\phi >= 1$) [4], namely:

$$\forall X \in \mathbf{X}, \quad \frac{u(f, X)}{u^*(\mathbf{X})} \leq \phi \tag{9}$$

By choosing ϕ slightly higher than 1, the virtual topologies satisfying (9) achieve near-to-optimal congestion. We next develop formula (9) in terms of MLU and MRU, respectively:

$$\forall X \in \mathbf{X}, \forall \text{ channel } (i,j), \ \sum_{s,d} \frac{x^{sd} \lambda_{ij}^{sd}}{\omega_{ij} \cdot c^L} \leq u^*(\mathbf{X}) \cdot \phi \tag{10}$$

$$\forall X \in \mathbf{X}, \forall \text{ node } i, \sum_{s,d: s \neq i} \ \sum_{j: \omega_{ij} > 0} \frac{x^{sd} \lambda_{ij}^{sd}}{c^R} \leq u^*(\mathbf{X}) \cdot \phi \tag{11}$$

Phase I is then formulated as the following LP problem and solved by testing if the objective is less than $u^*(\mathbf{X}) \cdot \phi$ where $u^*(\mathbf{X}) \cdot \phi \leq 1$:

Objective:

$$\max \left\{ \max\{ \sum_{s,d} \frac{x^{sd} \lambda_{ij}^{sd}}{\omega_{ij} \cdot c^L}, \sum_{s,d: s \neq i} \ \sum_{j: \omega_{ij} > 0} \frac{x^{sd} \lambda_{ij}^{sd}}{c^R} \} \mid \forall X \in \mathbf{X} \right\} \tag{12}$$

Variable: $\Omega = \{\omega_{ij}\}$

Subject to: $\forall X \in \mathbf{X}$,

- Traffic routing:

$$\lambda_{ij}^{sd} \in \{0, 1\} \text{ applying the SPF algorithm} \tag{13}$$

- Total flow on a lightpath:

$$\forall \text{ channel } (i,j), \quad \sum_{s,d} x^{sd} \lambda_{ij}^{sd} \leq \omega_{ij} \cdot c^L \tag{14}$$

- Lightpath routing:

$$\pi_{nm}^{ij} \text{ applies } k\text{-shortest-path algorithm} \tag{15}$$

- Number of channels on a link:

$$\forall \text{ link } (m,n), \quad \sum_{i,j} \pi_{mn}^{ij} \leq w \tag{16}$$

Note that the value of $u^*(\mathbf{X})$ is not known a priori. By default, $u^*(\mathbf{X})$ is computed separately, but this is a poor choice considering the repeated computation on the same instances. To avoid the redundancy, we maintain a variable called *best_result* during the search, such that $u^*(\mathbf{X})$ can be obtained on the run. Constraint (15) finds k shortest (link) paths between each node pair, and selects one for routing the lightpath. We set $k = 3$. Regarding the complexity, phase I has a total of $O(|V|^2)$ variables and $O(|V|^2)$ constraints.

C. Phase II – Traffic Routing with Optimized Power Usage

With the penalty envelope as a safeguard, phase II searches for an energy-minimized traffic routing during the normal operation, given common demand T and candidate virtual graphs from phase I.

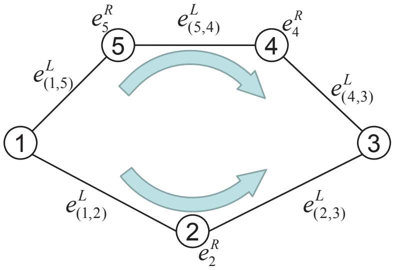

Fig. 4. Power consumption incurred by SP and non-SP routings.

While a SPF routing was assumed previously, traffic flows here are routed based on the minimization of power usage. As an example shown in Figure 4, there are two paths between nodes 1 and 3. The shortest path 1-2-3 consumes E_1, $e^L_{(1,2)} + e^R_2 + e^L_{(2,3)}$. The alternative path 1-5-4-3 consumes E_2, $e^L_{(1,5)} + e^R_5 + e^L_{(5,4)} + e^R_4 + e^L_{(4,3)}$. The alternative path is preferred if $E_1 \geq E_2$. This is possible because the bandwidth of a channel or router port is usually larger than a single flow. t^{13} may cost no power if it can be groomed into the already working components along 1-5-4-3. We use $\widehat{\Omega}$ to indicate the number of working lightpaths in Ω (see Table I).

One remaining issue is the convergence of the MLU/MRU bound. To account for the routing changes, we use $\widehat{\lambda}^{sd}_{ij}$ (see Table I) to measure the load increment produced by a non-SPF routing in phase II. From Constraints (10) and (11), one observes that $\widehat{\lambda}^{sd}_{ij}$ has to be smaller than $1-\phi\cdot u^*(\mathbf{X})$ to ensure MLU$<$ 100% and MRU$<$ 100%. Based on this observation, phase II is formulated as the following LP problem:

Objective:

$$\sum_{s,d:s\neq i} \sum_{j:\omega_{ij}>0} \beta \cdot t^{sd}\lambda^{sd}_{ij} + \sum_{i,j} e^L_{(i,j)} \cdot \widehat{\omega_{ij}} \tag{17}$$

Variables: $\Lambda = \{\lambda^{sd}_{ij}\}$

Subject to:

- Flow conservation at each node:

$$\forall s,d,i \in V, \quad \sum_j \lambda^{sd}_{ij} - \sum_j \lambda^{sd}_{ji} = \begin{cases} 1 & i = s \\ -1 & i = d \\ 0 & \text{otherwise} \end{cases} \tag{18}$$

- Total flow on a lightpath:

$$\forall \text{ channel } (i,j), \quad \sum_{s,d} x^{sd}\lambda^{sd}_{ij} \leq \omega_{ij} \cdot c^L \tag{19}$$

- MLU bound conservation:

$$\forall X \in \mathbf{X},\ \forall \text{ channel } (i,j),\ \sum_{s,d} \frac{x^{sd}\widehat{\lambda}^{sd}_{ij}}{\omega_{ij}\cdot c^L} \leq 1 - u^*(\mathbf{X})\cdot\phi \tag{20}$$

- MRU bound conservation:

$$\forall X \in \mathbf{X},\ \forall \text{ node } i,\ \sum_{s,d:s\neq i} \sum_j \frac{x^{sd}\widehat{\lambda}^{sd}_{ij}}{c^R} \leq 1 - u^*(\mathbf{X})\cdot\phi \tag{21}$$

In (17), β is effectively the marginal router power usage per unit traffic. Phase II has a total of $O(|V|^3)$ variables and $O(|V|^3)$ constraints for a sparse mesh virtual graph.

VII. PERFORMANCE EVALUATION

We compare four design models: 1) an energy-minimized design, *Energy-Min*, optimizing the objective function (1) only; 2) a congestion-minimized design, *MRU&MLU-Min*, optimizing the objective function (5) only; 3) an optimal design, *Optimal*, optimizing the formulation (8); and 4) our heuristic algorithm, *Heuristic*. Model 3) functions as a reference for the best solution it can achieve through model 4). For each model, we use the CPLEX software package to solve the corresponding LP optimization on a desktop with 3.0 GHz CPU and 2G memory. The performance is tested on the two backbone networks shown in Figure 5: 8-node 12-link NLR [13] and 14-node 21-link NSFNET [14]. The specifications of the two networks are summarized in Table II.

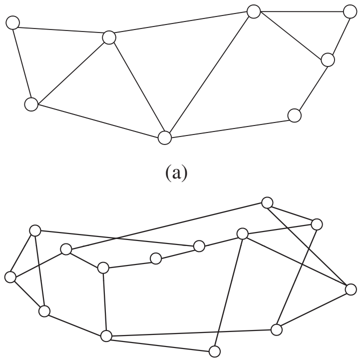

(a)

(b)

Fig. 5. Physical topology maps of (a) NLR and (b) NSFNET.

TABLE II
NETWORK SPECIFICATIONS

Network	w	c^L (Gbps)	c^R (Gbps)	e^L (Watts)	e^R (Watts)
NLR	8	5	120	75	200/15Gbps port
NSFNET	8	10	160	150	250/20Gbps port

We use the traffic model described in [15] to generate the predicted common demand. In particular, the required bandwidth between two nodes is proportional to the product of their populations given that PoP nodes are mostly located at major cities. To emulate traffic spikes, we create six representative

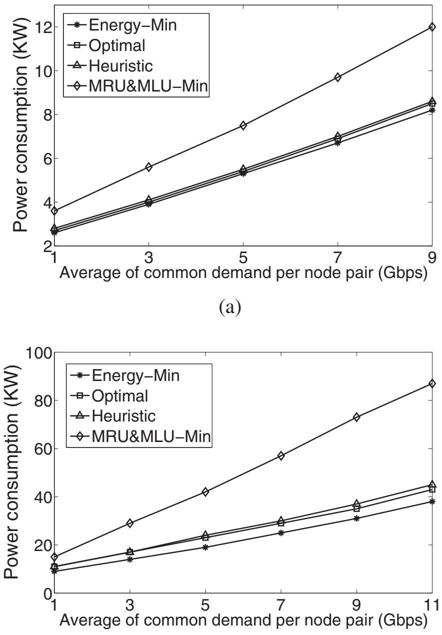

(a)

(b)

Fig. 6. Total power consumption. a) NLR, and b) NSFNET.

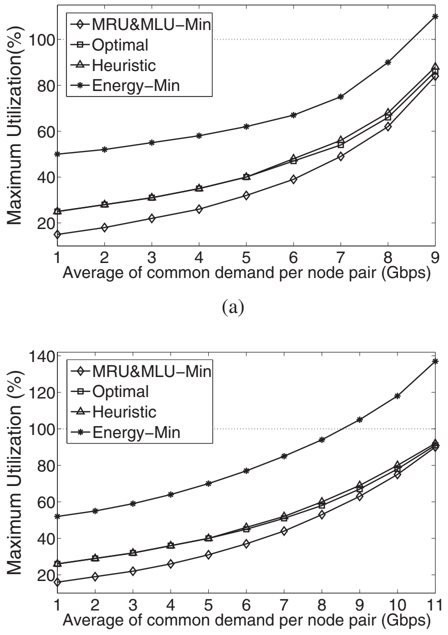

(a)

(b)

Fig. 7. The maximum utilization. a) NLR, and b) NSFNET.

traffic matrices. In each representative matrix, 10% of all node pairs are *randomly* selected to carry three times of their normal traffic rate. An average of the simulation results from 10 runs are shown. We set $\phi = 1.2$ throughout the simulation.

Figure 6 compares the resulting power consumption, E, of the design models. It is clear that *Energy-Min* consumes the least power and *MRU&MLU-Min* consumes the most. Four models keep the same power ranks in all tested networks. *Optimal* well tracks the *Energy-Min* curve with no more than 3% higher in NLR and no more than 11% higher in NSFNET. In both cases, *Optimal* is much superior to *MRU&MLU-Min* results. Also, the *Heuristic* is found to perform very closely to *Optimal*, thereby verifying the effectiveness of our heuristic algorithm on energy minimization.

Figure 7 compares the models in term of the maximum utilization under the occurrence of traffic spikes. As expected, *MRU&MLU-Min* yields the minimal value. *Energy-Min*, however, has much higher utilization than all other models. In general, the maximum utilization increases as the traffic rate between node pairs, showing nearly exponential growth. *Optimal* follows the trend of *MRU&MLU-Min*, and the two curves begin to merge at large traffic demands. *Heuristic* performs very closely to *Optimal*. *Energy-Min* model without considering congestion goes over the capacity limit as the traffic demand increases (e.g, $t^{sd} = 9$ Gbps in NSFNET), while *Optimal* and *Heuristic* always remain within bound.

We finally examine Figure 6 and Figure 7 together to see the overall performance of the design models. It is clear that an energy-minimized design suffers from severe link/node congestion when traffic fluctuates, and a congestion-minimized design consumes excessive energy. The proposed heuristic well matches the optimal solution. It balances the two metrics, achieving near optimal power consumption with the utilization slightly higher than the minimum. As shown in Table III, NSFNET network ($|V|$=14) requires 10^8 variables and 10^7 constraints, making the exact solution very expensive. Our heuristic uses only a small fraction of the time, and is thus well suited for real-sized networks.

TABLE III
COMPARISON OF COMPUTATION TIME ON NSFNET

Design Models	**# of variables**	**# of constraints**	**Computation time**
Optimal	$O(\|V\|^7)$	$O(\|V\|^6)$	22 hours
Heuristic	$O(\|V\|^3)$	$O(\|V\|^3)$	37 minutes

VIII. CONCLUSION

In this paper we introduce a new network design model by considering both energy savings and robustness to traffic spikes, which are “fundamental” challenges to ISP backbone networks. The proposed two-phase heuristic leads to close-to-optimal results on both two factors, while reducing the computation time to less than 40 minutes compared to 22 hours for the optimal algorithm of the simulated networks. Our model does not consider the power consumption of network cooling devices. It is known that heat dissipation has become a primary issue. Supplying sufficient cooling may cost several times more power than that delivered to routers. How to adapt our design to include cooling consumption is interesting for future work.

REFERENCES

[1] B. Mukherjee, *Optical WDM networks*. Springer Science, 2006.

[2] S. Huang, D. Seshadri, and R. Dutta, "Traffic grooming: A changing role in green optical networks," in *Proceedings of the 28th IEEE conference on Global telecommunications (GLOBECOM)*, 2009, pp. 1–6.

[3] M. Gupta and S. Singh, "Greening of the Internet," in *Proceedings of ACM SIGCOMM*, 2003, pp. 19–26.

[4] H. Wang, H. Xie, L. Qiu, Y. R. Yang, Y. Zhang, and A. Greenberg, "COPE: traffic engineering in dynamic networks," in *Proceedings of ACM SIGCOMM*, 2006, pp. 99–110.

[5] R. M. Krishnaswamy and K. N. Sivarajan, "Design of logical topologies: A linear formulation for wavelength-routed optical networks with no wavelength changers," *IEEE/ACM Transactions on Networking*, vol. 9, no. 2, 2001.

[6] K. Zhu and B. Mukherjee, "Traffic grooming in an optical WDM mesh network," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 1, pp. 122–133, 2002.

[7] G. Shen and R. S. Tucker, "Energy-minimized design for IP over WDM networks," *IEEE/OSA Journal of Optical Communications and Networking*, vol. 1, no. 1, pp. 176–186, 2009.

[8] J. Chabarek, J. Sommers, P. Barford, C. Estan, D. Tsiang, and S. Wright, "Power awareness in network design and routing," in *Proceedings of IEEE INFOCOM*, 2008, pp. 457–465.

[9] F. Vismara, V. Grkovic, F. Musumeci, M. Tornatore, and S. Bregni, "On the energy efficiency of IP-over-WDM networks," in *Proceedings of the IEEE Latin-American Conference on Communications*, 2010, pp. 1–6.

[10] C. Zhang, Y. Liu, W. Gong, J. Kurose, R. Moll, and D. Towsley, "On optimal routing with multiple traffic matrices," in *Proceedings of IEEE INFOCOM*, vol. 1, 2005, pp. 607–618.

[11] "Alcatel-Lucent WaveStar OLS 1.6T product specifications (2009)," (http://www.alcatel-lucent.com).

[12] G. Iannaccone, C. Chen-Nee, S. Bhattacharyya, and C. Diot, "Feasibility of IP restoration in a tier 1 backbone," *IEEE Journal on Networking*, vol. 18, no. 2, pp. 13–19, 2004.

[13] "National LambdaRail network (2006)," (http://www.nlr.net/).

[14] "NSFNET (1998)," (http://www.nsfnet-legacy.org/about.php).

[15] X. Ma, K. S., and K. Harfoush, "Towards realistic physical topology models for Internet backbone networks," in *Proceedings of the IEEE 6th International Symposium on High-Capacity Optical Networks and Enabling Technologies (HONET)*, 2009, pp. 36–42.