Green-PolyH: A Green Traffic Engineering Solution Over Uncertain Demands

Alejandro Ruiz Rivera and Kwan-Wu Chin
School of Electrical, Computer and Telecommunications Engineering
University of Wollongong, NSW, Australia
Email: ar506@uowmail.edu.au, kwanwu@uow.edu.au

Sieteng Soh
Department of Computing
Curtin University of Technology, WA, Australia
Email: s.soh@curtin.edu.au

***Abstract*—Green traffic engineering (TE) solutions play a central role in minimizing the energy expenditure of a network where they seek the minimal links/routers/switches required to support a given traffic demand. A key limitation, however, is that they are not designed to be robust against random traffic demands. To this end, this paper reports the first robust, green TE solution, called Green-PolyH, that considers demands defined by a given polyhedral set. Advantageously, Green-PolyH ensures all such demands do not cause the utilization of links to exceed a given threshold. Our experiments over well-known topologies show that savings above 80% are achievable whilst remaining robust to traffic changes.**

***Index Terms*—Green Networks, Energy Savings, Traffic Engineering, Hose Model, Polyhedra Set**

I. Introduction

Green TE approaches are relied upon by network operators during off-peak periods to reroute traffic with the aim of reducing the number of operating links/switches/routers. For example, the solution in [1] switches off links that are not part of any shortest path trees (SPTs) computed by routers. In [2], the authors aim to switch off the maximal number of links. Another example is [3] where routers are ranked either randomly, or according to their degree, number of flows or number of active neighbors. Routers are then shut down according to their rank or power consumption. Apart from that, researchers have also considered the impact on the Border Gateway Protocol (BGP) before shutting a link down [4].

All TE solutions must consider flow demands, which are succinctly captured by traffic matrices (TMs). With these TMs in hand, a TE solution can then devise a routing to optimize a given objective; e.g., minimize the maximum link load or reduce the number of operating network elements. However, as noted in [5], obtaining an accurate TM is expensive due to the volume of traffic. A key innovation that simplifies the representation of random traffic demands is a polyhedra set [6]; see Section II for details. Briefly, the set or polytope describes all possible demands that can be taken by source and destination pairs. In addition, the polytope is readily described in terms of inequalities and thus can be incorporated into a Linear Program (LP). Lastly, any convex combinations of the extreme points of the polytope are also valid TMs.

To date, no works have considered green TE over a polyhedra set. This is a critical gap because current green TE solutions, see [7], are not *robust* against random TMs. Consequently, they are unable to adapt to varying demands. Indeed, a new set of routers/links may need to be activated or deactivated every time demands change. In this respect, a straightforward solution is to allocate resources according to peak demands. However, doing so may yield many unnecessary links/routers/switches; i.e., there will be unnecessary energy expenditure. Conversely, too few resources cause congestion, high delays and/or packet loss. We note that existing TE works that consider polyhedra set, see [8] and references therein, do not consider the problem of minimizing the energy expenditure of a network.

Henceforth, this paper makes the following *contributions*. First, we formalize a novel problem that calls for the minimum resources to be used for a given polyhedra set. Second, we propose the first green TE solution that ensures all demands described by a polyhedra set are supported with the key constraint that no link utilization exceeds a given threshold. Advantageously, as the resulting solution is robust for all demands within the polyhedra set, a network operator does not need to recompute a new solution whenever the TM changes.

Next, we present our network model and define the polyhedra set formally. Then in Section III, we present the problem, followed by our approach in Section IV. Results from our experiments over two well-known topologies are presented in Section V. We conclude in Section VI.

II. Network Model

We model a connected network as a directed graph $G(V, E)$. The set V contains nodes/routers/switches and E is the set of edges. Each edge e has capacity c_e. Let $\Theta = \{(s,t) : s,t \in V, t \neq s\}$ be the set of commodities where source s has demand d_{st} for destination t. We will write the TM as $d \in \mathcal{R}^{|\Theta|}$ and d_{st} is a component of d; note, d is technically a vector but the term *traffic matrix* is ubiquitous in the literature. The possible values that d can take are governed by a polytope $\mathfrak{D}$ [6] that is defined as $\mathfrak{D} = \{d \in \mathcal{R}^{|\Theta|} : Ad \leq \alpha, d \geq 0\}$. Here $A \in \mathcal{R}^{\Gamma \times |\Theta|}$ and $\alpha \in \mathcal{R}^{\Gamma}$, where Γ is the number of constraints or inequalities.

As an example, consider a triangle topology with $V = \{1, 2, 3\}$ and the set of edges $E = \{(1,2),(2,1),(1,3),(3,1),(2,3),(3,2)\}$. Assume a special instance of the polyhedra set called the Hose model [9], where each node v has a total outgoing and incoming bandwidth that is denoted as C_v^+ and C_v^-, respectively. Moreover, we have $\Theta = \{(1,2),(2,1),(1,3),(3,1),(2,3),(3,2)\}$. In this example, a possible inequality is $d_{12} + d_{13} \leq C_1^+$, meaning the total outgoing demands from node 1 must not exceed C_1^+. Conversely, for incoming demands into node 2, we have $d_{12} + d_{32} \leq C_2^-$. In general, we have

$$\sum_{j \neq i, (i,j) \in \Theta} d_{ij} \leq C_i^+, \ \forall i \in V \quad (1)$$

$$\sum_{i \neq j, (i,j) \in \Theta} d_{ij} \leq C_j^-, \ \forall j \in V \quad (2)$$

We write P_{st} as the set containing the first K simple shortest paths for commodity (s,t); each of which is indexed by k. In other words, the k-th path is p_k^{st}. Note, K is an upper bound as some (s,t) pairs may have no more than K paths. We will use P_e to denote the set of paths that use link e. Hence, the total demand over a given link e is $L_e^d = \sum_{p \in P_e} B(p)$, where the function $B(p)$ returns the demand transmitted on path p given TM $d \in \mathfrak{D}$. Each link has utilization $u_e = L_e^d / c_e$ and it is bounded by γ. Lastly, we assume all links consume the same amount of energy; see [2].

III. The Problem

The problem at hand is to minimize the number of links used to route demands from the given set $\mathfrak{D}$ such that for *any* $d \in \mathfrak{D}$ (i) the demand for each commodity $(s,t) \in \Theta$ is routed over one path in P_{st}, and (ii) u_e of all link $e \in E$ is no more than γ; i.e., the utilization of each active link must be no more than γ for *any* TMs in $\mathfrak{D}$.

Consider topology (i) of Figure 1. Assume all links are undirected and have unit capacity. Also, let the required link utilization be no more than $\gamma = 0.8$. In addition, node-A has outgoing demands d_{AE} and d_{AF} that are constrained by the inequality $d_{AE} + d_{AF} \leq 0.8$. The total incoming traffic to both node E and F must be less than 0.8; i.e., $d_{AE} \leq 0.8$ and $d_{AF} \leq 0.8$. Given these inequalities, we thus have a polyhedron with extreme points $(0,0)$, $(0,0.8)$ or $(0.8,0)$. The aim is to switch off as many links as possible whilst supporting the demands from node-A. One possible solution is to switch off links (B,E) and (E,F) to yield topology (ii); a saving of 28.6%; two links out of seven have been switched off. Notice that the active links are able to support the said extreme points as well as any convex combinations of these points. The optimal solution, as depicted in Figure 1(iii), is to switch off three links, i.e., (A,C), (C,E) and (E,F); a saving of 42.9%.

We will now formalize the problem. Let $x_e \in \{0,1\}$ be a decision variable that is set to one if link e is active, and zero otherwise. With a slight abuse of notation, we will also use p_k^{st} to denote a binary decision variable that indicates whether

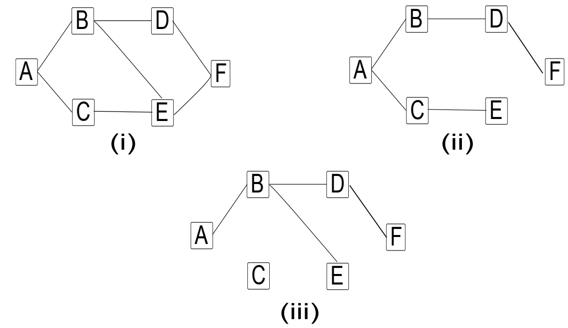

Fig. 1. (i) Original topology, (ii) a possible solution with two switched off links, and (iii) the optimal solution that supports the polyhedral set defined by the inequalities $d_{AE} + d_{AF} \leq 0.8$, $d_{AE} \leq 0.8$ and $d_{AF} \leq 0.8$.

the said path is chosen. We thus have,

$$\begin{array}{lll} \text{MIN} & \sum_{e \in E} x_e & \\ \text{subject to} & \sum_{k=1}^{|P_{st}|} p_k^{st} = 1, & \forall (s,t) \in \Theta \\ & x_e \geq p_k^{st} & \forall p_k^{st} \in P_e, \forall e \in E \\ & \max_{d \in \mathfrak{D}} \left\{ L_e^d \right\} \leq \gamma c_e, & \forall e \in E \end{array}$$

The first constraint ensures only one path for each commodity is chosen. The second constraint ensures x_e is one, i.e., a link is up, if there is a path using it. The third constraint ensures that for a given set of chosen paths or active links, the worst case TM d in $\mathfrak{D}$ does not cause the load of edge e to exceed γc_e. We remark that powering off switches can be considered by introducing a decision variable $s_j \in \{0,1\}$ for a switch j and setting it to zero when all its incident links are off. To ease exposition, we will ignore this straightforward extension.

Recall that each commodity in Θ has up to K paths. The search space in terms of paths is thus of size $K^{|\Theta|}$. Another key challenge is the number of extreme points in the polyhedra set $\mathfrak{D}$ grows exponentially with the number of demands; i.e., checking the third constraint is computationally expensive. Thus, it is not surprising that routing over polyhedral sets is NP-hard [10].

IV. Proposed Solution

We propose a heuristic algorithm, called Green-PolyH, that iteratively moves paths away from links with low utilization. In each iteration it removes a link with low utilization that is not in the final solution from the network. It also reroutes all paths that traverse the link. It then checks the utilization of the remaining links when fed demands from the given polyhedra set. If the resulting link utilization exceeds γ, the link is added into the final solution and Green-PolyH reverts back to the previous routing solution. Otherwise, the link is removed permanently.

It is important to note that the resulting solution, i.e., network topology, is valid for any demands within the polyhedra set. This means if the demands during an off-peak period are captured succinctly using a polyhedra set, then Green-PolyH produces a topology that remains valid for the duration of the period.

With the aid of Algorithm 1, we will now present the details of Green-PolyH. We will use the set Δ_t to record links that have been removed temporarily, and $\mathcal{A}$ is a set containing links included in the final solution. Initially, for each commodity, we route its demand over its shortest path; let R denote the set of paths for demands in Θ. Unused links are added into Δ_t and removed from the network temporarily; see lines 2 to 6. For each given routing R, Green-PolyH solves an LP called *LP-MaxUTIL()* to determine the total demand traversing a link (see line 9 to 12); the formulation for *LP-MaxUTIL()* will be explained later. The total demand of link e is stored in l_e; the set $\mathcal{L}$ is used to store the total demand of each link. We also note that in practice *LP-MaxUTIL()* is only applied on links that are traversed by at least one path.

Green-PolyH then checks the link utilization of all links. In particular, if the maximum utilization across all links is less than γ, then there is an opportunity to reroute commodities away from the least loaded link and thereby switch off said link. This is the goal of lines 13 to 21. First, in line-14, temporary links are removed permanently. This is reasonable because the current network $G(V, E)$ has sufficient capacity to handle all demands in $\mathfrak{D}$. Then in line-15, the current routing is saved; we will need to revert back to this routing if removing the link selected in line-16 results in at least one link with utilization that exceeds γ.

After that, the function *Reroute()* is called to reroute commodities as follows: (i) Select a link e^* in G, that is not in the set $\mathcal{A}$ with the *lowest* utilization. If all links in G are already in $\mathcal{A}$, then return "DONE", (ii) Reroute *all* commodities corresponding to paths in P_{e^*} onto another path that does not involve link e^*. If not successful, then add e^* into $\mathcal{A}$ and go back to Step (i). Otherwise, i.e., all paths in P_{e^*} have been rerouted, add e^* and any unused links into Δ_t and return the new routing R, $\mathcal{A}$, and Δ_t. Lastly, in line-20, it removes the link in Δ_t from $G(V, E)$ before restarting the *while* loop.

Green-PolyH may have removed a critical link; i.e., one that is required to ensure the worst case demand from $\mathfrak{D}$ does not cause link utilization to exceed γ. To this end, lines 22 to 27 address the scenario where one or more links have utilization beyond γ. First, Green-PolyH adds the links in Δ_t, so called critical links, to $\mathcal{A}$. This means in the next iteration, *Reroute()* will no longer consider these links. Then Green-PolyH restores these critical links, see line 24, and the previous routing solution; see line-26.

As mentioned, Line-10 of Algorithm 1 calls an LP solver to compute *LP-MaxUTIL*. Its aim, for a given routing R and a link e, is to determine the maximum aggregated demand that traverses e. Let $I^{st}_{e,R}$ be an indicator function that returns one if commodity (s, t)'s path, as determined by routing R, traverses link e. Then the following LP solves for a TM d in $\mathfrak{D}$ that maximizes the load of link e.

$$\underset{d \in \mathfrak{D}}{\text{MAX}} \sum_{(s,t) \in \Theta} I^{st}_{e,R} \times d_{st} \tag{3}$$

Note, it is possible that the value generated by (3) exceeds link e's capacity, i.e., $l_e > c_e$. This is not a concern because the goal is to identify whether the current routing R causes link e's utilization, i.e., u_e, to exceed γ.

Algorithm 1: Green-PolyH

input : $G(V, E)$, Θ, P_{st}, $\mathfrak{D}$

output: $\mathcal{A}$ – the set of active links

1 $\Delta_t = \mathcal{A} = R = \emptyset$

2 **for** $(s, t) \in \Theta$ **do**

3     $R = R \cup$ RouteShortest(G, P_{st})

4 **end**

5 Δ_t = GetUnusedLinks(G, R)

6 G = RemoveLinks(G, Δ_t)

7 **while** *true* **do**

8     $\mathcal{L} = \emptyset$

9   **for** $e \in E$ **do**

10         $l_e =$ `Solve` LP-MaxUTIL(G, e, R, $\mathfrak{D}$)

11         $\mathcal{L} = \mathcal{L} \cup l_e$

12   **end**

13   **if** $MAX\{\frac{l_e}{c_e} \mid l_e \in \mathcal{L}\} \leq \gamma$ **then**

14         $\Delta_t = \emptyset$

15         $R_{temp} = R$

16   [$\mathcal{A}$, Δ_t, R, $Code$] = Reroute(G, $\mathcal{A}$, $\mathcal{L}$, R)

17   **if** *Code ==*'*DONE*' **then**

18   Return $\mathcal{A}$

19   **else**

20   G = RemoveLinks(G, Δ_t)

21   **end**

22   **else**

23         $\mathcal{A} = \mathcal{A} \cup \Delta_t$

24   G = RestoreLinks(G, Δ_t)

25         $\Delta_t = \emptyset$

26         $R = R_{temp}$

27   **end**

28 **end**

We conclude this section with a few key facts.

Proposition 1. *Green-PolyH ensures all commodities in Θ remain connected at all times; i.e., the resulting graph induced by links in the set $\mathcal{A}$ is connected.*

Proof. In line-3, each (s, t) is routed on the shortest path. As the network is connected, the proposition is true. Consider an arbitrary iteration k where all $(s, t) \in \Theta$ are connected, and the utilization of all links is below γ. Assume link e_1 has the lowest utilization. Then *reroute()* either (i) successfully establishes an alternative path for all commodities traversing e_1, meaning all rerouted commodities remain connected, or (ii) commodities over e_1 cannot be rerouted and thus link e_1 is added into $\mathcal{A}$. As no links on paths traversing link e_1 have been removed, all commodities remain connected. Observe also that if a commodity has only one path, then all links on the path will eventually be included into $\mathcal{A}$. Lastly, assume case (i) and e_1 has been removed temporarily from G in iteration line-20, and at iteration $k + 1$, the utilization of one or more links exceed γ. Then as per line-24, Green-PolyH reverts back to the previous routing of iteration k, which by assumption connects all commodities. Also, Green-PolyH restores the link in Δ_t. Hence, the proposition for this case is also true. □

As noted in Section II, the poyhedra set $\mathfrak{D}$ is defined by Γ inequalities and $|\Theta|$ commodities. If we assume the Hose model [9], then for a given network, *LP-MaxUTIL()* contains $|\Theta|$ decision variables and $2|V|$ inequalities; each node has a constraint that bounds its aggregated incoming and outgoing demands. This means the size of LP solved by *LP-MaxUTIL()* is proportional to the network size and number of commodities. The number of times Green-PolyH calls *LP-MaxUTIL()* is stated in the next proposition.

Proposition 2. *LP-MaxUTIL() is called at most* $|E|^2$ *times.*

Proof. We start by showing that the *while* loop, i.e., line 7-28, repeats at most $|E|$ times because some links are removed from G in Line 6 and 20. In each iteration, a link is either added into $\mathcal{A}$ or removed permanently from G; the former occurs when rerouting is unsuccessful or a new routing causes high utilization. The latter happens if Green-PolyH successfully reroutes all paths from the link. Recall that Step (i) of *Reroute()* ignores links in $\mathcal{A}$. Hence, in both cases, in subsequent iterations, Green-PolyH ignores links in $\mathcal{A}$. Now, as *LP-MaxUTIL()* is carried out on a link-by-link basis, Green-PolyH thus calls *LP-MaxUTIL()* no more than $|E|^2$ times. □

V. EVALUATION

Our experiments are conducted in MATLAB. We use two well known topologies: Abilene (11 nodes, 28 links) and AT&T (25 nodes, 112 links). We present the results from only two experiments: *Exp-1* and *Exp-2*. In both experiments, without loss of generality, we will consider the Hose model [9], and set $C_i^+ = C_i^-$; we will denote this symmetric case as $C_i^\pm$. We note that the Hose model succinctly captures the variation in demands because C_i^+ and C_i^- can be respectively set to the outgoing and incoming capacity of a network interface card. In *Exp-1*, we study the impact of $C_i^\pm$ on energy savings, in terms of percentage of shutdown links, when $|\Theta| = 10$ and $|\Theta| = 50$. In *Exp-2*, we fix $C_i^\pm$ to a random value from the range $[0, \gamma]$ at the start of each experiment. We then study the impact of increasing number of (s,t) pairs. We run each experiment 20 times. Lastly, as per [11], we set the MLU (γ) to 80%.

Figure 2 presents the result from Exp-1. As expected, when the network load is low, i.e., $|\Theta| = 10$, more savings are observed; AT&T has a saving above 80% at this load but this decreases to 67% when $C_i^\pm$ approaches γ. A similar trend is observed for Abilene, with energy savings decreasing from 44% to 41.4%. When $|\Theta| = 10$, many links will have a low utilization, and thus, there are more opportunities to reroute flows. However, as the load of these ten (s,t) pairs increases, it becomes increasingly difficult to reroute a flow to reduce the number of operating links without causing one or more links to exceed the specified MLU. We now examine the case where $|\Theta| = 50$. As link utilization tends to be high, even when $C_i^\pm = 0.1$, there is little or no opportunity to reroute flows in order to reduce the number of active links. Hence, for all tested values of $C_i^\pm$, the percentage of shutdown links corresponds to the maximal number of links that can be switched off safely without impacting robustness.

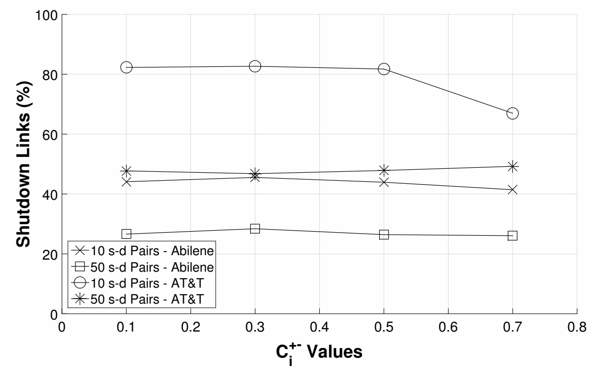

Fig. 2. Percentage of shut down links versus increasing load.

The result from Exp-2 is shown in Figure 3. For both topologies, energy savings reduce when network load increases. In particular, for the lowest network load, i.e., $|\Theta| = 10$, Green-PolyH is able to shut down 82% and 43.4% of the links for AT&T and Abilene, respectively. For the highest network load, i.e., $|\Theta| = 35$, energy savings reduce to 56% and 26% for AT&T and Abilene, respectively.

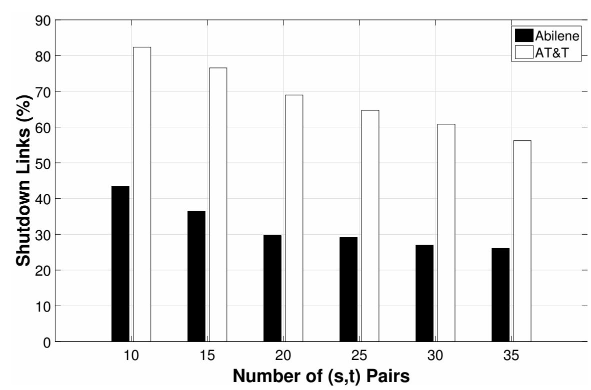

Fig. 3. Percentage of shut down links versus increasing (s,t) pairs or $|\Theta|$ values.

VI. CONCLUSION

We have presented the first green and robust TE solution that ensures active links have the capacity to support *any* demand from a given polyhedral set. This is significant because our solution considers random traffic matrices when shutting down links. Experimental results over two well-known topologies confirm the efficacy of our solution in terms of saving energy and ensuring the utilization of all links is below a given threshold.

References

[1] A. Cianfrani, V. Eramo, M. Listanti, M. Marazza, and E. Vittorini, "An energy saving routing algorithm for a green OSPF protocol," in *IEEE INFOCOM*. San Diego, CA: IEEE, Mar. 2010.

[2] M. Zhang, C. Yi, B. Liu, and B. Zhang, "GreenTE: Power-aware traffic engineering," in *IEEE ICNP*, Kyoto, Japan, Oct. 2010.

[3] L. Chiaraviglio, M. Mellia, and F. Neri, "Minimizing ISP network energy cost: Formulation and solutions," *IEEE/ACM Transactions on Networking*, vol. 20, no. 99, pp. 463–476, Apr. 2011.

[4] A. Ruiz-Rivera, K.-W. Chin, and S. Soh, "A novel framework to mitigate the negative impacts of green techniques on BGP," *Elsevier Journal of Network and Computer Applications*, vol. 48, no. 2, pp. 22–34, Feb. 2014.

[5] P. Tune and M. Roughan, "Internet traffic matrices: A primer," in *Recent Advances in Networking (Volume I)*, H. Haddadi and O. Banaventure, Eds. ACM SIGCOMM eBook, Aug. 2013.

[6] W. Ben-Ameur and H. Kerivin, "Routing of uncertain traffic demands," *Optimization and Engineering*, vol. 6, pp. 283–313, 2005.

[7] A. Bianzino, C. Chaudet, D. Rossi, and J.-L. Rougier, "A survey of green networking research," *IEEE Communications Surveys Tutorials*, vol. 14, no. 1, pp. 3 –20, 2012.

[8] C. Chekuri, "Routing and network design with robustness to changing or uncertain traffic demands," *ACM SIGACT News*, vol. 38, no. 3, pp. 106–129, Sep. 2007.

[9] N. G. Duffield, P. Goyal, A. Greenberg, P. Mishra, K. K. Ramakrishnan, and J. van der Merwe, "A flexible model for resource management in virtual private networks," *ACM/SIGCOMM Computer Communications Review*, vol. 29, no. 4, pp. 95–109, 1999.

[10] M. Minoux, "Robust network optimization under polyhedral demand uncertainty is np-hard," *Discrete Applied Mathematics*, vol. 158, no. 5, pp. 597–603, Mar. 2010.

[11] A. Nucci, N. Taft, P. Thiran, H. Zang, and C. Diot, "Increasing the link utilization in IP over WDM networks," in *SPIE ITCom*, Boston, MA, Jul. 2002, pp. 55–75.