

Disruption-Free Green Traffic Engineering with NotVia Fast Reroute

Ning Wang, Christos Michael and Kin Hon Ho

Abstract— We introduce a practical green traffic engineering technique in fast reroute (FRR) enabled IP networks which is able to avoid transient routing disruption caused by sleeping-mode reconfigurations during runtime. Following the description on the time-driven “delegate before sleep” operations based on the NotVia FRR platform, we also propose a simple optimization algorithm for selecting sleeping links and determining the sleeping duration according to dynamic traffic behaviors. The aim of the approach is to avoid traffic congestion due to reduced network capabilities during off-peak operation time.

I. INTRODUCTION

With global efforts on the migration towards future green ICT, the current practice of resource over-provisioning and “always-on” strategies in operational ISP networks may no longer be regarded as “cool” solutions in the future. One specific proposal to achieve energy efficiency is to reconfigure network elements such as routers or links to the sleeping mode when the traffic loading is low [1-4], for instance during off-peak hours. Specifically, customer flows are diverted to a set of alternate delivery paths, providing opportunities for “idle” devices to enter the sleeping mode for energy reduction.

It is important to note that topology changes (e.g. removal of links from the working topology) during runtime are generally deemed to be harmful in IP networks due to the routing re-convergence problem. Currently this issue has not been directly tackled by any energy-efficiency proposal involving sleeping-mode reconfigurations. The only work that mentioned this issue is [4], and the proposed approach is to rely on a hybrid IP+MPLS solution for fast traffic diverting purposes. In this letter we describe a simple scheme that enables seamless sleeping-mode reconfigurations in pure IP based network environments for practically enabling green traffic engineering (TE), but without causing any transient traffic disruption to end users. The basic idea is to periodically put a set of network links to the sleeping mode following a *time-driven* manner. An offline algorithm is proposed for computing the *Scheduled Sleeping Link Group (SSLG)* for which affected traffic can be seamlessly diverted on to alternative paths involving only active links during off-peak time. We propose to use existing IP Fast Reroute (FRR) techniques for such a purpose. Specifically, based on the given physical network topology and off-peak time traffic matrix (TM), a set of sleeping links can be computed *a priori* to form the SSLG. When the network

operation enters the *scheduled* off-peak time, customer traffic is first rerouted (delegated) onto the alternate paths enabled by FRR mechanisms before the links in the SSLG simultaneously go to the sleeping mode for energy saving. The routing switches back to the normal IP routing only when all the SSLG links wake up at the end of the off-peak time. We consider NotVia [5] as the underlying IP FRR technique for supporting path delegation, as it is the most sophisticated FRR scheme that is able to achieve 100% reroute coverage in any connected network topology. Such a feature is vital for disruption-free green TE, as it is able to support sleeping-mode reconfiguration on any arbitrary network device, when desired. It is important to note that the use of the NotVia technique in the proposed green TE does not interfere with its original functionality for fast failure recovery during peak time, as the SSLG is allocated with a dedicated NotVia address that is never activated during this period. On the other hand, in this letter it is not our aim to elaborate link failure protection during off-peak time with reduced network capability which is much more complicated.

In this letter we first describe how disruption-free green TE can be practically operated in NotVia-enabled IP networks. A greedy offline heuristic is then proposed for selecting pre-scheduled SSLG links to enter the sleeping mode. Meanwhile, the planned sleeping duration also needs to be determined by taking into account the reduced network capability against changing traffic behaviors. Finally, implications to network resilience will also be discussed.

II. GREEN TE IN NOTVIA-ENABLED NETWORKS

A. Basic Description

Before specifying the proposed *delegate-before-sleep* operation in green TE, we first describe the basic NotVia mechanism for failure protection. NotVia relies on special IP addresses assigned to each protected interface for enabling FRR. The semantics of a NotVia address is that “a packet addressed to a NotVia address must be delivered to the router advertising that address, not via the protected component with which that address is associated”. When a link failure occurs, the repairing router encapsulates the affected packets to a NotVia address for the protected interface. The actual diverting path towards the NotVia address (a.k.a. NotVia tunnel) is the alternate shortest path to reach the *next-hop* but not including the protected link. From this tunnel endpoint, the packets will follow the native IGP path towards the final destinations. In case of multiple link failures that can be predicated in advance, NotVia paths can be also computed according to the anticipated *Shared Risk Link Group (SRLG)* information, where customer traffic is rerouted without traversing any SRLG link.

Ning Wang and Christos Michael is / was with University of Surrey in U.K. (e-mail: n.wang@surrey.ac.uk, michaeldkaizoku@hotmail.com).

Kin-Hon Ho was with the Hong Kong Polytechnic University (e-mail: polyuroyho@gmail.com).

As far as our proposed disruption-free green TE is concerned, we compute a set of network links to form a Scheduled Sleeping Link Group (SSLG), based on which the corresponding NotVia delegation paths are pre-configured for enabling seamless traffic diversion during the transition period from normal network configuration to sleeping mode. An offline algorithm is responsible for determining a set of links to be included into the SSLG which is scheduled to enter the sleeping mode during a *pre-determined* off-peak period, e.g. during midnights and weekends when the traffic loading is expected to be low. Once the network operation enters the off-peak period, each delegating router, i.e. the head node of a link (or multiple links) belonging to the SSLG, is responsible for first diverting affected customer traffic onto the NotVia delegation path(s) *before* sleeping. Specifically, the delegating router identifies the destination prefixes that will be affected by the shutdown of the sleeping link(s), and then encapsulates the packets towards these destinations by using the NotVia address(es) with SSLG awareness. As such, these packets are diverted onto the NotVia tunnel(s) without traversing *any* of the SSLG links scheduled to enter the sleeping mode. Only when each delegating router makes sure that there is no more packet being sent onto the local SSLG link(s) to enter the sleeping mode, it can actively put it/them to sleep.

B. A Greedy Heuristic for Computing the SSLG

Before describing the algorithm for computing the SSLG during off-peak operation time, we first explain how traffic matrices (TMs) play a key role here. Today's operational networks are often periodically measured based on which distinct TMs can be derived, for instance at every 15-minute interval in the GEANT Network [6]. We analyzed the published GEANT TM dataset [7] across 7 consecutive days (trace starting from Monday noon), and we found that the maximum link utilization (MLU) for these 600+ TMs varies significantly (from 31% to 98%) during the period, as shown in Fig. 1. We can notice that the traffic pattern is regular on the daily basis except that the MLU during the weekend is much lower than that during weekdays. Our green TE objective is to compute an optimized SSLG to form a network topology with reduced capacity which can be applied during the operation period where the network loading (reflected by MLU) is expected to be low according to historically measured TMs.

Now we specify the offline algorithm for determining the SSLG according to the following input: (1) the physical network topology, (2) a specific traffic matrix T^* during the off-peak time, and (3) a pre-defined MLU threshold to avoid excessive link removals that may cause traffic congestion in the residual topology during off-peak time operation. The determination of T^* can be based on the historical measurement, and according to our greedy approach, we select the specific TM that produces the lowest MLU with the original full topology. The fundamental *Utilization Constraint (UC)* for computing the SSLG is that *any MLU resulted from the residual network topology after SSLG link removal during off-peak time should not exceed the "worst-case" MLU with the full topology during peak-time hours (pMLU)*. The pMLU metric refers to the highest MLU across the entire network operation period, for instance 98% in the GEANT network (shown in Fig. 1). The proposed greedy heuristic aims to put the

maximum number of links to sleeping mode as long as this UC is satisfied. Of course ISPs may apply other policies in setting the UC according to own considerations. The calculated residual topology is then used to determine the *duration* (window size) for the off-peak period in actual network operations. Specifically, historically measured traffic matrices are mapped onto this reduced working topology. For those consecutive TMs whose resulting MLU is below pMLU the corresponding time could be put into the off-peak period. We notice from Fig. 1 that, although the daily traffic pattern in GEANT follows similar style, it might not be the case that the resulting off-peak window size falls into exactly the same time everyday due to deviations day-by-day. As such, in order to ease the *time-driven* reconfiguration complexity in practice, it is also required to define a *unified* off-peak time duration on regular daily bases (e.g. the scheduled off-peak operation needs to start/finish exactly at the same time point every weekday). Generally, effective tweaking of the window size helps to cope with traffic uncertainty in real operational ISP networks.




Fig. 1. MLU performance in GEANT across 7 days

Input:

- Physical network topology $G = (V, E)$ with V being the *node* set and E being the *link* set;
- Traffic matrix T^* during off-peak time;
- Worst-case peak-time MLU (pMLU) as threshold.

Output:

- The SSLG scheduled to enter sleeping mode (L^s).

Step 1: Sort all network links in E in *ascending* order according to their utilizations based on T^* ; $L^s = \emptyset$.

Step 2: Consider removing the first link l in the ordered list by diverting the customer flows traversing l onto the alternative shortest path bypassing all the links in $\{l\} + L^s$.

Step 3: **If** the anticipated utilization of all links on the delegation path is still below or equal to pMLU, **and** the residual topology remains connected,

then remove link l from the topology; $L^s = L^s + \{l\}$.

Otherwise go to step 5.

Step 4: Update the projected utilization of the links along the delegation path, and reorder the link list in the remaining topology according to the updated utilizations.

Step 5: Go to Step 2 for the next link to be considered for removal until all the links have been examined.

Fig. 2. The proposed greedy algorithm for computing the SSLG

Fig. 2 shows the proposed offline algorithm for determining sleeping links to form the SSLG. The overall strategy can be described as follows. The initial SSLG is set to be empty. Then all the network links are sorted in ascending order according to

the mapping of T^* onto the physical network topology. Starting from the link with the least utilization, we consider removing it from the topology, and then update the utilization of all the links along the alternative shortest path without involving this link. If none of these links has its utilization exceeding pMLU, then we put this link to the SSLG. When this procedure is repeated for additional links to be potentially included into the SSLG, it is important to note that the corresponding alternative paths should not include those links that have already been put into the SSLG. Finally the NotVia delegation paths are actually configured based on the SSLG determined by the algorithm.

C. Network Resilience Considerations

We now discuss resilience issues during practical network operations. First of all, the configuration of the NotVia based *delegation* paths for the purposes of green TE does not affect the configuration of normal NotVia *protection* paths against network failures. SSLG-aware NotVia addresses used for green TE operations are separated from those associated with conventional SRLGs for multi-failure protection. As far as the latter is concerned, dedicated and even multiple SRLGs can be defined and allocated with separate NotVia addresses. During *peak-time* operation hours, upon normal link failures, the conventional NotVia protection tunnels can be activated according to the actual SRLG. If the failure cannot be restored before the beginning of the scheduled off-peak time, then the reduced topology should not be activated. In addition, during *off-peak* hours where the NotVia delegation tunnels are activated based on the SSLG, then upon unexpected network failures, NotVia protection tunnels cannot be activated for enabling FRR protection. The reason for both scenarios is that, traffic being carried over a NotVia path cannot be further diverted through another NotVia path in order to avoid forwarding loops caused by multiple levels of NotVia tunneling [5][8]. How to protect against failures during off-peak period will be our future research topic.

III. PERFORMANCE EVALUATIONS

We use the power consumption metric introduced in [2] in our performance evaluation based on the GEANT network data, which takes into account different numbers of amplifiers associated with each physical link. We first investigate the overall power reduction performance with the greedy approach in which the TM with the lowest MLU (31%) is used as T^* for computing the SSLG. According to our results, 28 out of 74 links can be included into the SSLG for sleeping subject to the UC, resulting in 44% energy reduction for that specific TM. If this computed topology is used during the off-peak time, then the resulting weekday duration for off-peak operation is 5 hours (covering 20 TMs) per day, plus the 35-hour weekend off-peak duration (excluding Saturday morning when the peak-time configuration is needed) for which the UC is satisfied. This configuration leads to overall 15.4% energy reduction across the entire week compared to the full-topology configuration.

Fig. 3 shows the actual MLU performance dynamics by using the greedy approach which leads to 28 links to be included in the SSLG, with daily scheduled peak time and off-peak time topology alternations across the week. The MLU

remains the same as shown in Fig. 1 during the determined peak time, while for the rest of the time the MLU performance is based on the off-peak time topology (in the dark area). As we can see, the MLU performance is always kept below the pMLU (98%). On the other hand, the off-peak MLU may increase in comparison to Fig. 1, which is the result of reduced network capabilities during the period. Although MLU is the selected TE metric for evaluating the performance of the proposed technique in this letter, other types of TE metrics such as network cost (e.g. the common one defined in [9]) can be also used for additional evaluation purposes.




Fig. 3. Actual MLU performance with time-driven SSLG reconfigurations

IV. CONCLUSION

We introduced a practical *delegate-before-sleep* green TE scheme in NotVia-enabled IP network environments, without incurring transient routing disruptions caused by runtime link sleeping-mode reconfigurations. The proposed simple greedy algorithm for computing the SSLG and its operation duration has offered potential power-saving gains according to our experiments based on the real GEANT network data without degrading of network performances. It will be our future research work to explore more intelligent approaches for further energy-saving performance enhancements.

ACKNOWLEDGEMENT

The research leading to these results has been performed within the UniverSelf project and received funding from the European Community's Seventh Framework Program (FP7/2007-2013) under grant agreement n° 257513.

REFERENCES

- [1] S. Nedeveschi, L. Popa, G. Iannaccone, "Reducing Network Energy Consumption via Sleep and Rate-adaptation", Proc. ACM NSDI 2008
- [2] L. Chiaraviglio, M. Mellia, F. Neri, "Energy-aware Backbone Networks: A Case Study", IEEE ICC 2009
- [3] W. Fisher, M. Suchara, J. Rexford, "Greening Backbone Networks: Reducing Energy Consumption by Shutting Off Cables in Bundled Links", Proc. ACM SIGCOMM Workshop on Green Networking, 2010
- [4] M. Zhang, C. Yi, B. Liu, B. Zhang, "GreenTE: Power-Aware Traffic Engineering", Proc. IEEE ICNP 2010
- [5] M. Shand, S. Bryant and S. Previdi, "IP Fast Reroute Using Not-Via Addresses," draft-ietf-rtgwg-ipfrr-notvia-addresses-07, April 2011
- [6] The GEANT Network: <http://www.geant.net>
- [7] The TOTEM Traffic Engineering Toolbox, <http://totem.run.montefiore.ulg.ac.be>
- [8] M. Menth, M. Hartmann, R. Martin, T. Cicic, A. Kvelbain "Loop-Free Alternates and Not-Via Addresses: A Proper Combination for IP Fast Reroute?", Computer Networks, Vol. 54, Issue 8, 2010, pp. 1300-1315
- [9] B. Fortz, M. Thorup, "Optimizing OSPF/IS-IS Weights in a Changing World", IEEE JSAC Vol. 20, Issue 4, 2002, pp. 756-767