

AMPLE: An Adaptive Traffic Engineering System Based on Virtual Routing Topologies



Ning Wang, University of Surrey

Kin Hon Ho, Hong Kong Polytechnic University

George Pavlou, University College London

ABSTRACT

Handling traffic dynamics in order to avoid network congestion and subsequent service disruptions is one of the key tasks performed by contemporary network management systems. Given the simple but rigid routing and forwarding functionalities in IP base environments, efficient resource management and control solutions against dynamic traffic conditions is still yet to be obtained. In this article, we introduce AMPLE — an efficient traffic engineering and management system that performs adaptive traffic control by using multiple virtualized routing topologies. The proposed system consists of two complementary components: *offline link weight optimization* that takes as input the physical network topology and tries to produce maximum routing path diversity across multiple virtual routing topologies for long term operation through the optimized setting of link weights. Based on these diverse paths, *adaptive traffic control* performs intelligent traffic splitting across individual routing topologies in reaction to the monitored network dynamics at short timescale. According to our evaluation with real network topologies and traffic traces, the proposed system is able to cope almost optimally with unpredicted traffic dynamics and, as such, it constitutes a new proposal for achieving better quality of service and overall network performance in IP networks.

INTRODUCTION

Traffic Engineering (TE) is an essential aspect of contemporary network management. Offline TE approaches aim to optimize network resources in a static manner, but require accurate estimation of traffic matrices in order to produce optimized network configurations for long-term operation (a resource provisioning period each time, typically in the order of weeks or even longer). However, these approaches often exhibit operational inefficiencies due to frequent and significant traffic dynamics in operational networks. Take the

published traffic traces dataset in the GEANT network as an illustration. The actual maximum link utilization (MLU) dynamics is substantial on a daily basis, varying from less than 40 percent during off-peak time to more than 90 percent in busy hours [2]! As such, using one single traffic matrix as input for offline computing a static TE configuration is not deemed as an efficient approach for resource optimization purposes in such dynamic environments.

Traffic engineering for plain IP-based networks (we will be referring to these as IGP-based networks, as is common in the literature since they route traffic based on the Interior Gateway Protocol, OSPF or IS-IS) has received a lot of attention in the research community [3, 4]. Existing IGP-based TE mechanisms are only confined to offline operation and hence cannot cope efficiently with significant traffic dynamics. There are well known reasons for this limitation: IGP-based TE only allows for *static* traffic delivery through native IGP paths, without flexible traffic splitting for dynamic load balancing. In addition, changing IGP link weights in reaction to emerging network congestion may cause routing re-convergence problems that potentially disrupt ongoing traffic sessions. In effect, it has been recently argued that dynamic/online route re-computation is to be considered harmful even in the case of network failures [5], let alone for dealing with traffic dynamics.

In recent years, the concept of virtual networks has received increasing attention from the research community, with the general spirit being to enable virtualized network resources on top of the same physical network infrastructure. Such resources not only include physical elements such as routers or links, but also soft resources such as logical network topologies through configurations that allow them to coexist gracefully. Our motivation differs from the existing proposals focusing on virtual network provisioning to support service differentiation, resource sharing or co-existing heterogeneous platforms [6]. Instead, we consider how multiple “equivalent” virtual network topologies, each having its own routing con-

An earlier version of this article has been published in [1].




Figure 1. Providing path diversity in the Abilene network topology.

figuration (such as IGP link weight setting), can be used for multi-path enabled adaptive traffic engineering purposes in IP-based networks. Multi-topology aware Interior gateway routing protocols (MT-IGPs) [7] are used as the underlying platform for supporting the coexistence of multiple virtual IGP paths between source-destination (S-D) pairs on top of the physical network infrastructure.

In our proposal we introduce AMPLE (Adaptive Multi-toPoLoGy traffic Engineering), a holistic system based on virtualized IGP routing topologies for dynamic traffic engineering. The fundamental idea behind this scheme follows the strategy of *offline* provisioning of multiple diverse paths in the routing plane and online spreading of the traffic load for dynamic load balancing in the forwarding plane, as advocated in [5]. The approach can be briefly described as follows. MT-IGPs are used as the underlying routing protocol for providing traffic-agnostic intra-domain path diversity between all source-destination pairs. With MT-IGP routing, customer traffic assigned to different virtual routing topologies (VRTs) follows distinct IGP paths according to the dedicated IGP link weight configurations within each VRT.

Figure 1 depicts an illustration of how path diversity can be achieved for S-D pairs in the Point-of-Presence (PoP) level Abilene network topology with three VRTs, by considering as an example from Sunny Vale to Washington. The i th number in the bracket associated with each link is the IGP weight assigned to it in the i th VRT. As illustrated in the figure, with each network link assigned distinct IGP link weights in the three VRTs, completely non-overlapping paths can be provisioned between the S-D pair. As such, the key task of the offline configuration is to compute MT-IGP link weights for providing maximum path diversity for every S-D pair. Once these link weights have been configured in the network, an adaptive algorithm in the forwarding plane performs traffic splitting ratio adjustment for load balancing across diverse IGP paths in short timescale (e.g. hourly or even more frequently) according to the monitored network and traffic conditions.

For example, if the link between Kansas City and Houston is highly loaded, some traffic originally carried through the green path (in VRT 1) can be shifted to the other two (i.e. the blue and pink paths in VRTs 2 and 3, respectively) by adjusting the traffic splitting ratio across the three VRTs at Sunny Vale. The ultimate goal is to intelligently adjust traffic assignment through splitting across multiple routing topologies at individual source PoP nodes in reaction to the monitored traffic conditions. In order to achieve this, the underlying MT-IGP link weights need to be carefully computed offline and set for maximizing path diversity, based on which adaptive traffic control is performed.

From a system point of view, AMPLE consists of two major components. The Offline Link Weight Optimization (OLWO) component focuses on the *static* dimensioning of the underlying network, with MT-IGP link weights computed for maximizing intra-domain path diversity across multiple VRTs. Once the optimized link weight configuration has been enforced onto the network, the Adaptive Traffic Control (ATC) component performs short timescale traffic splitting ratio adjustment for adaptive load balancing across diverse IGP paths in the engineered VRTs, according to the up-to-date monitored traffic conditions. Given the fact that traffic dynamics are both frequent and substantial in today's ISP networks, our proposed TE system offers a promising solution to cope with this in an efficient manner.

SYSTEM OVERVIEW

Figure 2 presents an overall picture of the proposed AMPLE TE system, with Offline MT-IGP Link Weight Optimization (OLWO) and Adaptive Traffic Control (ATC) constituting the key components. As previously mentioned, the ultimate objective of OLWO is to provision offline maximum intra-domain path diversity in the routing plane, allowing the ATC component to adjust at short timescale the traffic assignment across individual VRTs in the forwarding plane. A salient novelty is that the optimization of the MT-IGP link weights does not rely on the availability of the traffic matrix *a priori*, which plagues existing offline TE solutions due to the typical inaccuracy of traffic matrix estimations. Instead, our offline link weight optimization is only based on the characteristics of the network itself, i.e. the physical topology. The computed MT-IGP link weights are configured in individual routers, and the corresponding IGP paths within each VRT are populated in their local routing information bases (MT-RIBs). While OLWO focuses on static routing configuration in a long timescale (e.g. weekly or monthly), the ATC component provides complementary functionality to enable short timescale (e.g. hourly) control in response to the behavior of traffic that cannot be usually anticipated.

As shown in the figure, the input for ATC includes:

- The diverse MT-IGP paths according to the link weights computed by OLWO.
- Monitored network and traffic data such as incoming traffic volume and link utilizations.

At each short-time interval, ATC computes a new traffic splitting ratio across individual VRTs for re-assigning traffic in an optimal way to the diverse IGP paths between each S-D pair. This functionality is handled by a centralized TE manager who has complete knowledge of the network topology and periodically gathers the up-to-date monitored traffic conditions of the operating network. These new splitting ratios are then configured by the TE manager to individual source PoP nodes, who use this configuration for remarking the multi-topology identifiers (MT-IDs) of their locally originated traffic accordingly. The TE manager function can be realized as a dedicated server, but for robustness and resilience it can be implemented in a distributed replicated manner for avoiding the existence of a single point of failure. In the next section we present the detailed design of individual components in the AMPLE system.

COMPONENT SPECIFICATION

OFFLINE LINK WEIGHT OPTIMIZATION

First of all, a fundamental issue in OLWO is how to determine the definition of “path diversity” between PoPs for traffic engineering. Let’s consider the following two scenarios of MT-IGP link weight configuration. In the first case, highly diverse paths (e.g. end-to-end disjoint ones) are available for some PoP-level S-D pairs, while for some other pairs individual paths are completely overlapping with each other across all VRTs. In the second case, none of the S-D pairs have disjoint paths, but none of them are completely overlapping either. Obviously, in the first case if any “critical” link that is shared by all paths becomes congested, its load cannot be alleviated through adjusting traffic splitting ratios at the associated sources, as their traffic will inevitably travel through this link no matter which VRT is used. Hence, our strategy targets the second scenario by achieving “balanced” path diversity across all S-D pairs.

Toward this end, we define the binary metric of Full Degree of Involvement (FDoI) to evaluate the overall path diversity for a given MT-IGP link weight configuration. More specifically, the FDoI value for a link with respect to an S-D pair is set to 1 if this link is shared by the shortest IGP paths across all VRTs for that S-D pair; otherwise it is set to 0. Let’s take Fig. 1 as an example again. The FDoI value for the link from Sunny Vale to LA with regard to the S-D pair (Seattle, LA) is 1, as this link is part of all the shortest IGP paths between Seattle and LA across the three VRTs. In comparison, the FDoI value for the same link with regard to the S-D pair (Sunny Vale, Washington) is 0, as alternate routes are available via Denver in other VRTs. The optimization objective of OLWO is to minimize the sum of FDoI values across all network links with regard to all S-D pairs. If this sum is equal to 0, then no critical link is formed given the underlying MT-IGP link weights, which means that at least one source in the network will always be able to find alternative path(s) to bypass the over-loaded link given any single link congestion scenario.

Our solution is based on an offline optimization algorithm for maximizing path diversity




Figure 2. AMPLE system overview.

across multiple VRTs (see [1] for details), and our evaluation based on two operational networks shows good path diversity performance: only three VRTs are sufficient to avoid any critical link for the GEANT network topology, while the Abilene topology needs four VRTs to achieve the same goal. As we will show later, even without necessarily creating high path diversity for every S-D pair, there is a high chance of achieving near-optimal TE performance based on the MT-IGP link weight setting in OLWO.

NETWORK MONITORING

Network monitoring is responsible for collecting up-to-date traffic conditions in real-time and plays an important role for supporting the ATC operations. AMPLE adopts a hop-by-hop based monitoring mechanism that is similar to the proposal of [8]. The basic idea is that a dedicated monitoring agent deployed at every PoP node is responsible for monitoring:

- The volume of the traffic originated by the local customers toward other PoPs (intra-PoP traffic is ignored).
- The utilization of the directly attached inter-PoP links.

As shown in Fig. 3, this monitoring agent gathers data on the locally originated traffic volume from all the access routers (ARs) attached to customers at the PoP. Meanwhile the agent also collects the utilization of the directly attached inter-PoP links from individual backbone routers (BRs).

In a periodic fashion (e.g. hourly), the central TE manager polls individual monitoring agents within each PoP and collects their locally monitored traffic volume and link utilizations. These statistics are then used by the central TE manager for updating its maintained traffic engineering

Given the optimized MT-IGP link weights produced by OLWO, adaptive traffic control (ATC) can be invoked at short-time intervals during operation in order to re-optimize the utilization of network resources in reaction to traffic dynamics.




Figure 3. Network monitoring and ATC.

information base (TIB, to be introduced in the next section) and computing traffic splitting ratios for the next interval. Such a hop-by-hop based paradigm works efficiently in a TE system with a central manager. The main reason is that new traffic splitting ratios are computed by the TE manager who is able to have the global view of the network, enabling it to achieve a global optimum in traffic control.

ADAPTIVE TRAFFIC CONTROL

Given the optimized MT-IGP link weights produced by OLWO, adaptive traffic control (ATC) can be invoked at short-time intervals during operation in order to re-optimize the utilization of network resources in reaction to traffic dynamics. The optimization objective of ATC is to minimize the maximum link utilization (MLU), which is defined as the highest utilization among all the links in the network. The rationale behind ATC is to perform periodic and incremental traffic splitting ratio re-adjustments across VRTs based on traffic pattern “continuity” at short a timescale, but without necessarily performing a global routing re-optimization process from scratch every time. In this section, we present a lightweight but efficient algorithm that can be applied for adaptive adjustment of the traffic splitting ratio at individual PoP source nodes to achieve this goal. In a periodic fashion, the following two operations are performed:

- Measure the incoming traffic volume and the network load for the current interval as described in the previous section.
- Compute new traffic splitting ratios at individual PoP source nodes based on the splitting ratio configuration in the previous interval, according to the newly measured traffic demand and the network load for dynamic load balancing.

To fulfill the second task, a traffic engineering information base (TIB) is needed by the TE manager to maintain necessary network state based on which new traffic splitting ratios are computed. Figure 4 presents the structure of our proposed

TIB, which consists of two inter-related repositories, namely the Link List (LL) and the S-D Pair List (SDPL). The LL maintains a list of entries for individual network links. Each LL entry records the latest monitored utilization of a link and the involvement of this link in the IGP paths between associated S-D pairs in individual VRTs. More specifically, for each VRT, if the IGP path between an S-D pair includes this link, then the ID of this S-D pair is recorded in the LL entry. It is worth mentioning that this involvement information remains static after the MT-IGP link weights have been configured (static information is presented in black in Fig. 4, while dynamic information that needs to be updated periodically at short timescale is shown in red). On the other hand, the SDPL consists of a list of entries, each for a specific S-D pair with the most recently measured traffic volume from S to D. Each SDPL entry also maintains a list of subentries for different VRTs, with each recording the splitting ratio of the traffic from S to D, as well as the ID of the bottleneck link along the IGP path for that S-D pair in the corresponding topology.

During each ATC interval, the TIB is updated upon the occurrence of two events. First, upon receiving the link utilization report from the network monitoring component, the TE manager updates the link utilization entry in the LL and the ID of the bottleneck link for each S-D pair under each VRT in SDPL. Second, when the adaptive traffic control phase is completed and the new traffic splitting ratios are computed, the splitting ratio field in SDPL is updated accordingly for each S-D pair under each VRT.

ATC is performed based on the up-to-date data maintained in the TIB. We start the ATC algorithm description by defining the following parameters:

- $t(u,v)$ — traffic from the source PoP node u to the destination PoP node v .
- $\phi_{u,v}(r)$ — traffic splitting ratio of $t(u,v)$ at u on routing topology r , $0.0 \leq \phi_{u,v}(r) \leq 1.0$.

The algorithm consists of the following three

steps. We define an iteration counter y which is set initially to zero.

Step 1: Identify the most utilized link l_{\max} in the network, which can be simply achieved by visiting the updated LL in the TIB.

Step 2: For the set of S-D pairs whose traffic flows are routed through l_{\max} in *at least one but not all* the routing topologies (i.e. $FDoI = 0$), consider each one at a time and compute its new traffic splitting ratio among the VRTs until the first feasible one is identified (see details in the follow-up description). A feasible traffic flow means that, with the new splitting ratios, the utilization of l_{\max} can be reduced without introducing new hot spots with utilization higher than the original value. To support this operation, all feasible S-D pairs that meet the above requirement are identified from the entry of l_{\max} in the LL.

Step 3: If such a feasible traffic flow is found, accept the corresponding new splitting ratio adjustment. Increment the counter y by one and go to Step 1 if the maximum K iterations have not been reached (i.e. $y \leq K$). If no feasible traffic flow exists or $y = K$, the algorithm stops and the latest resulting values for the traffic splitting ratio are configured in the corresponding entry in the SPDL in order to be executed by individual source PoP nodes.

The parameter K controls the algorithm to repeat at most K iterations in order to avoid long running time. The value of K can be carefully determined by taking into account the trade-off between the TE performance and system complexity. In Step 2, the task is to examine the feasibility of reducing the load of the most utilized link by decreasing the splitting ratios of a traffic flow assigned to the routing topologies that use this link, and shift a proportion of the relevant traffic to alternative paths with lower utilization in other topologies. More specifically, the adjustment works as follows. First, a deviation of the traffic splitting ratio, denoted by δ where $0.0 < \delta \leq 1.0$, is taken out for trial. For the traffic flow $t(u,v)$ under consideration, let R^+ be the set of routing topologies in which the IGP paths from u to v traverse l_{\max} . The main idea is to decrease the sum of traffic splitting ratios on all the routing topologies in R^+ by δ and at the same time to increase the sum of the ratios on other topologies that do not use l_{\max} by δ . (We denote this set of topologies by R^- where $R^- = R/R^+$.) Specifically, for all the topologies in R^+ , which share a common link with the same (maximum) utilization, their traffic splitting ratios are evenly decreased. Hence, the new traffic splitting ratio for each routing topology in R^+ becomes:

$$\phi_{u,v}(r)' = \phi_{u,v}(r) - \delta / |R^+| \quad \forall r \in R^+$$

On the other hand, let μ_r be the bottleneck link utilization of the IGP path in routing topology $r \in R^-$. To obtain μ_r , the TE manager should first identify the ID of the bottleneck link along the IGP path between the associated S-D pair from the SDPL, and then refer to the LL to obtain its utilization. The traffic splitting ratio of each routing topology in R^- increases in an inverse proportion to its current bottleneck link utilization, i.e.




Figure 4. Traffic engineering information base structure: a) entry structure of LL; and b) entry structure of SDPL.

$$\phi_{u,v}(r)' = \phi_{u,v}(r) + \left(\frac{1 - \mu_r}{\sum_{r \in R^-} 1 - \mu_r} \times \delta \right) \quad \forall r \in R^-$$

The lower (higher) the bottleneck link utilization, the higher (lower) the traffic splitting ratio will be increased.

An important issue to be considered is the value setting for δ . If not appropriately set, it may either lead to slow convergence or overshoot the traffic splitting ratio, both of which are undesirable. On one hand, too large value of δ may miss the chance to obtain desirable splitting ratios due to the large gap between each trial. On the other hand, too small (i.e. too conservative) value of δ may cause the algorithm to perform many iterations before the most appropriate value of δ is found, thus causing slow convergence to the equilibrium. Taking this consideration into account, we apply an algorithm to perform an exponential increment of δ starting from a sufficiently small value. If this adjustment is able to continuously reduce the utilization of l_{\max} without introducing negative new splitting ratios on R^+ , the value of δ will be increased exponentially for the next trial until no further improvement on the utilization can be made or the value of δ reaches 1.0 (i.e. the maximum traffic splitting ratio that can be applied).

WORKING AS A WHOLE SYSTEM

After presenting the detailed information on individual components, we now briefly describe how they work in unison as a whole TE system. First, optimized MT-IGP link weights are configured on top of the underlying MT-IGP platform and remain static until the next offline OWLO cycle. During this period, ATC plays the major role for adaptively re-balancing the load according to the traffic dynamics in short-time intervals. As a bootstrap procedure, the initial traffic splitting is evenly distributed across VRTs, but this will be recomputed based on follow-up traffic monitoring results. In response to the periodic polling requests by the TE manager, the monitoring agents attached to individual PoP nodes report back the incoming traffic volume (from access routers) and inter-PoP link utilizations (from backbone routers). The TE manager accordingly updates the traffic volume between each S-D pair in the SDPL and link utilization information stored in the LL of the TIB. According to the obtained link utilization information,

the bottleneck link ID along the IGP paths between individual S-D pairs in each VRT is also updated in the SDPL. Based on the updated information, the TE manager computes the new traffic splitting ratio for each S-D pair across individual routing topologies. These new splitting ratios are configured in the SDPL and the TE manager, then instructs all the source PoP nodes within the network to use these new values for traffic splitting during the next interval. In addition, these values in the SDPL will also be used as the starting point for the future computation of the splitting ratios in the next interval. Once each source PoP node has received the new values for traffic splitting from the central TE manager, it enforces them by remarking the MT-ID values carried by the locally originated traffic packets in the new proportions across individual routing topologies. Such a MT-ID remarking operation follows the same style as the technique described in [9] for enabling IP fast reroute functions based on MT-IGP platforms.

EXPERIMENTAL RESULTS

In order to evaluate the performance of AMPLE, we use the real topologies and traffic traces from the GEANT and Abilene networks that are provided by the TOTEM Project [10]. We present results based on a seven-day long traffic traces dataset. Although the dataset in

[10] provides traffic traces measured every five minutes, for consistency with the GEANT scenario, we also use seven-day long traffic matrices at the interval of every 15 minutes. In this article we compare the following optimization methods:

Actual: The actual static link weight setting in the current operational networks. MT-IGP routing is not used.

Multi-TM: We use the TOTEM toolbox to compute a set of static link weights for multiple traffic matrices. The objective is to make the IGP TE robust to traffic demand uncertainty [3]. Specifically, the link weights are computed at the beginning of each day based on the sampled traffic matrices (one per hour) on the same day of the previous week. MT-IGP routing is not used.

AMPLE-n: Our proposed adaptive TE algorithm that runs based on n MT-IGP routing topologies with their link weights computed by the OLWO. The ATC operations are performed at 15 minute intervals according to the latest measured traffic conditions.

Optimal: As the baseline for our comparisons, we use the GNU Linear Programming Kit (GLPK) function in the TOTEM toolbox to compute the optimal MLU for each distinct traffic matrix associated with the given topologies.

Figure 5 plots the MLU versus the time intervals of traffic traces for the GEANT and Abilene networks. From the figure we can have an overall glance of the traffic dynamics pattern in both networks during the sampling period. In a further evaluation, Table 1 shows results of the following additional statistics that are derived from Fig. 5:

- **Average maximum link utilization (AMU):** The average value of the MLU across all the traffic traces during the seven-day period.
- **Highest maximum link utilization (HMU):** The highest value of the MLU across all the traffic traces during the period.
- **Proportion to near-optimal performance (PNO):** The percentage over all the traffic traces in which AMPLE can achieve near-optimal performance. We define here the meaning of near-optimality to be the MLU that is within 3 percent of the gap from Optimal.

An overall observation is that *AMPLE can substantially reduce the MLU for most of the traffic traces*. For example, in the GEANT network, the *Actual* link weight approach produces AMU that is 86 percent higher than that of the optimal value, whereas with AMPLE the value varies between 0.1 percent and 43 percent, depending on the number of routing topologies that are used. In general, the larger the number of routing topologies used, the closer to the optimal performance can be achieved. Similar results are also observed for the HMU performance.

For the PNO metric in Table 1, if AMPLE is based on two routing topologies, the value is only 13.1 percent but it still performs significantly better than all the other approaches. We can now start to see the practical usefulness of our approach for improving network utilization: When the number of routing topologies increases to three, the PNO boosts up to 78.3 percent. With 99.6 percent of all the traffic traces,




Figure 5. MLU comparison between schemes with seven days of traffic traces: a) the GEANT network MLU performance; and b) the Abilene network MLU performance.

AMPLE achieves near-optimal performance with four routing topologies. These results reveal that, for the GEANT network, AMPLE has a very high chance of achieving near-optimal TE performance under any scenario of traffic traces with four routing topologies. For the Abilene network, our experimental results reach conclusions that are similar to that of the GEANT network.

Another observation is that the *Multi-TM* approach does not achieve good performance in minimizing the MLU according to Fig. 5. There are two reasons for this. First, the ultimate objective of Multi-TM is to minimize the cost represented by a piece-wise linear function [3] rather than specifically on the MLU. Second, even if multiple traffic matrices with different pattern characteristics are considered in link weight optimization, unexpected traffic spikes may still introduce poor TE performance. This is especially the case in the Abilene scenario whose traffic pattern is less regular compared to the GEANT traffic pattern.

SUMMARY

In this article we have introduced AMPLE, a novel TE system based on virtualized IGP routing that enables short timescale traffic control against unexpected traffic dynamics using multi-topology IGP-based networks. The framework encompasses two major components, namely, Offline Link Weight Optimization (OLWO) and Adaptive Traffic Control (ATC). The OLWO component takes the physical network topology as the input and aims to produce maximum IGP path diversity across multiple routing topologies through the optimized setting of MT-IGP link weights. Based on these diverse paths, the ATC component performs intelligent traffic splitting adjustments across individual routing topologies in reaction to the monitored network dynamics at short timescale. As far as implementation is concerned, a dedicated traffic engineering manager is required, having a global view of the entire network conditions and being responsible for computing optimized traffic splitting ratios according to its maintained TE information base. Our experiments based on the GEANT and Abilene networks and their real traffic traces have shown that AMPLE has a high chance of achieving near-optimal network performance with only a small number of routing topologies, although this is yet to be further verified with traffic traces data from other operational networks when available. A potential direction in our future work is to consider a holistic TE paradigm based on AMPLE, which is able to simultaneously tackle both traffic and network dynamics, for instance network failures.

REFERENCES

- [1] N. Wang, K-H. Ho, and G. Pavlou, "Adaptive Multi-topology IGP Based Traffic Engineering with Near-Optimal Performance," *Proc. IFIP Networking 2008*.
- [2] S. Uhlig et al., "Providing Public Intradomain Traffic Matrices to the Research Community," *ACM Sigcomm Comp. Commun. Rev. (CCR)*, vol. 36, no. 1, Jan. 2006, pp. 83–86.

Optimization Method	GEANT (%)			Abilene (%)		
	AMU	HMU	PNO	AMU	HMU	PNO
Optimal	30.05	52.82	—	12.2	33.42	—
Actual	55.74	96.91	0	19.59	63.24	1.19
Multi-TM	48.56	104.15	0.44	53.2	230	0.15
AMPLE-2	42.9	94.61	13.08	18.61	60.96	64.14
AMPLE-3	31.95	60.36	78.34	12.36	33.44	88.69
AMPLE-4	30.08	52.88	99.56	12.4	49.6	97.77

Table 1. MLU performance statistics.

- [3] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS Weights in a Changing World," *IEEE JSAC*, vol. 20, no. 4, May 2002, pp. 756–67.
- [4] D. Xu, M. Chiang, and J. Rexford, "Link-State Routing With Hop-By-Hop Forwarding Can Achieve Optimal Traffic Engineering," *Proc. IEEE INFOCOM*, Apr. 2008.
- [5] M. Caesar et al., "Dynamic Route Computation Considered Harmful," *ACM Comp. Commun. Rev. (CCR)*, vol. 40, no. 2, Apr. 2010, pp. 66–71.
- [6] N. M. Mosharaf Kabir Chowdury and R. Boutaba, "A Survey of Network Virtualization," *Computer Networks*, vol. 54, issue 5, Apr. 2010, pp. 862–76.
- [7] P. Psenak et al., "Multi-Topology (MT) Routing in OSPF," RFC 4915, June 2007.
- [8] A. Asgari et al., "Scalable Monitoring Support for Resource Management and Service Assurance," *IEEE Network Mag.*, vol. 18, issue 6, Nov. 2004, pp. 6–18.
- [9] A. Kvalbein et al., "Multiple Routing Configurations for Fast IP Network Recovery," *IEEE/ACM Trans. Net.*, vol. 17, no. 2, 2009, pp. 473–86.
- [10] S. Balon et al., "Traffic Engineering and Operational Network with the TOTEM Toolbox," *IEEE Trans. Network and Service Management*, vol. 4, no. 1, June 2007, pp. 51–61, Project website (including GEANT/Abilene Network Topology and Traffic Dataset): <http://totem.info.ucl.ac.be/>

BIOGRAPHIES

NING WANG (n.wang@surrey.ac.uk) is a Lecturer at the Centre for Communication Systems Research (CCSR), University of Surrey in UK. He received his B.Eng. (Honors) degree from the Changchun University of Science and Technology, P.R. China in 1996, his M.Eng. degree from Nanyang University, Singapore in 2000, and his PhD degree from the University of Surrey in 2004 respectively. His research interests include Internet traffic engineering, network virtualization, QoS provisioning and content-aware networks.

KIN HON KO (royhokh@googlemail.com) was an Instructor at the Department of Computing, The Hong Kong Polytechnic University (PolyU). He holds a B.Sc. (Hons) in Computer Studies from City University of Hong Kong, a M.Sc. (Eng.) in Data Communications from the University of Sheffield, and a Ph.D. from University of Surrey. Before joining PolyU, he was a postdoctoral research fellow at the CCSR, University of Surrey, and a full-time teaching professional at City University of Hong Kong. His current research interests cover network operations and management, traffic engineering, dependable and secure network design, P2P/IPTV traffic management and Green Internet design.

GEORGE PAVLOU (g.pavlou@ee.ucl.ac.uk) holds the Chair of Communication Networks at the Dept. of Electronic & Electrical Engineering, UCL. Over the last 25 years he has undertaken and directed research in networking, network management and service engineering, having extensively published in these areas. He has contributed to ISO, ITU-T and IETF standardization activities and has been instrumental in a number of key European and UK projects that produced significant results. His current research interests include traffic engineering, content-based systems, autonomous networking, policy-based management and infrastructure-less wireless networks.