

Joint Optimization of Intra- and Inter-Autonomous System Traffic Engineering

Kin-Hon Ho, George Pavlou, Ning Wang and Michael Howarth, *Member, IEEE*

Abstract—Traffic Engineering (TE) involves network configuration in order to achieve optimal IP network performance. The existing literature considers intra- and inter-AS (Autonomous System) TE independently. However, if these two aspects are considered separately, the overall network performance may not be truly optimized. This is due to the interaction between intra- and inter-AS TE, where a good solution of inter-AS TE may not be good for intra-AS TE. To remedy this situation, we propose a joint optimization of intra- and inter-AS TE in order to improve the overall network performance by simultaneously finding the best egress points for inter-AS traffic and the best routing scheme for intra-AS traffic. Three strategies are presented to attack the problem, sequential, nested and integrated optimization. Our evaluation shows that, in comparison to sequential and nested optimization, integrated optimization can significantly improve overall network performance by being able to accommodate approximately 30%-60% more traffic demand.

I. INTRODUCTION

Traffic Engineering (TE) [1] is the set of techniques to optimize operational IP network performance by tactically routing traffic on paths that satisfy certain performance objectives. The task of TE is the following: given a network topology and an expected set of traffic demands (i.e. the traffic matrix), determine and configure the best routing strategy so that the overall network performance is optimized.

With the hierarchical structure of the Internet, traffic is routed *within* an Autonomous System (AS) or *between* ASes. Accordingly, the scope of TE covers intra- and inter-AS traffic management. For *intra-AS TE*, the operator of an AS controls traffic routing within the network by either setting the link weights of the corresponding IP routing protocol (typically OSPF or IS-IS) or establishing Label Switched Paths (LSPs) through MultiProtocol Label Switching (MPLS). A typical intra-AS TE optimization objective is to minimize the overall network cost or the maximum link utilization. On the other hand, *inter-AS TE* [2,3] aims to control traffic entering and

exiting an AS with optimization objectives such as load balancing over inter-AS links that are common congestion points [4,5]. For a particular AS, the network operator can control traffic exiting the AS by assigning the traffic to the best egress points. This is known as *outbound* inter-AS TE. Likewise, the network operator can also control traffic entering the AS by selecting the best ingress points, which is called *inbound* inter-AS TE. In current practice, the commonly used method to enforce inter-AS TE is by adjusting Border Gateway Protocol (BGP) route attributes such as *local-pref*, AS path length, MED, etc. Other methods such as inter-AS MPLS [6] are also applicable.

Internet Service Provider (ISP) networks carry both *intra-AS traffic*, for which the source and destination are attached to the network, and *inter-AS traffic*, for which the source, destination or both are outside the network. ISPs employ both intra- and inter-AS TE to optimize the routing of these types of traffic [7,8]. However, although work exists on both intra- and inter-AS TE, most of the existing literature deals with them *separately*: this may result in suboptimal overall network performance, primarily due to the interaction effect between inter- and intra-AS TE. More specifically, inter-AS TE can change the ingress and egress points of inter-AS traffic, thus causing the traffic to be routed on different ingress-to-egress paths within the network. This fundamentally changes the intra-AS traffic matrix, i.e. traffic load between each ingress and egress node pair. Such a change could therefore significantly affect the performance achieved by intra-AS TE.

If intra- and inter-AS TE are not jointly optimized, a sequential approach may be seen as the most reasonable solution. In this sequential approach, the inter-AS TE solution can become the input for intra-AS TE. In other words, the best egress/ingress points are first determined and then the intra-AS traffic matrix should be computed before running intra-AS TE. However, since the objectives and constraints of the intra-AS TE are not taken into account, the decisions made by inter-AS TE often do not provide a good input for the intra-AS TE, sometimes even leading to infeasible solutions. As a result, it is difficult to claim that a truly good overall TE solution has been achieved when each TE type is considered separately. In this paper, we propose the joint optimization of intra- and inter-AS TE as an effective means to achieve better overall TE solutions than the ones obtained through the sequential approach. More specifically, we investigate the following two challenges:

- *How should intra- and inter-AS TE be combined and how do we formulate their joint optimization problem?*

Manuscript received February 8, 2008; revised November 10, 2008; approved April 13, 2009. The associate editor coordinating the review of this paper and approving it for publication was Prof. Alberto Leon-Garcia.

Kin-Hon Ho is with the Department of Computing, The Hong Kong Polytechnic University, Hong Kong (e-mail: royhokh@googlemail.com).

George Pavlou is with the Department of Electronic and Electrical Engineering, University College London, Torrington Place, London, United Kingdom, WC1E 7JE (e-mail: g.pavlou@ee.ucl.ac.uk).

Ning Wang and Michael Howarth are with the Centre for Communication Systems Research, University of Surrey, Guildford, Surrey, United Kingdom, GU2 7XH (e-mail: n.wang@surrey.ac.uk, m.howarth@surrey.ac.uk).

- *How can we solve the joint TE optimization problem to achieve better overall network utilization?*

Our contributions in this paper are as follows. First of all, we explain the TE interaction effect that can lead to suboptimal overall network utilization for the sequential approach. Then, for the first challenge listed above, we formulate a bi-criteria joint optimization problem of intra- and inter-AS TE with an aim of optimizing their objective cost functions simultaneously. Since the interaction effect can be generally applied to any intra- and inter-AS TE approach, it is possible to formulate the joint optimization problem for each of the combinations of intra- and inter-AS TE approaches, i.e. intra-AS TE combined with inbound/outbound inter-AS TE. As the primary objective of this paper is to illustrate the benefits of the joint optimization of intra- and inter-AS TE, we consider *MPLS-based intra-AS TE* and *outbound inter-AS TE* in this paper.

For the second challenge above, we consider three strategies to solve the joint TE optimization problem, namely sequential, nested and integrated optimization. These strategies aim to obtain non-dominated solutions with respect to the intra- and inter-AS TE objectives. We evaluate the performance of these strategies by simulation using Rocketfuel [9] topologies and synthetic traffic matrices. Our simulation results show that the integrated optimization, which considers intra- and inter-AS TE simultaneously, achieves better overall network utilization than the other two approaches. This improvement allows the network to support approximately a 30%-60% increase in traffic demand. We believe that our work provides an insight into the interaction between intra- and inter-AS TE, enabling ISPs to further optimize the performance of their networks over the sequential approaches.

The rest of this paper is structured as follows. In the next section, we provide background on intra- and inter-AS TE. Section III explains the TE interaction effect. We formulate the joint intra- and inter-AS TE optimization problem in Section IV and present three strategies to solve it in Section V. In Sections VI and VII, we present our evaluation methodology and results for these strategies. We discuss some potential enhancement work to the joint intra- and inter-AS TE in Section VIII. Finally, Section IX and X present related work and our conclusion respectively.

II. BACKGROUND OF TRAFFIC ENGINEERING

Traffic engineering typically takes as input elements such as the traffic matrix and network topology, and then executes an algorithm to produce a traffic routing scheme that optimizes the overall network performance. We present below some important background for this work. For ease of presentation, Table 1 shows the notation used throughout the paper.

A. Types of Internet Traffic

According to [10], Internet traffic received by an AS can be classified into four types: i) *internal* traffic that travels from an ingress access link to an egress access links; ii) *transit* traffic that travels from an ingress peering link to an egress peering link; iii) *inbound* traffic that travels from an ingress peering

link to an egress access link, and iv) *outbound* traffic that travels from an ingress access link to an egress peering link. For our work, we define the following types of Internet traffic and use them throughout the paper:

- *local intra-AS traffic*: traffic destined to egress access links. This corresponds to internal and inbound traffic.
- *inter-AS traffic*: traffic destined for downstream ASes and whose egress peering points can be selected by inter-AS TE. This corresponds to outbound and transit traffic.
- *intra-AS traffic*: overall traffic that traverses the network. This corresponds to both local intra-AS and inter-AS traffic, i.e. the sum of the above two.

Table 1. Notation used in this paper

| NOTATION | DESCRIPTION |
|------------------|------------------------------------------------------------------------------------------------------------------|
| K | A set of downstream routing prefixes |
| I | A set of ingress points |
| J | A set of egress points (inter-AS links) |
| E | A set of intra-AS links |
| $t_{inter}(i,k)$ | Bandwidth demand of the inter-AS traffic flow at ingress point $i \in I$ destined to routing prefix $k \in K$ |
| $t_{loc}(i,j)$ | Bandwidth demand of the local intra-AS traffic flow between ingress point $i \in I$ and egress point $j \in J$ |
| $t_{intra}(i,j)$ | Bandwidth demand of the intra-AS traffic flow between ingress point $i \in I$ and egress point $j \in J$ |
| $T_{Inter}(I,K)$ | Inter-AS (ingress-to-prefix) traffic matrix consisting of all $t_{inter}(i,k)$ |
| $T_{Loc}(I,J)$ | Local intra-AS (ingress-to-egress) traffic matrix consisting of all $t_{loc}(i,j)$ |
| $T_{Intra}(I,J)$ | Intra-AS traffic matrix consists of all $t_{intra}(i,j)$ |
| $Out(k)$ | A set of egress points that has reachability to routing prefix k |
| C_{inter}^j | Capacity of egress point j (inter-AS link) |
| bw_{inter}^j | Residual bandwidth of C_{inter}^j |
| C_{intra}^l | Capacity of intra-AS link l |
| bw_{intra}^l | Residual bandwidth of C_{intra}^l |
| $x_{i,k}^j$ | A binary variable indicating whether inter-AS traffic flow $t_{inter}(i,k)$ is assigned to egress point j |
| $y_{i,j}^l$ | A binary variable indicating whether intra-AS traffic flow $t_{intra}(i,j)$ is assigned to intra-AS link l |
| $P_{i,j}$ | A set of candidate paths realizing intra-AS traffic flow $t_{intra}(i,j)$ |
| $w_{i,j,p}$ | A binary variable indicating whether path $p \in P_{i,j}$ is chosen to realize the traffic flow $t_{intra}(i,j)$ |
| $s(i,k)$ | A variable storing the egress point that has been assigned to $t_{inter}(i,k)$ |

B. Traffic Matrices

A Traffic Matrix (TM) represents a matrix of traffic load from one network point to another one over a particular time interval. In the **inter-AS traffic matrix**, each element $t_{inter}(i,k)$ represents the volume of inter-AS traffic that enters the network at ingress point i and is destined to routing prefix k that belongs to a remote downstream AS. The best egress point of each inter-AS traffic flow can be selected by inter-AS TE. We denote $s(i,k)$ the egress point assigned to $t_{inter}(i,k)$. On the other hand, in the **intra-AS traffic matrix**, each element $t_{intra}(i,j)$ represents the volume of traffic that enters the network at ingress point i and exits at egress point j . It is the sum of the local intra-AS and inter-AS traffic volume between each ingress and egress node pair:

$$t_intra(i,j) = t_loc(i,j) + \sum_{k \in K : s(i,k)=j} t_inter(i,k)$$

Figure 1 shows a network with ingress point i and egress points j and j' . We assume both egress points can reach routing prefixes k and k' . Given local intra-AS traffic flow $t_loc(i,j)$ and inter-AS traffic flows $t_inter(i,k)$ and $t_inter(i,k')$ with $s(i,k) = j$ and $s(i,k') = j'$, the elements of the intra-AS traffic matrix are $t_intra(i,j) = t_loc(i,j) + t_inter(i,k)$ and $t_intra(i,j') = t_inter(i,k')$.




Figure 1. Decomposition of traffic matrices

The intra- and inter-AS TM may be obtained through measurement or estimation. Intra-AS traffic can be measured [10] using for example, Cisco’s NetFlow. It can also be estimated from measured link load statistics [11,12]. For deriving inter-AS TM, the authors in [13] describe a methodology using Cisco’s NetFlow and BGP routing data. Alternatively, a methodology is proposed in [14] based on server logs from content delivery networks and packet level traces from large user sets.

C. Intra- and Inter-AS Traffic Engineering

The general intra-AS TE problem can be summarized as follows: given a network topology and an intra-AS traffic matrix, determine an appropriate set of OSPF/ISIS link weights or MPLS LSPs so as to optimize the network performance such as bandwidth consumption and load balancing in the network. Many techniques exist to solve this problem [15]. For example, Fortz and Thorup [16] proposed a tabu search technique to derive optimal link weight settings. Xiao et al [17] proposed a greedy heuristic for MPLS-based TE. Trimintzios et al. proposed a MPLS-based TE system for quality of services provisioning [18]. More intra-AS TE work can be found in [15,19,20] and references therein.

The general inter-AS TE problem [4] can be summarized as follows: given a network topology, an inter-AS traffic matrix and BGP routing prefixes, select ingress/egress points for the traffic so that the network performance is optimized. A common inter-AS TE objective is to balance the load over inter-AS links. Various outbound inter-AS TE algorithms have been proposed [4, 21,22]. In current practice, inter-AS TE is typically enforced by adjusting BGP route attributes such as *local-pref*, AS path length, MED, etc. The reader is referred to [23] for the explanation of these techniques. In addition, ongoing work in inter-AS MPLS [24] provides an alternative

method to enforce inter-AS TE. A comprehensive survey of intra- and inter-AS TE approaches and algorithms can be found in [25].

III. INTERACTIONS BETWEEN INTRA- AND INTER-AS TE

The existing literature considers intra- and inter-AS TE separately. However, as already mentioned, this may result in non-optimal overall network performance. In this section, we explain an interaction effect between intra- and inter-AS TE that would lead to suboptimal overall network performance.

A. Effect of Inter-AS TE on Intra-AS TE

The interaction is the effect of inter-AS TE on the performance that can be achieved by intra-AS TE. Although we only consider outbound inter-AS TE, this effect is also applicable to inbound inter-AS TE as both can influence intra-AS TE performance by changing the intra-AS traffic matrix.

Recall from Section II.B that the intra-AS TM is derived from both local intra-AS and inter-AS traffic. Inter-AS TE assigns egress points to inter-AS traffic in order to balance the load over inter-AS links. Consequently, the traffic will be routed on different ingress-to-egress paths according to the assigned egress points. This may make the current intra-AS TE solution undesirable for the new state of the network and thus require re-optimization. However, since the traffic load for each ingress-egress node pair is varied, it changes the current intra-AS traffic matrix. For example, referring to the scenario of Section II.B, if $s(i,k)$ was changed to j' , the two elements of intra-AS traffic matrix would then become $t_intra(i,j) = t_loc(i,j)$ and $t_intra(i,j') = t_inter(i,k) + t_inter(i,k')$. It can be seen that different egress point selections for the inter-AS traffic will result in different intra-AS TMs. Since the traffic matrix as an input to intra-AS TE changes, this may consequently cause intra-AS TE to produce solutions with different routing schemes and possibly suboptimal overall performance. It is thus the goal of this paper to investigate how to derive the best egress point selection for optimal inter-AS resource utilization as well as a good intra-AS traffic matrix that leads intra-AS TE to achieve optimal resource utilization within the network.

In fact, not only inter-AS TE affects intra-AS TE, but this is also true the other way round. Intra-AS TE can affect the performance that can be achieved by inter-AS TE due to the hot-potato routing. In our previous work, we explained this effect with an example. However, due to space limitations, we omit this part in this paper. Interested readers may be referred to [26] for more details.

B. The Need for Joint TE Optimization

The interaction effect has shown that inter-AS TE indeed can affect the performance that can be achieved by intra-AS TE. A recent study has also shown that the Internet bottleneck is approximately equally distributed on intra- and inter-AS links [5]. Therefore, a good overall TE solution should perform satisfactorily with respect to both intra- and inter-AS TE objectives. It is important that inter-AS TE aspects should be considered during intra-AS TE and vice versa in order to

improve the overall network performance. We therefore propose to operate on local intra-AS TM and inter-AS TM collectively. In the next section, we present a bi-criteria integer programming formulation for the Joint intra- and inter-AS TE optimization (Joint-TE) problem.

IV. JOINT INTRA- AND INTER-AS TE OPTIMIZATION

A bi-criteria optimization formulation is that one can express two notions that are of concern in defining what represents an optimal solution. Their objectives are typically expressed in a form of cost functions. In this paper, we formulate a bi-criteria Joint-TE problem by taking into account both intra- and inter-AS TE cost functions.

A. Cost Functions

We employ the network cost function proposed by Fortz and Thorup [16]. This is a piecewise linear function of link utilization, which imitates the response time of M/M/1 queues to represent the cost of network links. It has been widely used for the performance evaluation of many TE schemes and algorithms [15,20,27,28,29]. As explained in [30], the underlying principle of this function is that it results in small cost for sending traffic over a link with low utilization while the cost increases steeply when the link utilization approaches 100%; links are also penalized heavily when utilization goes over 100%. By using the piecewise linear cost function, the two objectives of minimizing bandwidth consumption and achieving load balancing are taken into account simultaneously. In this paper, we use the piecewise linear cost function for both intra- and inter-AS TE for consistency and generality, as well as for its ability to express the common key objectives of each TE type. Nevertheless, these cost functions may be different according to the operational objectives of particular ASes.

It is worth mentioning that the inter-AS TE objectives may also include policies between ASes instead of being solely based on performance metrics such as link utilization. For example, a customer AS may have a policy of minimizing or limiting the amount of inter-AS traffic to be forwarded to some of its provider ASes which incur a high charging rate on using their networks. In this case, we can deliberately reduce the capacity of the inter-AS links connecting to the ‘expensive’ ASes, and the cost function, which aims to minimize the amount of traffic traversing low capacity links for the sake of load balancing, is equivalent to minimizing the amount of traffic to be sent towards those expensive ASes. Thus, our cost function can also taken inter-AS policy into consideration to a certain extent. As far as multiple inter-AS policies are concerned, one may model (or quantize) these policies as a cost function, and then combine them with the performance-related objectives.

B. Problem Formulation of Bi-Criteria Joint-TE Optimization

The objective of the Joint-TE problem is to minimize both overall intra- and inter-AS costs:

$$\text{Minimize } \sum_{l \in E_{intra}} u_l \quad (1)$$

$$\text{Minimize } \sum_{j \in J} Y_j \quad (2)$$

Subject to

$$u_{intra}^l = \frac{\sum_{i \in I} \sum_{j \in J} y_{i,j}^l \alpha_{intra}(i,j)}{C_{intra}^l}, \quad "l \in E \quad (3)$$

$$u_{inter}^j = \frac{\sum_{i \in I} \sum_{k \in K} x_{i,k}^j \alpha_{inter}(i,k)}{C_{inter}^j}, \quad "j \in J \quad (4)$$

$$x_{i,j}^l \in \{0,1\}, \quad "(l \in E, i \in I, j \in J) \quad (5)$$

$$y_{i,j}^l \in \{0,1\}, \quad "(l \in E, i \in I, j \in J) \quad (6)$$

$$w_{i,j,p} \in \{0,1\}, \quad "(p \in P_{i,j}, i \in I, j \in J) \quad (7)$$

$$\sum_{p \in P_{i,j}} w_{i,j,p} = 1, \quad "(i \in I, j \in J) \quad (8)$$

$$y_{i,j}^l \leq w_{i,j,p}, \quad "(l \in P, p \in P_{i,j}, i \in I, j \in J) \quad (9)$$

$$\sum_{j \in J} x_{i,k}^j = 1, \quad "(i \in I, k \in K) \quad (10)$$

$$u_l = f(u_{intra}^l) \quad (11)$$

$$Y_j = f(u_{inter}^j) \quad (12)$$

$$f(m) = m, \quad 0 \leq m < \frac{1}{3} \quad (13)$$

$$f(m) = 3m - \frac{2}{3}, \quad \frac{1}{3} \leq m < \frac{2}{3} \quad (14)$$

$$f(m) = 10m - \frac{16}{3}, \quad \frac{2}{3} \leq m < \frac{9}{10} \quad (15)$$

$$f(m) = 70m - \frac{178}{3}, \quad \frac{9}{10} \leq m < 1 \quad (16)$$

$$f(m) = 500m - \frac{1468}{3}, \quad 1 \leq m < \frac{11}{10} \quad (17)$$

$$f(m) = 5000m - \frac{16318}{3}, \quad \frac{11}{10} \leq m \quad (18)$$

Equations (3) and (4) define the utilization of intra- and inter-AS links. Constraints (5)-(7) represent the standard integrality constraint to the three defined variables. Together with constraint (7), constraint (8) ensures that each intra-AS traffic flow $t_{intra}(i,j)$ is routed along a single LSP within the network. This is an optional constraint but we follow the assumption made in [31] in order to preserve scalability and to minimize complexity on network management by avoiding excessive LSPs to be managed and arbitrary traffic splitting, though such splitting could permit better network performance. Constraint (9) ensures that, whenever traffic flow $t(i,j)$ is assigned to intra-AS link l , then the path to which l is associated must have been selected. Together with constraint (5), constraint (10) ensures that only one egress point is selected for each inter-AS traffic flow and therefore no traffic flow can be split over multiple egress points. This is because only one single egress point can be selected for each remote prefix according to the standard BGP routing configuration [23]. The remaining constraints define the cost of each intra- and inter-AS link as a function of its utilization based on the piecewise linear cost function [16]. Given the lossless property of the links, an additional constraint is the flow conservation, which ensures that the total incoming

traffic at an intermediate node is equal to the total outgoing traffic from the node.

In principle, the Joint-TE problem is the combination of the intra- and inter-AS TE problems. The intra-AS TE problem formulation consists of objective function (1) as well as constraints (3), (6)-(9), (11) and (13)-(18). On the other hand, the problem formulation of inter-AS TE consists of objective function (2) as well as constraints (4), (5), (10) and (12)-(18). Hence, throughout the rest of this paper, when we refer to the problem of intra- or inter-AS TE, this corresponds to the problem with its own objective function and constraints. In addition, we use the term intra- and inter-AS cost to represent objective functions (1) and (2) respectively. In contrast, for the Joint-TE problem, both cost function (1) and (2) as well as constraints (3)-(18) are considered simultaneously.

C. Optimization Criteria

An optimal solution of the bi-criteria optimization problem (1) and (2) is that each of the two TE objectives attains an optimal value at the same time. However, in general, either it is not possible to find such optimal solutions or they do not exist. In other words, the two objective functions are conflicting. For example, the cost of inter-AS TE may be low but the cost of the corresponding intra-AS TE is high. Moreover, it is not possible to compare the two objective values mathematically and sensibly. For example, we cannot distinguish mathematically which is a better TE solution, (10,30) or (30,10), where (x,y) represents intra- and inter-AS costs respectively. However, we can observe that (10,30) is better than (20,30) or (10,40). On the other hand, the value of the cost function varies with the number of links and their utilization. It does not make sense to compare the two objective values when the number of intra- and inter-AS links and their capacities are different, as it is typically the case in real operational networks. Consequently, this leads us to find non-dominated solutions, which is a primary goal when solving a multi-criteria optimization problem. A solution is called *non-dominated* if there is no other solution that is strictly better in one of the objective functions, and has the same or better values in the others [32]. Thus, solution (10,30) in the above example is a non-dominated solution of (20,30) and (10,40).

There are multiple ways to identify non-dominated solutions. A commonly used method is to design a metric or cost function that combines both intra- and inter-AS TE objectives. However, it is often unclear how to determine the relative weights between the two objectives. An alternative approach, which we consider in this paper, is to search non-dominated solutions in such a way that the inter-AS cost remains at least near-optimal while substantially improving the intra-AS cost. In other words, optimizing inter-AS resource utilization is assumed to be more important than the intra-AS one. The rationale for this assumption is the following:

- intra-AS bandwidth over-provisioning has been employed by ISPs as an effective means to provide good quality service to all traffic in their IP backbone networks [33].

- inter-AS links are common points of congestion in the Internet [4,5]. This may primarily be due to the fact that the rapid growth of peer-to-peer traffic consumes the major part of inter-AS link bandwidth [34,35]. Moreover, an inter-AS link is relatively more difficult to extend than an intra-AS link due to time-consuming and complicated negotiations between two ASes. Therefore, the ASes need to carefully control traffic routed on their inter-AS links.

The method of predefining a lexicographic importance order is commonly used in solving multi-criteria optimization problems. It allows us to produce Joint-TE solutions that can be mathematically distinguished and sensibly compared.

D. Intra-AS and Outbound Inter-AS TE Algorithms

In this section, we introduce two algorithms for solving the intra- and inter-AS TE problems individually. These algorithms will be used by the strategies we propose in the next section as a key component to solve the Joint-TE problem. In this paper, we deliberately treat both intra- and inter-AS TE as black-boxes that we combine independently. Both use TE algorithms that are based on previously established techniques and can achieve near-optimal solutions. The intention of using a TE algorithm that produces near-optimal solutions is to attempt to minimize the possibility that any performance improvement is solely caused by a large performance gap between the optimal solution and the solution achieved by the algorithm.

We use the optimal aware heuristic proposed by Sridharan et al in [19] as the algorithm for solving our intra-AS TE problem. The algorithm solves a MPLS-based intra-AS TE problem using the piecewise linear cost function. The basic idea of the algorithm is that, in the beginning, it solves the Linear Programming (LP) formulation of the intra-AS TE problem with the relaxation of the intra-AS traffic splitting constraint (6) and (7) as follow:

$$y_{i,j}^l \in [0,1], \quad (l \in E, i \in I, j \in J) \quad (19)$$

$$w_{i,j,p} \in [0,1], \quad (p \in P_{i,j}, i \in I, j \in J) \quad (20)$$

This relaxation means that an intra-AS traffic flow can be split over multiple intra-AS paths for a destination in the network. In general, we can readily obtain an optimal solution for a LP problem by using mathematical methods such as Simplex and Interior Point methods [30] or deploying off-the-shelf optimization tools. However, since this optimal solution allows arbitrary intra-AS traffic splitting across multiple paths between any pair of nodes, it is not scalable and practical since excessive path status has to be maintained at each node in the network. It is also not a feasible solution to our intra-AS TE problem as the latter does not permit such arbitrary traffic splitting in the network given constraint (6)-(9). Thereafter, a greedy heuristic is performed to transform this infeasible optimal solution to a traffic non-splittable solution that is feasible for our intra-AS TE problem, while attempting to maintain its optimality during the transformation. The basic operation of the heuristic algorithm is first to sort all the traffic demands in descending order according to their traffic volume

and then, for each traffic flow in that ordered sequence, finds the minimum cost path such that the intra-AS cost is as close as possible to the optimal one that is obtained by solving the LP problem formulation. We call this intra-AS TE algorithm **INTRA-OPTIMAL-AWARE-ALG** throughout this paper. The reader is also referred to [19] for more details of the algorithm.

The idea of the optimal aware heuristic algorithm has also been used by prior TE work [20,36]. Their simulation results showed that this kind of algorithm can achieve near-optimal solutions. Hence, we also propose using it to solve the inter-AS TE problem. We call it **INTER-OPTIMAL-AWARE-ALG** throughout this paper. The algorithm works as follows:

Algorithm 1 INTER-OPTIMAL-AWARE-ALG

Step-1: A LP formulation of the inter-AS TE problem with relaxation of the inter-AS traffic splitting constraint (5) is solved, i.e.

$$x_{i,j}^l \in [0,1], \quad (l \in E, i \in I, j \in J) \quad (21)$$

The solution could attain an optimal inter-AS cost and the desired utilization of each inter-AS link that achieves this optimal cost. This solution, however, permits arbitrary splitting of inter-AS traffic over multiple egress points and is therefore infeasible for our inter-AS TE problem as dictated by constraints (5) and (10).

Step-2: Set the desired inter-AS link utilization as a capacity constraint. This constraint ensures that the total traffic on each inter-AS link does not exceed the desired utilization.

Step-3: Sort inter-AS traffic flows in descending order according to their traffic volume. Assign each traffic flow in that order to the egress point (i.e. inter-AS link) that has the lowest utilization while not violating its capacity constraint. If several such egress points exist, the selection tie-break is in order of maximum residual capacity and the visited sequence.

Step-4: If there exist unassigned inter-AS traffic flows, remove the capacity constraint on all the inter-AS links and re-run Step-3 for only these unassigned flows until all the traffic flows have been assigned.

The LP formulations (i.e. the first step) of **INTRA-OPTIMAL-AWARE-ALG** and **INTER-OPTIMAL-AWARE-ALG** were modeled in AMPL [37] and solved with the CPLEX optimization engine [38]. To enable us to show that these two algorithms produce near-optimal TE performance, we include their optimal values (i.e. by only solving the LP formulation of both intra- and inter-AS TE problem individually) in our experiments for baseline performance comparison.

V. STRATEGIES FOR SOLVING JOINT-TE OPTIMIZATION

Based on the problem formulation, assumptions and TE algorithms described in the preceding section, we present three strategies to solve the Joint-TE problem, namely sequential, nested and integrated optimization.

A. Sequential Optimization

Sequential optimization solves the Joint-TE problem *sequentially* where the optimal solution of one TE problem

becomes the input for the other one. It can be regarded as the most reasonable solution in the current use of intra- and inter-AS TE. As with the interaction effect explained in Section III, sequential optimization refers to the following approach:

Algorithm 2 SEQOPT-INTER-INTRA

Step-1: Egress points are assigned to inter-AS traffic by using the **INTER-OPTIMAL-AWARE-ALG** so as to optimize the inter-AS cost.

Step-2: Compute the intra-AS traffic matrix by taking into account the local intra-AS TM and the inter-AS TM with the assigned egress points.

Step-3: Solve the intra-AS TE problem by the **INTRA-OPTIMAL-AWARE-ALG** (described in Section IV.D) with the objective of minimizing the intra-AS cost.

This strategy is logical in the sense that intra-AS TE is not performed until both ingress and egress points of the inter-AS traffic have been determined. The advantage of sequential optimization is that different algorithm and analysis techniques can be applied to each of the TE problems. However, the sequential optimization approach does not consider intra-AS route optimization during inter-AS TE nor egress point selection during intra-AS TE.

B. Nested Optimization

The sequential optimization generates only one solution from inter-AS TE, and this is used as input to the intra-AS TE. However, this solution may not be a good input. In fact, for many optimization problems, there exists more than one optimal solution. Hence, there may exist inter-AS TE solutions that are not only nearly optimal but also good with respect to intra-AS TE. Thus, we seek for those inter-AS TE solutions and then input them to the intra-AS TE one at a time until the solution with the best overall TE cost is found.

We propose a nested optimization to implement the abovementioned idea. It can be regarded as an enhanced and *iterative* version of the sequential optimization. The algorithm proposed for the nested optimization in a similar sequence to **SEQOPT-INTER-INTRA**.

Algorithm 3 Nested Optimization

Step-1: A Genetic Algorithm (GA) is used to identify a set of lowest-cost inter-AS TE solutions. When the GA converges, all chromosomes have nearly identical lowest cost but they may have different solutions. In fact, in order to explore a larger solution searching space for intra-AS TE, these solutions should not be restricted to an identical cost. In this paper, the solution with inter-AS cost not exceeding the visited lowest cost by more than 0.001% is considered.

Step-2: For each of this set of inter-AS TE solutions, we recompute the intra-AS traffic matrix and perform intra-AS TE using **INTRA-OPTIMAL-AWARE-ALG**. During the TE optimization, the best and the worst visited solutions in terms of intra-AS cost are recorded. We call them **NESTEDBEST** and **NESTEDWORST**. The **NESTEDBEST** and **NESTEDWORST** solutions reflect respectively the extent to which the sequential

optimization solution can be further optimized and how worse could be the solution.

We modified our previously proposed GA [21] to solve the inter-AS TE problem. The GA has included a heuristic similar to Step-3 of **INTER-OPTIMAL-AWARE-ALG** to enhance the quality of the solution. In fact, although the number of candidate inter-AS TE solutions may be quite large, we were able to identify a large number of such solutions as alternative inputs for the intra-AS TE. We observed that many of those inter-AS TE solutions have significantly different egress point selection results, which lead intra-AS TE to produce solutions with various performance. As with the sequential optimization, the nested optimization allows different algorithm and analysis techniques to be applied to each of the TE problems while attempting to obtain better overall TE solutions.

C. Integrated Optimization

Integrated optimization aims to solve the Joint-TE problem by *simultaneously* optimizing the intra- and inter-AS TE objectives. We propose an integrated approach that requires as starting solutions an inter-AS and an intra-AS TE configuration with known egress points and ingress-to-egress paths. The starting solutions can be any quality, regardless of whether they are optimized by TE or not. The integrated approach then proceeds to enhance the quality of the starting solution using neighborhood search algorithm. The integrated approach guarantees that the produced solutions are no worse than the input solutions and in practice are much better.

1) Overview

The Neighborhood Search Algorithm (NSA) is widely regarded as an important tool to solve hard combinatorial optimization problems efficiently. The primary reasons for the widespread use of neighborhood search techniques in practice are their intuitive appeal, flexibility and ease of implementation, and their excellent empirical results [39].

The basic steps of NSA are as follows. Consider a current starting solution x . NSA explores the solution space by identifying the neighborhood of x , $N(x)$. The neighbors of x are solutions that can be obtained by applying a single local transformation (also called a *move*) on x . The best solution in the neighborhood is selected as the new current solution. This neighborhood searching iterates until the stopping criterion is satisfied. Finally the algorithm returns the best visited solution.

During neighborhood search, NSA can move the current solution to the best neighbor that either improves or worsens the quality of the solution. To avoid cycling, a specially designed memory list is used to store previously visited solutions or certain attributes of them for a certain number of iterations. A neighbor solution is rejected if it is already in the list. In order to make the neighborhood searching more effective, an intensification or diversification technique is used to force the algorithm to explore parts of the solution space

that have not been searched yet. Our NSA is outlined in Figure 2 and its fundamental components are described below.

2) Non-TE Starting Solutions

Starting solutions for inter- and intra-AS TE can be respectively obtained by randomly assigning egress points to the inter-AS traffic and then routing each intra-AS traffic flow on the shortest hop paths. They are regarded as non-TE solutions. Nevertheless, we will also evaluate the impact of using TE optimized starting solutions on network performance.

3) Neighborhood Structure and Search Strategy

Algorithm 4 Neighborhood Search Algorithm

Require: Starting intra- and inter-AS TE solutions
Notation: $\Psi_j(\Phi)$ represents the cost of inter-AS link j with residual bandwidth Φ
Notation: $\gamma_l(\Phi)$ represents the cost of intra-AS link l with residual bandwidth Φ

1. $iter \leftarrow 0$
2. **while** $iter < MAX_ITER$ /* stopping criteria */
3. $iter \leftarrow iter + 1$
4. **if** no significant cost improvement for a certain number of iterations **then**
5. perform intra-AS TE on the current solution /* diversification */
6. **for** each inter-AS traffic flow $t_inter(i,k)$
7. $f(i,k) \leftarrow s(i,k)$
8. $F \leftarrow \alpha Y_{s(i,k)}(bw_{inter}^{s(i,k)}) + \sum_{i \in P_{i,j}^{w_{i,j},p}=1} \alpha_i (bw_{intra}^i)$
9. $F' \leftarrow \alpha Y_{s(i,k)}(bw_{inter}^{s(i,k)} + t_inter(i,k)) + \sum_{i \in P_{i,j}^{w_{i,j},p}=1} \alpha_i (bw_{intra}^i + t_inter(i,k))$
10. $D \leftarrow F - F'$ /* delta is the saved cost */
11. **for** each $j \in Out(k)$ which does not constitute a move in the memory list
12. $W \leftarrow \alpha Y_j(bw_{inter}^j) + \sum_{i \in P_{i,j}^{w_{i,j},p}=1} \alpha_i (bw_{intra}^i)$
13. virtually add $t_inter(i,k)$ to $t_intra(i,j)$
14. virtually release the resource used by $t_inter(i,k)$ on i
15. re-compute intra-AS path z between ingress point i and egress point j
16. $W' \leftarrow \alpha Y_j(bw_{inter}^j - t_inter(i,k)) + \sum_{i \in P_{i,z}} \alpha_i (bw_{intra}^i)$
17. **if** $D \geq W' - W$ **then** /* if saved cost \geq added cost */
18. $D \leftarrow W' - W$ /* lines 16-20 aim to find the minimum added
19. $f(i,k) \leftarrow j$ cost, i.e. to find the maximum profit cost */
20. $sel_path \leftarrow z$
21. **end if**
22. **end for**
23. **if** $f(i,k) \neq s(i,k)$
24. stop the **for** loop in line 6 /* Implementation of FM */
25. **end if**
26. **end for**
27. **if** $f(i,k) \neq s(i,k)$
28. update resource utilization on intra- and inter-AS link with respect to the
29. new assignment
30. replace i by sel_path
31. add $t_inter(i,k)$ to $t_intra(i,f(i,k))$
32. $s(i,k) \leftarrow f(i,k)$
33. **end if**
34. **end while**

Figure 2. Neighborhood search algorithm for Joint TE optimization

We consider a neighborhood structure that is based on shifting inter-AS traffic to different egress points while at the same time rerouting the corresponding ingress-to-egress paths. Details of this local transformation are as follows:

We define the path cost to be sum of the cost of the inter-AS link and the cost of each link on the intra-AS route to which an inter-AS traffic flow has been assigned. In order to place greater importance on optimizing the inter-AS cost relative to the intra-AS cost, we introduce α as a factor with large value

to scale the inter-AS cost. Note that α itself has no particular meaning to the Joint-TE problem. It is an intermediary to identify all the non-dominated combinations of intra- and inter-AS costs. For each inter-AS traffic flow, we calculate the profit cost, i.e. the saved cost minus the added cost. The *saved cost* (line 9 in Figure 2) is the path cost of the traffic flow minus the path cost of the traffic flow that would have been removed. This saved cost reflects how much cost would have been reduced if the traffic flow was removed from the path. The *added cost* (line 16) is calculated for each potential egress point except for the one that is currently assigned to the inter-AS traffic flow. It is the cost of a new path towards the potential egress point that the traffic flow would have been assigned minus the path cost of the original path towards this egress point. The new path is the result of rerouting the original path taking into account the traffic flow (lines 13-15) and it is a minimum cost path that can be found by Dijkstra’s algorithm using the instantaneous intra-AS link cost as the routing metric. Consequently, the added cost reflects how much the cost would increase when the traffic flow is assigned onto a new egress point. The neighborhood search strategy specifies which solution in the neighborhood is chosen at each iteration. The following two methods are commonly used:

- *Best Method (BM)*: Compute the profit cost for each inter-AS traffic flow. Choose the one yielding the solution with the *highest* profit cost as the next move.
- *First Method (FM)*: Compute the profit cost for each inter-AS traffic flow. Choose the *first* one yielding a solution with positive profit cost.

It is of great importance for both the solution quality and the search efficiency. We have found in our experiments that BM can achieve approximately 5%-10% performance improvement over FM, but the computational complexity of BM is several orders of magnitude higher than the FM, which makes it impractical to use. We therefore decided to use FM in our NSA. Our finding is also consistent with the prior work that has evaluated the tradeoff between quality and efficiency of BM and FM [40].

4) Use of Memory List

The memory list is operated as a first-in-first-out queue. The first element in the list is removed and then a new solution is pushed into the tail of the list. As suggested in [39], the size of the list depends on the size and the characteristics of the problem. We define the size of the list to be a large value (100) in order to avoid looping. This number does not significantly affect the performance that can be achieved by the NSA because the number of potential traffic-to-egress-point assignments that are not in the memory list is still very large.

5) Diversification

The FM-based neighborhood search strategy always selects, among all the neighboring solutions, the one that yields the largest performance improvement over the current solution. As a result, the overall TE costs improve as algorithm iteration increases, and eventually the algorithm converges until no solution with further performance improvement can be found.

Since the nature of such neighborhood search strategy is greedy (i.e. always choosing the best only at each iteration), the NSA algorithm tends to get trapped into local optima easily. Therefore, we propose an intelligent diversification step in our NSA algorithm to drive the search to examine new regions of solution in order to avoid being trapped in local optima. The diversification step is that if there is no obvious improvement in the solution for a certain number of iterations, we modify the current solution by re-running intra-AS TE on it. We define the threshold of obvious improvement to be 10% and the number of iterations to be 500.

6) Stopping Criterion

Many stopping criteria can be developed depending on the nature of the problem being studied. The most common criterion, which is employed in this paper, is a maximum number of iterations (*MAX_ITER*). However, we do not arbitrarily select this number since the performance of the NSA is mainly dependent on how many times inter-AS traffic flows can be reassigned. Therefore, the maximum iteration number should be related to the number of inter-AS traffic flow. In our experiments, we found that setting the maximum iteration number to be 4 times the number of inter-AS traffic flows gives us a sufficiently good result.

7) The Overall Complexity of the Integrated Optimization Strategy

The worst-case overall time complexity of the NSA-based joint intra- and inter-AS TE is analyzed. We denote n and e the number of inter-AS traffic flows and the number of egress points. The NSA calculates the profit cost for each inter-AS traffic flow by evaluating each potential egress point (lines 10-20 in Fig. 2). The most time consuming step in this block is to find a new minimum cost route by using the Dijkstra’s algorithm (line 14). The time complexity of Fibonacci-heap implementation of the algorithm is $O(|V|\log|V|+|E|)$ where $|V|$ and $|E|$ are the number of nodes and links in the network [41]. Since the worst-case is to examine all the inter-AS traffic flows until the first positive profit cost solution is found, the whole step (lines 6-20) could take $O(n \cdot e \cdot (|V|\log|V|+|E|))$ time. The NSA then iterates until the maximum iteration number (*MAX_ITER*) has been reached (line 2). Therefore, the overall worst-case time complexity of the NSA can be summarized as $O(\text{MAX_ITER} \cdot n \cdot e \cdot (|V|\log|V|+|E|))$.

VI. EVALUATION METHODOLOGY

Having described the Joint-TE strategies, we present our methodology in this section for the evaluation of their performances.

A. Network Topology

We use Rocketfuel [9] Point-of-Presence (POP) level topologies in our experiments. For each topology, POPs correspond to cities. Some POPs have inter-AS links connected to other ASes and we call them border POPs. These border POPs are effectively egress points in our Joint-TE problem. For simplicity, we assume that each border POP is associated with one inter-AS link. Nevertheless, our Joint-TE problem formulation can also be directly applied to the case

where an egress border POP is associated with multiple inter-AS links connecting to different border POPs in the adjacent downstream ASes. In this case, from the local AS point of view, we can treat each border POP at the adjacent AS side as the traffic egress point in our Joint-TE problem. Thus, our Joint-TE problem formulation and algorithms can be also used in this case without any modification.

On the other hand, since the Rocketfuel data do not contain link bandwidth, we set the capacity of each intra- and inter-AS link to be 2.5Gbps and 622Mbps respectively. The rationale of using this capacity setting is based on the evaluation model used in [42]. In this model, the ratio of intra-AS link capacity to inter-AS link capacity is considered to be 4:1. The topology details are summarized in Table 2.

Table 2. Rocketfuel topologies used for evaluation

| AS # | AS Name | # POP nodes | # Intra-AS links | # Border POP nodes |
|------|-------------|-------------|------------------|--------------------|
| 1239 | Sprint (US) | 52 | 84 | 40 |
| 7018 | AT&T (US) | 114 | 148 | 41 |

B. Internet Routing prefixes

For scalability and stability, inter-AS TE can focus only on a small fraction of routing prefixes which are responsible for a large fraction of the traffic [2]. In this paper, we consider 200 such popular routing prefixes. Nevertheless, each of them may not merely represent an individual prefix but also a group of distinct routing prefixes that have the same set of candidate egress points [10] in order to improve network and TE algorithm scalability. Hence, the number of routing prefixes we consider could actually represent an even larger value.

Each border POP can be an ingress or egress point. In order to evaluate the effect of inter-AS TE on the performance of intra-AS TE, we consider the scenario where if a border POP receives a route advertisement towards routing prefix k from adjacent AS Y , then AS Y cannot inject traffic for k into it. This corresponds to multi-hop traffic [10] in which the traffic traverses the network instead of being directed to another egress link of the same border POP. As a result, we cannot assign all the routing prefixes on each border POP as route advertisements. Instead, for generality, we consider half of the routing prefixes randomly selected as route advertisements and the other half as inter-AS traffic in each border POP. We note that this routing prefix generation is a best effort attempt to model prefix distribution, as no synthetic model for the actual behavior of real networks can be found in the literature.

C. Traffic Matrices

We generate synthetic traffic matrices for our evaluation. We generate inter-AS traffic from each POP towards each of the considered routing prefixes. Note that if the POP is a border POP, it can only inject traffic heading towards the routing prefixes that have not been selected as route advertisements. Previous work has shown that intra- and inter-AS traffic are not uniformly distributed [43,44]. According to [45], AS traffic volumes are top-heavy and can be approximated by the Weibull distribution with shape parameter 0.2-0.3. We therefore generate the inter-AS TM following this distribution with the shape parameter 0.3.

We use the Gravity Model (GM) outlined in [46] to generate the local intra-AS TM. The GM approach was proposed based on the findings in [12]. Furthermore, following the suggestions in [44], we randomly classify 40% of POPs as “small”, 40% as “medium” and 20% as “big”. The amount of incoming traffic at a POP is proportional to its size.

D. Algorithm Parameters

For the GA in the nested optimization, we use the suggested values from previous GA research to achieve satisfactory effectiveness and convergence rate of the algorithm [47]. The population size is 200 and the probability of mutation is 0.01. We set the GA to produce maximum 200 distinct inter-AS TE solutions to compute **NESTEDBEST** and **NESTEDWORST**.

E. Performance Metrics

The following performance metrics are used to evaluate the Joint-TE strategies. For these metrics, lower values are better than high values.

- **Overall intra- and inter-AS cost:** these metrics capture the costs of objective function (1) and (2).
- **Total bandwidth consumption:** the required amount of bandwidth to accommodate all the traffic demands in the network, being the sum of the traffic loads over all the intra-AS links.
- **Maximum intra- and inter-AS link utilization:** the maximum intra-AS (inter-AS) link utilization is the maximum utilization on all the intra-AS (inter-AS) links in a network. Minimizing this value ensures that traffic is moved away from congested to less utilized links and is balanced over the links.

VII. EVALUATION RESULTS

This section presents and analyzes evaluation results on the performance metrics achieved by different Joint-TE strategies.

A. Evaluation of the Overall Inter-AS TE Cost and Maximum Link Utilization

We have evaluated the overall inter-AS TE costs achieved by all the strategies for both network topologies. We found that their results exhibit a common characteristic – their overall inter-AS TE costs are nearly identical as well as very close to the optimal solution that is obtained by solving the LP formulation of the inter-AS TE problem.

Figure 3 shows the overall inter-AS cost (y-axis) achieved by all the strategies as a function of the inter-AS traffic demand (x-axis) for the Sprint and the AT&T topologies. The curve denoted by **OPTIMAL** is the optimal solution to the LP formulation of the inter-AS TE problem (i.e. Step-1 of **INTER-OPTIMAL-AWARE-ALG**). However, as already mentioned, this is not feasible for our Joint-TE problem which does not permit splitting inter-AS traffic over multiple egress points. The shapes of the result curves follow the piecewise linear cost function. The figures show that all strategies have almost identical overall inter-AS costs. The inter-AS costs of sequential and nested optimization are similar because the **INTER-OPTIMAL-AWARE-ALG** and the GA use a similar heuristic. These algorithms effectively produce near-optimal

inter-AS costs (approximately within 2% from the optimal solution). The NSA of integrated optimization has also reached a similar inter-AS cost because a significant importance has been given to optimize inter-AS cost over intra-AS cost, as mentioned in Section IV.C. As a consequence, it is expected that those Joint-TE strategies will strive to obtain a near-optimal inter-AS TE cost and, hence, their costs are almost identical.

Figure 4 shows the corresponding maximum inter-AS link utilization of the solutions in Figure 3. On the whole, all the strategies obtain a similar utilization level although there are some small differences among them at some traffic demands. This mainly stems from the fact that the piecewise linear cost function gives the same cost to links with utilizations in the same piecewise linear segment, such as from




Figure 3. Overall Inter-AS cost (top: Sprint, bottom: AT&T)





Figure 4. Maximum inter-AS link utilization (top: Sprint, bottom: AT&T)

1/3 to 2/3 (as shown between any two dashed horizontal lines in the figures). In other words, such links are considered at same utilization level. In the figures, we see that all the strategies have utilization level within the same block at any traffic demand.

We conclude from Figures 3 and 4 that the inter-AS TE solutions achieved by all the strategies are nearly identical and near-optimal. Recall in Section IV.C that, in order to achieve an unambiguous comparison, our aim is to derive TE solutions that remain overall inter-AS cost near-optimal while substantially improving intra-AS cost. At this point, the objective of inter-AS TE has been achieved. In order to determine which strategies produce the best and the worst overall network performance, we proceed to evaluate their overall intra-AS costs.

B. Evaluation of the Overall Intra-AS TE Cost and Maximum Link Utilization




Figure 5. Normalized overall intra-AS cost (top: Sprint, bottom: AT&T)

In our intra-AS TE experiments, the intra-AS TM is the entire local intra-AS TM with an increasing amount of inter-AS traffic. Following the method of deriving a universal measure of congestion in [27], we normalize the resulting overall intra-AS cost by the cost derived from the hop-count based shortest path routing with infinite capacity on each intra-AS link under **SEQOPT-INTER-INTRA** which could be regarded as the most reasonable sequential approach. If the normalized cost is larger than one, it implies that the algorithm is performing as badly as if all flows were along shortest hop paths with loads matching the capacities.

Considering the general picture for the normalized overall intra-AS costs in Figure 5, we first see that **NESTEDWORST** has the worst overall intra-AS cost. Then comes the sequential optimization method and **NESTEDBEST** is better of them. The curve denoted by **OPTIMAL (SEQOPT)** is the optimal solution to the LP formulation of the intra-AS TE problem for **SEQOPT-INTER-INTRA**, which again is not a feasible solution to our Joint-TE problem as it allows arbitrary intra-AS traffic splitting over multiple paths between any pair of nodes. We see that the intra-AS cost of **SEQOPT-INTER-INTRA** is within 3-5% of the optimal solution, thus showing that the **INTRA-OPTIMAL-AWARE-ALG** can achieve performance very closing to the optimal solutions that allow arbitrary traffic splitting.

The above results are expected and can be explained as follows. First of all, the nested optimization finds the best and the worst intra-AS TE solution by evaluating many feasible inter-AS TE solutions, with each solution performing a sequential optimization. Hence, **NESTEDBEST** and **NESTEDWORST** could be regarded as the lower and the upper bound of the sequential optimization. As can be seen, indeed, the performance of **SEQOPT-INTER-INTRA** is always between **NESTEDBEST** and **NESTEDWORST**. The difference in performance between the sequential and the nested optimization sufficiently demonstrates that there indeed exist optimal inter-AS TE solutions that are far better and worse with respect to intra-AS TE. Since the fundamental characteristics of both sequential and nested optimization follow a sequential model that do not optimize intra- and inter-AS TE costs simultaneously, their performances are generally poor. This contrasts to the superior performance of the integrated optimization.






Figure 6. Maximum intra-AS link utilization (top: Sprint, bottom: AT&T)

We follow the methodology in [16] to quantify and compare the performance of different Joint-TE strategies. The comparison metric is the amount of traffic demand the network can cope with before it gets congested (i.e. before the normalized intra-AS cost reaches unity). Our experiments show that the integrated optimization allows the network to cope with approximately 30%-60% more traffic demand than the other non-integrated optimization approaches. This significant improvement implies that the intra-AS traffic matrices for the sequential and the nested optimization are not yet optimized to achieve better overall network performance.

Intra-AS cost can reflect the performance of maximum intra-AS link utilization and total bandwidth consumption. Turning our attention first to the maximum intra-AS link utilization, Figure 6 shows the utilization of solutions in Figure 5 achieved by the different strategies. The integrated optimization is the best strategy. It attempts to keep the maximum utilization below 100% to avoid the high cost penalty as the network load increases. As a consequence, the number of additional traffic flows the network can support before suffering congestion (which we define here as a maximum utilization above 100%) is approximately at least 30% and 60% more than the non-integrated optimization strategies for AT&T and Sprint topologies respectively.






Figure 7. Total bandwidth consumption (top: Sprint, bottom: AT&T)

Figure 7 shows that the total intra-AS bandwidth consumption achieved by the integrated optimization is lower than the other strategies. Together with the results in Figure 6, we observe that the integrated optimization has employed more intra-AS routes that are short and well load balanced in the network. This explains why integrated optimization can achieve better overall intra-AS cost than the other strategies.

C. Summary of Evaluation

From our evaluation of inter- and intra-AS costs, we see that the integrated optimization has successfully produced from non-TE starting solutions to the final solutions that have the same inter-AS cost as those obtained from inter-AS TE (INTER-OPTIMAL-AWARE-ALG and the GA) with an improved intra-AS cost. It is worth noticing that, in comparison to those intelligent intra- and inter-AS TE algorithms that produce near-optimal solutions in a decoupled mode, the overall performance improvement (30%-60%) of the integrated optimization is remarkable. An implication of this finding is to encourage ISPs to move towards an integrated TE approach that is aware of both intra- and inter-AS TE simultaneously.

D. Effect of Joint Optimization on Traffic Granularity

This section elaborates on why the integrated optimization achieves better intra-AS TE performance than the others by investigating the effect of different Joint-TE strategies on traffic granularity. Traffic granularity refers to the level of traffic aggregation that constrains the load balancing ability of traffic aware routing [19]. In this paper, we define coarse (or fine) granularity as the aggregation of traffic into large (or small) trunks in terms of traffic volume. It is known that the traffic granularity can affect TE performance: the smaller the traffic trunks, the greater the flexibility to achieve load balancing.

Since inter-AS TE selects optimal egress points, traffic is then routed onto different corresponding intra-AS routes. As a result, the traffic granularity of each intra-AS traffic flow varies, depending on the outcome of inter-AS TE. It should be emphasized that the distribution of coarse and fine granularity among intra-AS traffic flows is critical to achieving optimal network performance.

We adopt the method described in [19] to order intra-AS traffic flows as follows. We define the Fractional Assignment

(FA) $r_{i,j}$ (also called rank in [19]) of the intra-AS traffic flow between ingress point i and egress point j to be $\frac{t_intra(i,j)}{T}$ where T is the total volume of intra-AS traffic. The smaller the value of FA, the smaller the volume of the traffic flow. In order to compare the achieved traffic granularity among different Joint-TE strategies, we normalize the FA of all traffic flows, because different strategies may have different maximum values of FA for each traffic flow. The normalized FA $r_{i,j}^N$ of the traffic flow is therefore defined as $\frac{r_{i,j}}{r_{max}}$ where r_{max} is the maximum FA value of all traffic flows under all the Joint-TE strategies. To group traffic based on $r_{i,j}^N$, an FA class (a,b) is defined and the members of it consist of all the traffic flows for which $a \leq r_{i,j}^N \leq b$.

Figure 8 presents the distribution of traffic as a function of FA class. The y-axis shows the number of traffic flows that belong to a particular FA class. We have derived this traffic distribution for many different traffic volume scenarios, and found similar results. Thus, here we only consider the highest traffic volume scenario.

The figure shows that all the Joint-TE strategies, except from the integrated optimization, have very similar results. The integrated optimization strategy produces more small FA class traffic flows than the other strategies. Because of these small FA traffic flows (equivalent to finer grain routing in [19]), the NSA allows traffic flows to be assigned more evenly among the intra-AS links, resulting in a more even link utilization and hence load balancing.




Figure 8. Evaluation of intra-AS traffic ranking and inter-AS traffic distribution (top: Sprint, bottom: AT&T)




Figure 9. Inter-AS traffic distribution on intra-AS traffic flows (top: Sprint, bottom: AT&T)

On the other hand, the integrated optimization also produces a few large FA classes. We conjecture that, in order to achieve optimal performance, some good quality intra-AS paths (e.g. the shortest one with high capacity) attract a high volume of traffic, while injecting low volume of traffic onto low-quality paths for achieving load balancing.

Recall that the intra-AS traffic is the sum of local intra-AS traffic and inter-AS traffic. Since the local intra-AS traffic is given, we evaluate how different Joint-TE strategies distribute the inter-AS traffic among the FA classes. For each of our ten FA classes of Figure 8 we have calculated the volume of inter-AS traffic, and divided this by the total inter-AS traffic volume to obtain the percentage of inter-AS traffic in each FA class. In Figure 9, we show the distribution of this inter-AS traffic percentage as a function of FA class. The figure shows that the integrated optimization strategy gives a relatively even distribution of inter-AS traffic across the different FA classes: our conjecture is that it is this more even distribution that contributes significantly to the large FA classes of Figure 8.

From our evaluations of traffic granularity, we see that the reason why the integrated optimization strategy achieved better intra-AS network performance is that inter-AS traffic is *appropriately* merged onto intra-AS traffic flows, resulting in optimal traffic granularity that leads intra-AS TE to produce optimal network performance. Such an optimal inter-AS traffic merging can be achieved by taking intra-AS TE objectives into account during inter-AS TE, i.e. the Joint-TE optimization.

E. Optimized vs. Non-Optimized Starting Solutions for Integrated Optimization

In the preceding experiments, we have used the non-TE (non-optimized) starting solutions for the integrated optimization. In this section, we evaluate the performance of the integrated optimization using a good quality starting solution. Our hypothesis is that much better performance improvement can be obtained when an optimized starting solution is used. We use the solution obtained from **SEQOPT-INTER-INTRA** as optimized starting solution, which is achieved by the existing intra- and inter-AS TE algorithms. As previously shown, the solutions are near-optimal with respect to inter- and intra-AS TE costs when these TE are accomplished separately.

In line with the phenomenon observed in Section VII.A, the integrated optimization using non-optimized and optimized starting solutions produce very similar overall inter-AS costs, which can be regarded as identical. For intra-AS TE performance, Figure 10 shows that the integrated optimization using optimized starting solution can achieve much better intra-AS cost than the sequential optimization. However, it is only slightly better (less than 1%) than that using non-optimized starting solution. This refutes our hypothesis.

The significant performance improvement over the sequential optimization is expected due to the simultaneous TE optimization and that the integrated optimization guarantees the performance of its solution is at least no worse than the starting solutions. On the other hand, the small performance improvement over that using non-optimized starting solutions reflects that the quality of starting solution has not much influence on the quality of the final solution. That is simply because the optimized starting solution itself is not really optimized from the viewpoint of Joint-TE and its quality could be even far inferior from the optimal Joint-TE solution than a non-optimized starting solution. Thus, an optimal inter-AS TE solution can be a mediocre starting point with respect to intra-AS TE, but there are inter-AS TE solutions that are nearly optimal and are far better with respect to the intra-AS TE objectives. An implication of this finding is that the existing effort on devising optimal intra- and inter-AS TE algorithms, which are assumed used in a decoupled mode, may not be sufficient to achieve a truly optimized network performance due to the TE interaction and the fact that the overall network performance can be significantly improved through their joint optimization. The existing inter-AS TE approaches may not improve intra-AS TE performance even though an opportunity for such improvement exists. On the other hand, the existing intra-AS TE approaches may not be able to achieve a truly optimized performance without making it aware of inter-AS TE. In line with our proposal in Section VII.C, the integrated TE approach is an appropriate solution.

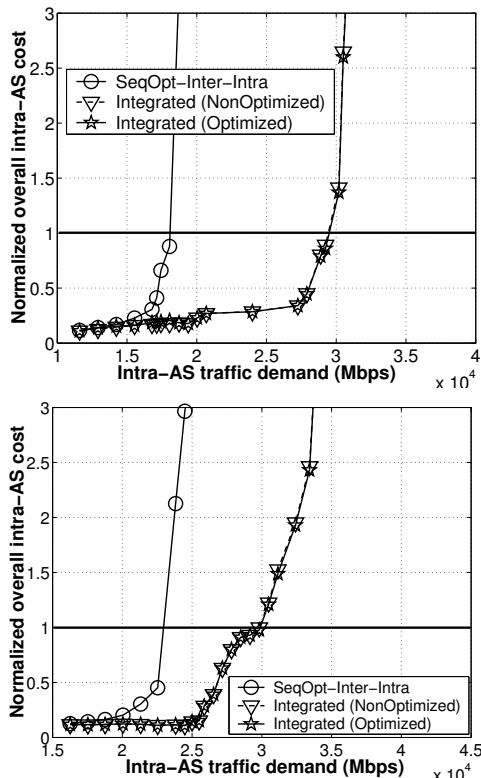


Figure 10. Optimized vs. Non-optimized starting solutions (top: Sprint, bottom: AT&T)

F. Convergence Behaviors of the Neighborhood Search Algorithm for Integrated Optimization

In this section, we analyze the convergence behaviors of the NSA proposed for the integrated optimization strategy. We show how the intra- and inter-AS costs improve iteration by iteration from an initial random TE solution. Figure 11 shows the overall inter-AS cost and the normalized overall intra-AS cost achieved by our NSA over a large number of iterations. As can be seen from the figures, since the neighborhood selection strategy is based on the First Method (FM), i.e. to always choose the solution which produces the largest performance improvement over the current solution, the overall TE costs gradually decrease as iteration increases, and eventually the NSA converges until no solution with further performance improvement can be found.

An observation in Figure 11 is that the normalized overall intra-AS cost does not follow a constant decreasing pattern. This is due to the fact that higher weighting is assigned to inter-AS TE (as discussed in Section IV.C and indicated by the use of α in our NSA), in which case the intra-AS TE objective is transiently sacrificed when the two TE objectives conflict with each other during the NSA evolution.




Figure 11. NSA convergence behaviors

VIII. FURTHER DISCUSSION AND FUTURE WORK

In this section, we discuss future enhancements that could make the Joint-TE more effective.

First of all, we have shown the effectiveness of the joint intra- and inter-AS TE, given that the traffic matrices are fairly accurate. However, in practice it can be hard to determine the traffic matrix accurately as it is susceptible to traffic or routing changes in upstream ASes. As a result, we would not always be able to achieve predictable TE performance when there is presence of a major deviation from the estimated traffic demand. A promising solution for this situation is to make TE robust to traffic uncertainty. Relevant work in this area has targeted robust intra-AS TE, most notably oblivious routing [48], valiant load-balancing [49] and two-hop routing [50]. The key objective of these research efforts is to engineer the network so that it is capable of handling a range of possible intra-AS traffic matrices with satisfactory network performance. We believe that relevant approaches could be brought in play and it would be an interesting research problem to make the Joint-TE robust against the uncertainty of intra- and inter-AS traffic matrices using robust or oblivious routing. In addition, we could investigate which one of these mechanisms gives the best performance for robust Joint-TE and also analyze their tradeoffs.

Another issue is that our Joint-TE does not rely on any data sharing or cooperation between ASes. It can be used by an AS by obtaining traffic matrices and BGP routing data that relate to its own network, which then become the inputs to our Joint-TE optimization. Therefore, the Joint-TE approach can be used independently by each individual AS. However, we believe that, in the case where individual ASes can cooperate with each other by disseminating their traffic and routing information, the overall TE performance in each of their

networks can be radically improved. Although it is difficult for competing SPs to agree on cooperation, we could possibly conceive this work under a federated environment where a SP owns and controls a number of individual ASes. Collaborative TE could help to make the traffic and routing behaviors of these ASes more predictable. This interesting issue is yet to be addressed in the literature and is worth being investigated.

On the other hand, since there is a misalignment between the objectives of intra- and inter-AS TE, counter-intuitive behavior can arise when individual autonomous ASes conduct selfish optimizations. We have shown by evaluation results the effectiveness of Joint-TE. It would also be interesting to investigate the interactions between competing SPs and their impact on network performance using analytical approaches such as game theory. A further step to this work is to analyze the interactions of Joint-TE between two independent ASes as they may often have misaligned TE objectives.

IX. RELATED WORK

Some recent work has investigated the interactions between intra- and inter-AS routing such as the dynamic and the disruption effects between BGP and hot-potato routing [51,52]. The authors in [53] have investigated the impact of BGP route changes on intra-AS traffic, which is related to our work. In addition, some work [7,54] has considered the interactions between BGP and IGP in its TE tools. However, none of these papers has either investigated the interactions between intra- and inter-AS routing from a TE perspective, or proposed any strategy or algorithm for their joint optimization.

Few attempts have been recently carried out towards this TE interaction. Our previous work proposes preliminary thoughts and approaches for the TE interaction [55]. Agarwal et al [56] evaluate the behavior of hot-potato routing during intra-AS link weight optimization. The key difference between our work and those previous works is that ours investigates several very different approaches for the TE interaction. These approaches covered rather complete and representative operations that can be considered by real-world network management. In addition, we propose bi-criteria algorithms for solving the TE interaction problem, which has not been investigated by previous works.

X. CONCLUSION

This paper has considered a joint intra- and inter-AS TE optimization scheme. We showed an interaction effect between intra- and inter-AS TE that can lead to suboptimal overall network performance. These interactions motivate the need for joint optimization of intra- and inter-AS TE in order to further optimize the overall network performance. We first formulated this joint TE optimization as a bi-criteria optimization problem. Then we presented three strategies, namely sequential, nested and integrated optimization, to solve it. Our experimental evaluations revealed that the integrated optimization, which solves intra- and inter-AS TE simultaneously, allows the network to accommodate approximately 30% - 60% more future traffic demand in

comparison to the other strategies that deal with intra- and inter-AS TE separately. The integrated optimization therefore provides a marked improvement on current industry practice towards the collective use of intra- and inter-AS TE.

REFERENCES

- [1] D. Awduche et al., "Overview and Principles of Internet Traffic Engineering," IETF RFC 3272, May 2002.
- [2] N. Feamster, J. Borkenhagen and J. Rexford, "Guidelines for Interdomain Traffic Engineering," ACM SIGCOMM Computer Communication Review, vol. 33, no. 5, October 2003, pp. 19-30.
- [3] B. Quoitin et al., "Interdomain Traffic Engineering with BGP," IEEE Communication Magazine, vol. 41, no. 5, May 2003, pp. 122-128.
- [4] T.C. Bressoud, R. Rastogi and M.A. Smith, "Optimal Configuration for BGP Route Selection," Proc. IEEE INFOCOM, 2003, pp. 916-926.
- [5] N. Hu et al., "Locating Internet Bottlenecks: Algorithms, Measurements, and Implications," ACM SIGCOMM Computer Communication Review, vol. 34, no. 4, October 2004, pp. 41-54.
- [6] R. Zhang and J.P. Vasseur, "MPLS Inter-Autonomous System (AS) Traffic Engineering (TE) Requirements," IETF RFC 4216, November 2005.
- [7] A. Feldmann et al., "NetScope: Traffic Engineering for IP Networks," IEEE Network, vol. 14, no. 2, March/April 2000, pp. 11-19.
- [8] Y.R. Yang et al., "On Route Selection for Interdomain Traffic Engineering," IEEE Network, vol. 19, no. 6, November/December 2005, pp. 20-27.
- [9] N. Spring, R. Mahajan and D. Wetherall., "Measuring ISP Topologies with Rocketfuel," IEEE/ACM Transactions on Networking, vol. 12, no. 1, February 2004, pp. 2-16.
- [10] A. Feldmann et al., "Deriving Traffic Demands for Operational IP Networks: Methodology and Experience", IEEE/ACM Transactions on Networking, vol. 9, no. 3, June 2001, pp. 265-280.
- [11] Y. Zhang et al., "Estimating Point-to-Point and Point-to-Multipoint Traffic Matrices: An Information-Theoretic Approach," IEEE/ACM Transactions on Networking, vol. 13, no. 5, October 2005, pp. 947-960.
- [12] A. Medina et al., "Traffic Matrix Estimation: Existing Techniques and New Directions," ACM SIGCOMM Computer Communication Review, vol. 32, no. 4, October 2002, pp. 161-174.
- [13] R. Teixeira et al., "Traffic Matrix Reloaded: Impact of Routing Changes," Proc. Passive and Active Measurement Workshop, 2005, pp. 251-264.
- [14] A. Feldmann et al., "A Methodology for Estimating Interdomain Web Traffic Demand," Proc. ACM Internet Measurement Conference (IMC), 2004, pp. 322-335.
- [15] S. Balon et al., "Traffic Engineering an Operational Network with the TOTEM Toolbox," IEEE Transactions on Network and Service Management, vol. 4, no. 1, June 2007, pp. 51-61.
- [16] B. Fortz and M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights," Proc. IEEE INFOCOM, 2000, pp. 519-528.
- [17] X. Xiao et al., "Traffic Engineering with MPLS in the Internet," IEEE Network Magazine, vol. 14, no. 2, March/April 2000, pp. 28-33.
- [18] P. Trimintzios et al., "Quality of Service Provisioning through Traffic Engineering with Applicability to IP-based Production Networks," Computer Communications, vol. 26, no. 8, May 2003, pp. 845-860.
- [19] A. Sridharan and R. Guerin., "The Impact of Traffic Granularity on robustness of Traffic Aware Routing," Technical Report, Multimedia and Networking Lab, University of Pennsylvanians, U.S, April 2004.
- [20] A. Sridharan, R. Guerin and C. Diot, "Achieving Near-Optimal Traffic Engineering Solutions for Current OSPF/IS-IS Networks," IEEE/ACM Transactions on Networking, vol. 13, no. 2, April 2005, pp. 234-247.
- [21] K.H. Ho et al., "Inter-Autonomous System Provisioning for End-to-end Bandwidth Guarantees," Computer Communications, vol. 30, no. 18, December 2007, pp. 3757-3777.
- [22] M. Amin et al., "A Closed-Loop Control Traffic Engineering System for the Dynamic Load Balancing of Inter-AS Traffic," to appear in the Journal of Network and Systems Management, 2009.

- [23] B. Halabi, *Internet Routing Architectures*, 2/E, Cisco Press.
- [24] J.P. Vasseur and A. Ayyangar, "Inter-area and Inter-AS MPLS Traffic Engineering", IETF Internet Draft, 2004.
- [25] N. Wang et al., "An Overview of Routing Optimization For Internet Traffic Engineering," *IEEE Communications Surveys & Tutorials*, vol. 10, no. 1, January 2008, pp. 36-56.
- [26] K.H. Ho et al., "Joint Optimization of Intra- and Inter-Autonomous System Traffic Engineering," *Proc. IEEE/IFIP Network Operations and Management Symposium (NOMS)*, 2006, pp. 248-259.
- [27] B. Fortz and M. Thorup, "Optimizing OSPF/IS-IS Weights in a Changing World," *IEEE Journal on Selected Areas in Communications*, vol. 20, no. 4, May 2002, pp. 756-767.
- [28] Jun Wang et al., "Edge-based Traffic Engineering for OSPF Networks," *Computer Networks*, vol. 48, no. 4, July 2005, pp. 605-625.
- [29] N. Wang, K.H. Ho and G. Pavlou, "Adaptive Multi-topology IGP Based Traffic Engineering with Near-Optimal Network Performance," *Proc. IFIP Networking Conference*, 2008, pp. 654-666.
- [30] M. Pioro and D. Medhi, *Routing, Flow, and Capacity Design in Communication and Computer Networks*, Morgan Kaufmann, 2004.
- [31] K. Kar, M. Kodialam and T.V. Lakshman, "Minimum Interference Routing of Bandwidth Guaranteed Tunnels with MPLS Traffic Engineering Applications", *IEEE Journal Selected Areas in Communications*, vol. 18, no. 12, December 2000, pp. 2566-2579.
- [32] R.E. Steuer, *Multi Criteria Optimization – Theory, Computation, and Application*, New York: Wiley, 1986.
- [33] C.N. Chuah, "A Tier-1 ISP perspective & observations of routing behavior", A talk, 2002.
- [34] S. Saroiu et al., "An Analysis of Internet Content Delivery Systems," *ACM SIGOPS Operating Systems Review*, vol. 36, special issue, winter 2002, pp. 315-327.
- [35] T. Karagiannis, P. Rodriguez and K. Papagiannaki, "Should Internet Service Providers Fear Peer-Assisted Content Distribution?," *Proc. ACM/USENIX IMC*, 2005, pp. 63-76
- [36] C.T. Chou, "Traffic Engineering for MPLS-based Virtual Private Networks," *Computer Networks*, vol. 44, no. 3, February 2004, pp. 319-333.
- [37] A Modeling Language for Mathematical Programming. Available at www.ampl.com.
- [38] ILOG CPLEX. Available at www.ilog.com.
- [39] F. Glover and M. Laguna, *Tabu Search*, Kluwer Academic Publishers, Norwell, MA, 1997.
- [40] J. Shen, F. Xu and P. Zheng, "A Tabu Search Algorithm for the Routing and Capacity Assignment Problem in Computer Networks," *Computers & Operations Research*, vol. 32, no. 11, November 2005, pp. 2785-2800.
- [41] T.H. Cormen et al., *Introduction to Algorithms*, MIT Press.
- [42] D. Goldenburg et al., "Optimizing Cost and Performance for Multihoming," *ACM SIGCOMM Computer Communication Review*, vol. 34, no. 4, October 2004, pp. 79-92.
- [43] W. Fang and L. Peterson, "Inter-AS Traffic Patterns and Their Implications", *Proc. IEEE Global Internet Symposium*, 1998, pp. 1859-1868.
- [44] S. Bhattacharyya et al., "POP-Level and Access-Link-Level Traffic Dynamics in a Tier-1 POP," *Proc. ACM Internet Measurement Workshop (IMW)*, 2001, pp. 39-53.
- [45] A. Broido et al., "Their Share: Diversity and Disparity in IP Traffic," *Proc. Passive and Active Measurement Workshop*, 2004, pp. 113-125.
- [46] A. Nucci et al., "IGP Link Weight Assignment for Operational Tier-1 Backbones," *IEEE/ACM Transactions on Networking*, vol. 15, no. 4, August 2007, pp. 789-802.
- [47] X.H. Lin, Y.K. Kwok and V.K.N. Lau, "A Genetic Algorithm Based Approach to Route Selection and Capacity Flow Assignment," *Computer Communications*, vol. 26, no. 9, June 2003, pp. 961-974.
- [48] D. Applegate and E. Cohen, "Making Routing Robust to Changing Traffic Demands: Algorithms and Evaluation," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, December 2006, pp. 1193-1206.
- [49] Z.S. Rui and N. McKeown, "Designing a Fault-Tolerant Network Using Valiant Load-Balancing," *Proc. IEEE INFOCOM*, 2008, pp. 2360-2368.
- [50] M. Kodialam et al., "A Versatile Scheme for Routing Highly Variable Traffic in Service Overlays and IP Backbones," *Proc. IEEE INFOCOM*, 2006, pp. 1-12.
- [51] R. Teixeira et al., "Dynamics of Hot-Potato Routing in IP Networks," *ACM SIGMETRICS Performance Evaluation Review*, vol. 32, no. 1, June 2004, pp. 307-319.
- [52] R. Teixeira et al., "Network Sensitivity to Hot Potato Disruptions," *Proc. ACM SIGCOMM*, 2004, pp. 231-244.
- [53] S. Agarwal et al., "The Impact of BGP Dynamics on Intra-Domain Traffic," *ACM SIGMETRICS Performance Evaluation Review*, vol. 32, no. 1, June 2004, pp. 319-330.
- [54] N. Feamster, J. Winick and J. Rexford, "Network-wide Prediction of BGP Routes," *IEEE/ACM Transactions on Networking*, vol. 15, no. 2, April 2007, pp. 253-266.
- [55] M. Howarth et al, *MESCAL deliverable D1.2*, January 2004. Available at www.mescal.org.
- [56] S. Agarwal, A. Nucci and S. Bhattacharyya, "Measuring the Shared Fate of IGP Engineering and Interdomain Traffic," *Proc. IEEE ICNP*, 2005, pp. 236-245.

Kin-Hon Ho is 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, being involved in European Union and United Kingdom computer networking research projects, 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 is a Professor of Communication Networks at the Department of Electronic and Electrical Engineering, University College London, United Kingdom where he coordinates the activities of the Networks and Services Research Lab. He holds a MEng in Engineering from the National Technical University of Athens, Greece, and MSc and PhD degrees in Computer Science from University College London, UK. He has been responsible for a number of European and UK research projects and industrial collaborations. His research interests focus on networking, network management and service engineering, including aspects such as network dimensioning, traffic engineering, quality of service management, policy-based systems, infrastructure-less wireless networks, autonomic networks and communications middleware. He has contributed to standardization activities in ISO, ITU-T, TMF, OMG and IETF.

Ning Wang is a Lecturer in Networks at the Centre for Communication Systems Research (CCSR), University of Surrey, United Kingdom. Before his appointment, he worked as a postdoc research fellow in CCSR on several EU IST research projects related to future Internet design and engineering. His major research interests include Internet routing algorithms and protocols, Quality of Service provisioning, traffic engineering, as well as peer-to-peer (P2P) and overlay networks.

Michael Howarth is a Lecturer in Networking at the CCSR, University of Surrey, United Kingdom. He holds a Bachelor's degree in engineering science and a D.Phil. in electrical engineering, both from Oxford University, and a M.Sc. in telecommunications from the University of Surrey. His research interests include IP Internet traffic engineering, security systems, protocol design and optimization, in fixed and mobile ad-hoc networks.