# Two Approaches to Internet Traffic Engineering for End-to-End Quality of Service Provisioning

Kin-Hon Ho, Michael Howarth, Ning Wang, George Pavlou and Stylianos Georgoulas

Centre for Communication Systems Research, University of Surrey, UK.

Email: {K.Ho, M.Howarth, N.Wang, G.Pavlou and S.Georgoulas}@eim.surrey.ac.uk

*Abstract*—**This paper considers Traffic Engineering (TE) for the provision of end-to-end quality of service (QoS) guarantees across multiple autonomous systems (ASes) in the Internet. We review an inter-AS QoS management model, and consider algorithms for offline inter-AS and intra-AS TE. We consider two approaches, decoupled and integrated, to allow these systems to work together to achieve optimal TE performance. We quantify through simulation the lower TE costs incurred by the integrated approach.**

*Keywords: Intra-AS and Inter-AS Traffic Engineering, End-to-end Quality of Service*

## I. INTRODUCTION

An important current topic of research is the provision of end-to-end Quality of Service (QoS) guarantees across the Internet. The Internet consists of many Autonomous Systems (ASes) managed by different Internet Service Providers (ISPs). In order to achieve end-to-end QoS, these ASes are required to work together. We use the term intra-AS QoS to refer to QoS provided within an AS, and the term inter-AS QoS to refer to QoS delivered by a combination of multiple ASes. Intra-AS QoS provisioning is relatively easy to accomplish because an ISP has full control of network resources within its own AS. However, inter-AS QoS requires ASes to collaborate to deliver a defined level of QoS across multiple ASes, and this is difficult because methods for collaboration are currently ill-defined. There is thus a need for approaches that manage both intra-AS and inter-AS QoS provisioning.

One approach, which forms the basis for the work we consider here, is to rely on Service Level Agreements (SLAs) in which ASes agree to provide QoS guarantees to other ASes. ISPs then perform Traffic Engineering (TE) to take advantage of these QoS guarantees and optimize the performance of their networks through traffic manipulation [1]. TE can be considered at two levels: intra-AS and inter-AS. The former focuses on the optimization of traffic and resource utilization performance within an AS, while the latter focuses on the optimization of inter-AS traffic exiting or entering an AS [2,12,16]. To assist ISPs in their end-to-end QoS provisioning, offline tools are needed that optimize both intra-AS and inter-AS TE objectives. A number of offline intra-AS TE tools and approaches have been proposed and evaluated, for example [3,4]. Inter-AS TE, however, is today commonly applied in a trial-and-error fashion [1]. Thus, developing a systematic offline inter-AS TE tool can improve the effectiveness of end-

to-end QoS provisioning and also avoid inter-AS misconfiguration. Such an offline inter-AS QoS TE tool, however, has not yet been developed.

In this paper we focus primarily on the inter-working between intra-AS and inter-AS TE: as described above, these are currently treated independently even though there exists a direct relationship between them. Specifically, they may affect each other's performance, because they optimize traffic and resource utilization performance at two different levels. For example, intra-AS resource utilization may need to be considered during inter-AS TE, because inter-AS traffic is routed within the AS and therefore consumes resources. Similarly, inter-AS routes may need to be tuned in order to achieve better intra-AS resource utilization. In this paper, we therefore explore the interaction between intra-AS and inter-AS TE by investigating the following operational challenge:

*If both offline intra-AS and inter-AS TE approaches are developed, how can they work together to achieve end-to-end QoS guarantees with optimal TE performance that takes account of both intra-AS and inter-AS optimization objectives?*

We propose two options for this interaction, namely the decoupled and integrated approaches. The decoupled approach performs a sequential optimization between intra-AS and inter-AS TE, while the integrated approach performs a joint optimization between them. We analyze both approaches in terms of TE performance. Our analysis gives ISPs insight into the relationship of intra-AS and inter-AS TE, to enable them to determine how offline TE systems can be designed in an effective way.

To the best of our knowledge, this paper is the first attempt to quantify the impact of end-to-end QoS provisioning through a combination of both intra-AS and inter-AS TE. The paper and our key contributions can be summarized as follows.

Section II reviews an Internet QoS management model for end-to-end QoS provisioning. This model is based on that developed by the European Union Information Society Technologies (IST) MESCAL project [5]. In Section III, we describe offline intra-AS and inter-AS TE systems, adapted from the MESCAL Internet QoS management model. In Section IV, we establish a direct relationship between intra-AS and inter-AS TE. We then describe two approaches to the interworking between intra-AS and inter-AS TE. We present a TE optimization problem and propose algorithms for both approaches in section V. In section VI, we present our performance comparison of the two approaches and describe

results from a simulation evaluation of both approaches. Finally, we conclude the paper in Section VII.

## II. THE QoS-BASED INTERNET

End-to-end QoS relies on provisioning at two levels: intra-AS and inter-AS.

### A. Intra-AS QoS Provisioning

Within an AS, an ISP has full control over its network resources for QoS provisioning. An ISP establishes SLAs with its customers that guarantee agreed QoS levels within the AS. The ISP then provisions and configures the network to ensure that customer traffic receives the expected QoS. To achieve this, routes that meet the QoS requirements of the customer traffic are found, and resources are allocated to the traffic.

### B. Inter-AS QoS Provisioning

One approach to Inter-AS QoS provisioning is a cascaded QoS model; this is the solution currently being investigated by MESCAL. The approach is based on two concepts: (1) the exchange of QoS and charging information between ASes, and (2) the establishment of SLAs between ASes to contract the negotiated end-to-end QoS guarantees. The key idea of this cascaded model is as follows.

An AS can offer guaranteed QoS levels to its own customer (upstream) ASes, including both destinations within its own AS and destinations to which the QoS is guaranteed by adjacent ASes. Such an offer specifies a remote destination(s), the QoS definition and a charge. Thus, for traffic whose destination is a downstream AS, the offer relies on the local QoS capabilities of the offering AS, supported by the SLAs established with its adjacent provider ASes. These SLAs, in turn, are based on the downstream AS' local QoS capabilities and any SLAs it has established with its adjacent provider ASes, and so on in a cascaded manner. In this cascaded model, an end-to-end SLA chain can be built, each SLA relying on SLAs between downstream ASes. At any point the QoS offered towards a specified remote destination reflects guarantees supported by SLAs.
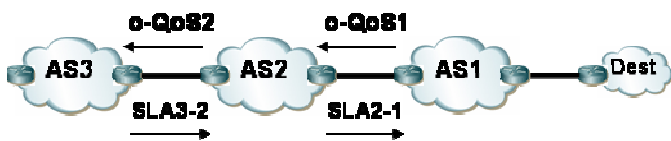


Figure 1.   Cascaded model

Figure 1 shows an example of the cascaded model. Let o-QoS1 be the QoS guarantee and charge offered by AS1 towards the destination 'dest'. AS2 receives an offer of this o-QoS1. We assume it decides to purchase the offered QoS: AS2 then establishes an SLA with AS1 (SLA2-1) to contract the detail of purchased QoS with the associated charge. Now AS2 has a QoS guarantee provided by AS1 for access to 'dest'. AS2 can in turn extend this QoS guarantee by concatenating its own QoS capability with SLA2-1, and then offering an extended QoS (o-QoS2) to AS3. Now o-QoS2 represents the QoS guarantee and charge from AS2 to the destination 'dest'.

AS3 receives o-QoS2 from AS2 and it in turn repeats the decision process, possibly purchasing the offered QoS and establishing SLA3-2. In summary, once offers from other adjacent ASes have been agreed as SLAs, an ISP may build new extended services upon cascaded existing ones. Further details of the MESCAL Internet QoS model can be found in [6]. As an example of QoS concatenation, if the QoS metric is delay, the concatenation operation is performed by addition.

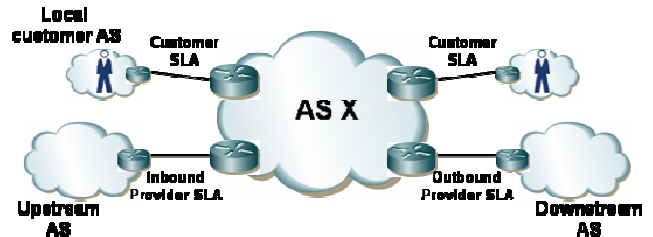### C. Structure of the QoS-based Internet



Figure 2.   Structure of the QoS-based Internet

We now show how a combination of SLAs can be used to form the QoS-based Internet. Figure 2 shows an AS in the QoS-based Internet, capable of providing end-to-end QoS guarantees by establishing SLAs with its local customers and adjacent providers. A *customer SLA* means that the ISP agrees to provide the contracted QoS guarantees to its local customers towards certain destinations (either within the AS or in downstream ASes). A *provider SLA* is divided into two categories: an *inbound provider SLA* is established with an upstream AS, and means that the ISP agrees to provide the contracted QoS guarantees to the traffic received from the upstream AS; an *outbound provider SLA* is established with the downstream AS, and means that the downstream AS provides the negotiated QoS guarantees towards destinations elsewhere in the Internet.

Since the ISP can utilize the resources within the AS and the resources contracted with its adjacent ASes, it wishes to know how to optimize the use of these resources to achieve end-to-end QoS provisioning in the most effective way. This can be resolved by Internet Traffic Engineering, and in the rest of this paper we focus on offline TE for network resource configuration.

## III. OFFLINE QoS-AWARE INTERNET TE SYSTEM

### A. Internet Traffic Types and Relationships

An offline TE system simultaneously examines many parameters such as traffic requirements and network resource constraints, and produces as output a set of network configurations (i.e. routing paths) that optimize the use of all network resources. Internet traffic received by an AS can be divided into four types: *Internal* traffic originates from and terminates within the same AS; *Transit* traffic originates from and terminates within other ASes; *Inbound* traffic originates from the other ASes and terminates within the AS; *Outbound* traffic originates from the AS and terminates within other ASes.

*Intra-AS traffic* involves traffic that has designated ingress and egress routers (e.g. internal traffic). Inbound, outbound and transit traffic also belong to intra-AS traffic once their ingress and egress routers are determined (these may be assigned by inter-AS TE as we will show later). On the other hand, *inter-AS traffic* involves the traffic that originates from or terminates within other ASes, and includes inbound, outbound and transit traffic. We now proceed to discuss in more detail systems that support intra-AS and inter-AS TE.

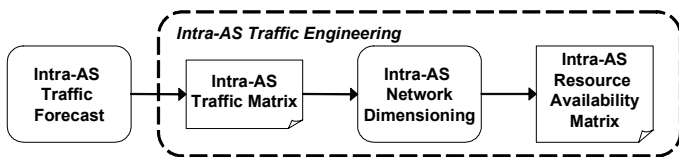## B. Offline QoS-Aware Intra-AS TE System



Figure 3. Offline intra-AS TE system

An offline intra-AS TE system has been proposed [4] which consists of the components shown in Figure 3.

*Intra-AS Traffic Forecast* (TF) forecasts the volume and QoS requirement of intra-AS traffic between every pair of ingress and egress points over a given time interval. The traffic forecast can be extrapolated from measurements, or estimated, or derived from existing / anticipated customer and inbound provider SLAs. The output of Traffic Forecast is an intra-AS Traffic Matrix (TM) which consists of a tuple (ingress router, egress router, QoS requirements) for each aggregate traffic flow. The Offline Intra-AS TE system takes the intra-AS TM as input and performs Intra-AS Network Dimensioning (ND) to achieve an effective QoS provisioning within the AS.

*Intra-AS Network Dimensioning* performs the selection of optimal routing paths for intra-AS traffic between ingress and egress routers so that QoS requirements within an AS are met while optimizing network resource utilization. The output of Intra-AS ND is an Intra-AS Resource Availability Matrix (RAM), which records the resources used within the AS and a network configuration that enforces the resulting routing paths.
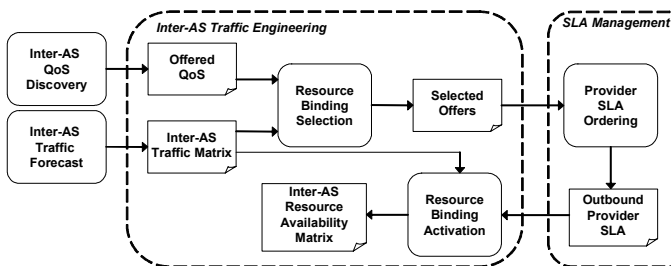
## C. Offline QoS-Aware Inter-AS TE System



Figure 4. Offline inter-AS TE system

Inter-AS TE focuses on the optimization of inter-AS traffic *exiting* or *entering* an AS. These are called *outbound* and *inbound* TE respectively. More specifically, outbound inter-AS TE focuses on an optimal assignment of outbound and transit

traffic onto inter-AS or egress resources. Inbound TE focuses on controlling inbound and transit traffic from adjacent ASes entering an AS through the best ingress points. In this paper we only focus on outbound TE: we therefore ignore inbound TE, since the ingress and egress points of inbound traffic are presumably determined. Outbound and transit traffic are inter-AS traffic because their egress points may be varied by outbound TE and are therefore the subject of this paper.

To ensure effective end-to-end QoS provisioning for inter-AS traffic exiting an AS, there are two issues to consider. First, an ISP needs to decide how much resource it should contract with adjacent ASes (in outbound provider SLAs) while minimizing cost. Second, the ISP has to decide the allocation and optimization of inter-AS resources by appropriately assigning inter-AS traffic. We now describe an offline inter-AS TE system (Figure 4) that primarily addresses the optimization of these two issues. This is a simplified version of the MESCAL TE functional architecture [6], allowing us to focus in this paper on our analysis of the integrated / decoupled approaches to TE.

*Inter-AS Traffic Forecast* (TF) produces the inter-AS Traffic Matrix (TM). As for intra-AS TF, the inter-AS traffic may be extrapolated from measurements, estimated or derived from SLAs. The inter-AS TM consists of a tuple (ingress router, remote destination, QoS requirements) for each flow.

*Inter-AS QoS Discovery* receives the set of advertised offered QoS guarantees from adjacent ASes, each offer specifying the QoS towards a remote destination together with a charge.

*Resource Binding Selection* (RBS) focuses on the selection of QoS advertisements. In general for a specific destination address an AS will receive multiple QoS advertisements from adjacent ASes, since the Internet has multiple connectivity. The AS therefore needs to decide which QoS advertisement it should select for optimum resource utilization and cost. RBS takes as inputs the inter-AS traffic demand from inter-AS TF and a set of QoS advertisements from inter-AS QoS Discovery. It computes the optimum set of resources (e.g. bandwidth) from each QoS advertisement. The purpose of RBS is to optimize towards business objectives by ensuring adequate provisioning of end-to-end QoS to meet forecast demand while minimizing the total charge in purchasing resources from adjacent ASes. The AS then needs to establish outbound provider SLAs with the corresponding adjacent providers to contract these selected resources. This is done by SLA management functions, which are out scope of this paper.

*Resource Binding Activation* (RBA) performs an optimal assignment of inter-AS traffic to the procured inter-AS resources. RBA is a successive step of RBS, and assumes that outbound provider SLAs have been established. The objective of RBA is to optimize inter-AS resource utilization so that no inter-AS resources are overloaded while meeting inter-AS traffic QoS requirements. Since in general an inter-AS traffic flow could be assigned to any one of multiple inter-AS resources, RBA can select the best inter-AS resource for forwarding the inter-AS traffic flow. Inter-AS resources are associated with egress points (egress routers or links) within an AS, so RBA can be viewed as an optimal egress point

selection for inter-AS traffic. The output of RBA is an Inter-AS RAM, which consists of the utilization of inter-AS resources and the outcomes of the optimal inter-AS traffic-resources assignment, including the egress points associated with the selected inter-AS resources for forwarding the inter-AS traffic to the adjacent ASes.

The proposed offline inter-AS TE system has been implemented using Java. We have proposed efficient heuristic algorithms and evaluated them for RBS and RBA in [7] and [8] respectively.

## IV.  INTRA-AS AND INTER-AS TE INTEROPERABILITY

Network operators currently may treat intra-AS and-inter-AS TE independently, even though a direct relationship exists between them. This therefore adversely affects the TE performance of (1) end-to-end QoS guarantees, and (2) intra-AS and inter-AS resource utilization. To achieve optimal TE performance, we need to understand this direct relationship between intra-AS and inter-AS TE, and explore their interaction so that they can work collectively. In this paper we propose two approaches for this interaction and evaluate them through simulation.

### A.  Traffic Engineering Performance

The objectives of TE performance can be classified into two categories: traffic-oriented and resource oriented [9]. Traffic-oriented objectives include aspects that enhance the QoS of traffic streams so as to meet the QoS requirement of the traffic. Resource-oriented objectives include aspects pertaining to the optimization of resource utilization so as to minimize congestion and cost and perform load balancing, etc.

### B.  Relationship between Intra-AS and Inter-AS TE

Although inter-AS TE only determines the best egress points for forwarding inter-AS traffic, the traffic itself has to be routed within an AS and therefore consumes resources. Selecting different egress points results in different routes within the AS, thereby yielding different effects on intra-AS traffic and resource utilization. For example, if Inter-AS TE selects an optimal egress point this may constrain intra-AS TE's solution space with either (1) sub-optimal intra-AS routes toward the egress point or (2) no feasible routes to satisfy the end-to-end QoS requirement of this flow or other traffic, thus failing to provide end-to-end QoS. Therefore, the objectives of intra-AS TE may need to be considered by inter-AS TE. Put another way, if some intra-AS traffic flows use overloaded routes towards a particular egress router, it may be desirable to change the egress routers of some of these flows in order to reduce traffic on the overloaded paths.

We therefore deduce that optimal TE performance with end-to-end QoS guarantees is unlikely to be achieved without the collaboration of both intra-AS and inter-AS TE.

### C.  Approaches to Intra-AS and Inter-AS TE optimization

We explore the interaction between intra-AS and inter-AS TE by investigating how best to get them to work collectively to achieve optimal TE performance. The scope of this interaction is between intra-AS TE and Inter-AS TE RBA, since both their objectives are to optimize traffic and resource utilization performance. We propose two potential approaches for this interaction, namely the decoupled and integrated approaches. Other options that are a combination of these two approaches are also feasible (and are being investigated within MESCAL) but for our initial analysis in this paper we focus on decoupled and integrated approaches to illustrate their respective benefits and drawbacks.

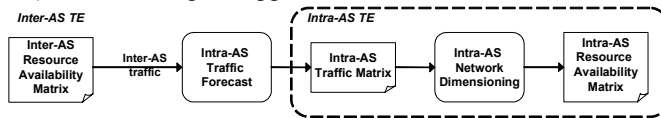#### 1)  The Decoupled Approach



Figure 5.   Decoupled approach

The decoupled approach divides offline intra-AS TE and inter-AS TE into two phases. Intra-AS and inter-AS TE retain their own optimization problems, objectives and algorithms. The two modules perform a sequential optimization, with only minimal interaction between them and limited parameter passing. Since inter-AS TE determines egress points for inter-AS traffic, and intra-AS TE engineers traffic that has designated ingress and egress routers, intra-AS TE cannot be performed until inter-AS TE has been accomplished. The first step therefore is for inter-AS TE to determine egress points for inter-AS traffic, and the second step is to optimize resource utilization within an AS using intra-AS TE (Figure 5). Solving the two optimization problems in the decouple approach could be NP-hard, e.g. intra-AS route selection and egress point selection problems can be respectively reduced to Integral Multicommodity Flow Problem and Generalized Assignment Problem which are NP-hard. The decoupled approach is essentially the general approach in current practice.

Upon the completion of inter-AS TE's RBA, inter-AS traffic can be effectively treated as intra-AS traffic since its ingress and egress routers are determined. The ingress and selected egress routers for inter-AS traffic recorded in inter-AS RAM are used by intra-AS TF to produce an intra-AS TM. This intra-AS TM, thus, not only involves local customer traffic (internal and inbound), but also outbound and transit traffic processed by inter-AS TE. Finally, the intra-AS TM is used by intra-AS TE to perform intra-AS optimal route selection.
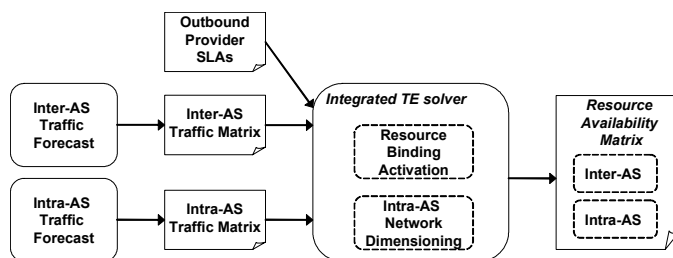
#### 2)  The Integrated Approach



Figure 6.   Integrated approach

In the integrated approach, inter-AS TE is performed in conjunction with intra-AS TE as a joint optimization of egress

point selection and intra-AS route selection. This constitutes a multi-objective optimization problem for both traffic and resource utilization (Figure 6). Solving the joint optimization problem is NP-hard due to the joint of two NP-hard problems.

The Integrated TE Solver is the optimization module, performing RBA and intra-AS ND simultaneously. It takes as inputs the outbound provider SLAs, and the intra-AS and inter-AS TM. The optimization algorithm in the integrated TE solver solves the joint optimization problem. Finally, the integrated TE Solver produces as output the RAM, consisting of both intra-AS and inter-AS RAM.

### D. Analysis of the Decoupled and Integrated Approaches

In this section, we analyze the decoupled and integrated approaches with respect to traffic engineering performance and provide heuristic algorithms for each approach. We also evaluate through simulation the behavior of each approach in the next section.

The decoupled approach can produce the best inter-AS TE performance because inter-AS TE is performed first and independently from the intra-AS TE. However, since inter-AS TE does not consider intra-AS resource utilization, two situations that have been mentioned in section IV.B may emerge. These two situations affect the traffic-oriented and resource oriented objectives of TE performance respectively. Since there is mutual dependency between intra-AS and inter-AS TE, the solution obtained by this approach may not be optimal.

In comparison with the decoupled approach, the integrated approach allows the optimization of intra-AS and inter-AS resource utilization simultaneously. For example, pre-selection can filter out candidate solutions that do not meet end-to-end QoS requirements, and a TE cost function can be employed to take into account the performance of both intra-AS and inter-AS resource utilization. Furthermore, a relative weighting could be introduced into the cost function to control the balance between the intra-AS and inter-AS TE objectives of resource utilization. As a result, the traffic engineering performance achieved by the integrated approach is usually satisfactory and can be flexibly controlled.
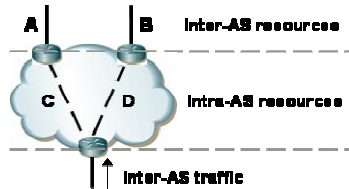


Figure 7.   Example of TE performance

Figure 7 illustrates the traffic engineering performance of the two approaches. Assume that the decoupled approach selects egress point A for an inter-AS traffic flow because this achieves the best inter-AS resource utilization; further assume that the utilization by selecting egress point A is just slightly better than the case when egress point B is selected. However, if the intra-AS route C towards the egress point A is sub-optimal, there will be a degradation of intra-AS resource utilization. In contrast, the integrated approach would select

egress point B since intra-AS route D achieves much better intra-AS resource utilization than that of route C without a significant degradation of inter-AS resource utilization. We believe that the integrated approach can provide a more flexible, controllable and effective TE solution.

## V.   PROBLEM FORMULATION

The optimization problem we consider in this paper is to optimize both intra-AS and inter-AS TE objectives while meeting traffic QoS requirements, using the decoupled and integrated approaches.

We use the following performance metrics as the optimization criteria to evaluate the decoupled and the integrated approaches: (i) Total network cost (the sum of intra-AS and inter-AS cost), (ii) total bandwidth consumption, and (iii) maximum intra-AS and inter-AS link utilization. The first metric captures the overall network cost. Overall intra-AS (respectively inter-AS) cost is defined as the sum of the cost of the intra-AS (inter-AS) links. Fortz and Thorup [13] propose a piecewise linear increasing function of link utilization which imitates the response time of M/M/1 queue to access the cost of intra-AS links. By using the piecewise linear cost function, two objectives of bandwidth usage and resource load balancing are taken into account simultaneously. These two objectives are related to our second and third performance metrics. In other words, the overall network cost is a function of both bandwidth consumption and link utilization.

In this paper, we adopt the piecewise linear function to quantify the cost of intra-AS and inter-AS links. Since inter-AS links are the bottleneck in the Internet [14], we assume that the cost of using them is a factor $\alpha$ times the cost of intra-AS links. We assume $\alpha=2$ as our initial evaluation in this paper. The impact of $\alpha$ on network performance will be evaluated in our future work.

The total bandwidth consumption is the amount of bandwidth needed to accommodate all traffic flows within an AS. It is calculated based on the bandwidth requirement of each traffic flow and the length of path on which the traffic flow has been assigned.

The utilization of a link is the amount of traffic on the link divided by its capacity. The maximum link utilization is the maximum utilization over all links in a network. Minimizing this objective ensures that traffic is moved away from congested to less utilized links and the distribution of traffic is balanced over the links [15].

For all three metrics, the lower values are preferred.

### A. Assumptions

We propose heuristic algorithms for the decoupled and integrated approaches. There are of course many possible algorithm or solution combinations for the two approaches. However, since this paper is not intended as a comparative study of these options, we will propose a classical greedy-based heuristic algorithm as the TE algorithm for the decoupled and integrated approaches. The proposed heuristic algorithms are similar to that proposed by Xiao [10] which has

been deployed in a real network system. The proposed algorithms for both approaches are very similar in order to accurately compare their TE performance. Although it might be appealing to test some more complex algorithms, the approach presented here is sufficient to illustrate the point of interest. For simplicity, but without loss of generality, we make the following assumptions for our algorithm and evaluation:

• Only outbound and transit traffic are considered.

• Bandwidth is considered as the QoS metric.

• The inter-AS resource objective to optimize is the inter-AS link utilization, and the outbound provider SLA is used as capacity constraint.

• Explicit routing is assumed and bandwidth constrained minimum cost routing algorithm is used for intra-AS route selection, where the cost is dynamically calculated for each considered traffic flow by the piece-wise linear cost function proposed in [13]. This not only minimizes resource consumption but also attempts to achieve load balancing within the network. The granularity of explicit paths is per-prefix.

• The AS under consideration has sufficient capacity to meet the end-to-end bandwidth requirements of all inter-AS traffic flows. Thus, the traffic-oriented TE objective can be negligible.

## B. Proposed Heuristic Algorithms

The greedy-based heuristic algorithm for the *decoupled* approach works as follows:

1. Sort all inter-AS traffic flows in a descending order according to their bandwidth requirement. Consider the first traffic flow in that order and assign it to a feasible inter-AS link that meets the bandwidth requirement while incurring the lowest cost. Update the inter-AS resource availability and repeat the selection for the next traffic flow until all the traffic flows have been considered.

2. Based on the sorted traffic flows produced by step 1, select a route that satisfies the bandwidth demand for the first traffic flow between the associated ingress and egress routers. Update network resource availability and repeat this route selection for the rest of the traffic flow in that sorted order.

The greedy-based heuristic algorithm for the *integrated* approach works as follows:

1. Sort all inter-AS traffic flows in a descending order according to their bandwidth requirement.

2. Consider the first ordered traffic flow, identify a set of inter-AS links together with their corresponding intra-AS routes that satisfy the bandwidth demand. Calculate a cost based on the utilization of each possible combination of inter-AS link and intra-AS route were the traffic flow to be assigned to them. Among these possibilities, select the one with the minimum cost. Update the inter-AS and intra-AS resource availability and repeat the selection for the next flow until all the traffic flows have been considered.

The primary difference between the algorithms proposed for the two approaches is that the decoupled approach algorithm divides egress router and intra-AS route selection into two successive phases while the integrated approach algorithm finds both egress router and intra-AS route simultaneously for each traffic flow by considering their conditions (i.e. cost).

## VI. PERFORMANCE EVALUATION

### A. Simulation Configuration

The simulation is based on 100-node topologies generated by BRITE [11] with node degree of 4. The number of border routers is set to 30% of the total network nodes. Note that inter-AS links can be ingress or egress links, and we only consider egress links for outbound TE in this paper. Without loss of generality, we assume that each border router is attached to a maximum of three egress links and the capacity of each egress link is randomly generated between 150 and 300 units. The capacity of each intra-AS link is randomly generated between 80 and 200 units. As inter-AS links are usually the bottleneck in the Internet, the total capacity of all intra-AS links should be larger than that of all inter-AS links.

Due to the fact that only a small fraction of prefixes are responsible for a large fraction of the traffic [12], we therefore consider 100 popular remote destinations which are uniformly and randomly distributed over all the border routers. The number of remote destinations that each border router can reach, specified in outbound provider SLAs, is randomly generated between 30 and 60 units, and these remote destinations are randomly distributed among all the egress links. The contracted bandwidth for a remote destination is randomly generated between 30 and 60 units.

For each aggregated inter-AS traffic flow, the remote destination and the ingress router are randomly generated. The bandwidth demand of each aggregated inter-AS flow is randomly generated between 1 and 40 units.

To ensure confident results, each simulation point takes an average value based on 10 trial runs.

### B. Computational Results and Analysis

Figure 8 and 9 show the inter-AS and intra-AS cost as a function of number of inter-AS traffic flows achieved by the decoupled and integrated approaches respectively. The inter-cost achieved by the two approaches is nearly identical. This is because the cost of using inter-AS links is higher than that of intra-AS links, so the inter-AS link utilization becomes a dominant factor in the selection decision in both approaches. It is possible that there are several inter-AS links that have very similar utilization, but the intra-AS routes connected to them may have different costs. In this case, the integrated approach can select the best combination of inter-AS links and intra-AS routes. We see in Figure 9 that the performance difference between the two approaches is primarily in their intra-AS cost.

The total network cost is defined as the sum of intra-AS and inter-AS cost. Since the inter-AS cost achieved by both approaches are nearly identical, the total cost will mainly
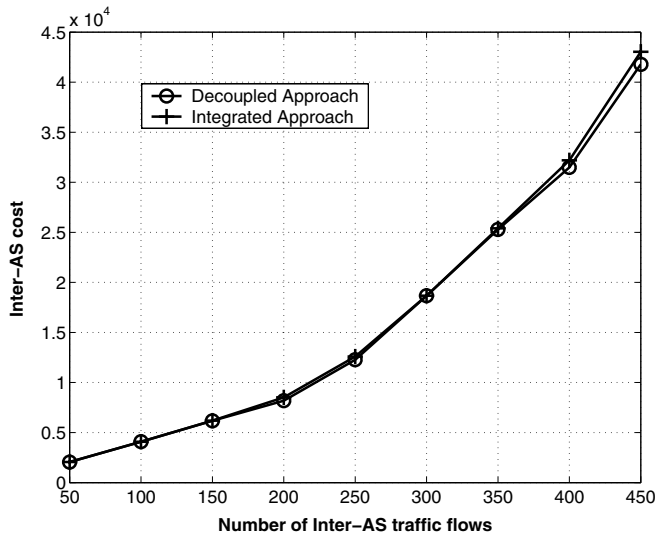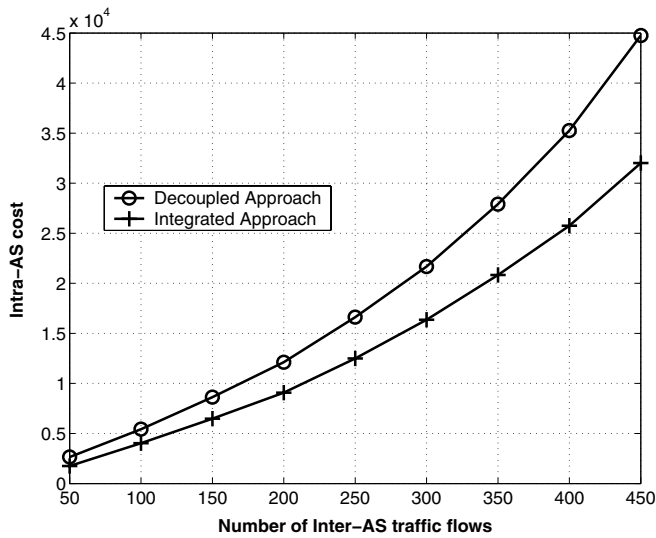
Figure 8.    Evaluation of inter-AS cost



Figure 10.  Evaluation of total bandwidth consumption



Figure 9.        Evaluation of intra-AS cost



Figure 11.  Evaluation of maximum inter-AS link utilization

depend on the intra-AS cost. Hence, the total cost achieved by the integrated approach is much lower than that achieved by the decoupled approach. This resembles the performance shown in Figure 9.

A major reason for the large intra-AS cost in the decoupled approach is due to the increase in bandwidth consumption and link utilization within an AS. Figure 10 shows that the integrated approach uses less bandwidth to accomplish its end-to-end QoS provisioning within the network than the decoupled approach. This is because, when choosing egress routers, the number of hops on the corresponding intra-AS routes has been considered as the selection criteria. The decoupled approach on the other hand may choose an egress router with the best inter-AS link utilization but at the expense of long intra-AS route towards the egress router, resulting in high bandwidth consumption.
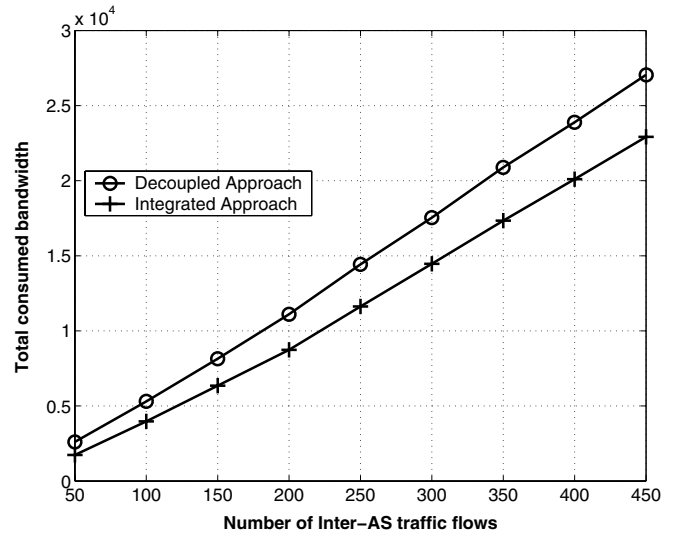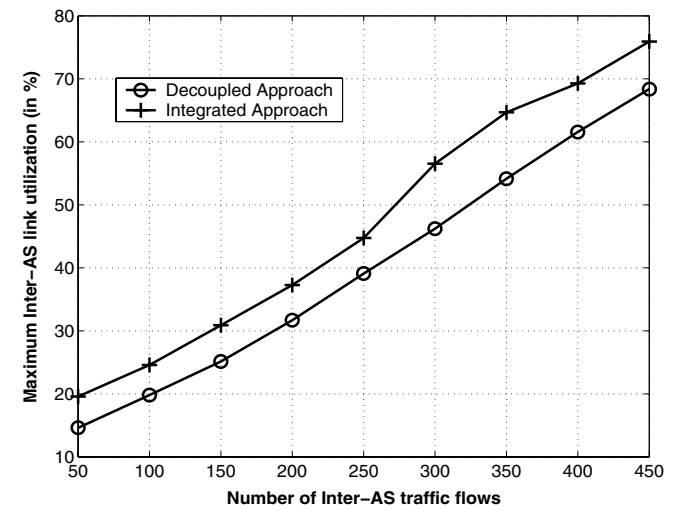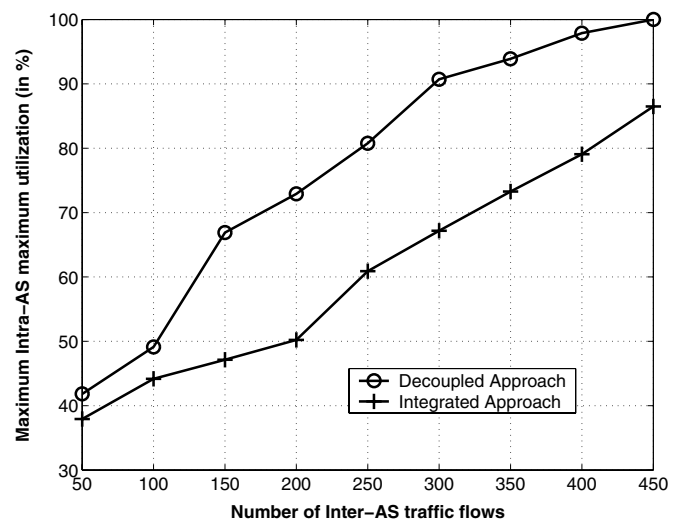


Figure 12.  Evaluation of maximum intra-AS link utilization

Although Figure 11 shows that the integrated approach has a slightly higher maximum inter-AS link utilization than the decoupled approach, both approaches incur nearly identical inter-AS costs, as shown in Figure 8. This may result partially from the piecewise linear cost function, which gives the same penalty to links with utilizations in the same block, such as between 1/3 and 2/3. In this case, such links are considered as at the same level of congestion. Based on the fact that both approaches result in nearly identical inter-AS TE performance, Figure 12 shows that the integrated approach exhibits the advantage of significantly reducing the maximum intra-AS link utilization, compared to the decoupled approach.

As the decoupled approach performs inter-AS TE prior to intra-AS TE, utilization performance on inter-AS link is good compared to that on intra-AS link. On the other hand, the integrated approach takes the balanced approach optimizing between intra-AS and inter-AS resource utilization, therefore the achieved inter-AS resource utilization may not be good as that achieved by the decoupled approach. Nevertheless, significant improvement in intra-AS utilization achieved by the integrated approach compared to the decoupled approach offsets this minor degradation in inter-AS resource utilization.

The performance differences on the total bandwidth consumption and resource utilization achieved by the decoupled and integrated approach in our experiments are consistent with our analysis given in section IV.D.

To compare the overall performance achieved by the decoupled and integrated approaches, our numerical experiments reveal that the integrated approach could save a significant amount of resource cost and achieve a good overall network resource performance, compared to the decoupled approach. Hence, we attempt to answer the question posed in the introduction section by introducing the integrated approach to achieve lower cost complete TE solution.

In fact, other factors can also affect the performance of the two approaches, such as the efficiency of algorithms, the definition of link cost function (linear, concave or discrete), network size and topology, etc. Further experiments are needed to understand their impact on traffic and resource utilization performance.

## VII. CONCLUSION

In this paper, we have described several mechanisms that contribute towards effective end-to-end QoS provisioning, including a cascaded end-to-end QoS management model and offline inter-AS TE systems. Our novel contribution in this paper has been to establish a direct relationship between intra-AS and inter-AS TE, and explore the interaction between them by proposing and analyzing both the decoupled and integrated approaches. We have shown through simulation how the integrated approach results in lower cost TE solutions with lower total consumed bandwidth.

We believe our work provides a starting point in the search for a complete understanding of this fundamental and interesting TE problem. Besides our experimental studies in this paper, there are still many research and practical problems. For example, the interaction between online intra-AS and inter-AS TE is a very important issue in the future.

### REFERENCES

[1] D. Awduche et al., "Overview and Principles of Internet Traffic Engineering," RFC 3272, May 2002.

[2] B. Quoitin et al, "Interdomain Traffic Engineering with BGP," *IEEE Communication Magazine*, May 2003.,

[3] A. Feldmann et al., "NetScope: Traffic Engineering for IP Networks," *IEEE Network*, vol. 14, no. 2, 2002, pp. 11-19.

[4] P. Trimintzios et al., "Service-driven Traffic Engineering for Intra-domain Quality of Service Management," *IEEE Network*, vol. 17, no. 3, May/June 2003, pp. 29-36.

[5] European Union IST MESCAL project. www.mescal.org.

[6] P. Flegkas et al., "Specification of business models and a functional architecture for inter-domain QoS delivery," MESCAL project deliverable D1.1, www.mescal.org.

[7] K.H. Ho et al., "An Incentive-based Quality of Service Aware Algorithm for Offline Inter-AS Traffic Engineering," *IEEE IPOM 2004*, Beijing, China, Oct 2004.

[8] K.H. Ho et al., "Multi-objective Egress Router Selection Policies for Inter-domain Traffic with Bandwidth Guarantees," *Proceedings IFIP Networking 2004*, pp. 271-283.

[9] D. Awduche et al., "Requirements for Traffic Engineering Over MPLS," RFC 2702, September 1999.

[10] X. Xiao et al., "Traffic Engineering with MPLS in the Internet," *IEEE Network*, vol. 14, no. 2, 2000, pp. 28-33.

[11] BRITE: Boston University Representative Internet Topology Generator. http://www.cs.bu.edu/brite/.

[12] N. Feamster et al., "Guildlines for Interdomain Traffic Engineering," *ACM SIGCOMM Computer Communications Review*, October 2003.

[13] B. Fortz and M. Thorup, "Internet Traffic Engineering by Optimizing OSPF Weights," *Proceedings 19th IEEE INFOCOM*, 2000, pp. 519-528.

[14] Akamai, "Internet Bottlenecks: the Case for Edge Delivery Services," Akamai White Paper, 1999.

[15] Y. Wang and Z. Wang, "Explicit Routing Algorithms for Internet Traffic Engineering," *Proceedings IEEE ICCCN* 1999, pp. 582-588.

[16] T.C. Bressoud, R. Rastogi and M.A. Smith, "Optimal Configuration for BGP Route Selection," *Proceedings 22nd IEEE INFOCOM*, 2003, pp. 916-926.