An Incentive-based Quality of Service Aware Algorithm for Offline Inter-AS Traffic Engineering

Kin-Hon Ho, George Pavlou, Michael Howarth and Ning Wang
Centre for Communication Systems Research, University of Surrey, UK
Email: {K.Ho, G. Pavlou, M. Howarth, N. Wang}@eim.surrey.ac.uk

Abstract—This paper focuses on incentive-based offline inter-AS traffic engineering with end-to-end Quality of Service (QoS) guarantees. We investigate a key inter-AS traffic engineering problem, the “egress router selection problem”. The objective is to select an egress router for each expected aggregate inter-AS traffic flow so that the required end-to-end QoS is provided and the capacity constraint of each inter-AS link is met while minimizing the total inter-AS transit cost. The problem is NP-hard and we propose a genetic algorithm to solve it. Simulation results show that our proposed approach performs better than conventional greedy-based approaches.

Keywords – *Offline Inter-AS Traffic Engineering, End-to-End Quality of Service.*

I. INTRODUCTION

The goal of Traffic Engineering (TE) is to optimize the performance of operational networks [1]; this is beneficial for both optimizing resource utilization and for supporting Quality of Service (QoS) guarantees. Traffic engineering can be applied to an Autonomous System (AS) in two levels in a hierarchical fashion: intra-AS and inter-AS TE. Traffic engineering has to be aware of QoS information in both levels in order to find routes that satisfy end-to-end QoS requirements. Intra-AS TE has attracted a lot of attention in recent years and many relevant solutions have been proposed. However, little work has been done to date on inter-AS TE. In this paper, we discuss the motivation, formulate the problem and propose a solution for inter-AS TE.

Inter-AS TE [2,3] focuses on the optimization with respect to traffic- and resource-oriented objectives [4] of inter-AS traffic *exiting* or *entering* an AS. We call these *outbound* and *inbound* inter-AS TE respectively. Since each AS is administratively independent and has no means to control other ASes, inter-AS TE may only be performed locally unless there is some cooperation or agreements between ASes.

In recent years, inter-AS TE had not drawn enough attention due to the absence of business incentives, lack of available data and immaturity of traffic engineering tools [5]. More recently, Internet Service Providers (ISPs) that operate ASes have started to acknowledge that the edges of their networks where they connect to adjacent ISPs ASes are the source of their greatest cost [5]. Three factors constitute this edge cost: (1) *resources on inter-AS links that connect to adjacent ASes are the key bottleneck in the Internet* [6]. Managing these resources so as to avoid overloading is clearly

critical both for adequate end-to-end best-effort services and for providing QoS guarantees. (2) *A charge has to be paid to adjacent ASes for purchasing transit services.* This gives ISPs an economic incentive to carefully select adjacent ASes for routing their traffic in order to maximize their profit by minimizing the charge paid for transit services. We call this charge the “inter-AS transit cost”. (3) *QoS-unaware inter-AS route selection prevents ISPs from generating more revenue by supporting end-to-end QoS guarantees for their customers.* Enabling and supporting the provision of QoS guarantees between ASes becomes an indispensable step to make inter-AS routing aware of QoS.

Considering the above factors, we believe that ISPs have an incentive to use traffic engineering techniques to optimize their edge costs by tactically controlling the traffic exiting the network both for minimizing their inter-AS transit cost and providing end-to-end QoS guarantees to their customers. These incentives drive outbound inter-AS TE.

Effective inter-AS TE requires both internal and external information i.e. information known by the ASes themselves and information obtained from adjacent ASes respectively. The required internal information is inter-AS link capacity, which represents a physical connectivity constraint and addresses edge cost factor (1) above. On the other hand, edge cost factors (2) and (3) can only be satisfied with external information; these are charging incurred by adjacent ASes and QoS information available to remote destinations. Charging information requires inter-AS TE to implement the business objective of minimizing the total inter-AS transit cost while QoS information enables QoS-aware inter-AS route selection for providing end-to-end QoS guarantees to customers. Since the Internet is large in scale and complex, a scalable QoS management model is needed to simplify the management of external information flows between ASes. We will describe the management model that underlies our work in the later section of this paper.

In this paper, we focus on incentive-based offline outbound inter-AS TE to optimize the edge cost of an AS. The outbound inter-AS TE problem can be translated to a problem of directing inter-AS traffic to the ‘best’ egress routers within an AS; we call this the “egress router selection” problem. As the egress routers connect adjacent ASes with inter-AS links, the egress router selection problem can also be viewed as the problem of selecting adjacent ASes for routing inter-AS traffic. Through the selected egress routers, traffic is forwarded to the

designated adjacent ASes over the corresponding inter-AS links, which in turn, forward the traffic to their adjacent ASes and so on, until the destination is reached. The egress router selection problem arises when an AS has multiple connections through different egress routers to adjacent ASes, so that a destination prefix is reachable through multiple egress routers. Selecting different egress routers can yield diverse effects on the inter-AS link utilization and the inter-AS transit cost because different charges and QoS are offered by different adjacent ASes. Addressing the egress router selection issue is important because appropriate selection benefits ASes by significantly improving their edge costs. Yet Inter-AS traffic engineering is today commonly applied in a trial-and-error only fashion [1].

The key objective behind our work is to migrate from trial-and-error towards a systematic approach to solve the egress router selection problem. Our goal can be summarized as: *Within an AS, select an egress router for each expected aggregate inter-AS traffic flow so that the required end-to-end QoS is provided and the capacity constraint of each inter-AS link is met while minimizing the total inter-AS transit cost.*

The algorithm we present here to solve this problem focuses only on bandwidth as the QoS metric; but it should be possible to extend it to accommodate other QoS metrics in a similar fashion. As such, QoS in the offered route refers to the maximum available bandwidth to the destination prefix and is associated with a charge per unit bandwidth. In our previous work [7,8], we solved the bandwidth guaranteed egress router selection problem with the objective of optimizing the AS's internal resource utilization. Since enterprises wish to satisfy business objectives such as minimizing the total cost of their infrastructure [9], we now extend our previous work by investigating the optimization of resources purchased from adjacent ASes so as to provision sufficient end-to-end bandwidth guarantees at the lowest cost. To the best of our knowledge, this is the first attempt at cost optimization toward business objectives through inter-AS traffic engineering while also providing end-to-end QoS guarantees for inter-AS traffic.

The rest of the paper is organized as follows. In section II we describe a scalable model to manage end-to-end QoS. In section III we present the problem formulation and algorithms for optimal egress router selection. In section IV we present the results of our analysis of the egress router selection algorithms. Finally, we conclude the paper in section V.

II. INTERNET QoS MODEL

A key assumption in this paper is that QoS will be deployed globally on the Internet. Our work is based on the MESCAL Internet QoS model [10] as the QoS provisioning framework.

The MESCAL Internet QoS model is based on two essential 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.

An AS will need to know the details of QoS guarantees offered by other ASes downstream (i.e. towards a given destination or set of destinations), and then purchases sufficient

QoS guarantees that the inter-AS traffic can be provided with the desired QoS. In addition, a usage-based charge (per unit of QoS guarantee) is associated with the offered QoS to reflect the cost of provision. Usage-based charging scheme is appropriate for pricing guaranteed services [11] and is incentive-compatible since it stimulates AS to provision end-to-end QoS efficiently by purchasing QoS guarantees from other ASes at the lowest cost.

In addition to purchasing QoS guarantees from adjacent ASes, an AS can in turn 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), a set of QoS parameters and a charge. Thus, for traffic whose destination is a downstream AS, the offer relies on the local QoS capabilities of the offering AS and also the SLAs established with its adjacent provider ASes. This SLA is called outbound provider SLA. These outbound provider SLAs, in turn, are based on the downstream AS' local QoS capabilities and any outbound provider 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 between any two ASes by concatenating the SLAs between corresponding ASes. Thus, at any point the offered QoS reflects guarantees from the offering AS towards a specified remote destination, potentially crossing multiple ASes.

Figure 1. MESCAL cascaded model

Figure 1 shows an example of the MESCAL cascaded model. Let o-QoS1 be the QoS associated with a charge offered by AS1 towards the destination 'dest'. AS2 receives an offer of this o-QoS1 information and let's assume it decides to purchase the offered QoS. AS2 then establishes an SLA with AS1 (i.e. SLA2-1) to contract the detail of purchased QoS with the associated charge. Now AS2 has some QoS guarantees provided by AS1 towards the destination. It can extend these QoS guarantees by concatenating its own QoS capability with SLA2-1, and then offer the extended QoS (i.e. o-QoS2) to AS3. As an example of QoS concatenation, if the QoS metric is delay, the concatenation operation is addition. Now o-QoS2 represents the QoS guarantees and a 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 are agreed as SLAs, an ISP may build new extended services upon existing ones.

An alternative model for end-to-end QoS provisioning is a global centralized broker. The cascaded model, in contrast: (1) makes possible to build scalable end-to-end QoS guarantees between any two ASes while only maintaining SLAs with adjacent ASes; (2) has backward compatibility with BGP, making inter-AS QoS deployment possible through extensions to BGP; and (3) retains privacy for all ASes regarding the details of their interactions.

Further details of the MESCAL Internet QoS model can be found in [10]. Since the end-to-end QoS is known by outbound provider SLAs, inter-AS traffic can be better controlled and tuned to enforce the QoS guarantees. This can be resolved by the egress router selection.

III. INCENTIVE-BASED EGRESS ROUTER SELECTION

Figure 2. Elements of a typical AS

Figure 2 shows a typical AS with all the elements required for the egress router selection. In an AS, each border router can be an ingress and/or egress router with inter-AS link(s) that connect to adjacent ASes. Inter-AS links can be ingress or egress links where inter-AS traffic is received from or passed to adjacent ASes respectively. Since this paper focuses on outbound inter-AS TE, the term “inter-AS link” in the rest of this paper only refers to egress link.

Each border router on an egress link receives a set of offered routes to remote destination prefixes. Each offered route consists of the tuple (prefix, QoS, charge) as described in our cascaded QoS management model. We assume that this AS has considered the best offered route towards each possible destination prefix at each border router if there are multiple route offerings, which may be determined by business policy. Thus, selecting an egress router implies that the corresponding attached egress link and the offered route are selected.

Individual Inter-AS traffic flows are aggregated at each border router (i.e. ingress router) according to destination prefix. Each aggregated inter-AS traffic flow is associated with a bandwidth demand which is the aggregated bandwidth demand over all the individual inter-AS traffic flows destined to the same destination prefix at the ingress router. The aggregated Inter-AS traffic may be measured, estimated or produced from SLAs and suitable extrapolation.

In the case where a destination prefix may be reached through the offered routes at multiple border routers (i.e. egress routers) over the corresponding associated inter-AS links, an opportunity arises to select the best egress router while optimizing the edge cost of the AS. Note that some types of ASes, such as tier 2 and 3, may have peering and transit connections to adjacent ASes towards certain destination prefixes. If traffic is directed to a peering connection, no charge is incurred (i.e. the charge equals to zero and the QoS represents the amount of local customer traffic to be exchanged). In contrast, if traffic is directed to a transit connection, charges are paid to the adjacent ASes. We define the total inter-AS transit cost to be the total charge an AS pays

TABLE I. NOTATION USED IN THIS PAPER

Notation	Description
K	A set of destination prefixes
I	A set of ingress routers
J	A set of egress routers
$t(i,k)$	The bandwidth demand of the aggregated inter-AS traffic flow at ingress router $i \in I$ destined to destination prefix $k \in K$
$Out(k)$	A set of egress routers that can reach destination prefix k
$NEXT_j$	A set of next hop addresses (addresses of border routers in downstream adjacent ASes) that connects to egress router $j \in J$
$f_k(j,n)$	True (1) / False (0); whether destination prefix k is routed through the inter-AS link connects between egress router j and next hop address $n \in NEXT_j$
$c_{inter}^{j,n}$	The capacity of the inter-AS link that connects between egress router j and next-hop address n
$p(k,j)$	The maximum available bandwidth of the offered route towards destination prefix k at egress router j
$Chg(k,j)$	Per unit bandwidth charge of $p(k,j)$
$x_{t(i,k)}^j$	True (1) / False (0); whether the traffic flow $t(i,k)$ has been assigned to the egress router j

for purchasing transit services. The information used to make egress router selection decision is the offered route tuple (prefix, QoS, charge) and the inter-AS link capacity.

Egress router selection can be realized by manipulating BGP attributes such as local preference, selective advertisement or policy-based routing. We also foresee that the solution can be realized by inter-AS MPLS [12]. But these approaches are outside the scope of this paper. as here we are only concerned with *off-line* inter-AS TE based on expected aggregate traffic demand, without focusing on enforcement of the egress router selection.

A. Problem Formulation

We formulate the egress router selection problem as an optimization problem. In table I, we summarize the notation used in the rest of this paper. The optimization objective of the egress router selection is to maximize the AS's profit by minimizing the total inter-AS transit cost

$$\text{Minimize} \sum_{i \in I} \sum_{k \in K} \sum_{j \in Out(k)} x_{t(i,k)}^j \cdot Chg(k,j) \cdot t(i,k) \quad (1)$$

subject to:

$$\sum_{i \in I} \sum_{k \in K} x_{t(i,k)}^j \cdot t(i,k) \cdot f_k(j,n) \leq c_{inter}^{j,n} \quad \forall (j,n) \quad \text{where } j \in J \text{ and } n \in NEXT_j \quad (2)$$

$$\sum_{i \in I} x_{t(i,k)}^j \cdot t(i,k) \leq p(k,j) \quad \forall (k,j) \quad \text{where } k \in K, j \in J \quad (3)$$

$$x_{t(i,k)}^j \in \{0,1\} \quad (4)$$

$$\sum_{j \in Out(k)} x_{t(i,k)}^j = 1 \quad \forall (i,k) \quad \text{where } i \in I, k \in K \quad (5)$$

Constraint (2) enforces the inter-AS link capacity constraint. This constraint not only avoids traffic exceeding the inter-AS link capacity but also contributes to balance the load over all the inter-AS links by specifying a desired link capacity. Constraint (3) enforces the bandwidth constraint for each offered route. Note that the bottleneck of end-to-end QoS provision is at the inter-AS level rather than the intra-AS level because there is likely to be sufficient local bandwidth to accommodate the traffic through over-provisioning within the AS. Therefore, constraint (3) ensures that end-to-end bandwidth guarantees can be provided to inter-AS traffic. Constraint (4) ensures the discrete variables assume binary values; constraint (5) ensures that only one egress router is selected for each traffic flow and inter-AS multi-path routing is not considered in this paper.

The NP-hardness of the egress router selection can be proved by reducing it to a well-known NP-hard problem – the Multi-Resource Generalized Assignment Problem (MRGAP) [13]. Due to the limited space available, we do not show our proof in this paper. Since the egress router selection is NP-hard, we propose a heuristic approach to solve it.

B. Genetic Algorithms

We purpose a genetic algorithm to solve the egress router selection problem. The steps involved in our GA for solving the egress router selection problem are as follows:

Step 1. Create a feasibility mapping table which maps all the feasible egress routers to each inter-AS traffic flow $t(i,k)$. An egress router j is feasible for traffic flow $t(i,k)$ if all the following constraints are satisfied:

- $j \in Out(k)$ (6)

- $t(i,k) \leq C_{inter}^{j,n}$ where n is the next-hop address for routing destination prefix k at egress router j (7)

- $t(i,k) \leq p(k,j)$ (8)

Constraint (6) ensures that the egress router acknowledges an offered route to the destination prefix of the traffic flow. Constraints (7) and (8) ensure that the bandwidth demand of the traffic flow does not exceed the capacity of the inter-AS link and the bandwidth of the offered route that are selected to route the traffic respectively; otherwise, the traffic flow may not receive the desired end-to-end bandwidth guarantees.

Traffic flow	1	2	...	g-1	g
Egress Router	ER1	ER2	...	ERm-1	ERm

Figure 3. Representation of an individual's chromosome

Step 2. Generate an initial population of C randomly constructed chromosomes. Figure 3 shows a representation of an individual chromosome which consists of g genes where g is the number of inter-AS traffic flows to be considered and each gene represents a traffic flow and egress router assignment. The identifier given to each traffic flow represents each

aggregated inter-AS traffic flow $t(i,k)$. Let $S_{t(i,k),c}$ represent the egress router assigned to traffic flow $t(i,k)$ in chromosome $c \in C$. Each of the initial chromosomes is generated by randomly assigning a feasible egress router to each traffic flow based on the feasibility mapping table. Note that an initial chromosome mapping table may not be a feasible solution since the capacity constraint (2) or (3) could be violated.

Step 3. Decode each chromosome to obtain its fitness value. The fitness of chromosome c is equal to the total inter-AS transit cost, given by

$$-\sum_{t(i,k)} Chg(k, S_{t(i,k),c}) \cdot t(i,k) \quad (9)$$

The negative sign reflects the fact that a lower charge is a better solution. If the chromosome contains an infeasible solution, a common approach is to penalize its fitness for the infeasibility. Instead of doing this penalization, we adopt the approach in [14] and associate an unfitness value for each chromosome. The unfitness value of chromosome c is the degree of infeasibility of the chromosome, which equals to the amount of violated capacity over all the inter-AS links and the offered bandwidth,

$$\sum_{j \in J} \sum_{n \in Next_j} \max \left\{ 0, \sum_{i \in I, k \in K, S_{t(i,k),c} = j} t(i,k) \cdot f_k(j,n) - C_{inter}^{j,n} \right\} + \sum_{k \in K} \sum_{j \in J} \max \left\{ 0, \sum_{i \in I, S_{t(i,k),c} = j} t(i,k) - p(k,j) \right\} \quad (10)$$

With the division of fitness and unfitness, chromosomes are evaluated in a two-dimensional plane, so the selection and replacement can direct the search toward feasible solutions.

Step 4. Select two parent chromosomes for reproduction. We use the pairwise tournament selection method. In a pairwise tournament selection, two individual chromosomes are chosen randomly from the population and the one that is fitter (i.e. highest fitness value) is selected for a reproductive trial. Two pairwise tournament selections are held, each of which produces one parent chromosome, in order to produce a child chromosome.

Step 5. Generate two child chromosomes by applying a simple one-point crossover operator on the two selected parents. The crossover point $p_{crossover}$ is randomly selected. The first child chromosome consists of the first $p_{crossover}$ genes from the first parent and the remaining $(n - p_{crossover})$ genes from the second parent. The second child chromosome takes the parent genes that have not been considered by the first child chromosome.

Step 6. Perform probabilistic mutation on each child chromosome. The mutation simply exchanges elements in two selected genes (i.e. exchanging assigned egress routers between two randomly selected traffic flow) without violating the constraints (6) – (8).

Step 7. The child chromosomes may be further improved by applying the following two problem-specific heuristic operators to improve their fitness and unfitness:

a) For each inter-AS traffic flow that has been assigned to an infeasible egress router with which either the capacity constraint (2) or (3) is violated, find a feasible egress router that incurs the lowest inter-AS transit cost. Denote $Diff_cost_{i(i,k)}$ the difference between the original cost and the new cost after the traffic flow has been assigned to the feasible egress router. Among those inter-AS traffic flows, select the one with the lowest $Diff_cost_{i(i,k)}$ and assign it to the corresponding selected feasible egress router. This operator is repeated at most H times where H is a parameter that optimizes the algorithm's performance.

b) For each inter-AS traffic flow, find a feasible egress router that produces the lowest inter-AS transit cost. If such egress router can be found, then reassign the traffic flow to it.

The heuristic operator (a) aims to reduce the unfitness value of the child chromosome by reassigning traffic flows from infeasible to feasible egress routers while keeping the total inter-AS transit cost as low as possible. The heuristic operator (b) attempts to improve the fitness of the child chromosome by reassigning traffic flows to egress routers with the lowest cost.

Step 8. Replace two chromosomes in the population by the improved child chromosomes. In our replacement scheme, chromosomes with the highest unfitness are always replaced. If there are no unfit solutions, the lowest fitness ones are replaced.

Step 9. Repeat step 4-8 until N_{cd} child chromosomes have been produced and placed in the population.

Step 10. Check if the termination criterion is met. The termination criterion is that either the average and the best fitness over all the chromosomes in two generations become the same or once the selected number of iteration, N_{it} , has been reached to avoid long convergence on the algorithm. Repeat step 4-9 until the termination criterion is met.

C. Alternative Algorithms

Since none of the published egress router selection algorithms address the objectives we wish to target, we compare the GA with a greedy heuristic that one imagines might be used by an ISP, which we refer to as Greedy-cost.

The Greedy-cost sorts all the inter-AS traffic flows in descending order based on their bandwidth demand and selects one at a time in that order. Among all the feasible egress routers that have sufficient resources to accommodate the traffic, select the one which incurs the least inter-AS transit cost. The unallocated resource on the corresponding selected inter-AS link and the offered bandwidth are then updated. This algorithm repeats for the next traffic flow until all traffic flows have been considered.

In addition to the Greedy-cost, the Random egress router selection is included as a baseline comparison. The Random egress router selection is similar to the Greedy-cost except that the selection is done at random. It may be viewed as the solution obtained from the trial-and-error inter-AS TE without using a systematic approach.

A. Simulation Configuration

We evaluate the three proposed egress router selection algorithms through simulation. The simulation software is written in Java. The simulation is based on 100-node topologies. The number of border routers is set to 30% of the total number of network nodes. According to the rank exponent of power-law properties [15], the number of inter-AS links of an N -node AS can be estimated by $2.5 \cdot N^{0.8} \cdot (N^{0.2} - 1)$ [16] when the rank exponent equals to -0.81. Note that inter-AS link can be ingress or egress links, and we only consider egress links in this paper. We assume that half of the inter-AS links are egress links and they are randomly distributed among all the border routers. 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.

Feamster [3] discovered that whilst a typical default-free routing table may contain routes for more than 100,000 prefixes, only a small fraction of prefixes are responsible for a large fraction of the traffic. Based on his finding, we consider 100 popular destination prefixes which are randomly distributed over all the border routers. The number of destination prefixes that each border router has been offered is randomly generated between 30 and 60, and the destination prefixes are randomly distributed among the egress links. The offered bandwidth of each destination prefix at each border router is randomly generated between 30 and 60, while the associated charge varies according to the simulation scenarios.

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

For the GA parameters, we adopt the suggested values from previous GA research to achieve satisfactory effectiveness and convergence rate of the algorithm [17]. The population size is 200, the value of H of the heuristic operator (a) is 40 since the egress router selection problem is highly constrained, $N_{cd} = 40$, the probability of mutation is 0.01 and $N_{it} = 150$.

B. Simulation Results

Figure 4 shows the total inter-AS transit cost of the Greedy-cost and the GA as a function of number of inter-AS traffic flows under a number of scenarios. The legend describes the name of the algorithm followed by the percentage of established peering connections.

The simulation consists of two major scenarios. The first scenario, namely all transit connections, consists of 0% peering connection established between the AS and its adjacent ASes. The charge is randomly generated between 1 and 10 (it could not be zero since all transit connections should have a non-zero transit cost). The figure shows that the GA performs better than the Greedy-cost as the number of inter-AS traffic flows increases. When the number of inter-AS traffic flows is small, the network has plenty of resources to accommodate these flows so that the two algorithms produce a similar egress router

Figure 4. Performance evaluation of algorithms

selection and total inter-AS transit cost. However, as the number of inter-AS traffic flows increases, network load increases and some resources reach their capacity limitations such that some traffic flows have been directed to the resources which incur high inter-AS transit cost. In this case, careful selection of egress routers is necessary in order to minimize the total inter-AS transit cost. This has been achieved by the GA.

On the other hand, the second scenario consists not only of transit connections but also some established peering connections. We evaluate three degrees of peering: 3, 6 and 9% of the total destination prefixes. Since a peering connection is free, the charge for the connection is equal to zero. The simulation data is based on the all transit connections scenario except that a designated number of destination prefixes are randomly selected as peering connection. Since our purpose is to merely evaluate the performance of inter-AS TE algorithms with some free-of-charge resources (i.e. peering), we follow the assumption in [18] of ignoring peering/transit restrictions.

The figure shows the anticipated result that the GA always performs better than the Greedy-cost in all the three peering scenarios. The increasing improvement as the number of inter-AS traffic flows increases can be explained by the reason given in the all transit connections scenario. Interestingly, the GA achieves greater improvement than the Greedy-cost as the degree of peering increases. This is because the available resources in peering connections do not incur any charges so that GA can more efficiently use them in order to further minimize the total inter-AS transit cost.

Table II shows the increasing improvement (in %) of the GA over the Greedy-cost and the Random selection as the total number of inter-AS traffic flows increases for the two considered simulation scenarios. In comparison to the Random selection, the performance of the GA is outstanding. This shows that a systematic inter-AS TE approach, especially using our proposed GA, is certainly important and valuable.

TABLE II. IMPROVEMENT OF THE GA OVER THE OTHER ALGORITHMS (IN %)

Number of inter-AS traffic	200	250	300	350	400	450	500
over Greedy-cost, with 0% peering	1.6	2.3	2.9	3.8	5.3	6	6.8
over Random with 0% peering	70.6	68.9	67.1	65.5	63.1	61.2	59
over Greedy-cost, with 3% peering	1.9	3.7	4.3	5	5.9	7.1	8
over Random with 3% peering	75.9	73.5	72	70.8	68.3	66.1	63.1
over Greedy-cost, with 6% peering	2.5	3.9	4.7	5.6	7.3	8.8	9.8
over Random with 6% peering	80.2	78.3	77	75	73	70.2	67.4
over Greedy-cost, with 9% peering	3.2	4.4	5.1	7.1	8.8	10.5	12
over Random with 9% peering	83.5	82.1	80.5	78.8	76.4	74.6	71

V. CONCLUSION

In this paper, we present and formulate an incentive-based offline inter-TE problem, namely egress router selection, with end-to-end QoS guarantees. We show the business objectives of an AS to perform an optimization on the egress router selection. We propose and develop a genetic algorithm to solve the problem, and show the results of an implementation embodying the algorithm on simulated network topologies and inter-AS traffic for all-transit and transit-peering scenarios. We conclude that the problem formulation is sound and our genetic algorithm shows marked improvement over alternatives in the two scenarios. Both the formulation and the solutions should be of great value to some ISPs, as they can get rid of the trial-and-error inter-AS route configurations and adopt our systematic approach to provide the required end-to-end QoS and to meet the capacity constraint of each inter-AS link while minimizing the total inter-AS transit cost. As future work, we plan to extend the egress point selection to support multiple classes of services.

REFERENCES

- [1] D. Awduche, A. Chiu, A. Elwalid, I. Widjaja and X. Xiao, "Overview and Principles of Internet Traffic Engineering," RFC 3272, May 2002.
- [2] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen and O. Bonaventure, "Interdomain traffic engineering with BGP," *IEEE Communication Magazine*, May 2003.
- [3] N. Feamster, J. Borkenhagen and J. Rexford, "Guidelines for Interdomain Traffic Engineering," *ACM SIGCOMM Computer Communications Review*, October 2003.
- [4] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell and J. McManus, "Requirement for Traffic Engineering Over MPLS," RFC2702, September 1999.
- [5] T.E. Monk, "Inter-domain Traffic Engineering: Principles and Case Examples", *The Internet Society's 12th Annual INET Conference*, Arlington, VA, USA, June 2002.
- [6] Akamai, "Internet Bottlenecks: the Case for Edge Delivery Services," Technical Report, Akamai White Paper, 1999.

- [7] K.H. Ho, N. Wang, P. Trimintzios, G. Pavlou and M. Howarth, "On Egress Router Selection for Inter-domain Traffic with Bandwidth Guarantees," in *Proceedings of IEEE HPSR*, Arizona, USA, April 2004.
- [8] K.H. Ho, N. Wang, P. Trimintzios and G. Pavlou, "Multi-objective Egress Router Selection Policies for Inter-domain Traffic with Bandwidth Guarantees," in *Proceedings of IFIP Networking 2004*, Athens, Greece, May 2004.
- [9] Aiber et al, "Autonomic Self-Optimization According to Business Objectives," in *Proceedings of the International Conference on Autonomic Computing (ICAC'04)*, May 2004.
- [10] MESCAL Deliverable D1.1. Available at <http://www.mescal.org>
- [11] B. Stiller et al, "Pricing and QoS," Quality of Future Internet Services Cost Action 263 Final Report, M. Smirnov et al (Eds.), Springer LNCS 2856, 2003, pp. 263-291.
- [12] J.P. Vasseur and A. Ayyangar, "Inter-area and Inter-AS MPLS Traffic Engineering," IETF Internet Draft, draft-vasseur-ccamp-inter-area-as-te-00.txt, February 2004.
- [13] B. Gavish and H. Pirkul, "Algorithms for the multi-resource generalized assignment problem," *Management Science*, vol 32, 1991, pp. 695-713.
- [14] P.C. Chu and J.E. Beasley, "A genetic algorithm for the set partition problem," Working paper. The Management School, Imperial College, London, England, 1995.
- [15] M. Faloutsos, P. Faloutsos and C. Faloutsos, "On Power-Law Relationships of the Internet Topology," in *Proceedings of ACM SIGCOMM'99* Cambridge, MA, USA, 1999.
- [16] M.El-Darieby, D. Petriu and J. Rolia, "A Hierarchical Distributed Protocol for MPLS Path Creation," in *Proceedings of IEEE ISCC'02*, Italy, July 2002.
- [17] 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, pp. 961-974, 2003.
- [18] J. Feigenbaum, C. Papadimitriou, R. Sami and S. Shenker, "A BGP-based Mechanism for Lowest-Cost Routing," in *Proceedings of the 21st ACM Symposium on Principles of Distributed Computing*, Monterey, California, July 2002, pp. 173-182.