

On Egress Router Selection for Inter-domain Traffic with Bandwidth Guarantees

Kin-Hon Ho, Ning Wang, Panos Trimintzios, George Pavlou and Michael Howarth

Centre for Communication Systems Research, University of Surrey, UK

Email: {K.Ho, N.Wang, P.Trimintzios, G.Pavlou, M.Howarth}@surrey.ac.uk

Abstract—As the Internet has grown in size and diversity of applications, the next generation is designed to accommodate flows that span over multiple domains with quality of service guarantees, and in particular bandwidth. In that context, a problem emerges when destinations for inter-domain traffic may be reachable through multiple egress routers. Selecting different egress routers for traffic flows can have diverse effects on network resource utilization. In this paper, we address a critical provisioning issue of how to select an egress router that satisfies the customer end-to-end bandwidth requirement while minimizing the total bandwidth consumption in the network.

I. INTRODUCTION

As the Internet has grown in size and diversity of applications, the next generation Internet is intended to accommodate flows with end-to-end Quality of Service (QoS) guarantees. To provide efficient end-to-end QoS guarantees, QoS routing and Traffic Engineering (TE) have become indispensable: the former selects a path that meets the QoS requirements while the latter optimizes resource utilization in order to be able to carry more traffic flows in the network. In the past decade, there has been a considerable amount of work on QoS routing and traffic engineering at the intra-domain level. However, only little attention has been given to the inter-domain problem. We consider that inter-domain QoS routing and traffic engineering should be addressed for the following reasons.

Inter-domain QoS: End-to-end QoS over the Internet includes both intra- and inter-domain QoS. Even though research in intra-domain QoS is mature, the lack of inter-domain QoS support hinders the deployment of end-to-end QoS. Thus, together with the current QoS-aware intra-domain routing, inter-domain QoS routing will facilitate an end-to-end QoS-based Internet, which will benefit both Internet Service Providers (ISPs) and their customers. The current inter-domain routing protocol, the Border Gateway Protocol (BGP), however, does not cater for QoS support.

Inter-domain TE: Inter-domain traffic engineering [1] concerns forwarding traffic entering or exiting a network based on some optimization objectives. One of the inter-domain traffic engineering problems is to direct the inter-domain traffic flows to the 'best' egress router within a domain towards certain destination prefixes; this we call the "egress router selection" problem. The problem arises when a domain has multiple connections to neighboring domains, so that a situation can emerge that a destination prefix is reachable

through multiple egress routers. Selecting different egress routers for traffic flows can have diverse effects on the network resource utilization. Addressing inter-domain TE is important because appropriate selection of egress routers for inter-domain traffic flows benefits ISPs by improving network resource utilization. Inter-domain traffic engineering, however, is commonly applied today in a trial-and-error only fashion.

Based on this reasoning, we aim to develop a systematic approach to solve the inter-domain TE problem with end-to-end QoS support, i.e. QoS guaranteed egress router selection.

Based on the assumption that most QoS requirements can be derived from bandwidth [2], our work focuses on bandwidth requirements. Thus, the problem we address becomes Bandwidth Guaranteed Egress Router Selection (BGERS).

Our goal is summarized as follows: *Given a customer traffic flow in an ISP network, select an egress router that satisfies the customer end-to-end bandwidth requirement while minimizing total bandwidth consumption in the network. Each customer traffic flow consists of a destination prefix that belongs to a remote domain and a bandwidth requirement. An egress router must be selected amongst egress routers that offer the guaranteed bandwidth to the destination prefix.*

Related work on inter-domain QoS routing is as follows. Bonaventure [3] focuses on how to distribute flexible QoS information by BGP in different network scenarios. Cristallo and Jacquet [4] propose a new attribute, the QoS_NLRI (Network Level Reachability Information), for the BGP UPDATE message to carry QoS information. Xiao [5] proposes a similar QoS extension to BGP to perform the bandwidth advertising and routing. On the other hand, research on egress router selection has only been done in the context of best-effort traffic. Bressoud [6] determines an optimal selection of outgoing links and associated border routers, where the selection optimizes the ISP's network resource utilization. The ISP, however, can only select an egress router based on prefix reachability and the egress link capacity information, without knowing whether the selection of egress routers can satisfy the end-to-end bandwidth requirement.

The key to solve the BGERS problem is support for traffic engineering information (e.g. bandwidth) within and between domains. In this paper, we propose a TE-enabled Internet architecture to achieve this, which includes traffic engineering extensions to both current intra- and inter-domain routing protocols. Our work extends the egress router selection problem presented in [6] by considering end-to-end bandwidth

guarantees. We propose three heuristic algorithms to solve the BGERS problem and evaluate their performance through simulation. To the best of our knowledge, this work is the first attempt at inter-domain traffic engineering using BGP policies to control inter-domain traffic flows with end-to-end bandwidth guarantees.

The rest of the paper is organized as follows. In section II we present a TE-enabled Internet architecture. In section III we formulate the BGERS problem and we propose heuristic algorithms to solve it. Section IV presents the evaluation of those algorithms through simulation. Finally, we conclude our work and discuss future research directions in section V.

II. TE-ENABLED INTERNET ARCHITECTURE

To address the BGERS problem, the TE-enabled Internet architecture requires that the current intra-domain and inter-domain routing protocols must be able to convey bandwidth information. We assume Traffic Engineering extensions to OSPF (OSPF-TE) [7] as the intra-domain routing protocol, which disseminates bandwidth information within the domain. Moreover, we assume Constrained Shortest Path First (CSPF) [8] with unit link cost to calculate a bandwidth constrained path between an ingress router and an egress router in the domain. The selected path is pinned and bandwidth is allocated on it. This can be done by establishing a Label Switched Path (LSP).

On the other hand, the lack of TE information support in the current BGP hinders the deployment of BGERS. To solve this deficiency, it is necessary to record bandwidth information in the BGP UPDATE message, which represents the ability of a domain to provide the route with such bandwidth. In [4], a new attribute, the QoS_NLRI, is proposed for this purpose. We assume that bandwidth information, which takes a single value, is conveyed through a similar attribute and call the extended BGP the Traffic Engineering extensions to BGP (BGP-TE).

The bandwidth information conveyed by BGP-TE is guaranteed by a Service Level Agreement (SLA) established between neighboring domains in a management time-scale. Each domain is configured based on established SLAs to make sure that sufficient bandwidth is provisioned for other domains. The outcome of this bandwidth provisioning is Bandwidth Capability (BC), which is the bandwidth that has been allocated to a path between an ingress router and an egress router within a domain towards certain destination prefixes. The bandwidth capability is advertised to the neighboring domains by BGP-TE.

The technical implication of signing an SLA with neighboring domains is the bandwidth capability binding: A domain binds its bandwidth capability to the bandwidth information advertised by the neighboring domains and uses the resulting binding as the basis for agreeing new SLAs with its customers. The bandwidth capability binding is unidirectional and is done by a simple algebraic method.

For a large scale Internet to provide bandwidth guarantees between edge domains, ISPs have to collaborate and provide transit services to other domains' traffic flows. The concatenation of SLAs between domains can ensure end-to-end bandwidth guarantees for end customers. In this context, the bandwidth information advertised by BGP-TE is the

unidirectional cascaded effect of bandwidth capability binding between each two domains along a BGP path. In other words, it is the concatenated bandwidth that is guaranteed starting from the downstream domain (i.e. the one which advertises the bandwidth information) until the destination domain.




Figure 1. Bandwidth capability binding

We give a small example of bandwidth capability binding in Figure 1. The example also shows how ISP 1 provides bandwidth guarantees to the traffic flows of customer 1 destined to customer 2, conforming to the customer SLA established with customer 1. We denote by BC_X the unidirectional bandwidth capability of ISP X towards destination prefixes that belong to customer 2. We assume that ISP 1 has established a provider SLA with ISP 2 for bandwidth guarantees to customer 2, thus ISP 2 will advertise BC_2 to ISP 1 through BGP-TE. When ISP 1 receives BC_2 , it performs bandwidth provisioning according to the customer SLA and binds BC_1 to BC_2 . This binding forms a unidirectional eBC (extended BC). The value of eBC is equal to the minimum of BC_1 and BC_2 . ISP 1 can then provide eBC bandwidth guarantees to the traffic flows of customer 1, conforming to the customer SLA. Each domain uses bandwidth information, provided by BGP-TE and OSPF-TE, to optimize its network resource utilization by selecting appropriate egress routers for inter-domain traffic flows with bandwidth guarantees.

III. BANDWIDTH GUARANTEED EGRESS ROUTER SELECTION

A. Overview




Figure 2. A general TE-enabled Internet architecture

With the TE-enabled Internet architecture, we can provision end-to-end bandwidth guarantees for inter-domain traffic flows.

In this section, we present the BGERs problem formulation and propose heuristic algorithms to solve it.

Figure 2 illustrates a general TE-enabled Internet architecture. For the ISP domain under consideration, we consider a set of border routers as well as a set of intra- and inter-domain links. An inter-domain link connects a border router of an ISP to a border router of the downstream domain. Each border router may connect to multiple inter-domain links. We assume that the ISP has established SLAs with its downstream domains for bandwidth guarantees and border routers in the ISP and downstream domains support BGP-TE. Through BGP-TE, each border router within the ISP receives route advertisements of destination prefixes associated with the bandwidth information from downstream domains. Each border router then selects the best route for each prefix based on the usual BGP decision process and distributes the route to other border routers within the domain through Internal BGP (IBGP) mesh between border routers. As a result, all the border routers within the domain have the same view on which border routers they can use to reach a specific destination prefix with an amount of guaranteed bandwidth. It is possible that a border router receives more than one route advertisement with a common prefix from other border routers through IBGP. Therefore, there is an opportunity to select a router among a number of egress routers for inter-domain traffic flows with bandwidth guarantees. The outcome of the BGERs can be realized by BGP policies such as using policy routing and manipulating BGP attributes. For the ISP's decisions on advertising bandwidth to its upstream domains (e.g. how much bandwidth is advertised and where to advertise) in order to extend its services, we consider this as the subject of inbound inter-domain TE, which is out scope of this paper.

We assume that the traffic matrix of customer flows is known through customer SLAs established in a management time-scale. Each customer traffic flow includes bandwidth requirement to a destination prefix and the ingress router where it enters the ISP domain. Moreover, individual customer traffic flows are aggregated at each ingress router according to their destination prefixes. In the rest of this paper, we refer the customer traffic flow as the one which is aggregated based on ingress router and destination prefix. The problem we address thus becomes: for each given customer traffic flow at each ingress router, select an egress router that satisfies the customer bandwidth requirement while minimizing total bandwidth consumption in the ISP's network.

B. Problem Formulation

We formulate the BGERs problem as an integer-programming problem. In table 1, we summarize the notation and the definitions used in the rest of this paper. Two objectives are addressed: satisfying the customer bandwidth requirement and minimizing the total bandwidth consumption.

The fundamental objective is to provide bandwidth guarantees to customer traffic flows by satisfying their bandwidth requirements. We define the following constraints to determine whether the bandwidth requirement of customer traffic flow $t(i, k)$ is satisfied:

1. There exists a feasible path p from the ingress router $i \in I$ to

TABLE 1. NOTATION USED IN THIS PAPER

Notation	Description
E	A set of intra-domain links
K	A set of destination prefixes
I	A set of ingress routers
J	A set of egress routers
$t(i, k)$	The aggregated bandwidth requirement of customer traffic flows destined to destination prefix $k \in K$ at ingress router $i \in I$
$Out(k)$	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 domains) that is connected to egress router $j \in J$
$f_i(j, n)$	True (1) / False (0); whether the prefix k can be reached through the inter-domain link between the egress router j and the next hop address $n \in NEXT_j$
C_{intra}^i	The capacity of intra-domain link $i \in E$
bw_{intra}^i	The current availability (unallocated bandwidth) on C_{intra}^i
$C_{inter}^{j, n}$	The capacity of the inter-domain link which is attached to egress router j and is connected to next-hop address n
$bw_{inter}^{j, n}$	The current availability (unallocated bandwidth) on $C_{inter}^{j, n}$
$p(k, j)$	The bandwidth advertised by BGP-TE on the egress router j to the destination prefix k after the BGP path selection
$bp(k, j)$	The current availability (unallocated bandwidth) on $p(k, j)$
$X_{t(i, k)}^j$	True (1) / False (0); whether the customer traffic flow $t(i, k)$ has been selected the egress router j
$Y_{t(i, k)}^i$	True (1) / False (0); whether the customer traffic flow $t(i, k)$ has consumed bandwidth on the intra-domain link i
$d(i, j, k)$	Number of hops of the feasible shortest path (found by CSPF) between the ingress router i and the egress router j for $t(i, k)$
$bw_{intra}^{p, i, j}$	The bottleneck bandwidth of intra-domain path p between the ingress router i and the egress router j , i.e. $\min_{i \in p} (bw_{intra}^i)$

the selected egress router $j \in J$ such that $bw_{intra}^{p, i, j} \geq t(i, k)$ (1)

2. $bp(k, j) \geq t(i, k)$ (2)

3. $bw_{inter}^{j, n} \geq t(i, k)$ where $n = FindInterdomainLink(j, k)$ (3)

The function $FindInterdomainLink(j, k)$ identifies a specific inter-domain link by giving the next hop address (i.e. the address of border router in the downstream domain) that can reach the destination prefix k through the egress router j . Constraint (1) is the intra-domain bandwidth constraint, which ensures that there exists a feasible path between the ingress router and the selected egress router, and the bottleneck bandwidth of the path is no less than the bandwidth requirement of the customer traffic flow. This path can be found by CSPF. Constraint (2) is the constraint of advertised bandwidth information at the selected egress router, which ensures that the advertised bandwidth is sufficient for the customer traffic flow. This implies that in each domain along the corresponding BGP path towards the destination there exists sufficient bandwidth for the customer traffic flow. Constraint (3) ensures that the inter-domain link at the selected egress router, which connects to the downstream domain reaching the destination prefix, has sufficient bandwidth for the customer traffic flow. If all three constraints are met, the customer bandwidth requirement is satisfied.

The objective of minimizing total bandwidth consumption translates to the problem of minimizing the total number of hops that a traffic flow must traverse in the network, i.e.

$$\text{Minimize } \sum_{k \in K} \sum_{i \in I} \sum_{j \in \text{Out}(k)} x_{(i,k)}^j \cdot d(i,j,k) \cdot t(i,k) \quad (4)$$

subject to:

$$\sum_{k \in K} \sum_{i \in I} x_{(i,k)}^j \cdot t(i,k) \cdot f_k(j,n) \leq c_{intra}^{j,n} \quad \forall (j,n) \text{ where } j \in J \text{ \& } n \in \text{NEXT}_j \quad (5)$$

$$\sum_{k \in K} \sum_{i \in I} y_{(i,k)}^j \cdot t(i,k) \leq c_{intra}^j \quad \forall j \in J \quad (6)$$

$$\sum_{i \in I} x_{(i,k)}^j \cdot t(i,k) \leq p(k,j) \quad \forall (k,j) \text{ where } k \in K, j \in J \quad (7)$$

$$x_{(i,k)}^j, y_{(i,k)}^j \in \{0,1\} \quad (8)$$

$$\sum_{j \in \text{Out}(k)} x_{(i,k)}^j = 1 \quad \forall (i,k) \text{ where } i \in I, k \in K \quad (9)$$

Constraint (5) is the capacity constraint for each inter-domain link; constraint (6) is the capacity constraint for each intra-domain link; constraint (7) is the capacity constraint for each advertised bandwidth towards the destination prefix; constraint (8) ensures the discrete variables to assume binary values; constraint (9) ensures that only one egress router is selected for each customer traffic flow.

The work in [6] has formulated the general egress router selection problem as a Generalized Assignment Problem (GAP) and proved that the problem is NP-complete. Due to the reason that our BGRS extends that of [6] and has the additional constraints (6) and (7) as an intra-domain and a cascaded inter-domain capacity constraint respectively, we consider it to be a variant of GAP, which is also NP-complete. Hence, we propose heuristic algorithms to solve it.

C. Heuristic Algorithms

We present three greedy-based heuristic algorithms, which use available information in different ways in order to achieve minimal total bandwidth consumption.

1) *Greedy-cost heuristic*: This sorts customer traffic flows in descending order based on their bandwidth requirements and selects one at a time in that order. The sorting aims to minimize the bandwidth consumption by initially assigning large traffic flows to the closest egress routers using the shortest paths.

Step 1 We evaluate each of the egress routers in the network individually to determine its feasibility for the customer traffic flow. We refer to this step as pre-selection. Pre-selection is based on the information such as bandwidth information advertised by BGP-TE, prefix reachability and intra- and inter-domain link available bandwidth. The egress router j is feasible if it satisfies both the following constraints:

$$1. j \in \text{Out}(k). \quad (10)$$

$$2. \text{The customer bandwidth requirement; i.e. the constraint (1)-(3)} \quad (11)$$

Constraint (10) ensures that by selecting the egress router j , the destination prefix k can be reached. Constraint (11) ensures that the bandwidth is adequate when selecting egress router j to satisfy the customer bandwidth requirement. Thus, pre-selection ensures that the feasible egress routers are able to satisfy customer bandwidth requirements.

Step 2 Among a set of feasible egress routers identified in step 1, we select an egress router with the minimum number of hops that the flow must traverse on the path from the ingress router to it. If there exist several such egress routers, the selection would tiebreak on the maximum bottleneck bandwidth of the intra-domain path and the inter-domain link.

Step 3 Once the egress router is selected, the corresponding selected intra-domain path is pinned and bandwidth is allocated on the pinned path, the corresponding selected inter-domain link and the advertised bandwidth capability in order to provide a bandwidth guarantee for the assigned customer traffic flow. RSVP-TE [9] can be used for bandwidth reservation if LSP is established.

Step 4 We consider the next customer traffic flow and repeat step 1 to step 4. The heuristic finishes when all the customer traffic flows have been considered.

2) *Greedy-penalty heuristic*: It is possible that assigning a customer traffic flow to an egress router in different orders results in different selection scenarios. For example, if we assign the customer traffic flow $t(i,k) = 2$ in the first place, we can assign it greedily to egress router j with $d(i,j,k) = 3$ (the total bandwidth consumed equals to 6). If we delay allocating it for a while, however, egress router j may not have sufficient bandwidth because its bandwidth has been allocated to other customer traffic flows and the considered customer traffic flow has to be assigned to egress router j' with $d(i,j',k) = 6$ (the total bandwidth consumed equals to 12). In this case, we have a penalty on the consumption of additional bandwidth (i.e. $12 - 6 = 6$) and we use plt to refer to this penalty value. A penalty-based algorithm aims to minimize the number of hops a flow must traverse by placing customer traffic flows in certain order according to plt . We propose a similar algorithm called Greedy-penalty heuristic as follows. Such an algorithm is also used to solve the Generalized Assignment Problem [10].

Step 1 For each unassigned customer traffic flow, we measure the desirability of assigning it to each feasible egress router that satisfies the constraint (10) and (11). The desirability is the total bandwidth consumed by the flow along the path between the ingress and the egress router (i.e. the number of hops times the requested bandwidth). In this case, the smaller the desirability, the better for the selection.

Step 2 Compute plt for each unassigned customer traffic flow, which is the difference between the desirability of the customer traffic flow's best and second best selection (i.e. the two egress routers which yield the smallest desirability). If there is only one feasible egress router to accommodate the customer traffic flow, we need to assign the customer traffic flow to it. Otherwise, this currently feasible egress router may become unfeasible afterwards, having been assigned to accommodate other customer traffic flows, which leads to insufficient bandwidth so that we would reject the customer traffic flow. In this case, we set plt to infinite.

Step 3 Among all unassigned customer traffic flows, the one yielding the largest plt is placed with its best selection (with the tiebreak decision as in the Greedy-cost heuristic). If multiple customer traffic flows have the same largest penalty, they are placed in the order of decreasing bandwidth requirement.

Step 4 Once the egress router is selected, the corresponding selected intra-domain path is pinned and bandwidth is allocated on the pinned path, the corresponding selected inter-domain link and the advertised bandwidth capability in order to provide a bandwidth guarantee for the assigned customer traffic flow. We iterate step 1 to step 4 until all the customer traffic flows have been considered.

3) *Greedy-random heuristic*: As with the Greedy-cost heuristic, this sorts the customer traffic flows in descending order based on their bandwidth requirements and selects one at a time in that order. It is identical to Greedy-cost heuristic except that the selection in step 2 is done at random, with uniform probability among all the feasible egress routers in the network. We consider this algorithm as the behavior of the current BGP for solving the BGERS problem. The current non-TE BGP will select an egress router with respect to bandwidth information completely at random.

IV. PERFORMANCE EVALUATION

A. Configuration

We evaluate the three proposed heuristic algorithms through simulation. The simulation results are based on 100-node transit domain topologies. The topologies are randomly generated by the method described by Waxman [11]. The set of ingress and egress routers are disjoint. We set the number of ingress routers to 30, whereas the number of egress routers is a variable, as we will evaluate some effects by changing its value between 10 and 30. Each egress router is attached to a maximum of two inter-domain links. We assume that the inter-domain resource is less than that of intra-domain resource. The capacity of each link within a domain is randomly generated between 400 and 500, and the capacity of each inter-domain link is randomly generated between 250 and 300.

Feamster [12] discovered that a typical default-free routing table may contain routes for more than 90,000 prefixes, but only a small fraction of prefixes are responsible for a large fraction of the traffic. Based on this finding, we consider 1000 routing prefixes. As these routing prefixes are usually popular destinations, we assume that each egress router can reach all of them. This set of routing prefixes is randomly distributed on the inter-domain link(s) of each egress router. Each routing prefix is advertised with available bandwidth randomly generated between 200 and 250.

For each customer traffic flow, the destination prefix and the ingress router are randomly generated and its bandwidth requirement is randomly generated between 10 and 40.

B. Performance Evaluation

Figure 3 presents the total bandwidth consumption as a function of the number of customer traffic flows under the three proposed greedy-based heuristic algorithms. This simulation is based on the scenario of 30 egress routers. The Greedy-penalty heuristic consumes less bandwidth than the others because it considers the penalties of all unassigned customer traffic flows and determines which of these flows, if assigned in the first place, can avoid consuming additional bandwidth. On the contrary, the Greedy-cost heuristic does not

take this into consideration and often results in a greater penalty in terms of consuming more bandwidth. As the Greedy-random heuristic randomly selects an egress router without considering any optimization, any efficient egress router selection algorithms should always outperform it.




Figure 3. Total bandwidth consumption




Figure 4. Bandwidth consumption difference between Greedy-cost and Greedy-penalty heuristic

In Figure 4, we show the difference of bandwidth consumption between the Greedy-cost and Greedy-penalty heuristics for a different number of egress routers. We study the bandwidth consumption difference under three traffic loads with 100% acceptance ratio at any considered number of egress routers: 50, 100 and 150 customer traffic flows. The bandwidth consumption difference is the total bandwidth consumption using the Greedy-cost heuristic minus the total bandwidth consumption using the Greedy-penalty heuristic. It is worthwhile to determine the improvement of bandwidth consumption when using the Greedy-penalty heuristic over the Greedy-cost heuristic.

When the number of traffic flows increases, the bandwidth consumption difference between the two heuristic algorithms increases. This can be explained by the case that, as traffic load to the egress routers increases, some egress routers do not have sufficient resource so that some customer traffic flows are directed to the "distance" egress router with possible great penalty in terms of consuming more bandwidth. It is the case where Greedy-penalty heuristic is used to avoid additional bandwidth consumption.

Something else that can be deduced from the figure is that as the number of egress routers increases, the bandwidth consumption difference decreases. This is the opposite effect to the previous one, with the aforementioned case occurs less

frequently as more capacity is added. As a result, the two heuristic algorithms are likely to have same selection for traffic flows and the performance of them tends to become identical.

From the above, we conclude that the Greedy-penalty heuristic provides significant performance improvement over the Greedy-cost one, under the situation where the network has a certain level of loading in order to take the advantage of penalty-based selection, and that no more than one egress router can preferentially accommodate most of the traffic flows while leaving the other egress routers barely selected. The latter situation is achievable due to the fact that resources are commonly distributed in the network for load balancing.




Figure 5. Bandwidth acceptance ratio for Greedy-penalty heuristic

For the rest of simulations, we continue to study the performance as the number of egress routers varies. As the Greedy-penalty heuristic outperforms the others, we only consider this one. Figure 5 shows the influence of the number of egress routers on the bandwidth acceptance ratio. The bandwidth acceptance ratio is the sum of bandwidths of accepted traffic flows over the sum of bandwidths of all the traffic flows. As the number of egress routers increases, the bandwidth acceptance ratio increases. This is due to the property that performance improves as more capacity, such as inter-domain link and advertised bandwidth capacity, is added by increasing the number of egress routers. It is also worthwhile to determine when the bandwidth acceptance ratio reaches a level of diminishing return.




Figure 6. Total bandwidth consumption vs. Number of egress routers for Greedy-penalty heuristic

To evaluate the influence of the number of egress routers on the total bandwidth consumption, we study the bandwidth consumption under three traffic loads as they were previously used: 50, 100 and 150 customer traffic flows. Figure 6 shows

the total network bandwidth consumption with a different number of egress routers. For all the traffic flows, as the number of egress routers increases, the total bandwidth consumption decreases. This is because, as the number of egress routers increases, the traffic flow can be directed to a "closer" router which results in reduced bandwidth consumption. This effect becomes more apparent when the number of traffic flows is large since the traffic load of each egress router is high, while adding additional egress routers can significantly improve the performance. On the contrary, this effect is less apparent when the number of traffic flows is small.

V. CONCLUSION

In this paper, we present the bandwidth guaranteed egress router selection problem and potential solutions in the context of a TE-enabled Internet architecture. The latter comprises traffic engineering extensions to the current intra-domain and inter-domain routing protocols. The objective is that, for each customer traffic flow, we select an egress router that satisfies the customer bandwidth requirement while at the same time we minimize the total bandwidth consumption in the network. We have developed three heuristic algorithms to solve the BGERs problem. Simulation results show that the Greedy-penalty performs better than the other two algorithms in terms of total network bandwidth consumption. We have also evaluated the influence of the number of egress routers on the total bandwidth consumption and bandwidth acceptance ratio. We found that the total bandwidth consumption decreases and the bandwidth acceptance ratio increases as the number of egress routers increases. As future work, we plan to extend the BGERs problem and solutions to accommodate other specific QoS metrics such as delay.

REFERENCES

- [1] B. Quoitin, S. Uhlig, C. Pelsser, L. Swinnen and O. Bonaventure, "Interdomain traffic engineering with BGP," *IEEE Communication Magazine*, May 2003.
- [2] G. Apostolopoulos, D. Williams, S. Kamat, R. Guerin, A. Orda and T. Przygienda, "QoS Routing Mechanisms and OSPF Extensions," RFC 2676, August 1999.
- [3] O. Bonaventure, "Using BGP to distribute flexible QoS information," Internet Draft <draft-bonaventure-bgp-qos-00.txt>, February 2001.
- [4] G. Cristallo and C. Jacquenet, "Providing Quality of Service Indication by the BGP-4 Protocol: the QoS_NLR attribute," Internet Draft <draft-jacquenet-qos-nlr-04.txt>, March 2002.
- [5] L. Xiao, K.S. Lui, J. Wang and K. Nahrstedt, "QoS Extension to BGP," in *Proceedings of IEEE ICNP* 2002.
- [6] T.C. Bressoud, R. Rastogi and M.A. Smith, "Optimal Configuration for BGP Route Selection," in *Proceedings of IEEE INFOCOM* 2003, San Francisco, March/April 2003.
- [7] D. Katz, K. Kompella and D. Yeung, "Traffic Engineering Extensions to OSPF Version 2," RFC 3630, September 2003.
- [8] E. Osborne and A. Simha, *Traffic Engineering with MPLS*, Cisco Press, July 2002.
- [9] D. Awduche, L. Berger, D.H. Gan, T. Li, V. Srinivasan and G. Swallow, "RSVP-TE: Extensions to RSVP for LSP Tunnels," RFC 3209, December 2001.
- [10] S. Martello and P. Toth, *Knapsack Problems: Algorithms and Computer Implementations*, John Wiley and Sons, New York, 1990.
- [11] B.M. Waxman, "Routing of multipoint connections," *IEEE Journal on Selected Areas in Communications*, 6(9) 1998, pp 1617-1622.
- [12] N. Feamster, J. Borkenhagen and J. Rexford, "Controlling the impact of BGP policy changes on IP traffic," AT&T Labs - Research, Technical Report HA173000-011106-02TM, 2001.