Traffic Engineering for Inter-domain Quality of Service

K.H. Ho, N. Wang, P. Trimintzios and G. Pavlou

Centre for Communication Systems Research, University of Surrey, UK

Abstract: We present a model for inter-domain traffic engineering with end-to-end quality of service support, focusing particularly on bandwidth guarantees. We address the problem of egress router selection, which involves the selection of the egress router for customer traffic towards the destination prefixes taking into account end-to-end bandwidth constraints. We also consider performance optimisation in order to minimise the bandwidth consumption and achieve load balancing. Due to the computational complexity of the problem, we propose two heuristics to solve it. The simulation results show that the heuristic that involves with coordination performs better than the one without coordination, in terms of average bandwidth consumption.

1. Introduction

The next generation networks are designed to accommodate a variety of traffic classes to provide diverse Quality of Service (QoS) guarantees. In that context, quality of service routing has become an indispensable step towards guarantees as it selects a path that meets the QoS requirements. Until recently, most researchers have been focusing on QoS routing at the intra-domain level, with only few attempts at the inter-domain problem. The next generation Internet must address end-to-end connectivity QoS guarantees. In order to provide QoS across the Internet, not only intra-domain routing protocols must support QoS, but also the inter-domain routing protocols should be able to convey QoS information such as available bandwidth, delay, etc. However, the current inter-domain routing protocol, the Border Gateway Protocol (BGP), does not convey any QoS information between domains. On the other hand, Traffic Engineering (TE) focuses on the performance optimisation of networks, in order to achieve efficiently utilised and load balanced network resources. Recall that intra-domain routing addresses an optimisation problem of finding a 'best' path between a source and a destination within a domain, whereas inter-domain routing addresses an optimisation problem of selecting egress routers toward destination prefixes. This paper aims to propose methods for optimising the egress router selection while meeting end-to-end QoS constraint, and more specifically bandwidth guarantees. Thus, the problem that this paper addresses becomes: for each customer traffic request, the network must select an egress router so that its bandwidth requirements are met while achieving some optimisation criterion such as minimising the network bandwidth consumption and improving load balancing. This paper's aim can be considered as inter-domain TE for end-to-end QoS support as traffic is strategically assigned to egress routers through which diverse inter-domain routes are established with the objectives of performance optimisation. Our work extends the ideas of [1] on egress router selection by supporting end-to-end bandwidth guarantees, which is very important for QoS provisioning.

2. Network model and problem formulation

Here we present our model for inter-domain TE with end-to-end QoS support. First, in order to support end-to-end QoS, the current intra-domain and inter-domain routing protocols must be QoS aware. We assume the existence of QOSPF [2] as the intra-domain QoS routing protocol, which disseminates bandwidth information. In order to make the current inter-domain routing protocol QoS aware, we consider extending BGP to convey QoS information between domains by introducing a new QoS attribute in the BGP UPDATE packet and we name this extension as QoS Extension to BGP (QBGP). The concept of QBGP was originally proposed in [3]. Our work assumes that the new QoS attribute enables BGP to convey bandwidth information to other domains. This bandwidth information is the maximum allowable bandwidth from one domain to another domain, which is guaranteed by a service level agreement (SLA) established between peering domains. This contracted bandwidth information is mapped to the new QoS attribute towards a specific destination prefix and through QBGP the attribute is advertised to a neighbouring domain only if a SLA has been established. The contracted bandwidth advertised may also be the result of cascading established SLAs between all domains along the path as recorded in the *AS_Path* attribute. The technical implication of signing the SLA is the binding of bandwidth capabilities between the involved parties. The domain that requests the SLA binds its bandwidth capability to the contracted bandwidth of the peering domain and uses it as the basis for agreeing new contracts with its customers. The operation of bandwidth capabilities binding is done by setting the upper limits of all the bandwidth capabilities of domains along the *AS_Path*. As a result, the advertisement of inter-domain bandwidth availability can be fulfilled by QBGP but the detailed implementation of this functionality is beyond the scope of this paper. Through cascading established SLAs between domains in the Internet, the bandwidth requirements of traffic across domains can be guaranteed. In addition, by employing QOSPF and QBGP to provide QoS within and between domain(s), end-to-end QoS support across the Internet can be achieved. Based on the aforementioned approach for supporting end-to-end QoS, we present the model for bandwidth guaranteed inter-domain traffic engineering. This model allows the Internet service provider (ISP) to flexibly select egress routers for customer traffic requests in order to allow network resource optimisation such as minimising bandwidth consumption and improving load balancing, with respect to network capacity constraints and end-to-end bandwidth constraints. Figure 1 illustrates a generic ISP network connectivity. For the ISP under consideration, we consider a set of border routers as well as a set of intra-domain and inter-domain links. An inter-domain link connects a border router of the ISP and a border router of the neighbouring domain. We assume that the ISP has established SLAs with its neighbouring domains for bandwidth guarantees. Through QBGP, each border router receives advertisements of destination prefixes associated with the contracted bandwidth from its neighbouring domains.

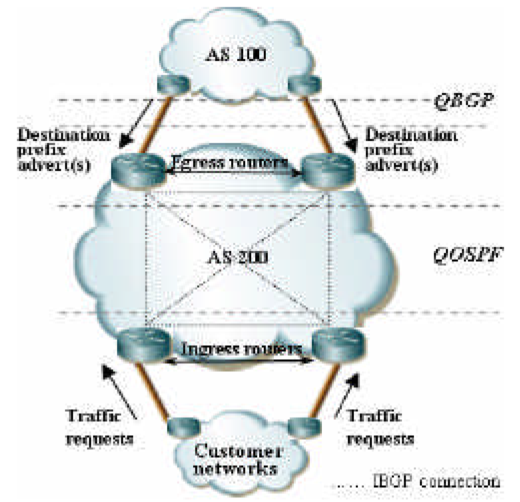

Figure 1. Generic network connectivity

These destination prefixes and the associated bandwidth information are then advertised to other border routers within the domain through fully meshed IBGP connections. As a result, all border routers have the same view on which border routers they can use to reach a specific destination prefix with an amount of guaranteed bandwidth. In figure 1, when an ingress router receives customer traffic requests with bandwidth requirements towards some remote destination prefixes, it looks up its BGP routing table and identifies a set of egress routers through which the destination prefixes can be reached. To select an egress router for each customer traffic request, the ingress router considers all those egress routers through which can reach the destination prefixes and meet customer bandwidth requests. This implies that the inter-domain routes towards the destination prefixes through the set of egress routers can support customer bandwidth requests. Moreover, not only the inter-domain route, but also the intra-domain path between the ingress and the egress router supports customer bandwidth demand. This is achieved by QOSPF with widest-shortest path selection algorithm [2]. Finally, among these eligible egress routers, which are able to meet the customer bandwidth requests, an egress router for each customer traffic request is selected, with the objectives of minimising network bandwidth consumption and load balancing the ISP network. The aforementioned egress router selection problem can then be formulated as an integer-programming problem with the following notations and definitions:

E – A set of intra-domain links H – A set of customers K – A set of destination prefixes

J – A set of egress routers c^{l}_{intra} – The capacity of intra-domain link $l \in E$

c^{j}_{inter} – The capacity of the inter-domain link attached to the egress router $j \in J$

$p(k,j)$ – The contracted bandwidth advertised on the egress router j to the destination prefix $k \in K$

$In(h)$ – A set of ingress routers from customer $h \in H$

$Out(k)$ – A set of egress routers that can reach the destination prefix k

$t(h,i,k)$ – Bandwidth demand from customer h through ingress router $i \in In(h)$ destined for destination prefix k

$d(i,j)$ – The number of hops between the ingress router i and the egress router j

$x^{j}_{(h,i,k)}$ – Indicate whether the request from customer h through ingress router i destined for destination prefix k has selected the egress router j.

$y^{l}_{(h,i,k)}$ – Indicate whether the request from customer h through ingress router i destined for destination prefix k has consumed bandwidth on the intra-domain link l

Our main objective of minimising overall network bandwidth consumption can be formulated as

$$Min \sum_{k\in K}\sum_{h\in H}\sum_{i\in In(h)}\sum_{j\in Out(k)} x^{j}_{(h,i,k)} \cdot d(i,j) \cdot t(h,i,k) \qquad (1)$$

We also consider minimising the maximum link utilisation as our second objective. These objectives are subject to the following constraints:

$$\sum_{k\in K}\sum_{h\in H}\sum_{i\in In(h)} x^{j}_{(h,i,k)} \cdot t(h,i,k) \le c^{j}_{inter} \quad \forall j \in J \qquad (2)$$

$$\sum_{k\in K}\sum_{h\in H}\sum_{i\in In(h)} y^{l}_{(h,i,k)} \cdot t(h,i,k) \le c^{l}_{intra} \quad \forall l \in E \qquad (3)$$

$$\sum_{h\in H}\sum_{i\in In(h)} x^{j}_{(h,i,k)} \cdot t(h,i,k) \le p(k,j) \quad \forall (k,j) \text{ where } k\in K, j \in J \qquad (4)$$

$$x^{j}_{(h,i,k)}, y^{l}_{(h,i,k)} \in \{0,1\} \qquad (5)$$

$$\sum_{j\in Out(k)} x^{j}_{(h,i,k)} = 1 \quad \forall (h,i,k) \text{ where } h\in H, i\in In(h), k\in K \qquad (6)$$

Constraint (2) is the capacity constraint for the inter-domain links; constraint (3) is the capacity constraint for the intra-domain links; constraint (4) is the capacity constraint for the advertised contracted bandwidth associated with destination prefixes; constraint (5) ensures the discrete variables to assume binary values; constraint (6) ensures that only one egress router is selected for each customer traffic request. Compared to [1], we have additionally considered constraint (3) and (4) as an intra-domain and a cascaded inter-domain capacity constraint respectively. The egress router selection problem is NP-Hard. Hence we propose two heuristics to solve it.

3. Heuristics for egress router selection problem

We propose two heuristics for the egress router selection problem. Both heuristics are based on a simple greedy method that assigns customer traffic requests to egress routers one at a time, selecting an egress router so that the customer bandwidth request is met and the network bandwidth consumption is minimised. The two heuristics are similar but they differ in the order in which an individual customer traffic request is considered. The first heuristic, Maximum Bandwidth Demand (MBD), is shown in figure 2a, while the second heuristic, Maximum Desirability Difference (MDD), is shown in figure 2b. The notations used follow the ones defined in the previous section.

Set $R = \{t(h,i,k)\}$, a set of customer traffic requests
Let b^{l}_{intra}, b^{j}_{inter}, $bp(k,j)$ be the available bandwidth of C^{l}_{intra}, C^{j}_{inter}, $p(k,j)$
Sort R into descending order according to requested bandwidth
While $R \neq \emptyset$ **do**
  Let $t(h,i,k)$ be the first element of R
  $min\_hops = \infty$
  $max\_bw = -1$
  $max\_egbw = -1$
  Select a set of egress routers,
    $EgSet = \{\forall j \in J \mid j \in Out(k),\ b^{j}_{inter} \geq t(h,i,k),\ p(k,j) \geq t(h,i,k)\}$
  While $EgSet\ != \emptyset$ **do**
    Let j be the first element of $EgSet$
    Compute a widest-shortest path, $path(h,i,k,j)$, between ingress router i and egress router j for $t(h,i,k)$
    Set $b(i,j)$ = MaxBottleneckBandwidth$(path(h,i,k,j))$
    If $b(i,j) \geq t(h,i,k)$ // intra-domain bandwidth constraint
      Set $d(i,j)$ = NumberofHops$(path(h,i,k))$
      If $d(i,j) < min\_hops \mid\mid d(i,j) == min\_hops\ \&\ b(i,j) > max\_bw \mid\mid d(i,j) == min\_hops\ \&\ b(i,j) == max\_bw\ \&\ b^{j}_{inter} > max\_egbw$
        $SelectedIntraPath = path(h,i,k)$
        $SelectedEgress = j$
        $min\_hops = d(i,j)$
        $max\_bw = b(i,j)$
        $max\_egbw = b^{j}_{inter}$
    $EgSet = EgSet / \{j\}$
  If $SelectedEgress \neq \emptyset$
    $b^{SelectedEgress}_{inter} = b^{SelectedEgress}_{inter} - t(h,i,k)$
    $\forall l \in SelectedIntraPath,\ b^{l}_{intra} = b^{l}_{intra} - t(h,i,k)$
    $bp(k,SelectedEgress) = bp(k,SelectedEgress) - t(h,i,k)$
    $\forall j \in Out(k) / \{SelectedEgress\},\ X^{SelectedEgress}_{(h,i,k)} = 0$
    $X^{SelectedEgress}_{(h,i,k)} = 1$
  else /* *leave the request without selecting any egress routers* */
    $\forall j \in Out(k),\ X^{SelectedEgress}_{(h,i,k)} = 0$
  $R = R / \{t(h,i,k)\}$

Figure 2a. Pseudocode of MBD

Set R, $R' = \{t(h,i,k)\}$, a set of customer traffic requests
Let b^{l}_{intra}, b^{j}_{inter}, $bp(k,j)$ be the available bandwidth of C^{l}_{intra}, C^{j}_{inter}, $p(k,j)$
While $R \neq \emptyset$ **do**
  $d^* = -\infty$
  While $R' \neq \emptyset$ **do**
    Let $t(h,i,k)$ be the first element of R'
    $CandidateEg = \emptyset$
    Select a set of egress routers,
      $EgSet = \{\forall j \in J \mid j \in Out(k),\ b^{j}_{inter} \geq t(h,i,k),\ p(k,j) \geq t(h,i,k)\}$
    While $EgSet \neq \emptyset$ **do**
      Let j be the first element of $EgSet$
      Compute a widest-shortest path, $path(h,i,k,j)$, between ingress router i and egress router j for $t(h,i,k)$
      Set $b(i,j)$ = MaxBottleneckBandwidth$(path(h,i,k,j))$
      If $b(i,j) \geq t(h,i,k)$ /* intra-domain bandwidth constraint */
        Set $d(i,j)$ = NumberofHops$(path(h,i,k,j))$
        $CandidateEg = CandidateEg \cup \{j\}$
        $w^{j}_{(h,i,k)} = d(i,j) \times t(h,i,k)$ /* desirability calculation */
      $EgSet = EgSet / \{j\}$
    If $size(CandidateEg) == 1$
      $d = -1$
    else /* calculate desirability difference = the difference between the smallest and the second smallest desirability of the request */
      $\forall j \in CandidateEg$, d = SecondMin$(w^{j}_{(h,i,k)})$ – Min$(w^{j}_{(h,i,k)})$
    If $d > d^*$
      $d^* = d$
      $SelectedReq = t(h,i,k)$
      $SelectedEgress = j$, Min$(w^{j}_{(h,i,k)} : \forall j \in CandidateEg)$
      $SelectedIntraPath = path(h,i,k,j)$
    $R' = R' / \{t(h,i,k)\}$
  If $SelectedReq \neq \emptyset$
    $b^{SelectedEgress}_{inter} = b^{SelectedEgress}_{inter} - SelectedReq$
    $\forall l \in SelectedIntraPath,\ b^{l}_{intra} = b^{l}_{intra} - SelectedReq$
    $bp(k,SelectedEgress) = bp(k,SelectedEgress) - SelectedReq$
    $\forall j \in Out(k) / \{SelectedEgress\},\ X^{SelectedEgress}_{(h,i,k)} = 0$
    $X^{SelectedEgress}_{(h,i,k)} = 1$
  else
    $\forall j \in Out(k),\ X^{SelectedEgress}_{(h,i,k)} = 0$
  $R = R / \{SelectedReq\}$

Figure 2b. Pseudocode of MDD

The first heuristic processes each request independently without coordination while the second heuristic, Maximum Desirability Difference (MDD), employs coordination and is based on the work of the Martello and

Toth [4]. It measures the desirability of assigning a customer traffic request to each egress router, denoted by $w^{j}_{(h,i,k)}$. We use the total amount of bandwidth consumed by the path between the ingress and the egress router, which is defined by the number of hops times the requested bandwidth, as the measurement of desirability. It iteratively considers all the unassigned customer traffic requests and determines the difference between the smallest and the second smallest desirability of the customer traffic request. The customer traffic request that has the maximum desirability difference is preferably considered over the others. If the customer traffic request has only one feasible egress router, it is given a higher priority. If several customer traffic requests have same maximum desirability difference, the selection would tiebreak on the maximum available bandwidth of intra-domain path and the maximum available bandwidth of inter-domain link associated with the egress router.

4. Simulation

In this section, we compare the performance of the two heuristics in terms of their average bandwidth consumption in the ISP network. The average bandwidth consumption is defined as the ratio of the total bandwidth consumed in the network over the total accepted bandwidth for all the customer traffic requests. We include the performance of a heuristic that randomly selects a candidate egress router. This method is used as a benchmark for comparison with the two proposed heuristics. For a given scenario, we consider a network of 100 nodes. The capacity of each intra-domain and inter-domain link is set to 500. We set both the number of egress and ingress routers to 30 and the number of prefixes to 1000. Each prefix has an average bandwidth request profile that is randomly generated between 1 and 20. We assume that customer traffic requests are aggregated by their destination prefixes. The destination prefix of each request is randomly generated and the requested bandwidth is, thus, set equal to the average bandwidth request profile of the destination prefix. Figure 3 shows the average bandwidth consumption as a function of the number of customer traffic requests. The MDD performs better than MBD, and they both outperform the Random heuristic in terms of average bandwidth consumption. The reason is that in MDD we give high priority to those customer traffic requests which have the greatest opportunity cost in terms of the difference of bandwidth consumption for the best and the second best egress router that have been chosen. Thus, the heuristic tries to minimise the possibility of the request being selected by the second best egress router, which consumes more network bandwidth. This reflects our main objective of minimising overall network bandwidth consumption. The second objective of minimising the maximum utilisation is achieved by selecting the egress router for which the associated path from the ingress router has the maximum available bandwidth if there are several such egress routers with the same number of minimum hops.

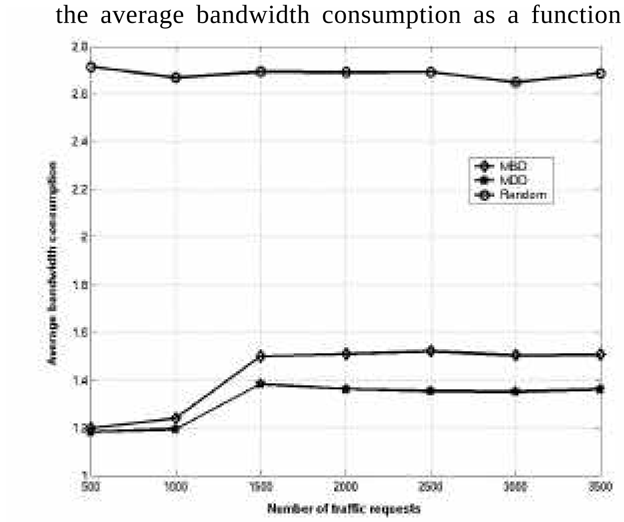

Figure 3. Performance comparisons

5. Conclusions

In this paper we presented an inter-domain traffic engineering model with end-to-end bandwidth support. The problem we addressed was the egress router selection: for each customer traffic request the ISP must carry, select an egress router for the request with the bandwidth requirement satisfied, while achieving some performance optimisation objectives such as minimising network bandwidth consumption and load balancing. We proposed two heuristics to solve the problem and compared their average bandwidth consumption.

References

[1] 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.

[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] G. Cristallo and C. Jacquenet, “An Approach to Inter-domain Traffic Engineering”, in *Proceedings of XVIII World Telecommunications Congress (WTC2002)*, Paris, France, September 2002.

[4] S. Martello and P. Toth, Knapsack Problems: Algorithms and Computer Implementations, John Wiley and Sons., New York, 1990.