We are IntechOpen, the world's leading publisher of Open Access books Built by scientists, for scientists

6,300

Open access books available

172,000

International authors and editors

190M

Downloads

Our authors are among the

154

Countries delivered to

TOP 1%

most cited scientists

12.2%

Contributors from top 500 universities

WEB OF SCIENCE™

Selection of our books indexed in the Book Citation Index
in Web of Science™ Core Collection (BKCI)

Interested in publishing with us?
Contact book.department@intechopen.com

Numbers displayed above are based on latest data collected.
For more information visit www.intechopen.com

Solving Inter-AS Bandwidth Guaranteed Provisioning Problems with Greedy Heuristics

Kin-Hon Ho¹, Ning Wang² and George Pavlou³

¹*Department of Computer Science, City University of Hong Kong,*

²*Centre for Communication Systems Research, University of Surrey,*

³*Department of Electronic and Electrical Engineering, University College London,*

¹*Hong Kong*

^{2,3}*United Kingdom*

1. Introduction

The current Internet consists of more than 26,000 Autonomous Systems (ASes) or domains, each being a network or group of networks managed by a single authority commonly known as Internet Network Provider (INP). With wide deployment of real-time multimedia applications in recent years, the emerging future-generation Internet is expected to provide end-to-end Quality of Service (QoS) guarantees across multiple ASes. In situations where stringent end-to-end QoS is required, ensuring that an adequate bandwidth is guaranteed by each AS along as the entire route in the Internet backbone is essential to achieve relevant performance targets (Zhang et al., 2004). Yet in practice, an AS is only capable of provisioning bandwidth guarantees within its own network. Hence, extending bandwidth guarantees beyond its boundary requires the AS to agree the supply of sufficient bandwidth from other ASes. This bandwidth supply is likely to be associated with a financial cost and therefore there is an economic incentive for an AS to carefully select its downstream provider ASes so as to minimize the cost of using that bandwidth.

Having purchased access to sufficient bandwidth from downstream ASes, the AS needs to utilize both this purchased bandwidth and its own network capacity in the most effective way in order to provide bandwidth guarantees for customer traffic. INPs thus need to optimize the utilization of these resources. Traffic Engineering (TE) is an effective technique to optimize IP operational network performance and subsequently improve network QoS capabilities (Awduche et al., 2002). INPs can thus use TE as an effective means for bandwidth guarantee provisioning while optimizing network resource utilization.

Concatenation of bandwidth guarantees between ASes makes it possible to provide an end-to-end guarantee between a source-destination pair in the Internet. These guarantees across ASes owned by different INPs require some level of agreement between themselves, usually summarized in a negotiated Service Level Agreement (SLA) at the AS level. An SLA is an agreement contracted between a customer AS and a provider AS that describes the characteristics of a service, specifying in particular the supported QoS and the associated cost. However, given that the Internet consists of thousands of ASes, SLA negotiation between ASes has to be carefully managed in an effective and scalable manner. In this chapter we adopt a cascaded negotiation model which allows ASes to build up end-to-end SLAs that provide end-to-end bandwidth guarantees. In this model, apart from route

Source: Advances in Greedy Algorithms, Book edited by: Witold Bednorz,
ISBN 978-953-7619-27-5, pp. 586, November 2008, I-Tech, Vienna, Austria

reachability information, each AS receives from adjacent downstream ASes a set of what we call bandwidth offers to designated remote AS destinations. If an AS decides to accept a bandwidth offer, an SLA is established between the two ASes. The AS can then in turn make bandwidth offers to its upstream (customer) ASes; these offers reflect both the AS' own resources and the SLAs established with the downstream ASes. The full set of SLAs enables all the ASes to support traffic with end-to-end bandwidth guarantees. However, the AS' tasks of making appropriate decisions on which bandwidth offers to accept, how much bandwidth to purchase and how to allocate the bandwidth among traffic aggregates are non-trivial. Inappropriate bandwidth offer selection or traffic assignment could result in respectively high cost or poor resource utilization. In order to obtain the best solutions, we propose a network dimensioning system that incorporates optimization modules that solve the two following problems:

- how to determine an appropriate amount of bandwidth to be purchased from each bandwidth offer so that the total cost of the bandwidth is minimized;
- given the knowledge of the available intra-AS bandwidth and the bandwidth purchased from downstream ASes, how to assign routes to the predicted traffic aggregates so that bandwidth demand is met while optimizing resource utilization.

We call these two problems the *Inter-AS Bandwidth Provisioning* and *Traffic Assignment* problems respectively. Our proposed network dimensioning system enables ASes to move from trial-and-error to a systematic approach for provisioning their end-to-end bandwidth guarantees. More specifically, we propose two efficient greedy heuristics to solve these optimization problems. It has been a long history that greedy heuristics are used for solving network optimization problems, such as traffic engineering (Sridharan et al., 2005), multicast routing (Shi & Turner, 2002) etc. Nevertheless, the optimization problems of end-to-end bandwidth guarantees provisioning across multiple ASes has not been addressed until recently, and in this chapter we will illustrate how greedy heuristics can gracefully solve these novel problems. The main contributions of this chapter can be summarized as follows:

- We propose a systematic network dimensioning system that can be used by ASes to achieve effective provisioning of end-to-end bandwidth guarantees. The network dimensioning system formulates two problems that respectively provide economic and engineering optimization, namely the inter-AS bandwidth provisioning and traffic assignment problems.
- We show that a heuristic approach can be used to solve the inter-AS bandwidth provisioning problem. To illustrate this, we use an efficient genetic algorithm embedded with two problem-specific greedy heuristics. Our proposed algorithm optimizes the bandwidth provisioning with 5%-30% and 75%-90% less cost than a conventional heuristic and a random-based algorithm respectively.
- We use a greedy-penalty heuristic algorithm to solve the traffic assignment problem. The proposed greedy-penalty heuristic results in 10% less total bandwidth consumption than a random-based algorithm.

2. Cascaded inter-AS negotiation model

The provision of end-to-end bandwidth guarantees requires each intermediate AS on the path from the source AS to the destination AS to guarantee the agreed bandwidth. However, this cannot be realized without first negotiating and agreeing SLAs among the ASes. Since the Internet is a collection of a large number of ASes, attention needs to be paid to how to manage such negotiation and SLA establishment in an effective and scalable manner. In this chapter,

we adopt a cascaded model, as proposed by the MESCAL project for negotiating QoS guarantees (e.g. bandwidth and delay) among ASes (Howarth et al., 2005).

The model is based on two concepts: (1) negotiation of bandwidth offers between ASes; (2) establishment of unidirectional SLAs between ASes for the agreed bandwidth. The key idea of the cascaded model is as follows. An AS offers bandwidth guarantees to its upstream ASes; each bandwidth offer specifies the reachable remote destination(s), the available bandwidth (e.g., maximum offered bandwidth) and a cost, for example, per unit of bandwidth. These destinations are either in customer ASes or reachable through downstream ASes. An upstream AS in general receives multiple bandwidth offers for any given destination, and has to decide which one to accept. Each accepted bandwidth offer is then established as a unidirectional SLA. The AS can then in turn make bandwidth offers to its upstream ASes, by combining its local bandwidth capabilities with the SLA. This process continues in a cascaded manner for further upstream ASes, and an end-to-end SLA chain can be built, with each SLA relying on the SLAs between downstream ASes.

Fig. 1. Illustration of the cascaded inter-AS negotiation model

Fig. 1 illustrates an example. Let $o\text{-BW1}$ be the bandwidth guarantee offered by AS1 towards destination 'dest'. AS2 receives this offer $o\text{-BW1}$. We assume that AS2 decides to accept the bandwidth offer: AS2 then establishes an SLA with AS1 ($SLA2\text{-}1$) for this bandwidth. Now AS2 has a bandwidth guarantee provided by AS1 for access to 'dest'. AS2 can in turn extend this bandwidth guarantee by concatenating its local bandwidth capability with $SLA2\text{-}1$, and then offering a bandwidth ($o\text{-BW2}$) to AS3. $o\text{-BW2}$ is the minimum of (a) the local bandwidth capability that AS2 is prepared to guarantee across its network and (b) $SLA2\text{-}1$. Now $o\text{-BW2}$ indicates the bandwidth guarantee from AS2 to destination 'dest'. AS3 receives $o\text{-BW2}$ from AS2 and it in turn repeats the decision process, possibly purchasing the offered bandwidth and establishing $SLA3\text{-}2$. In summary, once offers from other adjacent downstream ASes have been agreed as SLAs, an INP may build new extended services upon cascaded existing ones. Further details of the cascaded model can be found in (Howarth et al., 2005).

The cascaded model has several advantages: (1) it makes possible to build scalable end-to-end QoS guarantees between any two ASes while only maintaining SLAs with adjacent ASes; (2) it has backward compatibility with BGP, making inter-AS QoS deployment possible through extensions to BGP; (3) it retains privacy for all ASes regarding the details of their interactions.

The decision on which bandwidth offers to accept, and how to effectively utilize the established SLAs and the AS' intra-AS resources is non-trivial. In the next section, we propose a network dimensioning system, incorporating TE mechanisms, to solve this problem and make the best decisions.

3. Decomposition of the network dimensioning system

We consider two optimization problems, an economic and an engineering one, that need to be solved for provisioning end-to-end bandwidth guarantees. First, an AS needs to

determine the appropriate amount of bandwidth to be purchased from each adjacent downstream AS so that the total bandwidth cost is minimized. Second, given these available bandwidth resources defined in the SLAs and the local network’s bandwidth, the AS has to determine how to assign routes to the supported traffic in order to satisfy their bandwidth requirements while at the same time optimizing network resource utilization. We illustrate on Fig. 2 a decomposition of a network dimensioning system which consists of several components. We envisage this system as being offline and running infrequently as part of a resource provisioning cycle, e.g. in the order of weeks.

Fig. 2. Architecture of the network dimensioning system

3.1 Components of the network dimensioning system

The proposed network dimensioning system consists of the following components:

1. **Inter-AS Traffic Forecast** predicts inter-AS traffic in the network for a period of time and records this information in an inter-AS Traffic Matrix (TM). Each element in the inter-AS TM is the aggregate traffic load that enters the network at an ingress point[†] and is destined for a remote destination prefix. The TM entry is represented by the tuple

$$< ingress\ point, remote\ destination\ prefix, long-term\ average\ traffic\ demand >$$

Some known methods can be used to compute the traffic aggregate, such as the effective bandwidth approach (Guerin et al., 1991) if the mean and peak rates of the traffic are known.

The inter-AS TM is an important element for network and traffic engineering. Whilst an accurate inter-AS TM could be obtained through fine-grained flow-level traffic measurement this is not suitable for long term predictions (Awduche et al., 2002). Nevertheless, these problems have recently been addressed with a methodology that allows an inter-AS TM to be predicted through measurement (Teixeira et al., 2005) and estimation for web traffic (Feldmann et al., 2004). Alternatively, an inter-AS TM can be extrapolated from customer SLAs.

2. **Inter-AS Bandwidth Discovery** discovers bandwidth offers from adjacent downstream ASes through offline techniques, e.g. advertisement. A bandwidth offer is uniquely identified by a connection point at which the offer is provided. Bandwidth offers are provided by adjacent ASes, and so the connection point, or inter-AS link on which it is offered, uniquely identifies the adjacent AS. Each bandwidth offer specifies a maximum bandwidth towards a remote destination prefix and is associated with a cost, for example per unit of bandwidth. Each bandwidth offer is represented by the tuple

$$\langle \text{egress router, adjacent AS border router address, remote destination prefix,} \\ \text{maximum offered bandwidth, cost} \rangle$$

3. **Inter-AS Bandwidth Provisioning (IBP)** addresses the economic problem described in the beginning of this section. For the sake of service resilience and load balancing, an increasing number of ASes have multiple connections to adjacent downstream ASes. As a result, an AS may receive multiple offers to each destination prefix from different adjacent downstream ASes. The goal of IBP is to take as input the inter-AS TM and a set of bandwidth offers, and to produce as output a decision on which bandwidth offers to accept and the amount of bandwidth to be purchased from each of the accepted offers. Based on the IBP outcome, the AS will then establish SLAs (in this chapter called outbound provider SLAs) with the adjacent downstream ASes to contract the bandwidth guarantees. We assume that the establishment of outbound provider SLAs is performed by the component “provider SLA ordering”, a process whose details are outside the scope.
4. **Traffic Assignment (TA)** deals with the engineering problem described in the beginning of this section. The goal of TA is to take as input an inter-AS TM, a set of outbound provider SLAs that are established after the IBP phase, and the available bandwidth resources of the AS, i.e. intra- and inter-AS link capacities, and then to assign appropriate routes for the supported traffic so that the bandwidth requirements are met while optimizing network resource utilization. An assignment of the route includes selection of an outbound provider SLA, an inter-AS link and an intra-AS route for the supported traffic. The key output of the TA is a Traffic Assignment matrix that records the outbound provider SLAs, inter-AS links and intra-AS routes that have been selected for the supported traffic. Based on this matrix, an INP can implement the TA solution by configuring the network accordingly.

3.2 Inter-AS bandwidth overprovisioning

We can employ overprovisioning in the IBP phase. This implies that some network resources are left unused so as to protect the core backbone from failures and to accommodate some degree of traffic demand fluctuation (Nucci et al., 2005). Overprovisioning is also the current solution adopted by some INPs for QoS provisioning within their networks. For these reasons, we consider a certain amount of inter-AS bandwidth overprovisioning in this chapter. During the IBP phase, the AS should not merely purchase bandwidth that marginally accommodates the forecasted traffic demand, because the bandwidth guarantee may not be maintained if even a small traffic upsurge occurs. A solution to this is to purchase more bandwidth than the forecasted traffic demand in order to insure against such traffic fluctuations. This also provides a buffer against inter-AS link failures, which may cause traffic to be shifted from one outbound provider SLA to another.

The task of IBP is thus to decide an appropriate amount of bandwidth to be purchased from the adjacent downstream ASes by taking into account overprovisioning. To do so, we introduce an overprovisioning factor $f_{over} \geq 1.0$ to specify the degree of inter-AS bandwidth overprovisioning. In principle, this factor is determined by considering the network’s traffic characteristics and the target link utilization. However, since optimization of f_{over} is not the subject to be concerned, we assume that a single value is used to represent the optimal overprovisioning that has already been determined by the ASes. The concept of overprovisioning factor has also been used by other researchers, e.g. (Nucci et al., 2005). In this work, inter-AS bandwidth overprovisioning is implemented as follows. If $t(i,k)$ denotes the average demand of an inter-AS traffic flow aggregate, we define an inflated traffic flow, $\hat{t}(i,k) = t(i,k) \cdot f_{over}$.

4. Optimal inter-AS bandwidth provisioning

In this section and the next, we present the problem statement, formulation and algorithms of both the IBP and the TA problems.

Fig. 3. Elements of the inter-AS bandwidth provisioning

Fig. 3 illustrates an AS topology with the key elements of the IBP problem. A set of border routers is connected to adjacent ASes. An ingress (or egress) router is the border router that receives (or sends) traffic from (or to) an adjacent AS. Each border router is associated with one or more inter-AS links. Each bandwidth offer is associated with a single inter-AS link. Each border router may receive multiple bandwidth offers for a remote destination prefix from different adjacent downstream ASes through different attached inter-AS links, for example, the top left border router in Fig. 3. Each inter-AS traffic flow enters the AS through a designated ingress router. We define the total inter-AS bandwidth provisioning cost to be the total charge an AS pays for purchasing bandwidth from its adjacent downstream ASes. The inter-AS bandwidth provisioning problem can be summarized as follows:
Given a set of bandwidth offers from adjacent downstream ASes, an inter-AS traffic matrix and a physical network topology, determine an appropriate amount of bandwidth to be purchased from each bandwidth offer so that the total inter-AS bandwidth provisioning cost is minimized while respecting the capacity constraints of the inter-AS links.
In solving the IBP problem we assume that the inter-AS traffic is non-splittable. This method not only can determine the appropriate amount of bandwidth to be purchased but also ensures that each traffic flow will be accommodated by at least one SLA during TA without causing the traffic to be split.

Note that some types of ASes, such as tier 2 and 3, may have both peering and customer-provider connections with adjacent ASes. A peering connection between two ASes refers to the case where each AS carries a similar amount of customer traffic from the other AS for free. On the other hand, a customer-provider connection refers to the case where the provider charges the customer for carrying traffic across its network. The IBP description in Section 3 assumed that an AS has only customer-provider connections with its adjacent downstream ASes and that a cost is associated with each bandwidth offer. In fact, peering connections can also be considered by IBP. In this case, the cost of bandwidth is typically zero and the maximum bandwidth represents the agreed amount of traffic to be exchanged.

4.1 Inter-AS bandwidth provisioning problem formulation

We formulate IBP as an integer programming problem. Table 1 shows the notation used throughout this chapter. The objective of the IBP problem is to minimize the total IBP cost:

$$\text{Minimize } \sum_{i \in I} \sum_{k \in K} \sum_{oBw(k, j, n) \in Out(k)} x_{i,k}^{j,n} \cdot chg_k^{j,n} \cdot t'(i, k) \quad (1)$$

subject to:

$$\sum_{i \in I} \sum_{k \in K} x_{i,k}^{j,n} \cdot t'(i, k) \leq c_{inter}^{j,n} \quad \forall (j, n) \text{ where } j \in J, n \in NEXT_j \quad (2)$$

$$\sum_{i \in I} x_{i,k}^{j,n} \cdot t'(i, k) \leq MaxBw_k^{j,n} \quad \forall (k, j, n) \text{ where } k \in K, j \in J, n \in NEXT_j \quad (3)$$

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

$$x_{i,k}^{j,n} \leq y_k^{j,n} \quad \forall (i, k, j, n) \text{ where } i \in I, k \in K, j \in J, n \in NEXT_j \quad (5)$$

$$\sum_{oBw(k, j, n) \in Out(k)} x_{i,k}^{j,n} = 1 \quad \forall (i, k) \text{ where } i \in I, k \in K \quad (6)$$

$$\sum_{n \in NEXT_j} y_k^{j,n} \leq 1 \quad \forall (k, j) \text{ where } k \in K, j \in J \quad (7)$$

Constraint (2) ensures that no inter-AS link carries traffic exceeding its capacity. Constraint (3) ensures that no bandwidth offer carries traffic exceeding its maximum capacity. Constraint (4) ensures that the discrete variables assume binary values. Constraint (5) ensures that, whenever traffic flow $t'(i, k)$ is assigned to bandwidth offer $oBw_k^{j,n}$, then this bandwidth offer must have been selected. Constraint (6) ensures that only one bandwidth offer is selected for each inter-AS traffic flow. Hence, traffic splitting over multiple bandwidth offers is not considered. Constraint (7) ensures that only one of the bandwidth offers, which are advertised at a border router through different inter-AS links, is selected for each remote destination prefix. This constraint ensures the BGP rule that only one route toward a remote destination prefix is selected as the best route. This makes the IBP implementation easier through BGP configuration.

Notation	Description
<i>- General notation -</i>	
K	A set of destination prefixes
I	A set of ingress routers
J	A set of egress routers
f_{over}	Overprovisioning factor
$t(i,k)$	Bandwidth demand of an inter-AS traffic flow entering the AS at ingress router $i \in I$ towards destination prefix $k \in K$. It is considered by the TA problem
$\hat{t}(i,k)$	Inflated traffic flow $t(i,k)$. It is considered by the IBP problem
$Out(k)$	A set of bandwidth offers that has reachability to destination prefix k
$NEXT_j$	A set of next hop addresses (addresses of the border routers in adjacent downstream ASes) that is associated with egress router $j \in J$
$C_{inter}^{j,n}$	Capacity of the inter-AS link that connects egress router j to next-hop address $n \in NEXT_j$
$bw_{inter}^{j,n}$	Residual bandwidth of $C_{inter}^{j,n}$
C_{intra}^l	Capacity of intra-AS link l
bw_{intra}^l	Residual bandwidth of C_{intra}^l
<i>- Notation used in the IBP problem -</i>	
$oBw_k^{j,n}$	Bandwidth offer that is associated with destination prefix k and is advertised through the inter-AS link that connects egress router j to next-hop address n
$MaxBw_k^{j,n}$	Maximum bandwidth of the offer $oBw_k^{j,n}$
$Chg_k^{j,n}$	A charge per unit bandwidth for $oBw_k^{j,n}$
$x_{i,k}^{j,n}$	Variable indicating whether traffic flow $t'(i,k)$ is assigned to bandwidth offer $oBw_k^{j,n}$
$y_k^{j,n}$	Variable indicating whether the bandwidth offer $oBw_k^{j,n}$ is selected
<i>- Notation used in the TA problem -</i>	
$pSLA_k^{j,n}$	Outbound provider SLA of the bandwidth offer $oBw_k^{j,n}$
$pSLAC_k^{j,n}$	Contracted bandwidth specified in outbound provider SLA $pSLA_k^{j,n}$
$pSLABw_k^{j,n}$	Residual bandwidth of $pSLAC_k^{j,n}$
$dist_{i,j}^k$	Number of hops on the intra-AS route between ingress router i and egress router j towards destination prefix k
$P_{i,j}$	A set of feasible intra-AS routes between ingress router i and the egress router j to which the selected outbound provider SLA is associated.
$w_{i,k}^p$	Variable indicating whether path $p \in P_{i,j}$ is chosen to realize traffic flow $t(i,k)$
$z_{i,k}^{j,n}$	Variable indicating whether traffic flow $t(i,k)$ is assigned to outbound provider SLA $pSLA_k^{j,n}$
$\Upsilon_{i,k}^l$	Variable indicating whether traffic flow $t(i,k)$ is assigned to intra-AS link l

4.2 Modified inter-AS bandwidth provisioning problem

We assume that when multiple bandwidth offers towards the same remote destination prefix k are present at a given border router j (i.e. $\exists n \text{ } MaxBw_k^{j,n} > 0$), the AS has already determined the best one as a candidate bandwidth offer. Thus, each border router will consider at most one bandwidth offer towards each remote destination. The decision of selecting the best bandwidth offer might be based on business factors such as the relationships between ASes and the reputations of adjacent downstream ASes. As a result of this assumption, the variable $y_k^{j,n}$ of which bandwidth offer has been considered for each remote destination prefix k at each border router j is pre-determined and this satisfies constraint (7) since at most one bandwidth offer will be considered (i.e. $\sum_{n \in NEXT_j} y_k^{j,n} \leq 1$).

Therefore, constraint (7) is automatically enforced.

4.3 A Lower bound of the inter-AS bandwidth provisioning problem

We derive an approximated optimal solution of the IBP problem that can be obtained efficiently by relaxing some constraints. This approximated optimal solution is thus a lower bound of the IBP problem. A lower bound typically has better result than the optimal solution because some problem constraints are relaxed. However, due to the relaxation, it is not a valid solution to the problem. Nevertheless, the lower bound is a good approximation of an optimal solution for heuristic algorithms to compare their performance. We show the derivation of a lower bound for the IBP problem as follows.

We derive a lower bound by relaxing some IBP problem constraints. First of all, constraint (7) is automatically enforced by our assumption that each border router has only considered the best candidate bandwidth offer towards each remote destination prefix. Second, we relax the non-bifurcation integer constraint (4). In many practical situations, integer programming problems, which require all variables to be integers, are NP-hard. Instead, a linear programming problem that has only non-integer variables can be generally solved efficiently in the worst case. Therefore, we relax constraint (4) to

$$0 \leq x_{i,k}^{j,n} \leq 1, \text{ non-integer} \quad (8)$$

Finally, we find that a lower bound can be readily calculated by the following method if inter-AS link capacity constraint (2) is relaxed. Relaxation of a capacity constraint means that the constraint is simply ignored based on the assumption that capacity is large enough to accommodate the traffic.

Given $Pr_{low} = \min_{\forall k,j,n} Chg_k^{j,n}$ and $Pr_{high} = \max_{\forall k,j,n} Chg_k^{j,n}$, we can define

$$b_k^\psi = \sum_{j \in J} \sum_{n \in Next_j} MaxBw_k^{j,n} \mid Chg_k^{j,n} = \psi \quad \forall Pr_{low} \leq \psi \leq Pr_{high} \text{ and } k \in K \quad (9)$$

where $b_k^\psi \geq 0$ is the sum of maximum capacity of all the bandwidth offers to remote destination prefix k with a charge equal to ψ , and

$$d_k = \sum_{i \in I} t'(i, k) \quad \forall k \in K \quad (10)$$

where $d_k \geq 0$ is the sum of bandwidth demands of all the traffic flows to destination prefix k . For each traffic demand d_k towards remote destination prefix k , we first attempt to assign it to the lowest cost bandwidth offer. If the lowest cost bandwidth offer cannot entirely accommodate the traffic demand due to capacity limitation, then the residual demand will be assigned to the next lowest cost bandwidth offer. This traffic demand assignment iterates until the bandwidth offer with a particular cost can entirely accommodate the traffic demand. A lower bound is calculated based on the traffic assigned to each bandwidth offer and its associated cost. A lower bound, using the abovementioned method, can be calculated by

$$\sum_{k \in K} \sum_{\psi = Pr_{low}}^{Pr_{high}} \left\{ \text{Min} \left[\text{Max} \left(d_k - \sum_{\alpha = Pr_{low}}^{\psi-1} b_k^{\alpha}, 0 \right), b_k^{\psi} \right] \cdot \psi \right\} \quad (11)$$

For a particular cost ψ , the *max* function determines the residual traffic demand that has not been allocated to the bandwidth offers that have lower cost than the one being considered. The *min* function attempts to assign this residual traffic demand to the bandwidth offer with the cost currently being considered. The inner summation symbol considers all bandwidth offers toward a remote destination prefix with different costs. The outer summation symbol considers all the remote destination prefixes.

4.4 A genetic algorithm embedded with greedy heuristics

We propose an efficient Genetic Algorithm (GA) to obtain a near-optimal solution of the IBP problem. Genetic Algorithm is an algorithm that operates by the natural selection of 'survival of the fittest' (Holland 1975). It has been successful in solving many large-scale optimization problems. In order to making the proposed GA in solving the IBP problem more efficiently, we propose two problem-specific greedy heuristics embedded into the GA. To solve the IBP problem, we modify and extend the GA (Chu & Beasley, 1997) proposed for solving the Generalized Assignment Problem (Martello & Toth, 1990). The steps of our GA are as follows:

Step 1. Create a feasibility mapping table which maps all the feasible bandwidth offers to each inter-AS traffic flow. A bandwidth offer $oBw_k^{j,n}$ is feasible for an inter-AS traffic flow $t'(i,k)$ if the following constraints are satisfied:

$$oBw_k^{j,n} \in \text{Out}(k) \quad (12)$$

$$t'(i,k) \leq C_{inter}^{j,n} \quad (13)$$

$$t'(i,k) \leq \text{Max}Bw_k^{j,n} \quad (14)$$

Constraint (12) ensures that the remote destination prefix in the bandwidth offer matches the requested remote destination prefix of the traffic flow. Constraints (13) and (14) ensure respectively that the bandwidth demand of the traffic flow does not exceed the capacity of either the inter-AS link to which the bandwidth offer is associated or the maximum capacity of the bandwidth offer. These constraints, however, do not guarantee that constraints (2) or (3) are met for the entire chromosome.

Step 2. Generate an initial population of C randomly constructed chromosomes. Fig. 4 shows a representation of an individual chromosome which consists of T genes where T is the number of inter-AS traffic flows and each gene represents an assignment between a traffic flow and a bandwidth offer. The identifier given to each traffic flow represents each inter-AS traffic flow $t(i,k)$. Let $S_{t(i,k),c} = \langle k,j,n \rangle$ represent the bandwidth offer $oBW_k^{j,n}$ that has been assigned to traffic flow $t(i,k)$ in chromosome $c \in C$. Each gene of the initial chromosomes is generated by randomly assigning a feasible bandwidth offer to each traffic flow according to the feasibility mapping table created in step 1. Note that an initial chromosome may not be a feasible solution as capacity constraint (2) or (3) could be violated.

Traffic flow	1	2	...	T-1	T
Bandwidth offer	o-BW1	o-BW2	...	o-BWm	o-BWn

Fig. 4. Representation of an individual's chromosome

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

$$-\sum_{i \in I} \sum_{k \in K} Chg(s_{t'(i,k),c}) \cdot t'(i,k) \quad (15)$$

The negative sign reflects the fact that a solution with lower cost has higher fitness. We define $Chg(s_{t'(i,k),c}) \cdot t'(i,k)$ to be the IBP cost for the traffic flow $t(i,k)$. If the chromosome contains an infeasible solution, a common approach is to penalize its fitness for the infeasibility. Instead of this, we adopt the approach in (Chu & Beasley, 1997) and associate an unfitness value for each chromosome. The unfitness value of chromosome c is the degree of infeasibility of the chromosome, which equals the amount of violated capacity summed over all the inter-AS links and all the bandwidth offers,

$$\sum_{j \in J} \sum_{n \in Next_j} \left\{ 0, \sum_{i \in I, k \in K: S_{t'(i,k),c} = \langle k,j,n \rangle} t'(i,k) - C_{inter}^{j,n} \right\} + \sum_{k \in K} \sum_{j \in J} \sum_{n \in Next_j} \left\{ 0, \sum_{i \in I: S_{t'(i,k),c} = \langle k,j,n \rangle} t'(i,k) - MaxBW_k^{j,n} \right\} \quad (16)$$

With the separation of fitness and unfitness values, chromosomes can be evaluated in a two-dimensional plane, so the selection and replacement can direct the search towards feasible solutions by replacing highly unfit chromosomes with lightly unfit or entirely fit ones.

Step 4. Select two parent chromosomes for reproduction. We use the pairwise tournament selection method. In pairwise tournament selection, two individual chromosomes are chosen randomly from the population and the one that is fitter (higher 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_{co} is randomly selected. The first child chromosome consists of the first p_{co} genes from the first parent and the remaining $(n - p_{co})$ 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 a probabilistic mutation on each child chromosome. The mutation simply exchanges elements in two selected genes (i.e. exchange the assigned bandwidth offers between two randomly selected traffic flows) without violating constraints (12) – (14).

Step 7. The fitness and unfitness values of child chromosomes can be improved by applying the following two problem-specific heuristic operators:

- *Heuristic-A*: For each inter-AS traffic flow that has been assigned to an infeasible bandwidth offer such that either capacity constraint (2) or (3) is violated, find a feasible bandwidth offer that incurs the lowest IBP cost for the traffic flow. Denote $\Delta f(i,k)$ the difference between the original IBP cost induced by the traffic flow and the new IBP cost after the traffic flow has been reassigned to a feasible bandwidth offer. Among those inter-AS traffic flows, select the one with the lowest $\Delta f(i,k)$ and assign it to the corresponding selected feasible bandwidth offer. This heuristic operator iterates at most H times where H is a parameter that optimizes the algorithm's performance or stops when no inter-AS traffic flows have been assigned to infeasible bandwidth offers.
- *Heuristic-B*: For each inter-AS traffic flow, find a feasible bandwidth offer that produces the lowest IBP cost. If such a feasible bandwidth offer has been found, reassign the traffic flow to it.

Heuristic-A aims to reduce the unfitness value of the child chromosome by reassigning traffic flows from infeasible to feasible bandwidth offers while keeping the total IBP cost as low as possible. *Heuristic-B* attempts to improve the fitness of the child chromosome by reassigning traffic flows to feasible bandwidth offers with lower costs.

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 by the fitter child chromosomes. If no unfit solution exists, 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 GA termination criterion is met. The termination criterion is that either both the average and the best fitness over all the chromosomes in the two consecutive generations are identical or once the selected number of iterations, N_{itr} , has been reached in order to avoid excess algorithm execution time. Steps 4 - 9 iterate until the termination criterion is met.

5. Optimal traffic assignment

Let us assume that the bandwidth offers selected by the IBP (Section 4) have now been accepted and configured as a set of outbound provider SLAs. Given this set and the available bandwidth capacity within the AS, we now consider how to assign routes to the traffic so as to meet the traffic's bandwidth requirements. Fig. 5 shows that from the viewpoint of AS-1, a route to the destination can be decomposed into three parts: (1) the intra-AS route, (2) the inter-AS link and (3) the inter-AS route from the downstream AS (AS-2) to the destination AS (AS-3). Sufficient bandwidth must be provisioned in all parts of this route in order to satisfy the bandwidth demand. Once the outbound provider SLA is

known, the available bandwidth resource on any part of the route is known to the AS: the intra- and inter-AS links are owned by the AS and the available bandwidth from the downstream AS to the destination AS is guaranteed by the outbound provider SLA. As a result, the TA problem can be defined as follows:

Given a set of outbound provider SLAs, an inter-AS TM and a physical network topology, assign end-to-end routes to the supported traffic so that the bandwidth requirement is satisfied while optimizing network resource utilization. A route assignment includes the selection of an outbound provider SLA, an inter-AS link and an explicit intra-AS route from the ingress router to the egress router where the selected outbound provider SLA is associated.

We assume that explicit intra-AS routes are implemented by MPLS. In addition, there are many optimization criteria for network resource utilization, such as minimizing resource consumption or load balancing. For simplicity, the network resource utilization used in this chapter is a general metric, the total bandwidth consumed in carrying traffic across the network.

Fig. 5. Essential components for end-to-end bandwidth guarantee

5.1 Traffic assignment problem formulation

As with the IBP problem of Section 4, we formulate the TA problem as an integer-programming problem. The fundamental objective is to provide bandwidth guarantees to inter-AS traffic by satisfying their bandwidth demands. We define the bandwidth demand of an inter-AS traffic flow $t(i,k)$ to be met if the following three constraints are satisfied:

There exists at least one feasible path $f_{path} \in P_{i,j}$ from ingress router i to egress router j to which the selected outbound provider SLA is associated , i.e.

Min $bw_{intra}^l \geq t(i,k)$ (17)
 $\forall l \in f_{path}$

$bw_{inter}^{j,n} \geq t(i,k)$ (18)

$pSLABw(k,j,n) \geq t(i,k)$ (19)

Constraint (17) ensures that there exists at least one feasible path between the ingress point and the selected egress point, and the bottleneck bandwidth of the path is not less than the bandwidth demand of the traffic flow. Constraints (18) and (19) ensure that the inter-AS link and the outbound provider SLA respectively have sufficient bandwidth to accommodate the traffic flow.

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

$$\text{Minimize } \sum_{i \in I} \sum_{k \in K} \sum_{oBw(k, j, n) \in Out(k)} z_{i,k}^{j,n} \cdot dist_{i,j}^k \cdot t(i, k) \quad (20)$$

subject to:

$$\sum_{i \in I} \sum_{k \in K} z_{i,k}^{j,n} \cdot t(i, k) \leq c_{inter}^{j,n} \quad \forall (j, n) \text{ where } j \in J, n \in NEXT_j \quad (21)$$

$$\sum_{i \in I} \sum_{k \in K} \gamma_{i,k}^l \cdot t(i, k) \leq c_{intra}^l \quad \forall l \in E \quad (22)$$

$$\sum_{i \in I} z_{i,k}^{j,n} \cdot t(i, k) \leq pSLAC_k^{j,n} \quad \forall (k, j, n) \text{ where } k \in K, j \in J, n \in NEXT_j \quad (23)$$

$$z_{i,k}^{j,n}, \gamma_{i,k}^l, w_{i,k}^p \in \{0, 1\} \quad (24)$$

$$\sum_{oBw(k, j, n) \in Out(k)} z_{i,k}^{j,n} = 1 \quad \forall (i, k) \text{ where } i \in I, k \in K \quad (25)$$

$$\sum_{p \in P_{i,j}} w_{i,k}^p = 1 \quad \forall (i, k) \text{ where } i \in I, k \in K \quad (26)$$

$$\gamma_{i,k}^l \leq w_{i,k}^p \quad \forall (l, p, i, k) \text{ where } l \in p, p \in P_{i,j}, i \in I, k \in K \quad (27)$$

Constraints (21), (22) and (23) ensure that the total traffic assigned to the inter-AS link, the intra-AS link and the outbound provider SLA do not exceed their respective capacities. Constraint (24) ensures the discrete variables assume binary values. Constraint (25) ensures that only one outbound provider SLA is selected for each traffic flow. Constraint (26) ensures that each traffic flow $t(i, k)$ is routed along a single intra-AS route in order to preserve scalability and minimize network management complexity. Constraint (27) ensures that, whenever traffic flow $t(i, k)$ is assigned to intra-AS link l , then the path to which l is associated must have been selected. Moreover, given the lossless property of the links, an additional constraint that has not been presented is the flow conservation constraint which ensures that the traffic flowing into a node must equal the traffic flowing out of the node for any intermediate node.

5.2 A greedy heuristic algorithm for the traffic assignment problem

In comparing the two problems in the network dimensioning system, the complexity of the TA Problem is higher than the IBP problem, in terms of number of decision variables and constraints. In addition, the TA is performed more frequently than the IBP: network capacity expansion is usually less frequent than traffic engineering. Based on these reasons, the algorithm for solving the TA problem should be more *efficient* than the IBP algorithm. In general, a GA can produce a better performance but with higher time complexity than simple greedy-based heuristics. Due to the higher complexity of the TA problem, we do not consider using GA to solve the TA problem as we did for the IBP problem. Instead, we present a simple and efficient greedy heuristic algorithm to solve the TA problem, namely greedy-penalty heuristic.

Greedy-penalty heuristic: It is possible that the order in which traffic flows are assigned to outbound provider SLAs may produce different selection results. For example, if we take a traffic flow $t(i,k) = 2$, we might assign it greedily to some outbound provider SLA $pSLA_k^{j,n}$ with intra-AS distance $dist_{i,j}^k = 3$. In this case, the total bandwidth consumed equals 6. If on the other hand we allocate it later in the process, the outbound provider SLA may not have sufficient bandwidth because its bandwidth has been allocated to other traffic flows and the considered traffic flow might have to be assigned to another outbound provider SLA $pSLA_k^{j',n'}$, for example, with $dist_{i,j'}^k = 6$. As a result, the total bandwidth consumed equals 12. In this case, we have a penalty on the consumption of additional bandwidth (i.e. $12 - 6 = 6$) and we use *penalty* to refer to this 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 *penalty*. We propose a greedy-penalty heuristic algorithm that takes into consideration the penalty value. Such an algorithm has also been used to solve the GAP (Martello & Toth, 1990).

Step 1 For each unassigned traffic flow, we measure the desirability of assigning it to each feasible outbound provider SLA that satisfies constraint (19). The desirability is the total bandwidth consumed by the traffic flow along the intra-AS route between the ingress and the egress router with which the outbound provider SLA is associated (i.e. the number of intra-AS hops times the bandwidth demand). Intra-AS route computation is done by Constrained Shortest Path First (CSPF) (Osborne & Simha, 2002), which finds a route that is shortest in terms of hop while satisfying the bandwidth requirement. The smaller the desirability, the smaller amount of bandwidth to be consumed, and thus the better the selection.

Step 2 Compute *penalty* for each unassigned traffic flow, being the difference between the desirability of the traffic flow's best and second best selection (i.e. the two outbound provider SLAs which yield the smallest desirability). If there is only one feasible outbound provider SLA with sufficient spare capacity to accommodate the traffic flow, we need to set *penalty* to infinity and immediately assign the traffic flow to it. Otherwise, this outbound provider SLA may subsequently become unavailable, resulting in an invalid solution.

Step 3 Among all unassigned traffic flows, the one yielding the largest *penalty* is placed with its best selection. In other words, this traffic flow is assigned to the feasible outbound provider SLA that achieves the smallest desirability. If multiple traffic flows which have the same largest penalty exist, the one with the largest bandwidth demand is placed. If there are several such traffic flows, one is chosen randomly.

Step 4 Once the outbound provider SLA is selected, the requested bandwidth is allocated on the corresponding selected intra-AS route and the outbound provider SLA to establish an end-to-end bandwidth guaranteed route. We iterate step 1 to step 4 until all the traffic flows have been considered.

6. Performance evaluation

We evaluate the proposed GA and the greedy-penalty heuristic algorithms by simulation. The simulation software was written in Java. The computation was carried out on a laptop with an Intel Pentium Centrino 1.5GHz Processor with 512MB RAM. All the results presented in this chapter are an average of 50 different simulation trials.

6.1 Network model

We use a network topology generated by BRITE (Brite) with 100 nodes and average node degree of 4. These numbers were chosen to represent a medium to large INP topology. All intra-AS links are unidirectional and each has capacity of 500 units. Note that, since no realistic data is publicly available, we assume that the values of link capacity, bandwidth offers, and traffic demand are unitless. Therefore, these values that we use in this chapter may represent any specific value depending on the definition of the corresponding unit.

Among the 100 nodes, 30 nodes are randomly selected as border routers and the remaining nodes are core routers. In practice, each border router may connect with several inter-AS links to adjacent ASes. However, for simplicity, and without loss of generality, we abstract these inter-AS links into one. Thus, each border router is associated with one virtual inter-AS link which can logically represent one or multiple physical inter-AS links. Therefore, 30 virtual inter-AS links are considered and each has capacity of 500 units.

6.2 Bandwidth offer model

It is well known 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 (Feamster et al., 2003). Based on this finding, we consider 100 remote destination prefixes to be included in the bandwidth offers. In fact, each of them may not merely represent an individual prefix but also a group of distinct address prefixes that have the same end-to-end path properties, e.g. geographical location, offering AS and maximum available bandwidth. Hence, the hundred prefixes we considered could reflect an even larger number of prefixes.

In a network, each border router can be an ingress or egress point. Without loss of generality, we consider the network scenario where if a border router receives a bandwidth offer towards destination prefix k from adjacent AS Y , then AS Y cannot inject traffic for k into it. This corresponds to multi-hop traffic (Feldmann et al., 2001) in which the traffic traverses the network instead of being directed to another egress link of the same border router. We adopt this model in order to evaluate the TA objective of total bandwidth consumption in the network. As a result, we cannot assign all the destination prefixes on each border router as bandwidth offers. Instead, at each border router we randomly select half of these hundred destination prefixes as bandwidth offers and the other half as inter-AS traffic. In other words, we set the average number of distinct bandwidth offers advertised at each border router to be half of the number of prefixes. Furthermore, each border router can generate the number of traffic flows towards half of these prefixes that have not been selected for bandwidth offers. We note that this destination prefix generation process is just a best effort attempt to model prefix distribution, as no synthetic model for the actual behavior of prefix distribution in real networks was found in the literature. The remote destination prefixes associated with the bandwidth offers are randomly selected. The maximum capacity of each bandwidth offer is uniformly generated between 100 and 200 units. The charge associated with each bandwidth offer varies according to the simulation scenarios.

6.3 Traffic model

Ingress points and remote destination prefixes of the inter-AS traffic matrix are randomly generated. Previous work has shown that inter-AS traffic is not uniformly distributed (Fang

& Peterson, 1999). According to (Broido et al., 2004)), the AS traffic volumes are top-heavy and can be approximated by a Weibull distribution with shape parameter 0.2-0.3. We therefore generate the inter-AS TM with traffic demand following this distribution with the shape parameter 0.3. As previously mentioned, we do not allow traffic-prefix looping, so that if the AS receives a bandwidth offer towards remote destination prefix k from an adjacent AS, then this adjacent AS cannot inject traffic into the AS for k . The number of inter-AS traffic flows to be considered ranges from 500 to a maximum 1500.

As mentioned in Section 3.1, each inter-AS traffic flow is an aggregate of individual traffic flows that have identical ingress points and remote destination prefixes. Hence, the number of inter-AS traffic flows we considered does not reflect the exact total number of individual traffic flows. Instead, the number could represent more individual traffic flows. We assume that moderate overprovisioning is considered by the IBP and unless specified, $f_{over} = 1.25$ (i.e. 25% inter-AS bandwidth overprovisioning). Table 2 shows the number of traffic flows, their corresponding traffic volume and overall inter-AS link utilization. Note that the total traffic volume presented in the table has already taken into account the overprovisioning factor.

Number of traffic flows	Total Traffic volume	Overall inter-as egress link utilization (%)
500	4465	30%
625	5578	37%
750	6719	45%
875	7813	52%
1000	8915	60%
1125	10046	67%
1250	11142	74%
1375	12259	82%
1500	13402	90%

Table 2. Inter-AS traffic

6.4 Algorithm parameters

For the IBP’s GA parameters, we adopt the suggested values from previous GA research to achieve satisfactory effectiveness and convergence rate of the algorithm (Lin et al., 2003). The population size is 200, the value of H of the heuristic operator (a) is 200 since the IBP problem is highly constrained by two capacity constraints, N_{cd} is set to 50, the probability of mutation is 0.01 and N_{it} is set to 100.

6.5 Evaluation of the IBP algorithms

We compare the performance of our proposed GA described in Section 4.4 with the following alternatives:

Greedy-cost heuristic: The Greedy-cost heuristic sorts all the inter-AS traffic flows in descending order of bandwidth demand and selects one at a time in that order. From the bandwidth offers that have sufficient bandwidth to accommodate the given traffic flow, we select the one which incurs the least IBP cost. The flow is then allocated to this bandwidth offer and its corresponding inter-AS route. This step is repeated for the next traffic flow until all flows have been considered. One can imagine this heuristic might be a conventional algorithm used by INPs to solve the IBP problem.

Greedy-random heuristic: A greedy-random heuristic algorithm is included as a baseline comparison. The random heuristic algorithm is similar to the Greedy-cost heuristic except that the bandwidth offer selection of traffic flows is done at random. It may be viewed as the solution obtained by a trial-and-error or an ad hoc IBP approach.

6.5.1 Evaluation of the Total IBP Cost

The aim of the proposed GA is to achieve better and near-optimal IBP cost in comparison with the alternative algorithms. Hence, the main objective of the evaluation in this section is to quantify the effectiveness of the proposed GA over the alternative algorithms. Fig. 6 shows the total IBP cost achieved by the Greedy-cost and the GA as a function of inter-AS traffic flows. The performance of the Greedy-random heuristic is not presented in this figure since it has a significant performance gap from the other heuristics. Nevertheless, it is compared to the alternative algorithms in Table 3. The legend in the figure shows the names of the algorithms followed by the percentage of established peering connections as mentioned at the beginning of Section 4.

Fig. 6. Evaluation of the total inter-AS bandwidth provisioning cost

The figure presents the results of two practical scenarios, and we evaluate whether the proposed GA performs consistently well under these scenarios. The first scenario consists of all customer-provider connections. In other words, no peering connection (i.e. 0%) is established and the charge of each bandwidth offer is non-zero. We generate an integer uniformly between 1 and 10 to represent each cost. The figure shows that the GA has a lower total IBP cost at all numbers of inter-AS traffic flows. We conjecture that when the

number of inter-AS traffic flows is small, the inter-AS links and the bandwidth offers have relatively plenty of bandwidth to cover all the traffic, and so the GA and the Greedy-cost algorithm would give equivalent IBP results and costs. In contrast, as the number of inter-AS traffic flows increases, both the overall inter-AS link and bandwidth offer utilizations increase and some inter-AS links or bandwidth offers have even reached their capacity limits. In this case, some traffic flows may be assigned to other bandwidth offers which have higher costs. This evaluation shows that a careful selection of bandwidth offers is important in order to minimize the total IBP cost. This can be achieved by the GA.

In addition, the total IBP costs of the GA at all volumes of traffic flows are closer to the lower bound than the Greedy-cost heuristic. This shows that the GA is not only able to achieve a better cost than the Greedy-cost, but also able to achieve a near-optimal cost.

In the second scenario not only are customer-provider connections considered but also peering connections. We evaluate three levels of established peering connections: 3%, 6% and 9% of the total number of bandwidth offers. Simulation data presented in this scenario is as for the previous one except that a designated number of bandwidth offers is randomly selected as peering connections. In current Internet peering practice, most ASes will only accept on a peer link traffic from the peers' customers. Since our purpose is to merely evaluate the performance of the algorithms, we follow the assumption in (Feigenbaum et al., 2002) that general policy routing and peering/transit restrictions are ignored.

Fig. 6 shows that the GA performs better than the Greedy-cost at all degrees of peering connection and all number of inter-AS traffic flows. This is similar to the results of the 0% peering scenario. The GA has better total IBP costs than the Greedy-cost heuristic as the degree of peering connection increases. This is because more and more peering connections do not incur any charges, so that the GA can more effectively utilize the cost-free bandwidth in order to further minimize the total IBP cost. In general, this performance improvement not only applies to the second scenario where some peering connections exist but also applies to the 0% peering scenario where some exceptional low cost bandwidth offers exist.

Number of Inter-AS traffic flows	1000	1125	1250	1375	1500
Over Greedy-cost with 0% peering	3.33	5.0	5.92	8.67	12.75
Over Random with 0% peering	76.16	75.97	75.68	75.6	75
Over Greedy-cost with 3% peering	4.98	6.91	10.13	12.61	17.16
Over Random with 3% peering	83.66	83.08	83.06	81.95	81.38
Over Greedy-cost with 6% peering	7.71	10.6	14.3	18.01	24.0
Over Random with 6% peering	89.22	88.7	88.47	87.67	87
Over Greedy-cost with 9% peering	12.59	16.45	20.96	24.87	31.76
Over Random with 9% peering	92.7	92.41	91.98	91.47	90.85

Table 3. Performance improvement of the GA over the alternative algorithms (in %)

Table 3 shows the relative improvement of the GA over the Greedy-cost and the Greedy-random heuristic algorithms at all numbers of inter-AS traffic flows with different degrees of peering connection. By summarizing the table and considering a reasonably high traffic volume, the proposed GA has approximately 5%-30% and 75%-90% performance improvement over the Greedy-cost and the Greedy-random heuristics respectively under different scenarios. In comparison with the Greedy-random heuristic, the performance of

the GA is remarkable. This shows the importance and value of using systematic approaches, such as the proposed GA, over the trial-and-error and ad hoc approaches.

6.5.2 Evaluation of the proposed GA average running time

In Table 4 we provide the average running time of the GA. The average running time increases as the number of traffic flows increases. We can see that even for quite high numbers of traffic flows the running times are acceptable. These times are perfectly acceptable taking into account the timescale of the provisioning system operation.

Number of traffic flows	Average running time (secs)
500	36.6
1000	78.6
1500	150.4

Table 4. Average running time of the GA

6.5.3 Discussion of the IBP algorithms

The simulation study in this section has evaluated the performance of three IBP algorithms. Simulation results have firstly shown that the proposed GA is efficient and is able to achieve better total IBP cost than the random-based and the conventional heuristic algorithms. The relative total IBP cost improvement achieved by the GA over the Greedy-cost heuristic and the random-based algorithms are great, with 5%-30% and 75%-90% cost savings respectively. We conclude that the IBP solutions obtained by the proposed GA are good overall. This has an implication for INPs that a systematic approach could be developed to optimize the total IBP cost significantly.

6.6 Evaluation of the TA algorithms

The previous section evaluated the performance of the proposed IBP algorithms. Once the IBP phase is completed, an AS performs TA to optimize network resource utilization in order to provide end-to-end bandwidth guarantees for the supported traffic. In this section, we evaluate the performance of our proposed TA algorithms.

We assume that outbound provider SLAs are successfully established in line with the first scenario in the evaluation of IBP algorithms, i.e. the GA IBP outcomes with a linear cost function and all customer-provider connections (0% peering). These outbound provider SLAs are then the input to the TA problem. We consider the following three approaches for the TA problem, namely Cost-only, Cost-Performance and Performance-only approaches. The words “Cost” and “Performance” used in the names of these approaches mean that the ordered priorities of the algorithm optimization targets are on the total IBP cost and the total bandwidth consumption respectively.

Cost-only: Given an IBP solution produced by the GA, there are multiple solutions for assigning traffic to satisfy all the TA problem constraints. Any of these solutions can be selected as the solution of the Cost-only approach since it does not optimize the total bandwidth consumption in the network. We use the Random-TA heuristic algorithm, as shown in Fig. 7, to find a solution for the Cost-only approach.

Random-TA Heuristic Algorithm

Sort inter-AS traffic flows in decreasing order of bandwidth demand

For each traffic flow in that order

- Assign an egress point randomly to the traffic flow
- Establish a bandwidth constrained path between the ingress and egress point
- Update utilized resources

End For

Fig. 7. The random-TA heuristic

Cost-Performance: Given an IBP solution produced by the GA, the Cost-Performance approach takes the proposed greedy-penalty heuristic algorithm as the TA algorithm to optimize the total bandwidth consumption in the network.

Performance-only: The Performance-only approach does not use the IBP solution. Instead, it takes all the bandwidth offers (rather than the outbound provider SLAs) as input and uses the Greedy-penalty heuristic algorithm to solve the TA problem. The total IBP cost is then equal to the sum of the cost of each accepted bandwidth offer. Since the total IBP cost is calculated by taking overprovisioning into consideration, we approximate the total IBP cost of the Performance-only approach by multiplying its solution cost by f_{over} in order to compare it with the total IBP costs achieved by the other two approaches.

6.6.1 Cost vs. performance

We evaluate the proposed three TA approaches. We test the hypothesis that the Greedy-penalty heuristic algorithm can improve the total network bandwidth consumption.

Fig. 8. Normalized total inter-AS bandwidth provisioning cost

Fig. 8 shows the total IBP costs of all the TA approaches at three different volumes of inter-AS traffic flows: 500, 1000 and 1500. The total IBP costs are normalized by the cost of

the solution produced by the GA. The total IBP costs of the Cost-only and the Cost-Performance approaches are identical because they both use the IBP solution produced by the GA. In contrast, the total IBP cost of the Performance-only approach is on average 4 times higher than the others. This significantly higher cost results from neglecting the IBP optimization so that some expensive bandwidth offers are selected, although, as we can see below, using them can significantly improve the total bandwidth consumption in the network.

Indeed, although the Performance-only approach has a very high total IBP cost, Fig. 9 shows that its total bandwidth consumption is approximately half of the other two approaches. Nevertheless, because of its high total IBP cost, the Performance-only approach can be assumed impractical. This implies that there can be conflict between the IBP cost and bandwidth consumption. Therefore, we need a compromising solution that would balance the interests of these two metrics. The Cost-Performance approach attempt to achieve such solution as it has low IBP cost and low total bandwidth consumption compared to the Cost-only approach with the amount closer to the Performance-only approach. This reduced total bandwidth consumption reveals that the proposed Greedy-penalty heuristic algorithm has on average a 10% improvement over the Random-TA heuristic algorithm.

Fig. 9. Normalized total bandwidth consumption in the network

6.6.2 Evaluation of the greedy-penalty heuristic algorithm average running time

Table 5 provides the average running time of the proposed greedy-penalty heuristic algorithm. The average running time increases as the number of traffic flows increases. These running times are perfectly acceptable taking into account the timescale of the provisioning system operation. The computation time could have been much longer if GA was used due to its evolutionary process.

Number of traffic flows	Average running time (secs)
500	6.2
1000	22.1
1500	64.76

Table 5. Average running time of the greedy-penalty heuristic

6.6.3 Discussion of the TA approaches

The simulation described in this section has evaluated the performance of three TA approaches. Simulation results have shown that the proposed Greedy-penalty heuristic algorithm used by the Cost-Performance approach is efficient and is able to achieve on average 10% less total bandwidth consumption than the random-based algorithm used in the Cost-only approach. The performance difference between the Performance-only approach and the other two reveals that a trade-off exists between the IBP and the TA optimization. This trade-off has also been discussed in (Goldenberg et al., 2004) where primarily optimizing monetary cost can degrade network performance and vice versa. However, the determination of relative weights between cost and performance optimizations is far from trivial, particularly when the units of the two metrics have different scales. It is thus in many cases difficult to express in terms of weights the trade-off between the two metrics. Therefore, we assume that from business point of view, an AS considers the IBP cost optimization as more important than the TA performance optimization. Based on this assumption and our simulation study, we conclude that the Cost-Performance approach, which uses our proposed GA and the greedy-penalty heuristic algorithm, performs well both in terms of the total IBP cost and the total bandwidth consumption, in comparison with the Cost-only and the Performance-only approaches. The Cost-Performance approach can be used by INPs to achieve an effective provisioning of end-to-end bandwidth guarantees. Moreover, since the TA problem has dealt with the selection of inter-AS route and explicit intra-AS route within the network, the Cost-Performance approach could be effectively applied to BGP/MPLS virtual private network provisioning (Rosen & Rekhter, 1999), a subject which is attracting a great deal of attention.

6.6.4 Impact of inter-AS overprovisioning factor on bandwidth consumption

We evaluate the impact of overprovisioning factor on the total bandwidth consumption achieved by the three TA approaches. The results of this evaluation are based on 1500 inter-AS traffic flows. The values of the inter-AS bandwidth overprovisioning factor examined are 1.25, 1.5, 1.75 and 2.0. As the inter-AS available bandwidth increases, the outbound provider SLA capacity constraint becomes less restrictive to the TA problem. Thus, in this case, we expect that the total bandwidth consumption in the network can be further improved. Fig. 10 shows that the total bandwidth consumption decreases as the overprovisioning factor increases. This is because a large overprovisioning factor reduces the outbound provider SLA capacity constraint and therefore increases the solution space for the TA algorithm, enabling it to find a result with lower total bandwidth consumption. As expected, the Cost-Performance approach has lower total bandwidth consumption than the Cost-only approach at any considered value of the overprovisioning factor. The total bandwidth consumption of the Performance-only approach for all considered values of the

overprovisioning factor is identical because the approach does not consider IBP. Therefore, its performance is not affected by the overprovisioning factor. Fig. 11 shows the normalized total bandwidth consumption achieved by the three TA approaches. As the overprovisioning factor increases, the relative improvement of the Cost-Performance approach over the Cost-only approach slightly increases from approximately 11% to 13%.

Fig. 10. Total bandwidth consumption achieved by different f_{over}

Fig. 11. Normalized total bandwidth consumption achieved by different f_{over}

Fig. 11 also reveals that the performance differences among the three TA approaches are consistent and are insensitive to changes on the overprovisioning factor. The results presented in these figures have revealed the effect of IBP on the TA performance with a different overprovisioning factor. The results confirm our conjecture that as the overprovisioning factor increases, more bandwidth is available in outbound provider SLAs for the TA algorithms to further optimize the total bandwidth consumption.

7. Conclusion

In this chapter we have reviewed a cascaded negotiation model for negotiating and establishing SLAs for bandwidth guarantees between ASes, and a network dimensioning system to solve the inter-AS bandwidth provisioning and the traffic assignment problems systematically.

We formulated the inter-AS bandwidth provisioning problem as an integer programming problem and prove it to be NP-hard. An efficient genetic algorithm was proposed to solve the problem. Our simulation study shows that the genetic algorithm has a near-optimal total inter-AS bandwidth provisioning cost. This cost is approximately 5%-30% and 75%-90% less than the cost achieved by a conventional greedy heuristic algorithm and a random-based algorithm respectively under two customer-peering scenarios.

We formulated the traffic assignment problem as an integer programming problem and prove it to be NP-hard. An efficient greedy-penalty heuristic algorithm was proposed to solve the problem. Our simulation study showed that the greedy-penalty heuristic algorithm achieved on average 10% less total bandwidth consumption than the random-based TA heuristic algorithm.

Finally, we evaluated the effects of different overprovisioning factor values on the total bandwidth consumption. The more the inter-AS bandwidth is overprovisioned, the less the total bandwidth is needed to carry the supported traffic across the network.

A limitation of our work is performance robustness. In case where the derived traffic matrix deviates significantly from the real traffic demands or link failures happen, the performance of IBP and TA may be affected since these network conditions have not been taken into account during the optimization. As future work, we will make the IBP and TA problems robust to traffic demand uncertainty and link failures. Although this may result in trade-offs between performance and robustness, we attempt to achieve good and balance solutions with respect to these two metrics.

8. References

- Awduche, D. et al. (2002). Overview and Principles of Internet Traffic Engineering, *IETF RFC 3272*.
- BRITE: Boston University Representative Internet Topology Generator. Website <http://www.cs.bu.edu/brite/>.
- Broido, A. et al. (2004). Their Shares: Diversity and Disparity in IP Traffic, *Proceedings of Passive and Active Measurement Workshop*, pp. 113-125, Antibes Juan-les-Pins, April 2004, Springer.
- Chu, P.C. & Beasley, J.E. (1997). A Genetic Algorithm for the Generalized Assignment Problem. *Computers and Operations Research*, Vol. 24, No. 1, January, pp. 17-23.

- Fang, W. & Peterson, L. (1999). Inter-AS Traffic Patterns and Their Implications, *Proceedings of IEEE GLOBECOM*, pp. 1859-1868, Rio Je Janeiro, December 1999, IEEE.
- Feamster, N. et al. (2003). Guidelines for Interdomain Traffic Engineering. *ACM SIGCOMM Computer Communications Review*, Vol. 33, No. 5, October, pp. 19-30.
- Feigenbaum, J. et al. (2002). A BGP-based Mechanism for Lowest-Cost Routing, *Proceedings of ACM Symposium on Principles of Distributed Computing*, pp. 173-182, Monterey, July 2002, ACM.
- Feldmann, A. et al. (2001). Deriving Traffic Demands for Operational IP Networks: Methodology and Experience. *IEEE/ACM Transactions on Networking*, Vol. 9, No. 3, June, pp. 265-279.
- Feldmann, A. et al. (2004). A Methodology for Estimating Interdomain Web Traffic Demand, *Proceedings of ACM IMC*, pp. 322-335, Taormina, October 2004, ACM.
- Goldenberg, D.K. et al. (2004). Optimizing Cost and Performance for Multihoming, *Proceedings of ACM SIGCOMM*, pp. 79-92, Portland, August 2004, ACM.
- Guerin, R.; Ahmadi, H. & Naghshineh, M. (1991). Equivalent Capacity and Its Applications in High-speed Networks. *IEEE Journal on Selected Areas in Communications*, Vol. 9, No. 7, September, pp. 968-981.
- Holland, J.H. (1975). *Adaptation in Natural and Artificial System*, The University of Michigan Press, ISBN 0262581116.
- Howarth, M. et al. (2005). Provisioning for Inter-domain Quality of Service: the MESCAL Approach. *IEEE Communications Magazine*, Vol. 43, No. 6, June, pp. 129-137.
- Martello, S. & Toth, P. (1990). *Knapsack Problems: Algorithms and Computer Implementations*, Wiley-Interscience, ISBN 0471924202.
- Lin, X.H.; Kwok, Y.K. & Lau, V.K.N. (2003). A Genetic Algorithm Based Approach to Route Selection and Capacity Flow Assignment. *Computer Communications*, Vol. 26, No. 9, June, pp. 961-974.
- Nucci, A. et al. (2005). Increasing Link Utilization in IP over WDM Networks Using Availability as QoS. *Photonic Network Communication*, Vol. 9, No. 1, January, pp. 55-75.
- Osborne, E. & Simha, A. (2002). *Traffic Engineering with MPLS*, Cisco Press, ISBN 1587050315.
- Rosen, E. & Rekhter, Y. (1999). BGP/MPLS VPNs, *IETF RFC 2547*.
- Shi, S.Y. & Turner, J.S. (2002). Multicast Routing and Bandwidth Dimensioning in Overlay Networks. *IEEE Journal on Selected Areas in Communications*, Vol. 20, No. 8, October, pp. 1444-1455.
- Sridharan, A.; Guerin, R. & Diot, C. (2005). Achieving Near-optimal Traffic Engineering Solutions for Current OSPF/IS-IS Networks. *IEEE/ACM Transactions on Networking*, Vol. 13, No. 2, April, pp. 234-247.
- Teixeira, R. et al. (2005). Traffic Matrix Reloaded: Impact of Routing Changes, *Proceeding of Passive and Active Measurement Workshop*, pp. 251-264, Boston, March/April 2005, Springer.
- Zhang, R. et al. (2004). MPLS Inter-Autonomous System Traffic Engineering Requirements, *IETF RFC 4216*.

Greedy Algorithms

Edited by Witold Bednorz

ISBN 978-953-7619-27-5

Hard cover, 586 pages

Publisher InTech

Published online 01, November, 2008

Published in print edition November, 2008

Each chapter comprises a separate study on some optimization problem giving both an introductory look into the theory the problem comes from and some new developments invented by author(s). Usually some elementary knowledge is assumed, yet all the required facts are quoted mostly in examples, remarks or theorems.

How to reference

In order to correctly reference this scholarly work, feel free to copy and paste the following:

Kin-Hon Ho, Ning Wang and George Pavlou (2008). Solving Inter-AS Bandwidth Guaranteed Provisioning Problems with Greedy Heuristics, *Greedy Algorithms*, Witold Bednorz (Ed.), ISBN: 978-953-7619-27-5, InTech, Available from: http://www.intechopen.com/books/greedy_algorithms/solving_inter-as_bandwidth_guaranteed_provisioning_problems_with_greedy_heuristics

INTECH
open science | open minds

InTech Europe

University Campus STeP Ri
Slavka Krautzeka 83/A
51000 Rijeka, Croatia
Phone: +385 (51) 770 447
Fax: +385 (51) 686 166
www.intechopen.com

InTech China

Unit 405, Office Block, Hotel Equatorial Shanghai
No.65, Yan An Road (West), Shanghai, 200040, China
中国上海市延安西路65号上海国际贵都大饭店办公楼405单元
Phone: +86-21-62489820
Fax: +86-21-62489821

© 2008 The Author(s). Licensee IntechOpen. This chapter is distributed under the terms of the [Creative Commons Attribution-NonCommercial-ShareAlike-3.0 License](https://creativecommons.org/licenses/by-nc-sa/3.0/), which permits use, distribution and reproduction for non-commercial purposes, provided the original is properly cited and derivative works building on this content are distributed under the same license.

IntechOpen

IntechOpen