Multi-Perspective Optimization of GÉANT Inter-domain Traffic

Manuel Pedro[1,2], Edmundo Monteiro[2], and Fernando Boavida[2]
[1] Polytechnic Institute of Leiria, ESTG, Morro do Lena,
Alto do Vieiro, 2411-901 Leiria, Portugal
[2] University of Coimbra, Pólo II, Pinhal de Marrocos,
3030-290 Coimbra, Portugal
{macpedro, edmundo, boavida}@dei.uc.pt

Abstract

Inter-domain traffic engineering is a key issue when resource optimization across autonomous systems boundaries is concerned. Mapping inter-domain traffic flows into existing service level agreements is, in general, a complex problem, for which an optimization tool is needed. This work aims at demonstrating the advantages of using such a tool to perform off-line inter-domain traffic optimization, using a multi-perspective approach. This optimization approach was applied to GÉANT, the European Research and Educational Network. The paper presents optimization solutions from various perspectives, namely bandwidth, monetary cost and routing trustworthiness. Results show the costs of the original GÉANT traffic engineering solution and the costs of the optimized solutions side-by-side. This case study also shows the advantages of being able to select between the perspectives that best fit the domain management policies.

1. Introduction

The main purpose of inter-domain resource optimization is to map incoming inter-domain traffic flows into inter-domain network resources, in order to optimize the use of network resources across autonomous systems (AS) boundaries. Network resources usage is, in any case, conditioned by existing Service Level Specifications (SLSs) that, in turn, result from the Service Level Agreements (SLAs) established between each domain and its neighbors. For the purpose of this paper, the terms 'domain' and 'autonomous system' are synonyms.

In order to describe the inter-domain relationships of an autonomous system, one can use a simple model, as shown in Figure 1. An autonomous system is interconnected with other autonomous systems by means of its ingress and egress interfaces. For the purpose of this paper, an egress interface includes the egress router identification plus the AS number of the peer neighbor AS.

The service offerings between autonomous systems as well as their mutual responsibilities are described by means of Service Level Agreements. In general, each SLA defines a set of contractual, administrative and technical requirements. The latter are called Service Level Specifications. An SLS comprises several items or clauses, including identification, application scope, flow identification, traffic conformance, excess treatment, and performance guarantees.

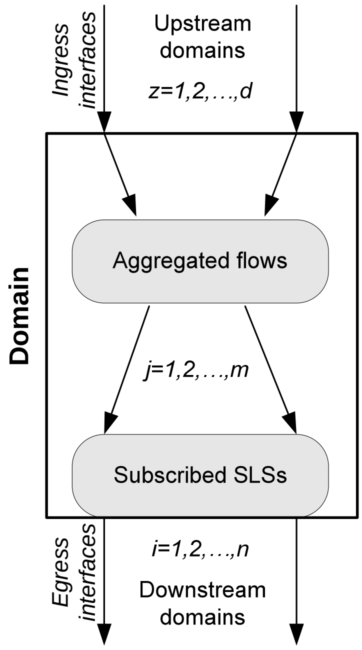

Figure 1. Inter-domain relationship model

In the context of the present work an SLS is characterized by an egress interface, a destination AS *d*, the corresponding maximum bandwidth requirements *b*, the monetary cost per unit of bandwidth *c*, and the route trustworthiness *r* associated with the SLS [3]. The monetary cost component reflects the monetary cost associated with the established SLA. On the other hand the routing trustworthiness reflects the intra-domain routing costs associated with the egress interface, and the inter-domain routing costs such as route quality, reliability and domain policies [3]. An SLS entry for a domain has the following format:

SLS entry = [egress interface, d, b, c, r]

On the other hand, a domain receives from upstream domains a collection of data flows towards other domains. Depending on the domain policy and on their common characteristics, such as destination, these flows may be aggregated. For the purpose of this work the flows' common characterization is the destination AS number d. This choice aimed the simplification of the high level flow's management tasks, without losing the main goal of destination prefixes. That is, an aggregated flow entry has the following format:

Aggregate flow entry = [ingress interface, d, a]

where *a* is the bandwidth requirement of the aggregated flow. The flow will be mapped into one of the existing SLSs. The appropriate selection of the SLSs for the inter-domain traffic flows benefits the domain by improving the network resources, maximizes the profits from a business point-of-view and, at same time, selects the most reliable routes according to internal and external information and business objectives [3]. The first benefit is reached through a correct bandwidth load-balancing, the second through a minimization of the costs, and the third through a high value of routing trustworthiness. In contrast, in current networks this task is executed in a trial-and-error fashion.

The objective of the work presented in this paper is to propose a multi-perspective optimization solution to inter-domain traffic engineering in the GÉANT network. To support our goal an optimization tool was used [3] on four months of data records from GÉANT.

In Section 2 of this paper an overview of related work is given. This is followed by a presentation of the optimization proposal and the optimization tool in Section 3. Section 4 presents and discusses the obtained results. The conclusions and guidelines for further work are presented in Section 5.

2. Related Work

Several studies on intra-domain resource optimization, such as [5] and [6], can be found in the literature. In the case of inter-domain, references [3] [7-11] constitute the framework for most of the current proposals.

The proposals [3] and [10] present multi-objective genetic algorithm based solutions. [10] contemplates monetary and bandwidth cost minimization, for traffic engineering of best effort traffic, and [3] extends the use of routing trustworthiness costs. None of these use real datasets to perform optimization.

Since the off-line inter-domain optimization tool proposed in [3] is the most complete, as far we know, it was our first choice to do this work.

To the best of the authors' knowledge, it is the first time an off-line inter-domain traffic optimization tool is applied to GÉANT network traffic data. On the other hand, it is also the first time that three different optimization perspectives are evaluated and compared on the basis of traffic data from a real network.

3. Proposal

In this section we describe the GÉANT datasets, and the methodology used to build the required dataset by the optimization tool. We also describe the selected optimization perspectives and the optimization tool's algorithm.

3.1. GÉANT Dataset

GÉANT is the pan-European backbone research network interconnecting Europe's national research and education networks (NRENs) of 34 countries [2]. GÉANT is composed of 23 routers interconnected by 38 links, and 53 links with others domains. All routers of GÉANT are border routers.

To support this work four months of dataset records from GÉANT were used. The dataset includes intra-domain and inter-domain traces of byte volume, for every 15-minutes window, and for every <source prefix, destination prefix>, and the daily BGP routing table dumps and its 15 minutes updates [1].

For the purpose of our work it was selected only one dataset per day, between 13:00 and 13.15, to characterize the entire day. This simplified dataset follows the patterns and the values presented in [1], apart from a scale factor, which simplifies extremely the decision process, by drastically reducing the processing times. The link load over the day can be found in [1].

3.2. Rebuilding the Inter-domain Dataset

In order to support our proposal we had to rebuild the following dataset from the GÉANT dataset: ingress inter-domain traffic matrix, the established Service Level Agreements, and the GÉANT mapping solution.

To estimate the required dataset, an inter-domain traffic matrix estimation tool was developed (Figure 2). This tool has as main purpose the reconstruction of the inter-domain traffic matrix, based on known inter-domain traffic volume in each node. Namely, it has to find the domain entry and exit points, and the intra-domain routes. To perform this task it receives the intra-domain traffic volume between the nodes, the routing table, the network topology, and the IGP weights, and follows the BGP route selection steps, and the IGP decision process.

On the other hand, to estimate the established SLSs, we need to know the maximum link bandwidth agreed with neighboring ASs, the bandwidth per destination AS, the bandwidth cost with each peer neighbor, and route trustworthiness for each route. Based on the available information [2], the maximum link bandwidth agreed between neighboring ASs was assumed to be equal to the corresponding link capacity. In the case of the maximum required bandwidth, for a given destination AS, it was defined as the average value, in a 15-minute window, of the maximum recorded value, during the 4 months of records. Regarding the link cost, we set a cost of 10 for commercial ASs, 5 for non-European and non-commercial ASs, and a cost of 1 for all European and non-commercial ASs. On the other hand, the route trustworthiness cost was defined as 100% minus the percentage of the number of route changes, during the four months observation period, for each destination AS.

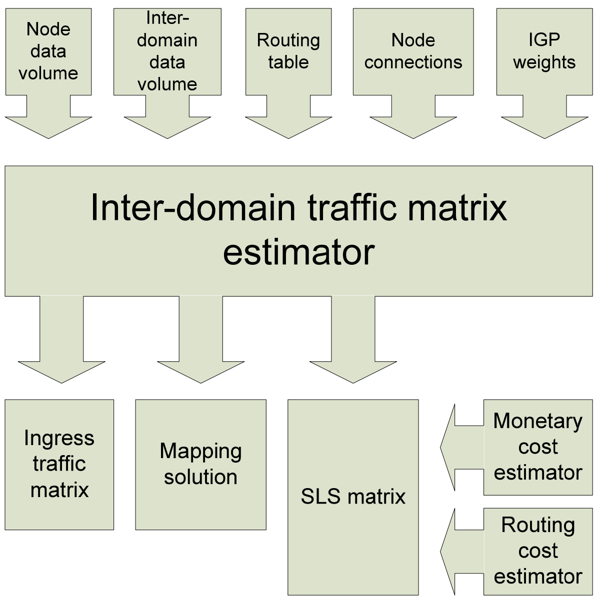

Figure 2. Traffic matrix estimator

3.3. Optimization Perspectives

The domain management decision process needs to take into account every optimization perspective before taking a decision. To simplify this process, the optimization can be performed using multiple perspectives. The perspectives considered by the used optimization tool are [3]:

- trade-off between all objectives
- minimization of the egress link bottleneck, or bandwidth cost
- maximization of the domain business profit
- maximization of the route trustworthiness

The problem can be expressed by three objective functions (1), (2), and (3) that represent respectively the total costs for bandwidth, monetary cost and routing. Formally, the problem can be stated as follows. Let $I = \{1,2,...,s\}$ be the set of SLSs and $J = \{1,2,..., f\}$ the set of aggregated traffic flows. For each SLS i there is a given resource capacity, expressed in terms of bandwidth, $b_i > 0$. For each $i \in I$ and each $j \in J$ there is a given set of costs, $B_{i,j} > 0$, for bandwidth, $C_{i,j} > 0$, for monetary, and $R_i > 0$, for routing, for assigning an aggregated traffic flow j to an SLS i. Additionally, $z_{i,j}$ is an indicator function that returns 1 if the traffic flow j is assigned to SLS i and 0 otherwise. The mathematical formulation is as follows:

$$y_1 = \sum_{i=1}^{s}\sum_{j=1}^{f} B_{i,j} \cdot z_{i,j} \quad (1)$$

$$y_2 = \sum_{i=1}^{s}\sum_{j=1}^{f} C_{i,j} \cdot z_{i,j} \quad (2)$$

$$y_3 = \sum_{i=1}^{s}\sum_{j=1}^{f} R_i \cdot z_{i,j} \quad (3)$$

$$\text{subject to } \sum_{j=1}^{f} a_{i,j} \cdot z_{i,j} \le b_i, \forall i \in I, \quad (4)$$

$$\text{with } \sum_{i=1}^{s} z_{i,j} = 1, \forall j \in J, \quad (5)$$

$$z_{i,j} \in \{0,1\}, \; \forall i \in I, \forall j \in J. \qquad (6)$$

The goal is to minimize the costs of (1), (2), and (3), where $B_{i,j}$, $C_{i,j}$, and R_i are respectively the cost functions for bandwidth, monetary, and routing. The capacity constraint (4) ensures that the total resource requirements of the traffic flows assigned to each SLS do not exceed the available capacity. The assignment constraint (5) guarantees that each traffic flow is assigned to exactly one SLS.

The cost functions that need to be minimized are:

$$B_{i,j} = \frac{1}{\left(b_i - b_j + 0.1\right)^2} \qquad (7)$$

$$C_{i,j} = c_i \cdot b_j \qquad (8)$$

$$R_i = 100 - r_i \qquad (9)$$

(7) is the bandwidth cost function used to measure the egress interfaces bottleneck, allowing the correct load-balancing in these interfaces, where b_i is the available bandwidth on egress interface *i* (the agreed SLS) for some destination and b_j is the bandwidth of the aggregate flow *j* (see Figure 1). The value 0.1 was added to the denominator in order to limit the value of B to 100.

On the other hand, (8) is the monetary cost function, which determines the amount to pay for using the established SLS *i*, by the aggregated flow *j*.

Lastly, (9) is the route trustworthiness objective function, which measures the aspects related to the egress links for the SLS *i*. For the purpose of this work, the route change history was considerate the trustworthiness parameter. The parameter *r* takes values between 0 (worst route) and 100 (the best route).

3.4. Optimization Algorithm

To perform the daily data optimization a multi-objective genetic algorithm tool [3] was used. The algorithm, belonging to the class of evolution strategies used in optimization, resembles the process of biological evolution, where each individual is described by its genetic code, called a chromosome. On the other hand each chromosome is composed of individual genes. In the problem in hand, a gene is the assignment of a single aggregate traffic flow to an SLS, and an individual (i.e., a chromosome) is a potential solution.

The tool returns an entire non-dominated Pareto front in a single run [4], i.e. a set of solutions where its components cannot be improved in terms of one objective without causing a simultaneous degradation in at least one of the other components.

The basic algorithm steps are presented in Figure 3. It starts with the creation of the initial generation, where the individuals are created randomly. Then, an evaluating step based on the proposed objective functions (1), (2), and (3), with costs (7), (8), and (9) respectively, follows. After that, and for a number of generations, a new generation of children is created that are compared with the corresponding generation of parents. From this comparison the better elements will compose the next generation of parents. The ranking step is done as proposed in [12].

```
Create the initial parent generation;
Evaluate the generation;
For a number of generations;
  Create the child generation;
  Evaluate both generations together;
  Rank both generations together;
  Replace worst parents with better
children;
End
```

Figure 3. Algorithm basic steps

The algorithm has a time complexity of $O(MN^2)$ where *M* is the number of objectives and *N* the size of the population.

4. Results

In this section we analyze the optimization results. We will begin by providing daily dynamics of the costs during the considered period. Then we will present the total costs gain compared with the costs of the original GÉANT mapping solution, extracted from the GÉANT dataset.

The results contemplate each optimization perspective, namely, the trade-off solution, the lowest bandwidth cost solution, the lowest monetary cost solution, and the best fitted solution in relation to routing trustworthiness.

All graphs have a ‘hole’ corresponding to six days for which there were no records in the GÉANT data.

4.1. Daily Costs

Figure 4 provides the evolution over time of the daily bandwidth costs of the original GÉANT solution, and the various optimization perspectives. We can see that GÉANT solution is always costly when compared to the optimized perspectives. The lower costs are achieved by the bandwidth perspective.

The periodic trends follow weekly patterns.

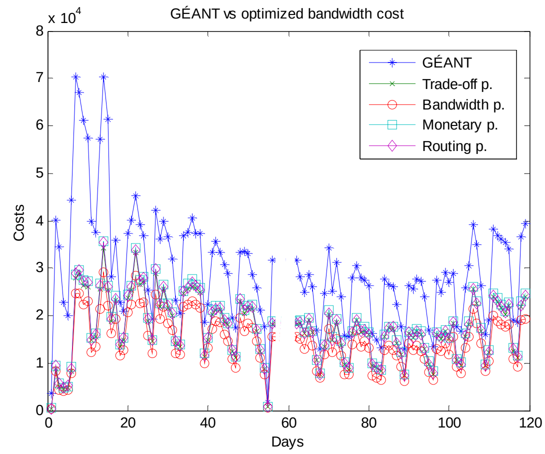

Figure 4. Evolution of bandwidth costs

Figure 5 shows the evolution over time of monetary costs for the original GÉANT solution, and for all other perspectives. Unlike the case of Figure 4, the cost of the GÉANT solution and the cost of the optimized perspectives are very close, although the latter are slightly lower.

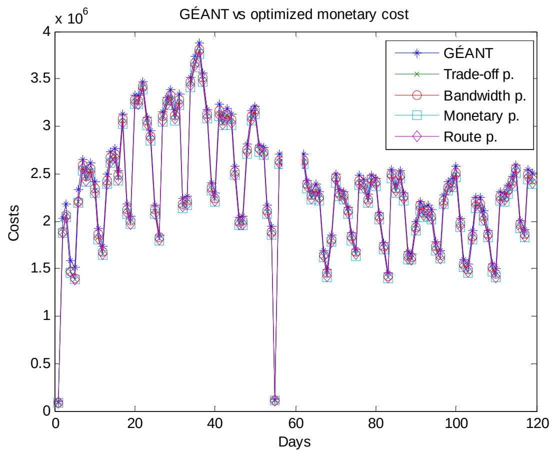

Figure 5. Evolution of monetary costs

Figure 6 provides the evolution over time of the daily routing costs for GÉANT solution and the optimized perspectives. The results follow the tendency already registered in Figure 4 with respect to the GÉANT solution. The lower costs are obtained with the routing perspective optimization.

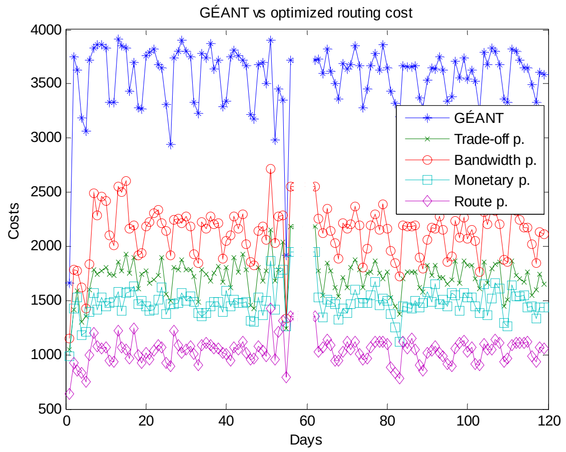

Figure 6. Evolution of routing costs

4.2. Gains

To get a global idea of the gains over the GÉANT solution, we have compared the various perspectives in terms of gain.

Figure 7 shows the percentage of gain, with respect to the original GÉANT solution, for bandwidth costs. One can say that the usage of the bandwidth perspective returns the best bandwidth gains for the domain (51%). On the other hand, for all perspectives the gains are always over 40%.

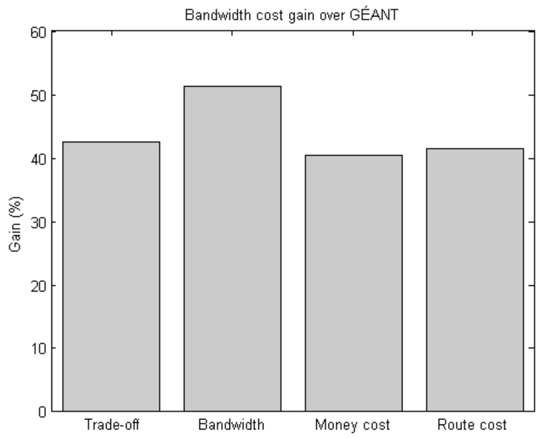

Figure 7. Total bandwidth gain

Figure 8 illustrates the percentage of gain, from the monetary point of view, with respect to the GÉANT solution. Here the values are not as high as in the latter case but, nevertheless, they are still positive for all perspectives. The monetary perspective has the best gain (4%) over the GÉANT solution.

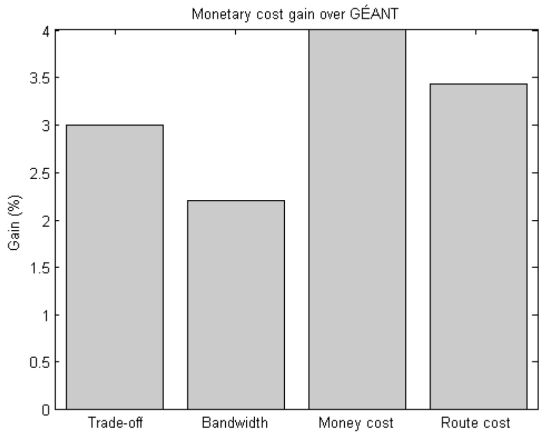

Figure 8. Total monetary cost gain

Lastly, Figure 9 shows the percentage of gain for the routing costs, with respect to the GÉANT solution. As we can see, there are considerable gains for all perspectives. The route cost solution is heading the gains with 56%. The remaining perspectives have gains of 14% for bandwidth and 40% for monetary costs.

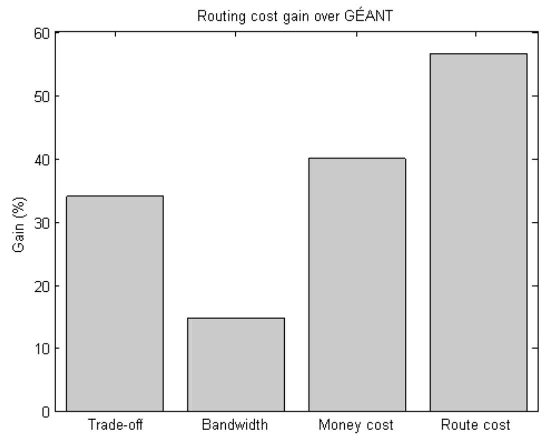

Figure 9. Total routing cost gain

5. Conclusion

Inter-domain resource optimization is one of the main challenges of current traffic engineering. This paper presented a GÉANT inter-domain optimization case study. The study was made using an optimization tool [3] and four months of GÉANT dataset records. Multi-perspective optimization solutions were shown, namely the trade-off perspective, the bandwidth perspective, the monetary cost perspective, and the routing cost perspective. The optimized solutions lead to high gains for every proposed perspective when compared with the GÉANT mapping solution.

As a general conclusion, one can say that the presented case study illustrates the advantages of using the inter-domain optimization tool [3] in a real network and the benefit of being able to select between the perspectives that best fit the domain management policies.

Further work will address the optimization of GÉANT-like networks from a QoS-aware perspective.

Acknowledgment

We are grateful to everyone that made possible the availability of GÉANT data, in particular Steve Uhlig from University of Louvain-la-Neuve.

This work was partially supported by the Portuguese Ministry of Science and High Education (MCES) and by European Union FEDER under programs POSC and PRODEP.

References

[1] S. Uhlig, B. Quoitin, S. Balon, and J. Lepropre, “Providing public intradomain traffic matrices to the research community”. ACM SIGCOMM Computer Communication Review, January 2006, 36 (1), 83-86.

[2] GÉANT [website]: http://www.geant2.net

[3] M. Pedro, E. Monteiro, and , F. Boavida. “An Approach to Off-line Inter-domain QoS-Aware Resource Optimization”. In Proc. of Networking 2006, (5th International IFIP-TC6 Networking Conference, Coimbra, Portugal, May 2006), Springer, 247-255.

[4] C. Coello, “An Updated Survey of GA-Based Multiobjective Optimization Techniques”. ACM Computing Surveys, 32(2), 2000, 109-143.

[5] M. Ericsson, M. Resende, P. Pardalos, “A genetic algorithm for the weight setting problem in OSPF routing”. Technical report, ATT Shannon Laboratory (2001).

[6] L. Buriol, M. Resende, C. Ribeiro, M. Thorup, “A memetic algorithms for OSPF routing”. In Proc. of 6th INFORMS Telecom (2002) 187-188

[7] K. Ho, N. Wang, P. Trimintzios, G. Pavlou, M. Howarth, "On Egress Router Selection for Inter-domain Traffic with Bandwidth Guarantees". In Proc. of IEEE Workshop in High Performance Switching and Routing (HPSR'2004), Phoenix, Arizona, USA (2004)

[8] T. Bressoud, R. Rastogi, M. Smith, "Optimal Configuration for BGP Route Selection". In Proceedings of IEEE INFOCOM' 2003, San Francisco (2003)

[9] M. Pioro, and D. Medhi, "Routing, Flow, and Capacity Design in Communication and Computer Networks". Morgan Kaufmann Series in Networking (2004)

[10] S. Uhlig, "A multiple-objectives evolutionary perspective to interdomain traffic engineering". International Journal of Computational Intelligence and Applications, Special Issue on Nature-Inspired Approaches to Telecommunications, World Scientific Publisher (2005)

[11] MESCAL project [website]: www.mescal.org

[12] K. Deb, A. Pratap., S. Agarwal, and T. Meyarivan, "A fast and elitist multi-objective genetic algorithm: NSGA-II". IEEE Transaction on Evolutionary Computation, 6(2), (2002). 181-197