

Comparing Traffic Engineering Objective Functions

Simon Balon^{*}, Fabian Skivée and Guy Leduc
Research Unit in Networking
University of Liège

{balon, skivee, leduc}@run.montefiore.ulg.ac.be

ABSTRACT

We compare and evaluate how well-known and novel network-wide objective functions for Traffic Engineering (TE) algorithms fulfil TE requirements. To compare the objective functions we model the TE problem as a linear program and solve it to optimality, thus finding for each objective function the best possible target of any heuristic TE algorithm. Our first results on the GÉANT network suggest that they can give quite different results.

Categories and Subject Descriptors: C.2.3 [Network Operations]: Network management

General Terms: Algorithms, Management, Performance

Keywords: Traffic Engineering, Objective function

1. INTRODUCTION

We consider the traffic engineering routing problem. We define this problem as follows. Given the topology of the network to be engineered and an estimate of the traffic matrix to be routed on it, the problem is to find a routing scheme that optimises the network, with the joint goal of good user performance and efficient use of network resources. The way classical algorithms fulfil this objective is not clear. Indeed, many algorithms try to minimise (or maximise) their home-made objective functions which are said (but not proven) to reflect traffic engineering objectives. The foundations of all these objective functions are related, but could lead to quite different results, as we see in our simulation.

Some in-depth reflection and comparison studies of traffic engineering objective functions are needed. To our knowledge, no such study exists. In many research papers, the quality of a new traffic engineering algorithm is evaluated regarding one specific objective function. If the algorithm obtains a good score, it is considered as good. But this is only meaningful if the objective function really reflects the traffic engineering goals. Furthermore, when analysing published papers it is difficult to figure out if the merits of a given TE method is due to its objective function or its heuristic algorithm. To fill this gap, we provide an objective comparison of many objective functions found in the literature.

^{*}Research Fellow of the Belgian National Fund for the Scientific Research (F.N.R.S)

2. CLASSICAL OBJECTIVE FUNCTIONS

We will first introduce some notations. A network is modelled by a directed graph, $G = (N, A)$ whose nodes and arcs represent routers and links. Each arc has a capacity $c(a)$. Traffic on the network is represented by a traffic matrix D whose component (s,d) is the value of the traffic demand from node s to node d .

Basically, the graph G and the traffic matrix D are the inputs of the problem. A traffic engineering algorithm has to find *good* paths between each pair of source and destination nodes to route corresponding traffic flows. Once the paths are chosen, we can associate with each arc a load $l(a)$ which is the total load on the arc, i.e. the sum over all demands of the amount of traffic sent over a . The utilisation of a link a is $u(a) = l(a)/c(a)$.

Typically, on-line algorithms have different objectives from off-line ones. On-line schemes usually try to minimise the probability of blocking future requests, while off-line ones try to minimise the load or the utilisation of the links, or try to maximise available bandwidth. To some extent, minimising the link utilisation (which is a relative measure) tends to maximise the available bandwidth (which is an absolute measure) on the links, thus also reducing the blocking probability of future requests. Clearly, these objectives are closely related, but no solid basis exists to choose one among all.

In [4], B. Fortz et al. try to find an optimal set of IGP weights so that classical shortest paths taking these modified metrics in consideration lead to a good routing scheme. A cost is associated with each link of the network. This cost (ϕ_a) is a piecewise linear function of the link load. The objective function they try to minimise is the sum over all links of this cost ($\phi = \sum_{a \in A} \phi_a$). *The idea behind ϕ_a is that it is cheap to send flow over an arc with a small utilisation. As the utilisation approaches 100%, it becomes more expensive, for example because we get more sensitive to bursts. (...) when the utilisation goes above 100% the penalty gets so high that this should never happen.*

In [5], Kodialam et al. introduce the concept of minimum interference routing. They propose an objective function which is a sum of the maxflow over all possible source-destination pairs on the residual topology: $\sum_{(s,t)} \theta_{st}$, where θ_{st} represent the maximum flow that can be sent from node s to node t in the residual network. Their on-line algorithm, called MIRA, is a heuristic that tries to maximise this objective function. Behind this function, the goal is to minimise the blocking probability of future requests, without information about them. If the maxflow between one source and one destination decreases, this means that the maximum request

that can be accepted between these two nodes decreases as well.

In [1], Blanchy et al. present an on-line traffic engineering algorithm to balance the load in the network. They propose the following objective function which is the sum of a load-balancing term and a shortest path term : $\sum_{a \in A} (u(a) - u_{mean})^2 + \alpha \sum_{a \in A} (u(a))^2$. It is interesting because *the combination of both terms will give more importance to the load-balancing term if the deviation is high enough to justify the detour, else it will let the "shortest path" term minimise the resources used. The weighted factor α allows to give more importance to one aspect or to the other.*

In [3], Elwalid et al. associate a cost with each link. They try to minimise the total cost which is the sum over all links of the link cost. They say that a natural choice for the link cost is the delay. If we take the delay to be the average delay of an M/M/1 queue, the cost of link a is assumed to be $Delay_a = \frac{1}{c(a)-l(a)}$. The network-wide cost function is thus defined as $Delay_1 = \sum_{a \in A} \frac{1}{c(a)-l(a)}$.

We introduce a new delay objective function: $Delay_2 = \sum_{a \in A} \frac{l(a)}{c(a)-l(a)}$. In this function, the delay of a link is weighted by the load of this link. This objective function reflects the average queuing + transmission delay of all the traffic of the network.

In [2], Degrande et al. propose to maximise an objective function which is composed of four terms : $C_F.F + C_T.T + C_B.B - C_U.U$ with F (airness), T (throughput), B (alance) and U (tilisation) the four performance parameters to be optimised and C_F , C_T , C_B and C_U the objective coefficients. We will consider only Balance and Utilisation which are resource oriented objectives. Balance is defined as 1 minus the maximum link utilisation: $B = 1 - u_{max}$. Network utilisation is defined as the sum of all link utilisations: $U = \sum_{a \in A} u(a)$.

3. SIMULATION

To compare all the objective functions, we will model the traffic engineering problem as a linear program (LP). In this formulation, all the flows can be arbitrarily split. For non-linear objective functions, we take a linear approximation. The simulation consists of solving to optimality the routing problem as defined in section 1, and this for all the presented objective functions. We made our simulations on the European research network (GÉANT). This network is composed of 23 routers and 38 bidirectional links. To build a realistic traffic matrix, we have collected netflow data on each interface of the network. We have aggregated this information to build a traffic matrix.

We will now present the variables we use to evaluate the quality of the routing solutions. Clearly, the maximum link utilisation (u_{max}) should be minimised. We also use the mean link utilisation (u_{mean}) and load (l_{mean}), and the standard deviation of utilisation (u_{stdv}). An important variable is also the mean link queuing + transmission delay of the network (del_{mean}). To reflect the residual available bandwidth of the network, we will also consider the total max flow ($\theta_{tot} = \sum_{(s,t)} \theta_{st}$) of the residual topology.

Table 1 presents the results of the problem we consider. $Degr_{(X,Y)}$ means Degrande's function with $C_B=X, C_U=Y$, $Blanchy_{(X)}$ means Blanchy's function with $\alpha=X$. We can see that all the objective functions are not equivalent. The lowest achievable value for u_{max} is given by $Degr_{(1,0)}$ which only optimises this variable. But the other variables for

Name	u_{max} %	u_{mean} %	u_{stdv} %	l_{mean} Mbps	del_{mean} $\frac{1}{Tbps}$	θ_{tot} Tbps
<i>Fortz</i>	45.2	7.5	9.8	270	0.77	4305
<i>MIRA</i>	58.4	7.0	10.0	280	0.75	4331
<i>Blanchy</i> ₍₀₎	<i>100.0</i>	<i>87.5</i>	3.8	<i>4837</i>	∞	<i>565</i>
<i>Blanchy</i> ₍₁₎	44.3	7.6	7.8	333	0.73	<i>4275</i>
<i>Blanchy</i> ₍₃₎	44.6	7.3	8.1	312	0.71	4297
<i>Delay</i> ₁	45.5	7.1	9.5	319	0.72	4303
<i>Delay</i> ₂	45.5	6.9	9.6	276	0.71	4323
<i>Degr</i> _(0,1)	46.5	6.9	9.7	274	0.71	4326
<i>Degr</i> _(1,0)	38.4	12.2	11.5	<i>494</i>	<i>1.29</i>	<i>4067</i>
<i>Degr</i> _(1,1)	38.4	6.9	9.7	274	0.71	4326

Table 1: Results on GÉANT network

this function are not very good. So we do not think that optimising only the max utilisation is a good idea. The lowest achievable value of the u_{mean} variable is given by the $Degr_{(0,1)}$. The best value of u_{stdv} is given by $Blanchy_{(0)}$. As we could predict other variables for the $Blanchy_{(0)}$ function are bad, because this function does not include traffic minimisation. The best value of the l_{mean} variable is given by $Fortz$. Many objective functions give good results for the del_{mean} variable. Bad objectives functions concerning this criterion are $Blanchy_{(0)}$ and $Degr_{(1,0)}$. $MIRA$ gives the optimum for the θ_{tot} variable. We notice that many other objective functions give values close to this optimum.

All the best values for each criterion are in bold and all the bad values are in italic. Considering these criteria on this topology, we can consider $Degr_{(1,1)}$ as the best objective function as it gives three best values and no bad value. But more tests should be performed in order to generalise this result.

4. CONCLUSION

We have shown how well-known network-wide objective functions reflect the requirements for Traffic Engineering. Our first results suggest that these are quite different and highlight the strength of some functions and the weakness of others. This preliminary study provides a first step towards an objective comparison of TE objective functions, allowing to choose the best one when designing a new TE algorithm.

Acknowledgment

This work has been partially supported by the Walloon Region (TOTEM project) and the European E-NEXT NoE.

5. REFERENCES

- [1] F. Blanchy, L. Mélon, and G. Leduc. An efficient decentralized on-line traffic engineering algorithm for MPLS networks. *Proc. of 18th ITC, 2003, Berlin, Germany*, 2003.
- [2] N. Degrande, G. V. Hoey, P. de La Vallée-Poussin, and S. V. den Busch. Inter-area traffic engineering in a differentiated services network. *J. Networks Syst. Manage.*, 11(4), 2003.
- [3] A. Elwalid, C. Jin, S. H. Low, and I. Widjaja. MATE: MPLS adaptive traffic engineering. In *INFOCOM*, 2001.
- [4] B. Fortz and M. Thorup. Internet traffic engineering by optimizing OSPF weights. In *INFOCOM*, 2000.
- [5] M. S. Kodialam and T. V. Lakshman. Minimum interference routing with applications to MPLS traffic engineering. In *INFOCOM*, 2000.