Routing Games for Traffic Engineering

Federico Larroca
TELECOM ParisTech
Paris, France
federico.larroca@telecom-paristech.fr

Jean-Louis Rougier
TELECOM ParisTech
Paris, France
rougier@telecom-paristech.fr

Abstract—Current data network scenario makes Traffic Engineering (TE) a very challenging task. The ever growing access rates and new applications running on end-hosts result in more variable and unpredictable traffic patterns. By providing origin-destination pairs with several possible paths, load-balancing has proved itself an excellent tool to face this uncertainty. In particular, mechanisms where routers greedily minimize a path cost function (thus requiring minimum coordination) have been studied from a game-theoretic perspective in what is known as a *Routing Game* (RG). The contribution of this paper is twofold. We first propose a new RG specifically designed for elastic traffic, where we maximize the total utility through load-balancing only. Secondly, we consider several important RGs from a TE perspective and, using several real topologies and traffic demands, present a thorough comparison of their performance. This paper brings insight into several RGs, which will help one in choosing an adequate dynamic load-balancing mechanism. The comparison shows that the performance gain of the proposed game can be important.

Index Terms—Traffic Engineering, Routing Games, Wardrop Equilibrium, Load balancing.

I. Introduction

As network services and Internet applications evolve, the traffic is becoming increasingly complex and dynamic. The convergence of data, telephony and television services on an all-IP network as well as user-mobility (which implies service-mobility) directly translates into a much higher variability and complexity of the traffic injected into the network. Moreover, the ever increasing access rates has seriously questioned overprovisioning as an economically viable solution. To cope with both the traffic increasing dynamism and the need for cost-effective solutions, a self-managing network architecture is required.

Dynamic load-balancing is emerging as a possible answer to the above issues [1]–[3]. In particular, mechanisms where each router greedily minimizes a certain cost function of its paths (thus requiring minimum coordination) have been intensively studied. This context constitutes an ideal case study for game theory, and is known as *Routing Game* in its lingo [4]. Different routing games have been defined, each with its specific path cost, and their equilibrium have been characterized. This characterization includes existence, uniqueness and optimality. With respect to the last aspect, different social cost functions (i.e. the cost of a given situation to all players as a whole) have been considered, and possible modifications to the path cost so that the greedy equilibrium and the social optimum coincide have been studied.

The first contribution of this paper is to propose a routing game specifically designed for elastic traffic. In this proposition, we aim at maximizing the total utility obtained by flows without changing their congestion control mechanism. This maximization is achieved instead by changing the amount of traffic routed through each path, thus changing the flows' mean rate. We also present a possible way of achieving the corresponding social optimum.

As a second contribution we provide a thorough comparison of three routing games (including our proposition), using two real topologies and several real traffic demands. The considered routing games are variants of the two most important ones: Congestion [4] and Bottleneck Games [5]. In particular, we analyze the social optimum of the three games, achieved by greedy users minimizing a modified path cost, measuring several important performance indicators for the three cases. This is an important study which will help network operators at the moment of choosing a possible dynamic load-balancing mechanism. As we shall see, our proposition is the one that obtains the best results.

The rest of the paper is organized as follows. In the following section we make the necessary definitions. In section III we describe the two other considered games. Our proposal is presented in section IV. The comparison between the games is left for section V. Finally, we conclude the paper in section VI.

II. Definitions

A. Network Model

The network is defined as a graph $G = (V, E)$. In the network there are a number of so-called *commodities*, indexed by $s = 1, .., S$, specified in terms of the triplet o_s, q_s and d_s; i.e. origin node, destination node and a certain fixed demand of traffic from the former to the latter. Each commodity s can use n_s paths connecting o_s to q_s (each noted as P_{si} for $i = 1, .., n_s$), and can distribute its total demand arbitrarily among them. Commodity s sends an amount d_{si} of its traffic through path P_{si}, where $d_{si} \geq 0$ and $\sum d_{si} = d_s$. This distribution of traffic induces the demand vector $d = (d_{si})$.

Given the demand vector, the total load on link l is then $\rho_l = \sum_s \sum_{i:l \in P_{si}} d_{si}$. The presence of this traffic on the link induces a certain "delay" given by the non-decreasing function $f_l(\rho_l)$. The total delay of path P (or its cost) is defined as $f_P = g(\{f_l(\rho_l)\}_{l \in P})$. The choice of g defines what kind of routing game we are considering. The case of $g(A) = \sum_{a \in A} a$

is called a *Congestion Game*, while $g(A) = \max_{a \in A} a$ results in a *Bottleneck Game*.

The function $f_l(\rho_l)$ does not necessarily represent the delay incurred by ρ_l. It could for instance give the cost for the network operator to route this amount of traffic through link l, or represent any other link-level performance measurement. It should be clear that the choice of $f_l(\rho_l)$ and what it represents almost dictates the choice of g.

Note that no explicit constraint on ρ_l was made. This is assumed to be implicitly present in the delay function. For instance, $f_l(\rho_l)$ goes to infinity as ρ_l reaches c_l (the link's capacity) and remains at infinity after this point.

B. Wardrop Equilibrium

The concept of the Wardrop Equilibrium was introduced in the context of transportation [6]. In it, commodities are assumed to be constituted by infinitely many agents, each controlling an infinitesimal amount of the demand. Each of these agents (or players) decides through which path to send its traffic. In this context the division d_{si}/d_s represents the portion of agents of commodity s that have P_{si} as their choice. If every agent acts selfishly, then the system will be at equilibrium when no agent can improve its delay by changing its path choice. Formally, a Wardrop Equilibrium is defined as follows.

Definition 1: A demand vector is a Wardrop Equilibria if for each commodity $s = 1 \ldots S$ and for each path P_{si} with $d_{si} > 0$ it holds that $f_{P_{si}} \leq f_{P_{sj}}$ for all P_{sj} with $j = 1, .., n_s$.

Once the functions $f_l(\rho_l)$ and g are defined, a *Social Cost* function ($SC(d)$) has to be chosen. This function measures the dissatisfaction of the commodities as a whole, and an optimum demand vector is one that minimizes it over all valid demands. The most typical example is the mean delay, defined as $\sum_{si} d_{si} f_{P_{si}}$. However, as we shall discuss, the optimum of the social cost function and the Wardrop Equilibrium are generally not the same.

III. RELATED WORK

A. Congestion Games

If each commodity generates packets of exponentially distributed size (with mean size equal to 1) as a Poisson process of intensity d_s, the network can be analyzed as a Jackson Network. The delay in each link can easily be calculated as $f_l(\rho_l) = 1/(c_l - \rho_l)$ [7]. This delay function, generally referred to as the *M/M/1 delay function*, has been used and studied in the past, both from a game theory perspective [8], as well as a Traffic Engineering one [1]. The social cost function in this kind of games is generally the mean delay:

$$SC(d) = \sum_{s=1}^{S} \sum_{i=1}^{n_s} d_{si} f_{P_{si}} = \sum_{l=1}^{L} \rho_l f_l(\rho_l) = \sum_{l=1}^{L} \frac{\rho_l}{c_l - \rho_l}$$

If each agent acts selfishly and seeks to minimize the delay he experiences, the resulting Wardrop Equilibrium will not minimize the mean delay, but will result in a local minimum of the so-called potential function [4]:

$$\Phi(d) = \sum_{l=1}^{L} \int_0^{\rho_l} f_l(x) dx \tag{1}$$

This means that, given a social cost function, we may modify the delay function so that the resulting Wardrop Equilibrium is the optimum demand vector. In this case, let us consider the following alternative delay function:

$$\widehat{f}_l(\rho_l) = \frac{c_l}{(c_l - \rho_l)^2} \tag{2}$$

In view of (1) the Wardrop Equilibrium of this new game (which we shall note as CG in the sequel) is a local minimum of:

$$\Phi(d) = \sum_{l=1}^{L} \int_0^{\rho_l} \frac{c_l}{(c_l - x)^2} dx = \sum_{l=1}^{L} \frac{\rho_l}{c_l - \rho_l} = SC(d)$$

which corresponds to the social optimum.

B. Bottleneck Games

In traffic engineering, a very important link-level performance indicator is the link utilization, defined as $u_l = \rho_l/c_l$. A link with a u_l close to one is a congested link with significant losses and queueing delays. A network-wide performance indicator is the maximum link utilization, which is arguably the most popular TE objective function ([2], [3] to name a few).

This is a natural context for a bottleneck game. The link and path delay are then defined as (we shall note the resulting game as BG):

$$f_l(\rho_l) = u_l = \frac{\rho_l}{c_l} \quad \text{and} \quad f_P = \max_{l \in P} \left\{ \frac{\rho_l}{c_l} \right\}$$

And the social cost function is simply the largest utilization:

$$SC(d) = \max_{P \in \mathcal{P}} \{f_P\} = \max_l \left\{ \frac{\rho_l}{c_l} \right\}$$

The authors of [5] studied general bottleneck games where the social cost function is the worst delay in the network, and showed that conditions exist under which the Wardrop Equilibrium is socially optimum. From a pure traffic engineering point of view, [2] previously proposed a greedy algorithm that converges to the social optimum.

IV. A NEW ROUTING GAME FOR ELASTIC TRAFFIC

A. Definition

Assume, for now, that demands are constituted of a fixed number of elastic or TCP flows (N_s flows for commodity s). If there are N_{si} flows in path P_{si}, the congestion control problem can be written as follows [9]:

$$\begin{aligned} \underset{x}{\text{maximize}} \quad & \sum_{s=1}^{S} \sum_{i=1}^{n_s} N_{si} U_{si}(x_{si}) \tag{3} \\ \text{subject to} \quad & \sum_{s} \sum_{i:l \in P_{si}} N_{si} x_{si} \leq c_l \end{aligned}$$

Where x_{si} is the rate obtained by each of the N_{si} flows and $U_{si}(x)$ is an non-decreasing, concave and continuous function. The above problem optimizes in the obtained rate, considering the N_{si}'s (i.e. routing) as given. However, to improve performance, we could jointly maximize in both N and x (adding the constraints $\sum N_{si} = N_s$ and $N_{si} \geq 0$). To implement such maximization, a first idea is that end-users decide their path. However, network operators generally do not share routing information with the end-users. On the other hand, path diversity is commonly used within a routing domain, where an ingress router can split traffic on a set of equal cost paths or disjoint MPLS tunnels. We thus propose to keep the separation between end-to-end congestion control (maximizing on x done by end-users) and routing (maximizing on N) but still try to solve (3). Some important aspects should be considered which we shall discuss in the following paragraphs.

The first obvious aspect is that routers do not know $U_{si}(x)$, and even if they knew it, we as network operators may not "like" it[1]. For this reason, we shall use an arbitrary $U(x)$ function which we believe convenient to operate the network (in our simulations $U(x) = \log(x)$).

The second aspect has to do with time-scales. Congestion control acts at the RTT timescale (generally in the order of ms) while routing does it at the seconds or minutes timescale. This means that we cannot consider the instantaneous values of x_{si}, but its temporal mean. Moreover, TCP flows have a certain finite life-time, thus N_s cannot be considered as static or given and a substitute should be used. We will first introduce the dynamic traffic model we use for this context, and then we will discuss how to address these issues.

We will suppose that each origin node o_s generates flows as a Poisson process of intensity λ_s. Each of these elastic flows have a random arbitrarily distributed workload (with mean ω_s) it has to transfer. After the transmission is finished, the flow disappears. Each flow is routed through path P_{si} with probability p_{si}, and it uses it throughout his lifetime. Under these conditions, the demand of the corresponding commodity can be written as $d_s = \lambda_s \omega_s$, and the demand vector as $d = (d_s p_{si})$.

The above described system has been thoroughly studied in the past. The mean throughput obtained by the flows traversing path P can be roughly approximated by the path's Available Bandwidth [10], [11]: $ABW_{P_{si}} = \min_{l \in P_{si}} \{c_l - \rho_l\} \approx x_{si}$. We shall then approximate x_{si} by $ABW_{P_{si}}$ in (3). The natural substitute of N_{si} is d_{si}, who plays the role of the amount of traffic using path P_{si}. This routing approximation to (3) results:

$$\underset{d}{\text{maximize}} \sum_{s=1}^{S} \sum_{i=1}^{n_s} d_{si} U \left(\min_{l \in P_{si}} \{c_l - \rho_l\} \right)$$

[1]For instance, we may consider that the bias against RTT of TCP Reno should not be re-enforced by routing.

Which is equivalent to:

$$\underset{d}{\text{minimize}} \sum_{s=1}^{S} \sum_{i=1}^{n_s} d_{si} \max_{l \in P_{si}} \{-U(c_l - \rho_l)\}$$

$$\text{subject to} \quad d_{si} \geq 0 \quad \sum_{i=1}^{n_s} d_{si} = d_s$$

This last problem may be seen as the following bottleneck game:

$$f_l(\rho_l) = -U(c_l - \rho_l) \qquad f_P = \max_{l \in P} \{f_l(\rho_l)\}$$

$$SC(d) = \sum_{s=1}^{S} \sum_{i=1}^{n_s} d_{si} f_{P_{si}} \tag{4}$$

However, as we mentioned in the last subsection, the Wardrop Equilibrium of this routing game will not necessarily coincide with the social optimum. In the following subsection we define an alternative game to achieve this optimum (as in subsection III-A).

B. Achieving the Social Optimum

In this section we present a mechanism to achieve the social optimum of the bottleneck game specified in (4). Minimization of the social cost function can be formulated by introducing the auxiliary variable t_{si} (which at optimality will be equal to $f_{P_{si}}$ and thus may be seen as a continuous function of d):

$$\underset{d,t}{\text{minimize}} \sum_{s=1}^{S} \sum_{i=1}^{n_s} d_{si} t_{si} \tag{5}$$

$$\text{subject to} \quad d_{si} \geq 0 \quad \sum_{i=1}^{n_s} d_{si} = d_s$$

$$t_{si} \geq -U(c_l - \rho_l) \, \forall s, i \, \forall l : l \in P_{si}$$

Necessary conditions on a local minimum can be obtained from the KKT conditions [12]. The Lagrangian function associated to problem (5) may be written as:

$$L(d, t, \nu, \lambda, \theta) = \sum_{s=1}^{S} \sum_{i=1}^{n_s} d_{si} t_{si} + \sum_{s=1}^{S} \nu_s \left(\sum_{i=1}^{n_s} d_{si} - d_s \right)$$

$$- \sum_{s=1}^{S} \sum_{i=1}^{n_s} \lambda_{si} d_{si} - \sum_{s=1}^{S} \sum_{i=1}^{n_s} \sum_{l:l \in P_{si}} \theta_{sil} (t_{si} + U(c_l - \rho_l))$$

Let d^* be a local optimum of problem (5). The KKT conditions state that there exist unique Lagrange multiplier vectors ν^*, $\lambda^* \geq 0$ and $\theta^* \geq 0$ such that:

$$t_{si}^* + \nu_s^* - \lambda_{si}^* + \sum_{l:l \in P_{si}} \widehat{\theta}_l = 0 \tag{6}$$

$$d_{si}^* - \sum_{l:l \in P_{si}} \theta_{sil}^* = 0 \tag{7}$$

$$\lambda_{si}^* d_{si}^* = 0 \tag{8}$$

$$\theta_{sil}^* (t_{si}^* + U(c_l - \rho_l^*)) = 0 \tag{9}$$

where

$$\widehat{\theta}_l = \sum_{s=1}^{S} \sum_{i:l \in P_{si}} \theta_{sil}^* U'(c_l - \rho_l^*) \tag{10}$$

Putting equation (6) and (8) together yields:

$$\widehat{f}_{P_{si}}^* = f_{P_{si}}^* + \sum_{l:l \in P_{si}} \widehat{\theta}_l = \begin{cases} -\nu_s^* & \text{if } d_{si}^* > 0. \\ -\nu_s^* + \lambda_{si}^* & \text{if } d_{si}^* = 0 \end{cases}$$

Then, for a given commodity s, all paths that are used at optimality have the same modified delay $\widehat{f}_{P_{si}}$ ($-\nu_s$ does not change over paths), which is smaller or equal than that of paths that are not used (remember that $\lambda_{si} \geq 0$). This means that a necessary condition for a demand vector to be a local minimum of (5) is to constitute a Wardrop Equilibrium of a routing game with path delay $\widehat{f}_P = f_P + \sum_{l:l \in P} \widehat{\theta}_l$.

Equation (9) indicates that θ_{sil}^* is zero for all links that are not the bottleneck of path P_{si}. If we assume that paths have only one bottleneck, and using equation (7), θ_{sil} is simply d_{si} in P_{si}'s bottleneck link and 0 elsewhere. We shall note UM (as in Utility Maximization) the routing game that uses $\widehat{f}_P$ as its path delay.

As we mentioned, the optimal value of t_{si} is a continuous function of d. So, problem (5) may be seen as minimizing a continuous function over a compact set. Its minimum, which given the previous observation is always achieved, is necessarily a Wardrop Equilibrium of the modified game, guaranteeing the existence of the latter. However, a Wardrop Equilibrium is not necessarily a minimum. For this last aspect, sufficient conditions may be derived that although difficult to prove in the general case, are verified in all the examples we evaluated (see Proposition 3.3.2 in [12]).

V. EVALUATION

In this section we present a comparative study between the three routing games, where each agent greedily chooses the path with the smallest modified delay function, so that the Wardrop Equilibrium results in the social optimum. The equilibrium was obtained through the greedy load-balancing algorithm described in [3]. The comparison will be made in two real networks, with several real demands, calculating for each of these demands two performance indicators: ABW_P and link utilization (u_l), whose importance we have already discussed. We could consider other performance indicators, such as queueing delay or path's propagation delay. However, calculation of the former depends heavily on the assumed traffic model, and we shall suppose that the latter has already been taken into account by the operator in the paths' choice.

Before presenting the performance analysis, we will overview three examples that illustrate the differences between the three games and help us get some insight and intuition.

A. Illustrative Examples

The first example we will consider is the simplest: a single commodity with two paths of the same capacity. However, one of the paths is constituted of two links, while the other
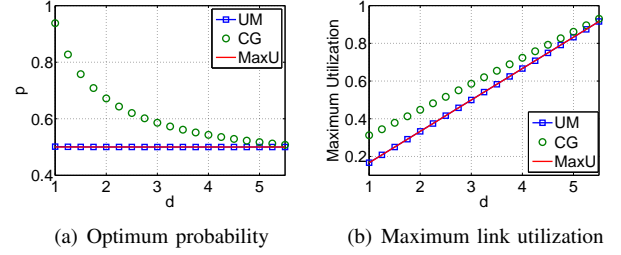

(a) Optimum probability   (b) Maximum link utilization

Fig. 1. The single-commodity two-paths case (a longer path subcase).

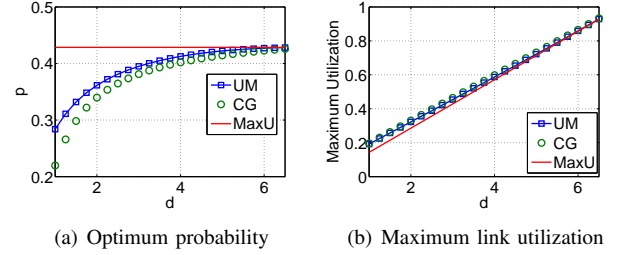(a) Optimum probability   (b) Maximum link utilization

Fig. 2. The single-commodity two-paths case (a narrower path subcase).

of only one. In particular, links' capacity will be 3.0. In Fig. 1(a) we can see p (the optimum traffic portion routed through the shortest path) as a function of d (the demand generated by the commodity) for the three different games: CG, BG and UM. Naturally, for both UM and BG the optimum is always 0.5. However, in light loads CG uses the shortest path almost exclusively, and as load increases it will need to use the longest one to avoid congesting links. Although this last behavior could seem more reasonable, if the network operator configured this paths is because they are relatively equivalent. Moreover, it significantly increases the maximum link utilization, as can be seen in Fig. 1(b).

The next example is the same as the previous one, except that both paths have the same length and a capacity of 3.0 and 4.0. In Fig. 2(a) we show p (the optimum traffic portion routed through the narrowest path) as a function of d for the three different games. Clearly p is always $3/7$ for BG. However, p changes with d for both the other two games this time. If the total demand d is small enough, the narrowest path is left unused since the obtained performance is inferior. As it can be seen in Fig. 2(b) the difference in the maximum utilization is not as significant as before.

The third example illustrates some fairness issues that are important to highlight. In Fig. 3(a) we can see the considered network. In it, all link capacities are equal and all commodities have the same demand d_0, except for commodity 1 that generates d and is the only one to have more than one path to choose from. We will consider $c_l = 5.0 \; \forall l$, $d_0 = 2$ and we will study the optimum p as we vary d. It is relatively simple to verify that the optimum for CG and BG is $p = 0.5$ independently of d. On the other hand, UM enforces fairness at a path level. This means that commodity 1 takes into account that the upper path "disturbs" two other paths while the lower one disturbs only one, resulting in more traffic being sent
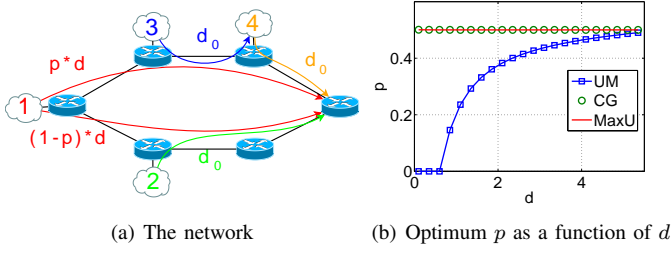
(a) The network   (b) Optimum p as a function of d

Fig. 3. An example illustrating the fairness issue.

(a) UM vs CG   (b) UM vs BG

Fig. 4. ABW_P for UM, CG and BG in the Abilene network

through the latter (see Fig. 3(b)). So, while d is relatively small and the *ABW* is enough, commodity 1 uses only the lower path. If any of these conditions is not true, p will rapidly go to 0.5, but always privileging better conditions on the upper path.

B. Abilene Network

In this section we present the performance analysis for the Abilene [13] network, a well-known academic network which consists of 12 nodes and 15 bidirectional links all with the same capacity. The topology comes as an example in the TOTEM toolbox [14] and we used 388 demands (spanning a complete week) from dataset X11 of [15]. The paths we used in this case were constructed by hand, trying to give commodities as much path diversity as possible, but limiting the hop count.

We will first present the path's ABW (ABW_P) results. For each demand we calculated the weighted mean ABW_P, where the weight of path P_{si} is d_{si}. This average provides us with a rough idea of the performance as perceived by traffic. A good value of this average indicator could however hide some pathological cases where some portions of traffic obtain a bad performance. That is why we also measured the 10% quantile and the minimum ABW_P. The comparison will be made by dividing the value obtained by UM by the one obtained by the other games.

In Fig. 4 we can see the boxplots of the ABW_P indicators. We first note that the weighted mean of ABW_P is always bigger in UM than in CG and BG, being generally between 1-2% and at most 5%. No conclusive results can be obtained from the quantile ABW_P. In the minimum ABW_P the results are clearer. UM achieves a minimum ABW_P that is generally between 6-13% (and can be as big as 18%) bigger than CG. As expected, the best results in this aspect are obtained by BG, although its gain over UM is not so important.

We now turn our attention to the link utilization (u_l) results. For each demand we calculated the mean, 90% quantile and maximum utilization on the network for each of the considered games. The comparison will be made by making the difference between the indicator obtained by BG and the other two games. Fig. 5 shows these results. Both the mean and the quantile do not present any substantial difference between the three games. It is in the maximum utilization that we can see a clearer distinction between them, where as expected BG always obtains the best results. However, and in concordance
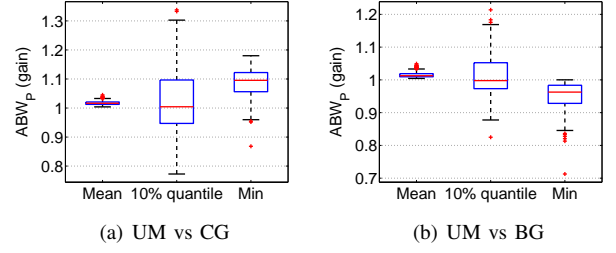
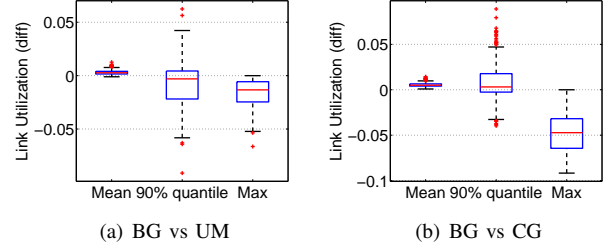

(a) BG vs UM   (b) BG vs CG

Fig. 5. u_l for BG, UM and CG in the Abilene network

with the ABW_P indicators, its gain over UM is smaller than over CG, the former being generally 1-2% and the latter between 3-7%. This results are related with the first example considered in section V-A, where the maximum link utilization was bigger for CG due to its reluctance to use longer paths.

C. Géant Network

The second case scenario is Géant [16]. This academic network connects 23 nodes using 74 unidirectional links, with capacities that range from 155 Mbps to 10 Gbps. The topology and demands (477 in total, covering a three week period) were obtained from TOTEM's webpage [14], [17]. In this case paths were constructed by a shortest path algorithm, where we used the inverse of the capacity as the link's weight. For each commodity we computed two paths. The first is simply the shortest path, we then prune the network of the links this path uses, and compute the second shortest path.

Results for the ABW_P in this case can be seen in Fig. 6. This time, the results of both UM and CG are more similar, where the mean and quantile ABW_P is somewhat bigger for UM than CG, and the minimum is relatively bigger for CG than UM. However, the results of the comparison between UM and BG are clearly in favor of the former. The mean ABW_P is generally 6-7% bigger, going as high as 11%. Good results are also obtained on the quantile. With respect to the minimum ABW_P, the results are logically better for BG, but the difference is not significant.

Fig. 7 show that the results for the link utilization are also very similar between UM and CG. The difference between the two and BG is not very significant, specially in the mean and quantile. With respect to the maximum utilization, BG obtains only subtle improvements over the rest. As we illustrated in the second example of section V-A, when capacities are different and paths are similar in length (as is generally the case for
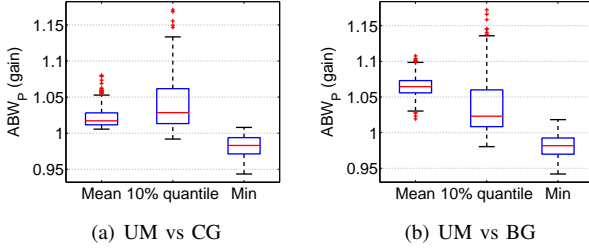
(a) UM vs CG   (b) UM vs BG

Fig. 6. ABW_P for UM, CG and BG in the Géant network
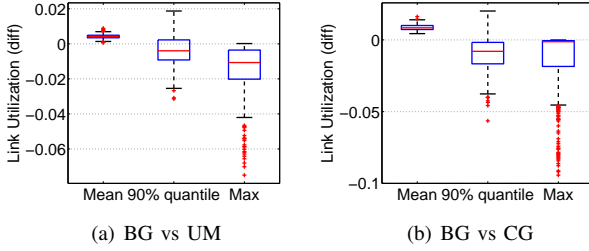

(a) BG vs UM   (b) BG vs CG

Fig. 7. u_l for BG, UM and CG in the Géant network

Géant) the difference in maximum utilization between the games is small.

VI. CONCLUDING REMARKS

The contribution of this paper is twofold. We first proposed and studied a routing bottleneck game, which defined the link's delay as a utility function of the available bandwidth. Such routing game was designed with elastic flows in mind, with the idea that TCP is in charge of the path's resource sharing, and routing continue to maximize utility by indirectly changing flow's obtained rate through their path choice. We then compared three routing games (including ours) from a traffic engineering perspective, a study that we believe will be helpful for network operators at the time of choosing a dynamic traffic engineering mechanism.

Along with our game (noted as UM), we considered a classical congestion game where the link's delay is a modified *M/M/1 delay function* (CG), and a bottleneck game where link's delay is its utilization (BG). From our study, conducted over two real networks along with several real traffic demands, some conclusions can be drawn. Firstly, performance as perceived by traffic (measured as the ABW) is always better in UM than both CG and BG. More specifically, the improvement over CG is generally not very big in mean, but may be important, specially in the worst ABW_P. This difference is originated in the implicit unfairness among commodities of the social cost function of CG. With respect to BG, UM obtains significantly better performance when links' capacity are not similar. Secondly, results on link utilization are very similar between UM and BG games. CG obtains similar results in the mean and quantile link utilization. However, due to its reticence to use longer paths, maximum link utilization can be significantly bigger in CG than the rest of the games.

All in all, it seems like UM is the most balanced game, in the sense that it generally outperforms the rest, and when it does not the difference is not important. Unfortunately, it comes with a price. The modified routing game through which the social optimum was obtained, requires to distinguish and measure the traffic generated by each commodity (remember that in section IV-B we defined θ_{sil} as either d_{si} or 0). Moreover, a mechanism so that a link can learn whether it constitutes a path's bottleneck has to be designed (θ_{sil} is not zero only if it is the bottleneck of path P_{si}). Although such mechanisms may be difficult to implement in a packet switched environment, they are relatively easy in an MPLS network (where, for instance, measuring d_{si} is straightforward).

We have shown that the Wardrop Equilibrium of the proposed game coincides with the minimum of the social cost for all the cases we considered. However, a proof in the general case constitutes an interesting and necessary future work. Another interesting question is what happens when we are interested in queueing delays. CG has been proposed as an answer, but it relies heavily on the assumed model. We are currently investigating how to design a measurement-based load balancing mechanism that converges to the social optimum, but where the delay function is not assumed a priori.

REFERENCES

[1] A. Elwalid, C. Jin, S. Low, and I. Widjaja, "MATE: MPLS adaptive traffic engineering," *INFOCOM 2001*, vol. 3, pp. 1300–1309, 2001.
[2] S. Kandula, D. Katabi, B. Davie, and A. Charny, "Walking the tightrope: responsive yet stable traffic engineering," in *ACM SIGCOMM '05*, 2005, pp. 253–264.
[3] S. Fischer, N. Kammenhuber, and A. Feldmann, "Replex: dynamic traffic engineering based on wardrop routing policies," in *CoNEXT 2006*, pp. 1–12.
[4] E. Altman, T. Boulogne, R. El-Azouzi, T. Jiménez, and L. Wynter, "A survey on networking games in telecommunications," *Comput. Oper. Res.*, vol. 33, no. 2, 2006.
[5] R. Banner and A. Orda, "Bottleneck routing games in communication networks," *Selected Areas in Communications, IEEE Journal on*, vol. 25, no. 6, pp. 1173–1179, August 2007.
[6] J. Wardrop, "Some theoretical aspects of road traffic research," *Proceedings of the Institution of Civil Engineers, Part II*, vol. 1, no. 36, pp. 352–362, 1952.
[7] L. Kleinrock, *Queueing Systems*. Wiley-Interscience, 1975.
[8] V. Mazalov, B. Monien, F. Schoppmann, and K. Tiemann, "Wardrop equilibria and price of stability for bottleneck games with splittable traffic," in *WINE*, 2006, pp. 331–342.
[9] R. Srikant, *The Mathematics of Internet Congestion Control*. Birkhäuser Boston, 2003.
[10] T. Bonald, L. Massoulié, A. Proutière, and J. Virtamo, "A queueing analysis of max-min fairness, proportional fairness and balanced fairness," *Queueing Syst. Theory Appl.*, vol. 53, no. 1-2, pp. 65–84, 2006.
[11] T. Bonald and A. Proutière, "On performance bounds for balanced fairness," *Perform. Eval.*, vol. 55, no. 1-2, pp. 25–50, 2004.
[12] D. P. Bertsekas, *Nonlinear Programming*. Athena Scientific, 1999.
[13] "The Abilene Network," http://www.internet2.edu/network/.
[14] "TOTEM: TOolbox for Traffic Engineering Methods," http://totem.info.ucl.ac.be/.
[15] Yin Zhang, "Abilene Dataset," http://www.cs.utexas.edu/~yzhang/research/AbileneTM/.
[16] "Géant Topology Map," http://www.geant.net.
[17] S. Uhlig, B. Quoitin, J. Lepropre, and S. Balon, "Providing public intradomain traffic matrices to the research community," *SIGCOMM Comput. Commun. Rev.*, vol. 36, no. 1, pp. 83–86, 2006.