

Routing Games with Elastic Traffic

Peter Key and Alexandre Proutiere
Microsoft Research Cambridge, UK

{Peter.Key, Alexandre.Proutiere} @ microsoft.com

ABSTRACT

In this paper, we introduce and investigate a novel class of multipath routing games with elastic traffic. Users open one or more connections along different feasible paths from source to destination and act selfishly— seeking to transfer data as fast as possible. Users only control their routing choices, and once these choices have been made, the connection rates are elastic and determined via congestion control algorithms (e.g. TCP) which ultimately maximize a certain notion of the network utility. We analyze the existence and the performance of the Nash Equilibria (NEs) of the resulting routing games.

1. INTRODUCTION

Our motivation is a fundamental routing and allocation problem, common to data communication networks, multipath routing, distributed systems and load balancing. We formalize the problem as a routing game. Consider a data network of m links with capacities $C = (C_1, \dots, C_m)$. n selfish users share the resources of these links (where $n \geq 2$). To transmit data, user i may open at most b_i connections from source to destination, and for each such connection, it has to choose one route from a set \mathcal{R}_i of possible routes - A route is just a subset of the set of links. Denote by s_i a pure routing strategy for user i , i.e. a set of routes, and let $s = (s_1, \dots, s_n)$ denote the vector representing the strategies of all users. Let \mathcal{S} denote the set of allowed strategies. Under strategies $s \in \mathcal{S}$, $n_{ij}(s)$ is the number of connections user i opens on route $j \in \mathcal{R}_i$. Strategies s are feasible if for any user i , $\sum_{j \in \mathcal{R}_i} n_{ij}(s) \leq b_i$.

Once users have chosen the routes on which they open connections, the network capacities are shared via congestion control algorithms (e.g. TCP). To model the way these algorithms share the resources, we use Kelly's optimization framework [1], where the achieved rates on the various routes maximize some notion of network utility. When users open several connections, we distinguish two types of congestion control: (i) uncoordinated control, where each user controls the rates on the different connections independently, (ii) coordinated control, where these rates are jointly controlled. In the former case, if r_{ij} denotes the rate achieved on route $j \in \mathcal{R}_i$, then the rate allocation solves the following social welfare maximization problem:

$$\begin{aligned} \max \quad & \mathcal{W}_u(s) = \sum_i \sum_j n_{ij}(s) U(r_{ij}) - \Gamma(r), \\ \text{s.t.} \quad & r_{ij} \geq 0, \quad n_{ij}(s) = 0 \implies r_{ij} = 0, \end{aligned} \quad (1)$$

where U is the concave and increasing utility function, $r = (r^1, \dots, r^m)$ represents the aggregate rate of connections on

the various links $r^l = \sum_i \sum_{j \in \mathcal{R}_i: l \in j} n_{ij}(s) r_{ij}$, and where Γ is a convex penalty function that captures the network costs, e.g. for (uncorrelated) capacity constraints

$$\Gamma(r) = \begin{cases} 0 & \text{if } r^l \leq C_l, \forall l \\ \infty & \text{otherwise.} \end{cases} \quad (2)$$

In the case of coordinated control, the rates solve a global social welfare problem:

$$\begin{aligned} \max \quad & \mathcal{W}_c(s) = \sum_{i=1}^n U(\sum_j n_{ij}(s) r_{ij}) - \Gamma(r), \\ \text{s.t.} \quad & r_{ij} \geq 0, \quad n_{ij}(s) = 0 \implies r_{ij} = 0. \end{aligned} \quad (3)$$

Denote by $r(s) = (r_{ij}(s), i \in [n], j \in \mathcal{R}_i)$ the solution of (1) or (3). We restrict our attention to the wide class of α -fair congestion control algorithm corresponding to utility function $U(\cdot) = (\cdot)^{1-\alpha}/(1-\alpha)$, for $\alpha > 0$ and $\alpha \neq 1$.

We assume that each user i acts strategically, choosing a strategy s_i (a set of routes) that maximizes their global rate $\sum_{j \in \mathcal{R}_i} n_{ij}(s) r_{ij}(s)$, and we analyze the resulting routing game. We are interested in the existence and performance of Nash Equilibria (NE). In particular we are interested in comparing the performance of NE to the socially optimal routing choices s^* solving:

$$\begin{aligned} \max_{s \in \mathcal{S}} \quad & \max \mathcal{W}(s), \\ \text{s.t.} \quad & r_{ij} \geq 0, \quad n_{ij}(s) = 0 \implies r_{ij} = 0, \end{aligned} \quad (4)$$

where $\mathcal{W} = \mathcal{W}_u$ for uncoordinated congestion control and $\mathcal{W} = \mathcal{W}_c$ for coordinated control. Let $w^* = \mathcal{W}(s^*)$. The notion of a Price of Anarchy is a convenient way to quantify the efficiency of Nash Equilibria, and is defined as follows. When the utility function considered takes positive values only, i.e. when $\alpha < 1$, (resp. negative values only, i.e., when $\alpha > 1$), the price of anarchy is for pure NEs:

$$\max_{s \in \text{PNE}} \frac{w^*}{w(s)}, \quad (\text{resp. } \max_{s \in \text{PNE}} \left| \frac{w(s)}{w^*} \right|),$$

where PNE denotes the set of pure NEs.

2. NON-ATOMIC GAMES

In this section, we analyze routing games when the population of users n is large (the size of the network being fixed). Formally, we consider a sequence of systems indexed by n , and we categorize users into a finite set \mathcal{C} of classes. Users of class c may open at most b_c connections on routes from a finite set \mathcal{R}_c . These users represent a proportion $\beta_c > 0$ of the total population of users, and we define $n_c = \beta_c \times n$. Now we

let n be arbitrarily large, and study the corresponding routing games. First, one can easily show that pure NEs exist. Then their performance is given in the following proposition.

THEOREM 1. (i) *With uncoordinated control, the price of anarchy tends to 1 when $n \rightarrow \infty$ for any α -fair allocation with $\alpha < 1$. For $\alpha > 1$, it tends to w_1^*/w_2^* , where w_1^* and w_2^* are respectively the maximal values in the following optimization problems:*

$$\max_{c \in \mathcal{C}} \sum_{c \in \mathcal{C}} (b_c \beta_c)^\alpha \frac{(\sum_{j \in \mathcal{R}_c} r_j)^{1-\alpha}}{1-\alpha}, \quad \text{s.t. } r_j \geq 0, \forall j, r^l \leq C_l, \forall l,$$

$$\max_{c \in \mathcal{C}} \sum_{c \in \mathcal{C}} \beta_c^\alpha \frac{(\sum_{j \in \mathcal{R}_c} r_j)^{1-\alpha}}{1-\alpha}, \quad \text{s.t. } r_j \geq 0, \forall j, r^l \leq C_l, \forall l,$$

(ii) *With coordinated control, the price of anarchy tends to 1 when $n \rightarrow \infty$.*

The proofs of all results in this paper can be found in [2].

3. ATOMIC GAMES

We now consider a fixed and *finite* population of users. In general it is difficult to analyze the corresponding atomic routing game, even when all routes consist of a single link only, i.e., in so-called *parallel link networks*. In such networks, the rate allocation does not depend on the choice of the utility function, both for uncoordinated and coordinated control. When uncoordinated control is used, the rate achieved by user i on a connection to link j is $r_{ij}(s) = C_j/n_j(s)$, where $n_j(s) = \sum_i n_{ij}(s)$ is the total number of connections on link j . When coordinated control is used, the allocation is more complicated to describe, but remains independent of the chosen utility function, see [3].

Parallel networks have an alternative interpretation, applicable to load-balancing questions: we may interpret each link as a server, where users may connect to different servers. As shown in the following example, pure NEs are not guaranteed to exist in the case of uncoordinated control. Possible non-existence of NEs is basically due to the fact that multipath routing games are not congestion games: one may show that in general, they admit no potential function.

EXAMPLE 1. *Consider a 2-link parallel homogeneous network whose resources are shared by users 1 and 2, where each link has unit capacity. User 1 can only open one connection on either link, while user 2 can open three connections and chose where to direct them. That is, $b_1 = 1, b_2 = 3, \mathcal{R}_1 = \mathcal{R}_2 = \{1, 2\}$. The allocation is independent of the rate control used by the users, and each link splits is the rate equally among the connections offered to it. For example, if user 1 creates one connection on link 1, while user 2 creates 2 connection on link 1 and one to link 2, which we write as $\{(1, 0), (2, 1)\}$ then user 1 receives $1/3$ while user 2 receives $5/3$. It is straightforward to show that this an ϵ -NE [4], where $\epsilon = 1/6$. $\epsilon > 0$ is equivalent to the existence of an improvement cycle (corresponding to the cycle $\{\{(1, 0), (2, 1)\}, \{(0, 1), (2, 1)\}, \{(0, 1), (1, 2)\}, \{(1, 0), (1, 2)\}\}$, which since the game is finite, implies the game is not an ordinal potential game [5] (and hence not a potential game).*

We now restrict our attention to parallel link networks where all links have unit capacity. The following two propositions characterize the existence and the performance of pure NEs in these networks. Let $b = \max_i b_i$.

THEOREM 2. (i) *For uncoordinated control, pure NEs exist when $\mathcal{R}_i = \mathcal{R}$ for all i , in the case of all α -fair allocations.* (ii) *With coordinated control, pure NEs exist for all α -fair allocations.*

THEOREM 3. *Assume that $m \leq n$, and for all $i, \mathcal{R}_i = \mathcal{R}$ and $b_i = b$.*

(i) *With uncoordinated control, the price of anarchy is equal to 1 for α -fair allocations with $\alpha < 1$, and when $\alpha > 1$, it is equal to:*

$$\frac{(nb - \lfloor \frac{nb}{m} \rfloor m)(\lfloor \frac{nb}{m} \rfloor + 1)^\alpha + (m - nb + \lfloor \frac{nb}{m} \rfloor m) \lfloor \frac{nb}{m} \rfloor^\alpha}{(n - \lfloor \frac{nb}{m} \rfloor m)(\lfloor \frac{nb}{m} \rfloor + 1)^\alpha + (m - n + \lfloor \frac{nb}{m} \rfloor m) \lfloor \frac{nb}{m} \rfloor^\alpha}.$$

(ii) *With coordinated control, the price of anarchy is 1.*

4. DISCUSSION

In multipath routing games with elastic traffic, we have seen that when the population is large, pure NEs exist and the price of anarchy is strictly greater than 1 only when considering uncoordinated α -fair congestion control with $\alpha > 1$. In such a scenario (with $\alpha > 1$), the inefficiency of NEs stems from the fact that a user has an incentive to open as many connections as possible, which increases network congestion and hence decreases the social welfare.

When the population is finite, pure NEs do not necessarily exist. However, in the case of parallel link networks, we have been able to provide conditions under which pure NEs exist. In such cases, we observe that once again, one has to pay the price of users selfishness only in the case of uncoordinated α -fair congestion control with $\alpha > 1$ (due to a similar phenomenon).

Some important issues remain to be solved. For example, it would be interesting to study the performance mixed Nash Equilibria. We believe that the price of anarchy is then no longer always equal to 1, even when the population of users is large. Some preliminary results [3] indicate that the price of anarchy can remain bounded in the case of coordinated control, but can be arbitrarily large when uncoordinated control is used (just as in traditional routing games [4]).

Characterizing the existence and the performance of NEs may be seen as a *static* analysis of multi-path routing games, whereas in reality users would dynamically adapt their routing choices as a function of the achieved rates, and leave the system once the transfer of their data flow has been completed. We reserve for future work the study of the resulting dynamic routing games.

5. REFERENCES

- [1] F. Kelly. Charging and rate control for elastic traffic. *European Trans. Telecommun.*, 8:33-37, 1997.
- [2] P. Key, A. Proutiere. Routing Games with Elastic Traffic. MSR Technical Report, 2009.
- [3] P. Key, L. Massoulié, and D. Towsley. Path selection and multipath congestion control. *In proc. of IEEE Infocom*, 2007.
- [4] N. Nisan, T. Roughgarden, E. Tardos, and V.V. Vazirani (Eds). Algorithmic Game Theory. *Cambridge University Press*, 2007.
- [5] M. Voorneveld and H. Norde. A characterization of Ordinal Potential Games *Economics Letters*, 19:235-242, 1997.