Congestion Games with Capacitated Resources*

Laurent Gourvès[1], Jérôme Monnot[1], Stefano Moretti[1], and Nguyen Kim Thang[2]

[1] LAMSADE, CNRS UMR 7243, Université Paris Dauphine, France
[2] IBISC, Université d'Evry Val d'Essonne, France

Abstract. We extend congestion games to the setting where every resource is endowed with a capacity which possibly limits its number of users. From the negative side, we show that a pure Nash equilibrium is not guaranteed to exist in any case and we prove that deciding whether a game possesses a pure Nash equilibrium is **NP**-complete. Our positive results state that congestion games with capacities are potential games in the well studied singleton case. Polynomial algorithms that compute these equilibria are also provided.

1 Introduction

The players of a *congestion game* interact by allocating bundles of resources from a common pool [18]. This type of games leads to well studied models for analyzing strategic situations including routing [9], network design [3] and load balancing [8]. They are a prominent model for resource sharing among uncoordinated selfish users.

Significant interest has been addressed over the last years to the analysis of practical congestion problems in the Internet. Data delays and losses due to *data congestions*, or the network collapse as a consequence of exceeding the *data flow capacity* of some links or nodes, has long been a real problem for the Internet [4]. Several policies have been proposed to control congestion, in order to regulate and improve the availability of broadband access to the Internet. *Priority rules*, for instance, have been adopted to regulate the users who enter into the network, with the objective to prevent congestion and to obtain a *Quality of Service* (QoS) that otherwise would not be available to users [5]. A classical example of priorities of users is provided by the access categories of the IEEE 802.11e standard, that was developed in order to offer QoS capabilities to Wireless Local Area Networks (WLANs) [15].

Congestion games [18] can only partially model the practical situation described above. In order to catch other realistic factors like capacities of resources and the different priority of users on the network, a more sophisticated model is required.

For this purpose, we introduce the class of *congestion games with capacitated resources*, where each resource is associated both with a *capacity* level,

* This work is supported by French National Agency (ANR), project COCA ANR-09-JCJC-0066-01.

representing the maximum number of users that such a resource may simultaneously *accommodate*, and with an *ordering* on the users, prescribing the *priority* of accommodation of the users. Given a certain profile of players' strategies, the cost of utilization of a resource for the players which have that resource in their strategy and which are accommodated on it, is a function of the number of players using it in that profile (as in the case of classical congestion games), whereas the cost of players having that resource in their strategy, but which are not accommodated, is prohibitive (supposed infinite).

In this paper we investigate the following questions: Do congestion games with capacitated resources always admit a pure strategy Nash equilibrium (NE in short) in any case as it holds for classical congestion games? If not, is it difficult to decide if an instance possesses a pure NE? Can we identify natural classes of instances admitting a pure NE? Are there polynomial (or more efficient) algorithms that build a pure NE for classes containing such an equilibrium?

2 Models and Notations

A *strategic (cost) game* is a tuple $\langle \mathcal{N}, (\Sigma_i)_{i\in\mathcal{N}}, (c_i)_{i\in\mathcal{N}} \rangle$, where $\mathcal{N} = \{1, \cdots, n\}$ is a finite set of *players*; Σ_i is a non-empty set of *pure strategies* for each player $i \in \mathcal{N}$; $c_i : \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{R}$ is an *individual cost function* specifying players i's cost $c_i(\sigma) \in \mathbb{R}$ for each strategy profile $\sigma = (\sigma_i)_{i\in\mathcal{N}} \in \Sigma_1 \times \cdots \times \Sigma_n$ and each $i \in \mathcal{N}$.

Using conventional notations, we denote by $\Sigma = \Sigma_1 \times \cdots \times \Sigma_n$ the *set of strategy profiles* or *strategy space* and we denote a *strategy profile* σ by (σ_i, σ_{-i}) if the choice of player i needs stressing. The *strategy space* Σ is *symmetric-strategy* if $\Sigma_1 = \Sigma_2 = \ldots = \Sigma_n$.

A *pure strategy Nash equilibrium* (or simply *pure Nash equilibrium*, NE in short) is a pure strategy profile $\sigma \in \Sigma$ such that, for all players $i \in N$, and all pure strategies $s_i \in \Sigma_i$, it holds that $c_i(\sigma) \leq c_i(s_i, \sigma_{-i})$. We only deal with pure strategies in this article so we often omit the word "pure".

For some given strategy profile, a *better move* of a player is a unilateral deviation such that his cost decreases strictly. If such a better move exists, we say that the corresponding player is *unhappy*, otherwise he is *happy*. In this setting a NE is a strategy profile where all players are happy. The *better-response dynamic* is the process of repeatedly choosing an arbitrary unhappy player and let him make an arbitrary better move. A *potential game* is a game in which, for any instance, the better-response dynamic always converges [17]. Such a property is typically shown by a potential function argument.

2.1 Congestion Models and Games

Rosenthal [18] defines a *congestion model* as a tuple $\langle \mathcal{N}, \mathcal{R}, (\Sigma_i)_{i\in\mathcal{N}}, (d_r)_{r\in\mathcal{R}} \rangle$ where $\mathcal{N} = \{1, \ldots, n\}$ is the set of players; $\mathcal{R}$ is a finite set of m *resources*; $\Sigma_i \subseteq 2^{\mathcal{R}}$ is the set of pure strategies of player i, for each $i \in \mathcal{N}$; $d_r : \{0, 1, \ldots, n\} \to \mathbb{R}^+$ is a *delay function* associated with resource r, for each $r \in \mathcal{R}$. This function depends on the number of players using resource r, denoted by $n_r(\sigma)$ or simply n_r when the context is clear. The interpretation is that every player of a resource r incurs a cost of $d_r(n_r)$ (with the convention that $d_r(0) = 0$). Delay functions are sometimes supposed monotone (e.g. [9]) but we do not make this restriction in this paper.

Given a congestion model $\langle \mathcal{N}, \mathcal{R}, (\Sigma_i)_{i\in\mathcal{N}}, (d_r)_{r\in\mathcal{R}} \rangle$, an associated *congestion game* is defined as a strategic cost game $\langle \mathcal{N}, (\Sigma_i)_{i\in\mathcal{N}}, (c_i)_{i\in\mathcal{N}} \rangle$ where for each $\sigma \in \Sigma$ and $i \in \mathcal{N}$, $c_i(\sigma) = \sum_{r\in\sigma_i} d_r(n_r(\sigma))$. Better-response dynamic always converges in congestion games because every better move decreases Rosenthal's potential function $\sum_{r\in\mathcal{R}} \sum_{i=1}^{n_r} d_r(i)$ [18].

An important subclass of congestion games is the class of *singleton congestion games* (also known as *parallel-link games*) in which every player's strategy consists of a single resource [1, 8, 10–12, 14, 16].

2.2 Congestion Games with Capacitated Resources

This section describes the model introduced and studied in this paper. Given a congestion model $\langle \mathcal{N}, \mathcal{R}, (\Sigma_i)_{i\in\mathcal{N}}, (d_r)_{r\in\mathcal{R}} \rangle$, we also assume that every resource $r \in \mathcal{R}$ has a *capacity* κ_r – an integer between 1 and n – which is the maximal number of players that can use resource r. Moreover, every resource r is associated with a linear order $\mathtt{pos}_r : \mathcal{N} \to \{1, \dots, n\}$, where $\mathtt{pos}_r(i) = t$ means that player i is in the t-th position of r ($\mathtt{pos}$ is strict total). We say that a player i has a *higher priority* than player j at resource r iff $\mathtt{pos}_r(i) < \mathtt{pos}_r(j)$. Notice that $\mathtt{pos}_r(i)$ is defined even if r does not appear in the strategy space of player i.

Let $N_r(\sigma)$ be the set of players using resource r in the strategy profile σ. A player $i \in N_r(\sigma)$ is *accommodated* by r iff the number of players in $N_r(\sigma)$ having a position lower than $\mathtt{pos}_r(i)$ is strictly smaller than the capacity of resource r, i.e., $|\{j \in N_r(\sigma) : \mathtt{pos}_r(j) < \mathtt{pos}_r(i)\}| < \kappa_r$. The *delay* $d_r(\sigma)$ of a resource r in profile σ is defined as $d_r(\min\{n_r(\sigma), \kappa_r\})$. The *delay* $d_r^i(\sigma)$ of player $i \in N_r(\sigma)$ on resource r is:

$$d_r^i(\sigma) = \begin{cases} d_r(\min\{n_r(\sigma),\, \kappa_r\}) & \text{if } i \text{ is accommodated,} \\ +\infty & \text{otherwise.} \end{cases} \tag{1}$$

A *congestion game with capacitated resources* (*capacitated congestion game* in short) is a strategic cost game where the *cost* of a player i in profile σ is defined as $c_i(\sigma) = \sum_{r\in\sigma_i} d_r^i(\sigma)$.

Note that capacitated congestion games follow the original congestion model of Rosenthal [18] when the resources are not overcrowded. When the capacity of a resource is exceeded, the game shares similarities with the player-specific model of Milchtaich [16] since we distinguish between accommodated and non accommodated players. However congestion games with capacitated resources are neither a refinement nor an extension of player-specific congestion games.

In congestion games with capacitated resources, a profile is a *Nash equilibrium* if the following conditions hold:

– no player, accommodated by *every* resource in his current strategy, can unilaterally deviate and decrease his cost;
– no player, not accommodated by at least one resource in his current strategy, can unilaterally deviate and incur a finite cost.

We say that a resource r is *saturated* if $n_r(\sigma) \geq \kappa_r$. We say that a player i is *displaced* by another player j in the following situation: i is accommodated by a resource r which is not used by j, j deviates so that r is in his new strategy and i is not accommodated by r anymore whereas j is (of course $\texttt{pos}_r(j) < \texttt{pos}_r(i)$).

3 Related Works

Various aspects of congestion games were investigated. The existence of pure NE, the convergence of better-response dynamic and the computation of equilibria are interleaved questions studied in [9, 14, 6, 2]. Computing a pure NE of a congestion game is a **PLS**-complete problem, even if strategies are symmetric. Nevertheless there are important subclasses for which a NE can be built in polynomial time, by the use of dedicated algorithms or simply via better response dynamic (see [19] for a survey).

Many extensions of the congestion model introduced in Rosenthal [18] have been studied in the literature of strategic games. *Player-specific congestion games*, have been introduced in [16] with the objective to model congestion situations where the delay of each resource in $\mathcal{R}$ depends not only on the number of players using that resource but also on the player's identity itself. The delay of a player $i \in \mathcal{N}$ on resource $r \in \mathcal{R}$ is a function $d_r^i : \mathbb{N} \to \mathbb{R}^+$.

A generalization of this model are (player-specific) *congestion games with priorities*, which have been introduced in [1] with the objective to model situations where each resource can assign priorities to the players, and players with a higher priority can displace *all* players with a lower priority. Every resource $r \in \mathcal{R}$ is associated with a map (not necessarily a bijection) $\pi_r : \mathcal{N} \to \{1, \ldots, |\mathcal{N}|\}$. Several players can *allocate* a resource r (those players form a set $N_r(\sigma)$) but only those with highest priority π_r are assigned to r. This latter subset of assigned players is denoted by $\hat{N}_r(\sigma)$.

Formally, for each strategy profile $\sigma \in \Sigma$ and each $r \in \mathcal{R}$ such that $N_r(\sigma) \neq \emptyset$, let $\hat{N}_r(\sigma) = \arg\max_{i \in N_r(\sigma)} \pi_r(i)$ be the set of players *assigned* to resource r. The delay incurred by an assigned player $i \in \hat{N}_r(\sigma)$ is $d_r^i(|\hat{N}_r(\sigma)|)$. Players in $N_r(\sigma) \setminus \hat{N}_r(\sigma)$, who are not assigned to resource r, incur an infinite delay.

Although there are some similarities between the congestion model with capacities introduced in this paper and the one with priorities introduced by [1] (e.g., the possibility to displace players with lower priority on a certain resource), in general, these two models generate well distinct strategic cost games.. Contrasting with the model discussed in this paper, Ackermann *et al* [1] suppose that there is no capacity on the resources, two players may have the same priority with respect to a given resource and two players with distinct priorities on a resource r can not be both assigned to r.

Finally, the notion of capacity in systems with congested resources has been considered in [7] (see also references therein). Nevertheless, capacitated congestion games and the model in [7] are different. In our setting, we consider a finite number of atomic players and resources have an order on the users, whereas in [7], players are non-atomic and resources are not endowed with an order.

4 Contribution and Organization

Our goal is two-fold: (i) characterize the existence of a NE in capacitated congestion games; and (ii) efficiently compute an equilibrium if it exists.

First, we consider capacitated congestion games in general. We prove that a capacitated congestion game always admits a NE if it consists of two resources; moreover, this equilibrium can be computed in linear time. Besides, a game with three resources (and more) does not necessarily possess a NE. This negative result holds even if the game is symmetric-strategy and all players' strategies except one are singleton. From a computational aspect, deciding whether a game, even symmetric-strategy and consisting of two players, has a NE is shown to be **NP**-complete. The results are presented in Section 5.

Next, we consider singleton capacitated congestion games. We show that the game is a potential game so it always admits a NE. The proof is based on a new geometrical approach of potential argument, which could be seen as a generalization of a dominant potential function in higher dimension. We believe that the approach would be useful in proving the existence of NE in other games and is of independent interest. In computational aspect, the better-response dynamic converges to a NE in at most $O(n^4 m)$ strategy changes (recall that n and m are the number of players and resources, respectively). Additionally, we give a more efficient algorithm to compute a NE when the game is symmetric-strategy. The results are presented in Section 6.

5 General Strategies

We begin with a simple symmetric-strategy game which does not admit a NE. There are two players, three resources x, y and z, and the priorities are the same for the three resources (priority is always given to the first player). The strategy space of the players is $\{\{x\}, \{y, z\}\}$. Resource x has capacity 1 and $d_x(1) = 2$. Resource y has capacity 2 and $d_y(1) = 3$ while $d_y(2) = 0$. Resource z has capacity 1 and $d_z(1) = 0$. The game is illustrated in Figures 1 and 2.

Notice that the example possesses some minimal characteristics for the existence of a NE: a game with one player obviously admits a NE and Theorem 1 states that capacitated congestion games defined on two resources always admit a NE. Moreover the instance falls into restricted cases which often make the existence of a NE likely: strategies are symmetric source-target paths of a directed network, delays are monotone and priorities on the resources are identical.

Theorem 1. *Every capacitated congestion game defined on two resources possesses a pure Nash equilibrium. Moreover, a NE can be computed in linear time.*

	$\{x\}$	$\{y,z\}$
$\{x\}$	2, $+\infty$	2, 3
$\{y,z\}$	3, 2	0, $+\infty$

Fig. 1. A 3-resource 2-player symmetric-strategy capacitated congestion game without any pure Nash equilibrium.

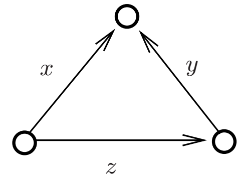

Fig. 2. The corresponding network where each arc is a resource.

Proof (Sketch of proof). We prove that Algorithm 1 outputs an equilibrium σ. Denote by r and s the resources. Observe that players with strategy space $\{\{r\},\{r,s\}\}$, $\{\{s\},\{r,s\}\}$ and $\{\{r\},\{s\},\{r,s\}\}$ cannot prefer to play $\{r,s\}$ over $\{r\}$ or $\{s\}$, in any profile, as the delay of every resource is non-negative. Hence, we can reduce the strategy space of those players to be $\{\{r\}\}$, $\{\{s\}\}$ and $\{\{r\},\{s\}\}$, respectively. The action of the players having only one strategy in their (reduced) strategy space is obviously known. Denote by $\hat{\mathcal{N}}$ the players whose (reduced) strategy space is $\{\{r\},\{s\}\}$.

Algorithm 1 2-resource

Input: a set $\mathcal{N}$ of players, two resources r, s
Output: A pure Nash equilibrium σ

```
N̂ ← ∅
If a player i has only one strategy in his reduced strategy space then assign him to
  that strategy, else let σ_i ← r and N̂ ← N̂ ∪ {i}
Rename players in N̂ such that pos_s(1) < pos_s(2) < ··· < pos_s(n̂) where n̂ = |N̂|
Let N̂_∞ and N̂_f be the set of players in N̂ with infinite cost and finite cost under
  the current profile σ, respectively
for i = 1 to n̂ do
  If i ∈ N̂_∞ and c_i(s, σ_{-i}) < c_i(σ) then σ_i ← s
end for
for i = 1 to n̂ do
  if i ∈ N̂_f and c_i(s, σ_{-i}) < c_i(σ) then
    σ_i ← s
    if i displaces a player j ∈ N̂ then
      σ_j ← r
    end if
  end if
end for
return profile σ
```

First, we show an invariant that at anytime, the algorithm maintains the property that no player of $\hat{\mathcal{N}}$ placed on s can or wants to move to r.

The property is clearly true before the first *for* loop. During the first *for* loop, no player who has moved from r to s has incentive to return back to r because he would get an infinite cost. For the second *for* loop, we prove the invariant by induction. The base case (before entering to the loop) is straightforward. We analyze a step by considering three subcases:

- Resource s is saturated before i moves and the deviation implies that a player $j' \notin \hat{\mathcal{N}}$ is displaced. In this case, the deviation does not incentivize a player $j \in \hat{\mathcal{N}}$ placed on resource s to move. Indeed j's cost is $d_s(\kappa_s)$ before and after i's deviation. After his deviation, i's cost is $d_s(\kappa_s)$ which is strictly smaller than his previous cost. Moving to r is not profitable to j.
- Resource s is saturated before i moves and the deviation implies that a player $j \in \hat{\mathcal{N}}$ is displaced. Observe that j cannot belong to $\hat{\mathcal{N}}_f$ because the loop follows the total order of priorities on s. The algorithm assigns j to r so that his cost is either equal to $+\infty$ or equal to the cost previously incurred by i. Then, the number of players on s remains unchanged. No player from $\hat{\mathcal{N}}$ placed on resource s has incentive to move, since otherwise the player can do it before the exchange of i and j, contradiction to the induction hypothesis.
- Resource s is not saturated before i moves and the deviation implies that at least one player $j \in \hat{\mathcal{N}}$ wants to unilaterally move to r. Players i and j have the same finite cost. By moving to r, player j would get either $+\infty$ or exactly the cost incurred by i before his deviation, contradiction.

The property holds at the end of the two phases. Now observe that a player $i \in \hat{\mathcal{N}}$ placed on r either has been displaced from s at some step or has had the opportunity to switch to s during the second loop but did not (could not) do so. Hence, those players are happy on resource r. The profile σ is then a pure Nash equilibrium. The algorithm is clearly linear in n.

When the number of resources is unbounded, the problem becomes much harder.

Proposition 1. *Deciding whether a symmetric-strategy capacitated congestion game has a NE is* **NP***-complete, even with two players.*

Proof (Sketch of proof). We reduce Partition — a **NP**-complete problem [13] — to the symmetric-strategy capacitated congestion game. In Partition, given n integers $\{a_1, \ldots, a_n\}$ such that $\sum_{j=1}^{n} a_j = 2B > 6$ and $0 < a_j < B$, one has to decide whether a subset $J \subseteq \{1, \ldots, n\}$ such that $\sum_{j \in J} a_j = B = \sum_{j \notin J} a_j$ exists.

Given an instance of Partition, we construct a capacitated congestion game with two players where the resources are the arcs of a network G and the players' strategies are all paths from a common source s to a common target t, see Figure 3. For arc e_0, $\kappa_{e_0} = 2$, $d_{e_0}(1) = B + 2$ and $d_{e_0}(2) = 0$. For arcs e_j and e'_j where $1 \le j \le n$, $\kappa_{e_j} = \kappa_{e'_j} = 2$, $d_{e_j}(1) = a_j$, $d_{e_j}(2) = B + 2$, and $d_{e'_j}(1) = 0, d_{e'_j}(2) = B + 2$. For arc e'_{n+2}, $\kappa_{e'_{n+2}} = 2$ and $d_{e'_{n+2}}(1) = 2$, $d_{e'_{n+2}}(2) = 0$. For arcs e_{n+1} and e'_{n+1}, their capacities are $\kappa_{e_{n+1}} = \kappa_{e'_{n+1}} = 1$ and player 1 has higher priority than player 2 in both arcs. Moreover, the delay functions are $d_{e_{n+1}}(1) = B$, $d_{e'_{n+1}}(1) = B - 1$.

One can show that the instance of Partition has a feasible solution iff the game defined on G admits a NE. □

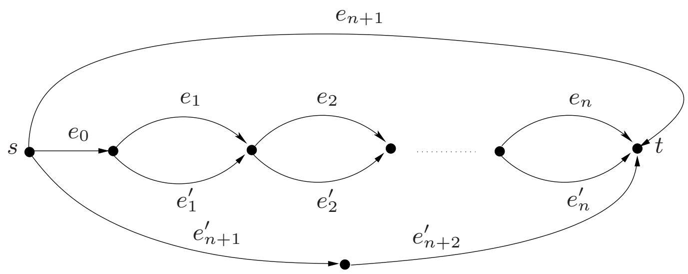

Fig. 3. The network associated with an instance of PARTITION.

6 Singleton Strategies

In this section, we are interested in studying the existence of NE and efficient algorithms to compute a NE in singleton capacitated congestion games. First, we present intuitively our approach in proving the existence of a NE.

Starting point Consider the following dominant order $\prec'$. Let $A = \{a_1 \leq \ldots \leq a_k\}$ and $B = \{b_1 \leq \ldots \leq b_k\}$ be two sets of k real-value elements that are named in increasing order. We say that $A \prec' B$ if there exists an index $1 \leq \ell \leq k$ such that $a_i = b_i$ for all $1 \leq i < \ell$ and $a_\ell < b_\ell$. This order is well-defined and has been used in proving the existence of Nash equilibria (for example [8]). We interpret this order in a geometrical view. For each set A and B, map all elements to points on a real line where the coordinate of a point equals the value of its corresponding element. For $u \in \mathbb{R}$, let A_u and B_u be the number of points corresponding to elements in A and B with coordinate smaller than or equal to u, respectively. Then, the order $\prec'$ could be equivalently defined as follows: $A \prec' B$ if for the smallest $u \in \mathbb{R}$ such that $A_u \neq B_u$, it holds that $A_u < B_u$. In fact, the smallest $u \in \mathbb{R}$ such that $A_u \neq B_u$ is a_ℓ where ℓ is the index in the former definition.

As we have seen, the dominant order could be geometrically interpreted as a one-dimension order. Taking this geometrical approach, we prove the existence of NE by designing a two-dimension order. Intuitively, the two dimensions are due to the nature of the game where the cost of a player depends on the resource delay and the priority of the player on the resource.

Theorem 2. *Singleton capacitated congestion games are potential games. Moreover, the better-response dynamic necessarily converges in $O(n^4 m)$ strategy changes.*

Proof. First, we give some definitions which are useful in the proof.

For each profile σ, a function $\texttt{rank}_\sigma : \mathcal{R} \to \mathbb{N}$ is defined as follows. If resource r is saturated[3] then $\texttt{rank}_\sigma(r) = \max\{\texttt{pos}_r(j) : \sigma_j = r, j \text{ is accommodated}\}$. Otherwise, $\texttt{rank}_\sigma(r) := n + 1$.

[3] A resource r is saturated if $n_r(\sigma) \geq \kappa_r$.

We define a function f that maps each profile σ to a multiset of points in $\mathbb{R}^+ \times \mathbb{N}$. Each resource r in profile σ is associated with the multiset $f(r,\sigma)$ of points $(d_r(1), n+1); (d_r(2), n+1); \ldots; (d_r(t_r(\sigma)-1), n+1)$ and $(d_r(t_r(\sigma)), \texttt{rank}_\sigma(r))$ where $t_r(\sigma) := \min\{n_r(\sigma), \kappa_r\}$. The multiset $f(\sigma) := \cup_{r \in \mathcal{R}} f(r,\sigma)$. An illustration of $f(\sigma)$ is given in Figure 4.

For a value $u \in \mathbb{R}^+$, to every profile σ we define the multiset $\sigma_u := \{(a,b) \in f(\sigma) : a \leq u\}$. Moreover, denote by $|\sigma_u|$ the cardinal of σ_u and $\|\sigma_u\| := \sum_{(a,b)\in\sigma_u} b$. By the definition, $|\sigma_u|$ is the number of points corresponding to profile σ which are on the left of the line $x = u$ and intuitively $\|\sigma_u\|$ is the total height of these points.

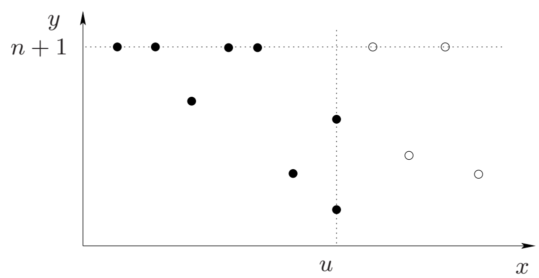

Fig. 4. An illustration of $f(\sigma)$, black filled dots if in σ_u.

Now we define a partial order $\prec$ on profiles. Formally, two profiles ν and σ satisfy $\nu \prec \sigma$ if for the smallest $u > 0$ such that $(|\sigma_u|, \|\sigma_u\|) \neq (|\nu_u|, \|\nu_u\|)$ we have $|\sigma_u| < |\nu_u|$, or $|\sigma_u| = |\nu_u|$ but $\|\sigma_u\| > \|\nu_u\|$. Intuitively, we can interpret this order as follows. Two profiles ν and σ satisfy $\nu \prec \sigma$ if for the smallest $u > 0$ such that $(|\sigma_u|, \|\sigma_u\|) \neq (|\nu_u|, \|\nu_u\|)$, either (1) the half-space on the left of the line $x = u$ contains more points of ν than those of σ; or (2) if they are equal, the total height of such points in ν is smaller than that of σ.

Now we can prove that after a better move of some player i from resource r in profile σ to a resource s, resulting in profile ν, we get that $\nu \prec \sigma$. Note that $f(\sigma)$ and $f(\nu)$ only differ on some points corresponding to resources r and s. In the following, we consider only these points. Let u be the cost of player i after the move, which equals $d_s(t_s(\nu))$ — the delay of resource s in profile ν. (Note that player i is accommodated by resource s in profile ν as he has taken a better move.)

Consider the set of points corresponding to resource r in $f(\sigma)$ and $f(\nu)$. If i has unbounded cost in profile σ (meaning that i is not accommodated), then $f(r,\sigma) = f(r,\nu)$. If i is accommodated in profile σ then either $f(r,\sigma) = f(r,\nu) \cup (d_r(\sigma), \texttt{rank}_\sigma(r))$ in case $n_r(\sigma) \leq \kappa_r$, or $f(r,\sigma) = f(r,\nu) \setminus (d_r(\kappa_r), \texttt{rank}_\sigma(r)) \cup (d_r(\kappa_r), \texttt{rank}_\nu(r))$ in case $n_r(\sigma) > \kappa_r$. However, as i has taken a better move,

$d_r^i(\sigma) = d_r(\sigma) > u$. Hence, restricting to points with first coordinate smaller than or equal to u, $f(r,\sigma) = f(r,\nu)$.

Consider the set of point corresponding to resource s in $f(\sigma)$ and $f(\nu)$. If s is unsaturated before the move of i then $f(s,\nu) = f(s,\sigma) \cup (d_s(\nu), \texttt{rank}_\nu(s))$ $= f(s,\sigma) \cup (u, \texttt{rank}_\nu(s))$. If s is saturated before the move of i then $f(s,\nu) = f(s,\sigma) \cup (u, \texttt{rank}_\nu(s)) \setminus (u, \texttt{rank}_\sigma(s))$.

Therefore, for any $u' < u$, $(|\sigma_{u'}|, \|\sigma_{u'}\|) = (|\nu_{u'}|, \|\nu_{u'}\|)$. Moreover, if s is unsaturated before the move of i, $|\sigma_u| < |\nu_u|$. Otherwise, $|\sigma_u| = |\nu_u|$ but $\texttt{rank}_\nu(s) < \texttt{rank}_\sigma(s)$, so $\|\nu_u\| < \|\sigma_u\|$. Hence, $\nu \prec \sigma$, i.e., after each better move, a new profile is $\prec$-smaller than the previous one. In conclusion, the game is a potential game.

Now we bound the number of strategy changes to reach an NE from arbitrary profile in the better-response dynamic. Let σ be an arbitrary profile. By the definition of order $\prec$, there are at most nm values of u that we have to consider. Moreover, for each u, $0 \leq |\sigma_u| \leq n$ and $0 \leq \|\sigma_u\| \leq n(n+1)$. Hence, there are at most $O(n^4 m)$ couples $(|\sigma_u|, \|\sigma_u\|)$ (where σ is a profile) which are $\prec$-different. Thus, from an arbitrary profile, the better-response dynamic converges to a NE in at most $O(n^4 m)$ strategy changes. □

In the following, we consider singleton capacitated congestion games with additional property of symmetry on players' strategy sets. We give an algorithm to compute a NE that is more efficient than the better-response dynamic by exploiting that property.

Theorem 3. *A NE in a symmetric-strategy, singleton capacitated congestion game can be computed in* $\min\{n, \kappa\}$ *strategy changes and the overall time complexity of the algorithm is* $O(\min\{n^2 m, \kappa^2\})$*, where* $\kappa = \sum_{r \in \mathcal{R}} \kappa_r$.

Proof. We show that Algorithm 2 computes a NE.

First consider the case $n \geq \sum_{r \in \mathcal{R}} \kappa_r$. By the algorithm, at the end of the *while* loop, all resources become saturated with delays $d_{r_1}(\kappa_1) \leq \ldots \leq d_{r_m}(\kappa_m)$. Next, κ_{r_1} first players according to $\texttt{pos}_{r_1}$ are assigned to resource r_1, then κ_{r_2} first players according to $\texttt{pos}_{r_2}$ among the remaining players are assigned to resource r_2 then so on. Finally, assign all remaining players to resource r_m. The outcome is a NE because: (1) a player assigned to a resource r_j cannot displace other player assigned to a resource $r_{j'}$ where $j' < j$; (2) a player assigned to a resource r_j cannot decrease his cost by moving to other resource $r_{j'}$ where $j' > j$.

Now, consider the case $n < \sum_{r \in \mathcal{R}} \kappa_r$. In this case, every player is accommodated to some resource. Suppose a player i, assigned to resource r in profile σ, has incentive to deviate to resource s resulting in profile σ'.

If i's deviation displaces some player i' then we get a contradiction. Indeed, $d_r(n_r(\sigma)) = c_i(\sigma) > c_i(\sigma') = c_{i'}(\sigma) = d_s(n_s(\sigma))$ and $\texttt{pos}_s(i) < \texttt{pos}_s(i')$ hold. However, the algorithm fills resource s before resource r (steps 8 to 12 of the algorithm) and player i should have been assigned to s instead of player i'.

Algorithm 2 Symmetric-strategy, singleton capacitated congestion games.

Input: Set $\mathcal{N}$ of n players, $\texttt{pos}_r$ and κ_r for all $r \in \mathcal{R}$
Output: An equilibrium σ
1: $n_r \leftarrow 0$ for all $r \in \mathcal{R}$
2: $\hat{n} \leftarrow \min\{n, \kappa\}$ where $\kappa = \sum_{r \in \mathcal{R}} \kappa_r$.
3: **while** $\hat{n} > 0$ **do**
4: Find r^* and k_{r^*} such that $d_{r^*}(k_{r^*}) = \min\{d_r(k_r) : n_r < k_r \leq \min\{n_r + \hat{n}, \kappa_r\}, r \in \mathcal{R}\}$.
5: $\hat{n} \leftarrow \hat{n} - (k_{r^*} - n_{r^*})$
6: $n_{r^*} \leftarrow k_{r^*}$
7: **end while**
8: Rename resources so that $d_{r_1}(n_1) \leq d_{r_2}(n_2) \leq \ldots \leq d_{r_m}(n_m)$
9: **for** $j = 1$ **to** m **do**
10: Assign to resource r_j the first n_j players $S \subset \mathcal{N}$ according to $\texttt{pos}_{r_j}$
11: $\mathcal{N} \leftarrow \mathcal{N} \setminus S$
12: **end for**
13: Assign all remaining players in $\mathcal{N}$ to an arbitrary resource, for example resource r_m.
14: **output** the current assignment σ.

Assume i does not displace anyone when deviating. We have indeed $d_r(n_r(\sigma)) = c_i(\sigma) > c_i(\sigma') = d_s(n_s(\sigma)+1)$. Consider the moment at which n_r is modified for the last time (line 6 of the algorithm). Let k_r and k_s be the number of players already assigned to resource r and s at that time, respectively. By the algorithm, n_r is modified because $d_r(k_r) = d_r(n_r(\sigma))$ is minimum among other choices. Besides, observe that at that time, $\hat{n} \geq (n_s(\sigma) - k_s) + 1$ since later, the algorithm will set $n_s(\sigma)$ as the number of players (who are different to i) on resource s. Therefore, resource s and $n_s(\sigma)+1$ is a candidate for the choice of the algorithm in line 4. Thus, $d_r(n_r(\sigma)) \leq d_s(n_s(\sigma)+1)$ — contradiction. Hence, every player in σ is happy, meaning that it is a NE. By the algorithm, the number of strategy changes is obviously $\min\{n, \kappa\}$ and the time complexity is dominated by the *while* loop which needs at most $O(\min\{n^2 m, \kappa^2\})$ operations. □

7 Conclusion

In the paper, we have assumed that each capacitated resource r is endowed with a linear order $\texttt{pos}_r$, indicating which players are accommodated when the resource is overcrowded. We believe that different and equally relevant ways to determine who is accommodated exist, and the existence of a NE should be investigated. For instance, an interesting open question is to know the computational complexity of symmetric-strategy capacitated congestion games with increasing delay functions. On a *dynamic* perspective, for instance, it would be interesting to study a model where the priorities of users depend on their timing of using resources (for routing problems, this could represent the arrival time to the starting node of an edge). On the other hand, in this perspective, dropping the assumption of priorities represented by linear orders could generate the technical problem of coordinating users asking for the same resource at the same time (on this issue, see the discussion about *timestamp games* in [10]).

References

1. H. Ackermann, P. Goldberg, V. Mirrokni, H. Röglin, and B. Vöcking. A unified approach to congestion games and two-sided markets. *Internet Mathematics*, 5(4):439–457, 2008.
2. H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6), 2008.
3. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
4. S. Bauer, D. Clark, and W. Lehr. The evolution of internet congestion. In *TPRC 2009, 37th research conference on Communication, Information and Internet policy*, 2009.
5. A. Campbell, C. Aurrecoechea, and L. Hauw. A review of QoS architectures. *ACM Multimedia Systems Journal*, 6:138–151, 1996.
6. S. Chien and A. Sinclair. Convergence to approximate nash equilibria in congestion games. In *SODA*, pages 169–178, 2007.
7. J. Correa, A. Schulz, and N. Stier-Moses. Selfish routing in capacitated networks. *Math. Oper. Res.*, 29(4):961–976, 2004.
8. E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibrium in load balancing. *ACM Transactions on Algorithms*, 3(3), 2007.
9. A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *STOC*, pages 604–612, 2004.
10. B. Farzad, N. Olver, and A. Vetta. A priority-based model of routing. *Chicago J. Theor. Comput. Sci.*, 1, 2008.
11. D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The structure and complexity of nash equilibria for a selfish routing game. *ICALP*, pages 785–785, 2002.
12. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing nash equilibria for scheduling on restricted parallel links. In *STOC*, pages 613–622. ACM, 2004.
13. M. Garey and D. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company, New York, 1979.
14. S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *AAAI*, pages 489–494, 2005.
15. S. Mangold, S. Choi, P. May, O. Klein, G. Hiertz, and L. Stibor. IEEE 802.11e Wireless LAN for quality of service. In *Proc. European Wireless*, volume 18, pages 32–39, 2002.
16. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
17. D. Monderer and L. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
18. R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.
19. B. Vöcking. Congestion games: Optimization in competition. In *ACiD Workshop*, volume 7 of *Texts in Algorithmics*, pages 9–20, 2006.