

Congestion Games with Capacitated Resources

Laurent Gourvès · Jérôme Monnot ·
Stefano Moretti · Nguyen Kim Thang

Published online: 26 March 2014
© Springer Science+Business Media New York 2014

Abstract The players of a congestion game interact by allocating bundles of resources from a common pool. This type of games leads to well studied models for analyzing strategic situations, including networks operated by uncoordinated selfish users. Congestion games constitute a subclass of potential games, meaning that a pure Nash equilibrium emerges from a myopic process where the players iteratively react by switching to a strategy that diminishes their individual cost. With the aim of covering more applications, for instance in communication networks, 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.

Keywords Congestion games · Pure nash equilibrium · Potential function · Algorithms

1 Introduction

In congestion games [32], a finite set of *players* may use *resources* from a common pool and the delays associated with the use of each resource is a function only of the

This work is supported by French National Agency (ANR), project COCA ANR-09-JCJC-0066-01. A preliminary version of this work appeared in the proceedings of SAGT 2012 [19].

L. Gourvès (✉) · J. Monnot · S. Moretti
LAMSADE, CNRS UMR 7243, Université Paris Dauphine, Paris, France
e-mail: laurent.gourves@dauphine.fr

N. K. Thang
IBISC, Université d'Evry Val d'Essonne, Évry, France

number of players using it. The cost of a player is the sum of the delays associated with each resource in his strategy choice, given the choices of the other players. The existence of (pure) *Nash equilibria* [31] for congestion games can be proved by constructing an exact *potential function* [30, 32], and it can be also proved that each game with an exact potential function is isomorphic to a congestion game ([30]; see [37] for an introduction to more general results on *potential games*).

This type of games are well studied models for strategic situations including routing [12], network design [3] and load balancing [11]. It is also 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 [6]. 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) [28]. But the interest for priority policies is not limited to the Internet. In many countries, priority-based mechanisms have increasingly been used to promote an efficient use of the highway. Think for instance to *High Occupancy Vehicle (HOV)* lanes, which provide priority to vehicles with more than two or three occupants, like shuttle vans, buses, car pools etc. Sometimes, and at less congested hours, low occupancy vehicles are also allowed to use priority lanes on payment of a toll aimed to contribute to cover system implementation costs. However, many transportation specialists claim that HOV lanes would not substantially solve the problem of traffic congestion, while they would extremely penalize low occupancy vehicles. Therefore, they argue for more sophisticated priority mechanisms which may improve the efficient exploitation of the highway traffic capacity [22]. Priority rules are also used in railway traffic regulation when the traffic demand exceeds the capacity of rails: for example, priority may be provided to trains increasing the revenues of the rail network provider, like in Germany [10], or to passenger trains over freight trains, like in Italy [27], etc.

Congestion games [32] 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, 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 $\mathcal{N} = \{1, \dots, n\}$ is a finite set of *players*; Σ_i is a non-empty set of *pure strategies* for each player $i \in \mathcal{N}$; $c_i : \Sigma_1 \times \dots \times \Sigma_n \rightarrow \mathbb{R}$ is an *individual cost function* specifying players i 's cost $c_i(\sigma) \in \mathbb{R}$ for each strategy profile $\sigma = (\sigma_i)_{i \in \mathcal{N}} \in \Sigma_1 \times \dots \times \Sigma_n$ and each $i \in \mathcal{N}$.

Using conventional notations, we denote by $\Sigma = \Sigma_1 \times \dots \times \Sigma_n$ 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 $\Sigma_1 = \Sigma_2 = \dots = \Sigma_n$.

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 $i \in \mathcal{N}$, 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 dynamics* 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 dynamics always converges [30]. Such a property is typically shown by a potential function argument.

2.1 Congestion Models and Games

Rosenthal [32] 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 $\mathcal{N} = \{1, \dots, n\}$ 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 $i \in \mathcal{N}$; $d_r : \{0, 1, \dots, n\} \rightarrow \mathbb{R}^+$ 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. [12]) 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 dynamics always converges in congestion games because every better move decreases Rosenthal’s potential function $\sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r} d_r(i)$ [32].

In general, a *potential function* $\Phi : \Sigma \rightarrow \mathbb{R}$ maps the incentives of all players into one function whose local optima correspond to the pure NE of the game. The class of games admitting a potential function is known as *potential games* [30].

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, 11, 13, 15, 16, 21, 29].

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 $\text{pos}_r : \mathcal{N} \rightarrow \{1, \dots, n\}$, where $\text{pos}_r(i) = t$ means that player i is in the t -th position of r (pos is strict total). We say that a player i has a *higher priority* than player j at resource r iff $\text{pos}_r(i) < \text{pos}_r(j)$. Notice that $\text{pos}_r(i)$ 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 $\text{pos}_r(i)$ is strictly smaller than the capacity of resource r , i.e., $|\{j \in N_r(\sigma) : \text{pos}_r(j) < \text{pos}_r(i)\}| < \kappa_r$. 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 [32] 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 [29] and with its generalization in [1]. However congestion games with capacitated resources are neither a refinement nor an extension of player-specific congestion games (this is discussed in detail in Appendix 1).

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 $\text{pos}_r(j) < \text{pos}_r(i)$).

3 Related Works

Various aspects of congestion games were investigated. The existence of pure NE, the convergence of better-response dynamics and the computation of equilibria are interleaved questions studied in [2, 7, 12, 21]. 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 dynamics (see [20, 36, 37] for surveys). Another important direction is the quality of equilibria compared to a social optimal profile (the so-called *price of anarchy* [23] and *price of stability* [3]) for the specific class of congestion games [8].

Many extensions of the congestion model introduced in Rosenthal [32] have been studied in the literature of strategic games. A non exhaustive list contains player-specific and weighted congestion games [29], potential games [30], local-effect games [26], congestion games with priorities [1] and graphical congestion games [5].

Player-specific congestion games [29] have been introduced 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 $d_r^i : \mathbb{N} \rightarrow \mathbb{R}^+$. This means that the delays of the same resource for two distinct players may be different.

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 (not necessarily a total order), and players with a higher priority can displace *all* players with a lower priority. Several players can *allocate* a resource but only those with highest priority are *assigned* to it. The congestion on a resource depends on the number of players assigned to it. The players who allocate a resource without being assigned to it incur an infinite delay.

Singleton congestion games with priorities (without player specific delay functions) are potential games, and a pure NE can be found in polynomial time [1]. Congestion games with priorities extend (*with ties*) the *two-sided markets model* introduced in [17], which is used to analyze markets where different kinds of agents are matched to another (e.g., students and colleges, client and server sides, readers and writers, firms and workers etc.). In the model of *two-sided markets with ties* adopted by [1], several players may propose to be allocated to a resource, but only the most preferred ones are assigned to that resource. In contrast, in many real examples of two-sided matching models, the assignments are subject to the limitations imposed by capacities, and not to the highest priorities, as it has been proposed in [1]. Consider for instance the two-sided market represented by the *college admissions problem* [17, 33], where colleges need to select sets of students, and students may rank schools according to their preferences. Since colleges have limited enrolling capacities, admissions are usually regulated by priorities (e.g., colleges may rank students based on their individual academic records). In this context, an assignment is

said to be stable if, whenever a student prefers another college to his assignment, the enrolling capacity of that college is already saturated by students with higher priority.

There are similarities between the model introduced in this article and the model of [1], e.g. the possibility to displace players with lower priority on a certain resource. Nevertheless 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 cannot be both assigned to r (see Appendix 1 for further details).

The notion of capacity in systems with congested resources has been considered in [9] (see also references therein). Nevertheless, capacitated congestion games and the model in [9] are different. In our setting, we consider a finite number of atomic players and resources have an order on the users, whereas in [9], players are non-atomic and resources are not endowed with an order.

A recent work is conducted on the price of anarchy and stability in a network with capacities on its links [14]. In this paper, the authors distinguish feasible and unfeasible strategy profiles (all capacity constraints are satisfied or not). A Nash equilibrium exists each time a feasible solution exists. Hence, no priority rule is needed in their work.

Other models have been presented in the literature with the objective to analyze the effects of players' priority on congestion. *Priority-based selfish routing models*, where nonatomic agents may have different priorities on the links of a network, have been presented in [13]. In this framework, different *priority schemes* [13] on the links have been considered, and the bounds of the corresponding *price of anarchy* [24] are studied and compared.

Finally, note that most of the results concerning the existence of Nash equilibria presented in this paper deal with *singleton congestion games*, which are situations where players are allowed to use only one single resource. Such games, also known as *parallel-link games*, are very well studied in the literature of congestion problems [1, 11, 13, 15, 16, 21]. We point out to the fact that the parallel-link network topology has received increasing attention in recent years in the context of QoS and the related implementation of priority policies in data networking [25, 35, 38].

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. Our main results are summarized in Table 1.

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

Table 1 Summary of our main results

	General strategies		Singleton strategies
	2 resources	3 resources	
\exists NE	yes	no	yes
Polynomially	yes	–	yes

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 dynamics converges to a NE in at most $O(n^4m)$ 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 Figs. 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.

Making further restrictions, like assuming non-decreasing delays, does not guarantee the existence of a NE either, as it is shown by the 2-player symmetric-strategy

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

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

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



game presented in Fig. 3. In this game, there are four resources $x, y, z,$ and w . The strategy space of the players is $\{\{x, y\}, \{y, z\}, \{x, z\}, \{w\}\}$. Resources x and y have capacity 1 and $d_x(1) = d_y(1) = 0$. Resource z and w have capacity 2 and $d_z(1) = 1$ and $d_w(1) = 2$, while $d_z(2) = 3$ and $d_w(2) = 4$. We have $\text{pos}_x(1) < \text{pos}_x(2)$ and $\text{pos}_y(1) > \text{pos}_y(2)$. The linear orders of z and w have no influence on the game since the resources can accommodate the two players.

Next result states that a pure NE is guaranteed to exist for two resources (it is not true for 3 resources by the example of Fig. 1).

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

Proof Denote by r and s the resources. Observe that a player with strategy space $\{\{r\}, \{r, s\}\}$, cannot prefer to play $\{r, s\}$ over $\{r\}$, in any profile, as the delay of a resource is non-negative. Hence, we can reduce his strategy space to $\{\{r\}\}$, meaning that he is happy whenever he plays $\{r\}$. With the same arguments, a player with strategy space $\{\{s\}, \{r, s\}\}$, is happy whenever he plays $\{s\}$. Similarly, the players with strategy space $\{\{r\}, \{s\}, \{r, s\}\}$ can have a reduced strategy space of $\{\{r\}, \{s\}\}$; we denote by \mathcal{N} these players.

	$\{x, y\}$	$\{y, z\}$	$\{x, z\}$	$\{w\}$
$\{x, y\}$	$+\infty$	1	$+\infty$	2
$\{y, z\}$	0	$+\infty$	3	2
$\{x, z\}$	$+\infty$	3	$+\infty$	2
$\{w\}$	0	1	1	4

Fig. 3 A 4-resource 2-player symmetric-strategy capacitated congestion game with non-decreasing delays and without any pure Nash equilibrium

We are going to build a strategy profile σ . The players who are not in $\hat{\mathcal{N}}$ are assigned to the single strategy of their (reduced) strategy space, and their assignments are not subsequently modified. For the players in $\hat{\mathcal{N}}$, we use Algorithm 1 to determine if they play $\{r\}$ or $\{s\}$.

Algorithm 1 2-resource

Input: a set \mathcal{N} of players, two resources r, s

Output: A pure Nash equilibrium σ

```

1:  $\hat{\mathcal{N}} \leftarrow \emptyset$ 
2: If a player  $i$  has only one strategy in his reduced strategy space then assign him
   to that strategy, else let  $\sigma_i \leftarrow r$  and  $\hat{\mathcal{N}} \leftarrow \hat{\mathcal{N}} \cup \{i\}$ 
3: Rename players in  $\hat{\mathcal{N}}$  such that  $\text{pos}_s(1) < \text{pos}_s(2) < \dots < \text{pos}_s(\hat{n})$  where
    $\hat{n} = |\hat{\mathcal{N}}|$ 
4: Let  $\hat{\mathcal{N}}_\infty$  and  $\hat{\mathcal{N}}_f$  be the set of players in  $\hat{\mathcal{N}}$  with infinite cost and finite cost under
   the current profile  $\sigma$ , respectively
5: for  $i = 1$  to  $\hat{n}$  do
6:   if  $i \in \hat{\mathcal{N}}_\infty$  and  $c_i(s, \sigma_{-i}) < c_i(\sigma)$  then  $\sigma_i \leftarrow s$ 
7: end for
8: for  $i = 1$  to  $\hat{n}$  do
9:   if  $i \in \hat{\mathcal{N}}_f$  and  $c_i(s, \sigma_{-i}) < c_i(\sigma)$  then
10:      $\sigma_i \leftarrow s$ 
11:     if  $i$  displaces a player  $j \in \hat{\mathcal{N}}$  then
12:        $\sigma_j \leftarrow r$ 
13:     end if
14:   end if
15: end for
16: return profile  $\sigma$ 

```

The algorithm starts by assigning all members of $\hat{\mathcal{N}}$ to r . Then $\hat{\mathcal{N}}$ can be partitioned in two sets $\hat{\mathcal{N}}_f$ and $\hat{\mathcal{N}}_\infty$ corresponding to accommodated and non accommodated players, respectively. The members of $\hat{\mathcal{N}}_\infty$ are examined by decreasing priority on resource s . A player in $\hat{\mathcal{N}}_\infty$ is assigned to s only if he is better off on s . After, the members of $\hat{\mathcal{N}}_f$ are examined by decreasing priority on resource s . A player in $\hat{\mathcal{N}}_f$ is assigned to s only if he is better off on s , and if the deviation causes the displacement of a player $j \in \hat{\mathcal{N}}$ then j is immediately moved to r . We are going to prove that the resulting strategy profile must be a pure NE.

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 on i . The base case (before entering 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 since it implies a worse cost ($+\infty$ or i 's previous cost).
- 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 with 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 $+\infty$ (which cannot be an improvement) or exactly the cost incurred by i before his deviation, contradicting the fact that i has decreased his cost.

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 not fixed, the problem becomes much harder as it is stated in the next proposition.

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

Proof We will reduce PARTITION — a NP-complete problem [18] — to the symmetric-strategy capacitated congestion game. In PARTITION, given n integers $\{a_1, \dots, a_n\}$ such that $\sum_{j=1}^n a_j = 2B > 6$ and $0 < a_j < B$, one has to decide whether a subset $J \subseteq \{1, \dots, n\}$ 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 Fig. 4. For




Fig. 4 The network associated with an instance of PARTITION

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 $1 \leq j \leq n$, $\kappa_{e_j} = \kappa_{e'_j} = 2$, $d_{e_j}(1) = a_j$, $d_{e_j}(2) = B + 2$, 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.

Let J be a subset of $\{1, \dots, n\}$ such that $\sum_{j \in J} a_j = B = \sum_{j \notin J} a_j$. Consider the following profile σ : $\sigma_1 = \{e_0\} \cup \{e_j : j \in J\} \cup \{e'_j : j \in \{1, \dots, n\} \setminus J\}$ and $\sigma_2 = \{e_0\} \cup \{e'_j : j \in J\} \cup \{e_j : j \in \{1, \dots, n\} \setminus J\}$. Each player has cost B . One can verify that no player can reduce his cost by unilaterally switching to another path. Therefore, σ is a NE.

Conversely, assume that σ is a NE of the game. First, we prove the following claim.

Claim 1 e_0 is the unique arc which is shared by two players □

Proof of claim Suppose that an arc $e \neq e_0$ is shared by the two players. If $e \in \{e_1, \dots, e_n\} \cup \{e'_1, \dots, e'_n\}$ then the cost of each player is at least $d_e(2) = B + 2$. In this case, one player can switch to arc e_{n+1} and get a cost of B . Otherwise, $e \in \{e_{n+1}, e'_{n+1}, e'_{n+2}\}$. As $\kappa_{e_{n+1}} = \kappa_{e'_{n+1}} = 1$ and player 1 has higher priority than player 2 in both arcs, the cost of player 2 is $+\infty$. However, player 2 can choose a path consisting of arcs in $\{e_1, \dots, e_n\} \cup \{e'_1, \dots, e'_n\}$ and incur a finite cost.

Now suppose that e_0 is not shared by the two players. Then, $e_0 \notin \sigma_i$ for some player $i \in \{1, 2\}$; the other player has index $3 - i$. If arc e_0 is used by player $3 - i$ then his cost is at least $d_{e_0}(1) = B + 2$. In this case, player $3 - i$ can deviate to the path $\{e_{n+1}\}$ or $\{e'_{n+1}, e'_{n+2}\}$, which is different from player i 's strategy, and have a cost of at most $B + 1$. Consider the case that arc e_0 is not used by player $3 - i$. By the previous argument, no edge different from e_0 is shared, so the players do not use the same path. If player 2 chooses (e_{n+1}) then he has a cost of B while player 1 must choose (e'_{n+1}, e'_{n+2}) with a cost of $B + 1$. Now, player 1 can switch to path (e_{n+1}) and reduce his cost to B because $\text{pos}_{e_{n+1}}(1) < \text{pos}_{e_{n+1}}(2)$. If player 2 chooses (e'_{n+1}, e'_{n+2}) then he has a cost of $B + 1$ while player 1 must choose (e_{n+1}) with a cost of B . Now, if player 1 changes his path by taking (e'_{n+1}, e'_{n+2}) , he gets a cost of $d_{e'_{n+1}}(1) + d_{e'_{n+2}}(2) = B - 1$ because $\text{pos}_{e'_{n+1}}(1) < \text{pos}_{e'_{n+1}}(2)$. In each case, we get a contradiction. Hence, e_0 is the unique arc shared by both players. □

By the claim, e_0 is shared by the players. If one player has a cost strictly larger than B , then he can decrease it by choosing e_{n+1} . Therefore, the player's costs in σ is at most B . Moreover, the arcs in their paths must belong to $\{e_0\} \cup \{e_1, \dots, e_n\} \cup \{e'_1, \dots, e'_n\}$. The total delays of arcs without e_0 in their paths is at least $\sum_j a_j = 2B$ by the definition of the delay functions and the fact that they do not share any arc except e_0 . Therefore, both have cost B . Hence, the set $J := \{a_j : e_j \in \sigma_1\}$ is a feasible solution to the 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.

6.1 Starting Point

Consider the following dominance order \prec' . Let $A = \{a_1 \leq \dots \leq a_k\}$ and $B = \{b_1 \leq \dots \leq b_k\}$ 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 [15]). 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 dynamics necessarily converges in $O(n^4m)$ strategy changes.*

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

For each profile σ , a function $\text{rank}_\sigma : \mathcal{R} \rightarrow \mathbb{N}$ is defined as follows. If resource r is saturated¹ then $\text{rank}_\sigma(r) = \max\{\text{pos}_r(j) : \sigma_j = r, j \text{ is accommodated}\}$. Otherwise, $\text{rank}_\sigma(r) := n + 1$.

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 $(d_r(1), n + 1); (d_r(2), n + 1); \dots; (d_r(t_r(\sigma) - 1), n + 1)$ and $(d_r(t_r(\sigma)), \text{rank}_\sigma(r))$ 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 Fig. 5.

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.

¹Recall that a resource r is saturated if $n_r(\sigma) \geq \kappa_r$.




Fig. 5 An illustration of $f(\sigma)$, black filled dots if in σ_u

Now we define a partial order $<$ on profiles. Formally, two profiles ν and σ satisfy $\nu < \sigma$ 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 $\nu < \sigma$ 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 show 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 $\nu < \sigma$. 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.)

Intuitively, a player i moves to some resource if his cost is improved, meaning that the delay $d_s(t_s(\nu))$ at the new resource is smaller than his previous cost. Then either the number of points with first coordinate at most $d_s(t_s(\nu))$ is larger; or if that number remains the same, it must be the case that i displaces some other player with lower priority. In any case, the potential function decreases. We are now proving the claim formally.

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 $f(r, \sigma) = f(r, \nu) \cup (d_r(\sigma), \text{rank}_\sigma(r))$ in case $n_r(\sigma) \leq \kappa_r$, or $f(r, \sigma) = f(r, \nu) \setminus (d_r(\kappa_r), \text{rank}_\sigma(r)) \cup (d_r(\kappa_r), \text{rank}_\nu(r))$ 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 $f(s, \nu) = f(s, \sigma) \cup (d_s(\nu), \text{rank}_\nu(s)) = f(s, \sigma) \cup (u, \text{rank}_\nu(s))$. If s is saturated before the move of i then $f(s, \nu) = f(s, \sigma) \cup (u, \text{rank}_\nu(s)) \setminus (u, \text{rank}_\sigma(s))$.

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 $\text{rank}_\nu(s) < \text{rank}_\sigma(s)$, so $\|\nu_u\| < \|\sigma_u\|$. Hence, $\nu < \sigma$, i.e., after each better move, a new profile is $<$ -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 dynamics. 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 $O(n^4m)$ couples $(|\sigma_u|, \|\sigma_u\|)$ (where σ is a profile) which are \prec -different. Thus, from an arbitrary profile, the better-response dynamics converges to a NE in at most $O(n^4m)$ strategy changes. \square

In the following, we consider singleton capacitated congestion games with the additional property of symmetry on the players' strategy sets. We give an algorithm to compute a NE that is more efficient than the better-response dynamics 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 $O(\min\{n^2m, \kappa^2\})$, where $\kappa = \sum_{r \in \mathcal{R}} \kappa_r$.*

Proof Consider Algorithm 2 which consists of two phases: (i) determine a load n_r (ie. a number of players) for every resource r , (ii) identify the n_r players to be assigned to resource r .

Algorithm 2 Symmetric-strategy, singleton capacitated congestion games.

Input: Set \mathcal{N} of n players, 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 $d_{r_1}(n_1) \leq d_{r_2}(n_2) \leq \dots \leq d_{r_m}(n_m)$
 - 9: **for** $j = 1$ **to** m **do**
 - 10: Assign to resource r_j the first n_j players $S \subset \mathcal{N}$ according to 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 .
-

The load of every resource is 0 at the beginning of the first phase. The algorithm iteratively places some loads which eventually add up to \hat{n} , defined as the minimum between the total number of players and the maximum number of players that can

be accommodated. Each iteration of the first phase is to find the least costly way to augment the load.

Once the loads on the resources are known, the algorithm starts the second phase. Every resource r is assigned exactly n_r players. The algorithm takes the most appealing resource r (r has the lowest delay when it accommodates n_r players) and assigns to it the n_r players with highest priority on r . The assignment on r is unchangeable and the algorithm proceeds greedily with the remaining resources and players until \hat{n} players are assigned. Finally, if $\hat{n} < n$ then all resources are saturated and the players that remain unassigned are put on r_m .

We are going to prove that the resulting strategy profile is a pure 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 $d_{r_1}(\kappa_1) \leq \dots \leq d_{r_m}(\kappa_m)$. Next, κ_{r_1} first players according to pos_{r_1} are assigned to resource r_1 , then κ_{r_2} first players according to 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 another 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 another 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 $\text{pos}_s(i) < \text{pos}_s(i')$ 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' .

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. \square

7 Conclusion

Though capacitated congestion games do not admit a pure NE in general, we have identified two subcases for which the existence of a pure NE is guaranteed : when there are only two resources and for singleton strategies. As a future work, it would be interesting to have a characterization of the instances which admit a pure NE, those for which a potential function exists, those for which a polynomial algorithm

can return a pure NE and those for which deciding the existence of a pure NE is NP-complete. Some factors like the number of resources, the symmetry of the players' strategy spaces, the monotonicity of the delay functions and the (non) homogeneity of the priority orders seem to be relevant. Concerning the existence of a pure NE, the examples of Section 5 give an insight into this characterization but many cases remain open. For example one can mention the instances with non-decreasing delays and homogeneous priority orders. If the players' strategy spaces can be asymmetric then a pure NE is not guaranteed to exist (see Appendix 2). For symmetric strategy spaces we believe that a pure NE exists but it is an open problem. Concerning the complexity of deciding if an instance admits a pure NE, Proposition 1 captures the symmetric 2-player case with homogeneous priority orders. We believe that other reductions can cover other cases.

We have assumed in the paper that each capacitated resource r is endowed with a linear order pos_{S_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 [13]).

Besides, it would be interesting to consider the model of two-sided markets in which the ones in both sides have dynamic priority orders on the ones of the other side. For example, certain matches between agents of the two sides could be a priori forbidden, or could determine some extra costs (like in the two-sided market represented by patients and hospitals, where patients should support further costs to switch from public hospitals to private ones). Finally, natural restriction/extension of our model would be singleton capacitated congestion games with linear delays and the study of instances with weighted players.

Acknowledgments We thank the anonymous referees for their comments which lead to significant improvements of the manuscript.

Appendix 1

Next discussion explains how the model introduced and studied in this article differs from player-specific congestion games [29].

Consider a congestion model with three players $\mathcal{N} = \{1, 2, 3\}$, two resources $\mathcal{R} = \{r, s\}$, a symmetric strategic space $(\Sigma_i = \{\{r\}, \{s\}\})_{i \in \mathcal{N}}$, and the delay of

Table 2 A congestion game with priorities

Players 1 \ 2	{r}	{s}
{r}	(∞, 2, 2)	(∞, 1, 1)
{s}	(1, 2, 2)	(∞, 1, 1)

Player 3: {r}

Players 1 \ 2	{r}	{s}
{r}	(∞, 1, 1)	(1, 2, 2)
{s}	(∞, 1, 1)	(∞, 2, 2)

Player 3: {s}

a resource equals the number of players using it. Consider the following player-specific congestion game with priorities. On both resources, players 2 and 3 have higher priority on player 1, precisely, $\pi_r(2) = \pi_s(2) = \pi_r(3) = \pi_s(3) = 2$ and $\pi_r(1) = \pi_s(1) = 1$. The corresponding congestion game with priorities is given in Table 2. Note that each triple (\cdot, \cdot, \cdot) represents the costs of players where the costs of player 1, 2 and 3 correspond to the first, second and third element of the triple.

Now, on the same congestion model, consider all possible capacitated congestion games and the two strategy profiles $\sigma = (\{r\}, \{r\}, \{r\})$ and $\sigma' = (\{r\}, \{r\}, \{s\})$. Whatever pos functions of resources r and s , if players' costs in profile σ are $(\infty, 2, 2)$ then $\kappa_r = 2$. On the other hand, if players' costs in profile σ' are $(\infty, 1, 1)$, then it implies that $\kappa_r = 1$. Therefore it cannot be possible to obtain the outcomes corresponding to σ and σ' of games shown in Table 2.

Besides, consider a congestion game with capacitated resources in which $\kappa_r = 2$, $\kappa_s = 1$ and $\text{pos}_r(3) < \text{pos}_r(2) < \text{pos}_r(1)$, $\text{pos}_s(3) < \text{pos}_s(2) < \text{pos}_s(1)$. We have that the costs of players in σ and σ' are $(\infty, 2, 2)$ and $(2, 2, 1)$, respectively. A congestion game with priorities in which players' costs are $(\infty, 2, 2)$ necessarily gives the highest priority to both players 2 and 3. Therefore, the outcomes corresponding to σ' is $(\infty, 1, 1)$. Hence, we have proved that neither of the two classes is included in the other.

However, a capacitated congestion game with two players and the same priorities over all resources can be seen as a player-specific congestion game [29]. Denote by i and j the two players of a capacitated congestion game with $\text{pos}_r(i) < \text{pos}_r(j)$ and delay d_r for every resource r . It suffices to consider a player-specific congestion game with delay functions d_r^i and d_r^j such that $d_r^i = d_r$ for every resource r , $d_r^j = d_r$ for every resource r with capacity equal to 2, and $d_r^j(1) = d_r(1)$ and $d_r^j(2) = +\infty$ for every resource r with capacity equal to 1.

Appendix 2

No pure NE exists in the next example. Every resource's delay is non-decreasing and all resources have the same priority order. Nevertheless the players have different strategy spaces.

There are two players (row and column) and three resources x, y and z . The strategy set of the row and column players are $\{\{x\}, \{y\}\}$ and $\{\{x, y\}, \{z\}\}$, respectively. Resource x has a capacity of 2, $d_x(1) = 1$ and $d_x(2) = 3$. Resource y has a capacity

of 1 and $d_y(1) = 0$. Resource z has a capacity of 1 and $d_z(1) = 2$. Priority is always given to the column player.

	$\{x, y\}$	$\{z\}$
$\{x\}$	3	2
$\{y\}$	1	2
	$+\infty$	0

References

- Ackermann, H., Goldberg, P., Mirrokni, V., Röglin, H., Vöcking, B.: A unified approach to congestion games and two-sided markets. *Internet Math.* **5**(4), 439–457 (2008)
- Ackermann, H., Röglin, H., Vöcking, B.: On the impact of combinatorial structure on congestion games. *J. ACM* **55**(6) (2008)
- Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, E., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. *SIAM J. Comput.* **38**(4), 1602–1623 (2008)
- Bauer, S., Clark, D., Lehr, W.: The evolution of internet congestion. In: *TPRC 2009, 37th Research Conference on Communication, Information and Internet Policy* (2009)
- Bilò, V., Fanelli, A., Flammini, M., Moscardelli, L.: Graphical congestion games. *Algorithmica* **61**(2), 274–297 (2011)
- Campbell, A., Aurrecochea, C., Hauw, L.: A review of QoS architectures. *ACM Multimedia Systems J.* **6**, 138–151 (1996)
- Chien, S., Sinclair, A.: Convergence to approximate nash equilibria in congestion games. In: *SODA*, pp. 169–178 (2007)
- Christodoulou, G., Koutsoupias, E.: The price of anarchy of finite congestion games. In: *STOC*, pp. 67–73 (2005)
- Correa, J., Schulz, A., Stier-Moses, N.: Selfish routing in capacitated networks. *Math. Oper. Res.* **29**(4), 961–976 (2004)
- Czerny, A.I., Mitusch, K., Tanner, A.: Priority rules versus scarcity premiums in rail markets. *WHU Otto Beisheim School of Management - Working Paper Series in Economics* **10**(3) (2010)
- Even-Dar, E., Kesselman, A., Mansour, Y.: Convergence time to nash equilibrium in load balancing. *ACM Trans. Algoritm.* **3**(3) (2007)
- Fabrikant, A., Papadimitriou, C., Talwar, K.: The complexity of pure nash equilibria. In: *STOC*, pp. 604–612 (2004)
- Farzad, B., Olver, N., Vetta, A.: A priority-based model of routing. *Chic. J. Theor. Comput. Sci.*, 1 (2008)
- Feldman, M., Ron, T.: Capacitated network design games. In: Serna, M. (ed.) *Algorithmic Game Theory - 5th International Symposium, SAGT 2012, Barcelona, Spain, 22–23 October 2012. Proceedings of Lecture Notes in Computer Science*, vol. 7615, pp. 132–143. Springer (2012)
- Fotakis, D., Kontogiannis, S., Koutsoupias, E., Mavronicolas, M., Spirakis, P.: The structure and complexity of nash equilibria for a selfish routing game. *Theor. Comput. Sci.* **410**(36), 3305–3326 (2009)
- Gairing, M., Lücking, T., Mavronicolas, M., Monien, B.: Computing nash equilibria for scheduling on restricted parallel links. In: *STOC*, pp. 613–622. ACM (2004)
- Gale, D., Shapley, L.S.: College admissions and the stability of marriage. *Am. Math. Mon.* **69**(1), 9–15 (1962)
- Garey, M., Johnson, D.: *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, New York (1979)
- Gourvès, L., Monnot, J., Moretti, S., Nguyen Kim, T.: Congestion games with capacitated resources. In: Serna, M. (ed.) *Algorithmic Game Theory - 5th International Symposium, SAGT 2012, Barcelona, Spain, 22–23 October 2012. Proceedings of Lecture Notes in Computer Science*, vol. 7615, pp.204–215. Springer (2012)

20. Gourvès, L., Moretti, S.: *Progress in Combinatorial Optimization*. Combinatorial optimization problems arising from interactive congestion situations. ISTE Ltd, Wiley (2011)
21. Jeong, S., McGrew, R., Nudelman, E., Shoham, Y., Sun, Q.: Fast and compact: a simple class of congestion games. In: AAAI, pp. 489–494 (2005)
22. Konishi, H., Mun, S.: Carpooling and congestion pricing: Hov and hot lanes. *Reg. Sci. Urban Econ.* **40**(4), 173–186 (2010)
23. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. In: 16th Annual Symposium on Theoretical Aspects of Computer Science, pp. 404–413 (1999)
24. Koutsoupias, E., Papadimitriou, C.: Worst-case equilibria. *Comput. Sci. Rev.* **3**(2), 65–69 (2009)
25. Lee, K.D., Leung, V.C.M.: Utility-based rate-controlled parallel wireless transmission of multimedia streams with multiple importance levels. *IEEE Trans. Mob. Comput.*, 81–92 (2009)
26. Leyton-Brown, K., Tennenholtz, M.: Local-effect games. In: Lehmann, D.J., Muller, R., Sandholm, T. (eds.) *Computing and Markets*, volume 05011 of Dagstuhl Seminar Proceedings. Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany IBFI, Schloss Dagstuhl, Germany (2005)
27. Lulli, G., Pietropaoli, U., Ricciardi, N.: Service network design for freight railway transportation: the italian case. *J. Oper. Res. Soc.* **62**(12), 2107–2119 (2011)
28. Mangold, S., Choi, S., May, P., Klein, O., Hiertz, G., Stibor, L.: IEEE 802.11e Wireless LAN for quality of service. *Proceedings of European Wireless*, vol. 18, pp. 32–39 (2002)
29. Milchtaich, I.: Congestion games with player-specific payoff functions. *Games and Economic Behavior* **13**(1), 111–124 (1996)
30. Monderer, D., Shapley, L.: Potential games. *Games and Economic Behavior* **14**, 124–143 (1996)
31. Nash, J.: Equilibrium points in n-person games. *Proceedings of the National Academy of Sciences of the United States of America* **36**(1), 48–49 (1950)
32. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *Int. J. Game Theory* **2**, 65–67 (1973)
33. Roth, A.E.: The college admissions problem is not equivalent to the marriage problem. *J. Econ. Theory* **36**(2), 277–288 (1985)
34. Serna, M. (ed.) *Algorithmic Game Theory - 5th International Symposium, SAGT 2012, Barcelona, Spain, 22–23 October 2012*. *Proceedings of Lecture Notes in Computer Science*, vol. 7615. Springer (2012)
35. Su, X., De Veciana, G.: Predictive routing to enhance qos for stream-based flows sharing excess bandwidth. *Comput. Netw.* **42**(1), 65–80 (2003)
36. Vöcking, B.: Congestion games: optimization in competition. In: *ACiD Workshop, Text in Algorithmics*, vol. 7, pp. 9–20 (2006)
37. Voorneveld, M.: *Potential Games and Interactive Decisions with Multiple Criteria*. PhD thesis, Tilburg University (1999)
38. Zhao, Z., Willman, B., Weber, S., de Oliveira, J.C.: Performance analysis of a parallel link network with preemption. In: 40th Annual Conference on Information Sciences and Systems, pp. 271–276 (2006)