

A Unified Approach to Congestion Games and Two-Sided Markets

Heiner Ackermann, Paul W. Goldberg, Vahab S. Mirrokni, Heiko Röglin, and Berthold Vöcking

Abstract. Congestion games are a well-studied model for resource sharing among uncoordinated selfish players. Usually, one assumes that the resources in a congestion game do not have any preferences regarding the players that can access them. In typical load-balancing applications, however, different jobs can have different priorities, and jobs with higher priorities get, for example, larger shares of processor time. We extend the classical notion of congestion game and introduce a model in which each resource can assign priorities to the players, and players with higher priorities can displace players with lower priorities. Not only does our model extend classical congestion games, it can also be seen as a model of two-sided markets with ties. Hence it unifies previous results for these two classical models.

We prove that singleton congestion games with priorities are potential games. Furthermore, we show that every player-specific singleton congestion game with priorities possesses a pure Nash equilibrium that can be found in polynomial time. Finally, we extend our results to matroid congestion games, in which the strategy spaces of the players are matroids over the resources.

1. Introduction

In a *congestion game*, there is a set of players who compete for a set of resources. Each player has to select a subset of resources that she wishes to access. The delay of a resource depends on the number of players accessing that resource, and every player is interested in accessing a subset of resources with small total delay. Congestion games are a well-studied model for resource sharing among uncoordinated selfish agents. They are widely used to model routing [Fabrikant

et al. 04], network design [Anshelevich et al. 04], load balancing [Even-Dar et al. 03, Goemans et al. 04], and dynamic frequency assignment for WLANs [Fischer et al. 07]. One appealing property of congestion games is that they are *potential games* [Rosenthal 73]. This means that there exists a potential function assigning a value to every state that decreases on going from one configuration to another by letting one player play a better response. In particular, this implies that every congestion game possesses a *pure Nash equilibrium* and that myopic players eventually reach a Nash equilibrium by iteratively playing better responses.

One drawback of the standard model of congestion games is that resources do not have any preferences with respect to the players. In typical load-balancing applications, however, different jobs can have different priorities, and depending on the policy, jobs with a low priority are stopped or slowed down when jobs with higher priorities are present. We introduce *congestion games with priorities* to model the scenario in which a job can prevent jobs with lower priorities from being processed. In our model, each resource can partition the set of players into classes of different priorities.

As long as a resource is accessed only by players with the same priority, these players incur a delay depending on the congestion, as in standard congestion games. But if players with different priorities access a resource, only players with the highest priority incur a delay, which depends on the number of players with this priority, and players with lower priorities incur an infinite delay. Intuitively, they are displaced by the players with the highest priority. This model is applicable if every player controls a stream of jobs rather than only a single one. In the latter case, it might be more reasonable to assume that jobs with lower priorities incur a large but finite delay.

Motivated by the application of congestion games to load balancing and dynamic frequency assignment for WLANs, we mainly consider congestion games in which each player has to choose exactly one resource to access, namely one server on which her job is to be processed. Such *singleton congestion games*, or congestion games on *parallel links*, have been studied extensively in the literature [Even-Dar et al. 03, Fotakis et al. 02, Gairing et al. 04, Ieong et al. 05]. We show that singleton congestion games with priorities are potential games, implying that uncoordinated players who iteratively play better responses eventually reach a pure Nash equilibrium. If the priorities assigned to the players coincide for every resource, then we even obtain polynomial-time convergence to a Nash equilibrium. *Player-specific congestion games* are introduced in [Milchtaich 96] as an extended class of congestion games in which every player can have her own delay function for every resource. Milchtaich shows that player-specific singleton congestion games are no longer potential games but that they possess pure

Nash equilibria that can be computed in polynomial time. We show that also in player-specific singleton congestion games with priorities, pure Nash equilibria exist and can be computed efficiently.

Interestingly, our model of player-specific congestion games with priorities extends not only congestion games but also the well-known model of *two-sided markets*. This model was introduced in [Gale and Shapley 62] to model markets on which different kinds of agents are matched to one another, for example men and women, students and colleges [Gale and Shapley 62], interns and hospitals [Roth 84], and firms and workers. Using the same terms as for congestion games, we say that the goal of a two-sided market is to match players and resources (or markets). In contrast to congestion games, each resource can be matched to only one player. A payoff is associated with each player–resource pair, and players are interested in maximizing their payoffs.

Hence the payoffs implicitly define a preference list with respect to the resources for each player. Additionally, each resource has a preference list of the players that is independent of the profits. Every player can *propose* to one resource, and if several players propose to a resource, only the most-preferred player is *assigned* to that resource and receives the corresponding payoff. In this way, every set of proposals corresponds to a bipartite matching between players and resources. A matching is *stable* if no player can be assigned to a resource from which she receives a higher payoff than from her current resource given the proposals of the other players. It is shown in [Gale and Shapley 62] that stable matchings always exist and can be found in polynomial time. Since the seminal work of Gale and Shapley there has been a significant amount of work in studying two-sided markets. See for example, [Knuth 76], [Gusfield and Irving 89], or [Roth and Sotomayor 90].

Just as in many situations it is unrealistic to assume that in congestion games the resources have no preferences with respect to the players, it is often unrealistic in two-sided markets to assume that the preference lists of the resources are strict. Our model of player-specific congestion games with priorities can also be seen as a model of *two-sided markets with ties*, in which several players can be assigned to one resource. If different players propose to a resource, only the most-preferred ones are assigned to it. If the most-preferred player is not unique, several players share the payoff of the resource. Such two-sided markets correspond to our model of congestion games with priorities, except that players are now interested in maximizing their payoffs instead of minimizing their delays, which does not affect our results. Two-sided markets with ties have been extensively studied in the literature [Gusfield and Irving 89, Iwama et al. 99].

In these models, ties are somehow broken, i.e., despite ties in the preference lists, every resource can be assigned to at most one player. Hence, these

models differ significantly from our model. One application of our model is markets into which different companies can invest: as long as the investing companies are of comparable size, they share the payoff of the market, but large companies can utilize their market power to eliminate smaller companies completely from the market. Player-specific congestion games and two-sided markets are the special cases of our model in which all players have the same priority or distinct priorities, respectively. In the following, we use the terms *two-sided markets with ties* and *player-specific congestion games with priorities* interchangeably.

We also consider a special case of *correlated two-sided markets with ties* in which the payoffs of the players and the preference lists of the resources are correlated. In this model, every resource prefers to be assigned to players who receive the highest payoff when assigned to it. We show that this special case is a potential game. Variants of correlated two-sided markets without ties have been studied in the context of content distribution in networks and distributed caching problems [Fleischer et al. 06, Goemans et al. 04, Mirrokni 05]. These markets have also been considered for discovering stable geometric configurations with applications in VLSI design [Hoffman et al. 06]. Our result implies that variants of the uniform distributed caching game with bandwidth constraints (defined in [Mirrokni 05, Fleischer et al. 06]) are potential games.

Additionally, we consider congestion games with priorities with strategy spaces that do not consist of singleton sets only. In particular, we consider player-specific congestion games with priorities in which the strategy space of each player is a matroid over the resources. For this case, we show that pure Nash equilibria exist and can be computed in polynomial time, extending a result for player-specific congestion games without priorities [Ackermann et al. 09]. These games can also be seen as many-to-one two-sided markets with ties. Many-to-one two-sided markets are well studied in the economics literature [Fleiner 03, Kelso and Crawford 82, Kojima and Ünver 08]. It is shown in [Kelso and Crawford 82] that if the preference list of every player satisfies a certain *substitutability property*, then stable matchings exist. It is proved in [Kojima and Ünver 08] that in this case, from every matching there exists a better-response sequence to a stable matching of polynomial length. This substitutability property is satisfied if the strategy spaces of the players are matroids. The crucial difference between our model of *many-to-one markets with ties* and the models considered in the economics literature is that in those models, every player specifies a ranking on the power set of the resources. This ranking is fixed and does not depend on the current matching. In our model with ties, however, players do not have fixed rankings but rankings that depend on the current matching.

2. Preliminaries

In this section, we define the problems and notation used throughout the paper.

2.1. Congestion Games

A *congestion game* Γ is a tuple $(\mathcal{N}, \mathcal{R}, (\Sigma_i)_{i \in \mathcal{N}}, (d_r)_{r \in \mathcal{R}})$, where $\mathcal{N} = \{1, \dots, n\}$ denotes the set of players, \mathcal{R} the set of resources, $\Sigma_i \subseteq 2^{\mathcal{R}}$ the strategy space of player i , and $d_r: \mathbb{N} \rightarrow \mathbb{N}$ a delay function associated with resource r . By m we denote $|\mathcal{R}|$, and we denote by $S = (S_1, \dots, S_n)$ the *state of the game* in which player i plays strategy $S_i \in \Sigma_i$. For a state S , we define the *congestion* $n_r(S)$ on resource r by $n_r(S) = |\{i \mid r \in S_i\}|$, that is, $n_r(S)$ is the number of players sharing resource r in state S . Every player i acts selfishly and wishes to play a strategy $S_i \in \Sigma_i$ that minimizes her individual delay, which is defined as $\sum_{r \in S_i} d_r(n_r(S))$. We call a state S a *pure Nash equilibrium*¹ if given the strategies of the other players, no player can decrease her delay by changing her strategy. It is shown in [Rosenthal 73] that every congestion game possesses at least one Nash equilibrium by considering the potential function $\phi: \Sigma_1 \times \dots \times \Sigma_n \rightarrow \mathbb{N}$ with $\phi(S) = \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r(S)} d_r(i)$. A congestion game is called a *singleton* if each strategy space Σ_i consists only of sets with cardinality one. The current state S of a singleton congestion game can be written as $S = (r_1, \dots, r_n)$, meaning that player i currently accesses resource r_i .

2.2. Player-Specific Congestion Games

Player-specific congestion games are congestion games in which every player i has her own delay function $d_r^i: \mathbb{N} \rightarrow \mathbb{N}$ for each resource r . The delay of player i is then computed with respect to the functions d_r^i .

2.3. Player-Specific Congestion Games with Priorities

We define player-specific congestion games with priorities to be a generalization of player-specific congestion games in which each resource r assigns a *priority* or *rank* $\text{rk}_r(i)$ to every player i . For a state S , let $\text{rk}_r(S) = \max_{i: r \in S_i} \text{rk}_r(i)$. We say that player i *accesses* resource r if $r \in S_i$, and we say that player i is *assigned* to resource r if $r \in S_i$ and $\text{rk}_r(i) = \text{rk}_r(S)$. We define $n_r^*(S)$ to be the number of players that are assigned to resource r , that is, the number of players i with $r \in S_i$ and $\text{rk}_r(i) = \text{rk}_r(S)$. The delay that an assigned player i incurs on r is $d_r^i(n_r^*(S))$. Players who access a resource r but are not assigned to it incur an infinite delay on resource r . Congestion games with priorities but without

¹In this paper, the term *Nash equilibrium* always refers to a *pure* Nash equilibrium.

player-specific delay functions are defined in the same way, except that instead of player-specific delay functions d_r^i there is only one delay function d_r for each resource r . We say that the priorities are *consistent* if the priorities assigned to the players by different resources coincide.

2.4. Two-Sided Markets

A *two-sided market* consists of two disjoint sets $\mathcal{N} = \{1, \dots, n\}$ and \mathcal{R} with $|\mathcal{R}| = m$. We use the terms *players* and *agents* to denote elements from \mathcal{N} , and we use the terms *resources* and *markets* to denote elements from \mathcal{R} . In a two-sided market, every player can be *matched* to one resource, and every resource can be matched to one player. We assume that with every pair $(i, r) \in \mathcal{N} \times \mathcal{R}$, a payoff $p_{i,r}$ is associated and that player i receives payoff $p_{i,r}$ if she is matched to resource r . Hence, the payoffs describe implicitly for each player a preference list with respect to the resource. Additionally, we assume that every resource has a strict preference list with respect to the players, which is independent of the payoffs. Each player $i \in \mathcal{N}$ can *propose* to a resource $r_i \in \mathcal{R}$. Given a *state* $S = (r_1, \dots, r_n)$, each resource $r \in \mathcal{R}$ is matched to the *winner of r* , which is the player whom r ranks highest among all players $i \in \mathcal{N}$ with $r = r_i$. If i is the winner of r , she gets a payoff of $p_{i,r}$. If a player proposes to a resource won by another player, she receives no payoff at all. We say that S is a *stable matching* if none of the players can unilaterally increase her payoff by changing her proposal given the proposals of the other players. That is, for each player i who is assigned to a resource r_i , each resource r from which she receives a higher payoff than from r_i is matched to a player whom r prefers to i .

2.5. Two-Sided Markets with Ties

We define a *two-sided market with ties* to be a two-sided market in which the preference lists of the resources can have ties. Given a vector of proposals $S = (r_1, \dots, r_n)$, we say that a player $i \in \mathcal{N}$ is matched to resource $r \in \mathcal{R}$ if $r = r_i$ and if there is no player $j \in \mathcal{N}$ such that $r = r_j$ and j is strictly preferred to i by r . For a resource r , we denote by $n_r(S)$ the number of players proposing to r and by $n_r^*(S)$ the number of players that are matched to r . We assume that every player i has a nonincreasing payoff function $p_r^i: \mathbb{N} \rightarrow \mathbb{N}$ for every resource r . A player i who is matched to resource r receives a payoff of $p_r^i(n_r^*(S))$. Also for two-sided markets with ties, we call a state S a *stable matching* if none of the players can increase her payoff given the proposals of the other players.

2.6. Correlated Two-Sided Markets with Ties

In *correlated two-sided markets with ties*, the preferences of players and resources are correlated. We assume as well that the preference lists of the resources are chosen according to the payoffs that are associated with the pairs from $\mathcal{N} \times \mathcal{R}$. That is, a player $i \in \mathcal{N}$ is preferred to a player $j \in \mathcal{N}$ by resource $r \in \mathcal{R}$ if and only if $p_{i,r} > p_{j,r}$. Due to this construction, if two players i and j are both matched to a resource r , then the payoffs $p_{i,r}$ and $p_{j,r}$ must be the same. We denote this payoff by $p_r(S)$, and we assume that it is split among the players that are matched to r . The payoff that a player receives who is matched to r is specified by a function $q_r(p_r(S), n_r^*(S))$ with $q_r(p_r(S), 1) = p_r(S)$ that is nonincreasing in the number of players matched to r .

2.7. Player-Specific Matroid Congestion Games with Priorities

In a player-specific *matroid* congestion game with priorities, each strategy space Σ_i must be the set of bases of a matroid over the set of resources. A set system $(\mathcal{R}, \mathcal{I})$ with $\mathcal{I} \subseteq 2^{\mathcal{R}}$ is said to be a matroid if $X \in \mathcal{I}$ implies $Y \in \mathcal{I}$ for all $Y \subseteq X$ and if for every $X, Y \in \mathcal{I}$ with $|Y| < |X|$ there exists an $x \in X$ with $Y \cup \{x\} \in \mathcal{I}$. A *basis* of a matroid $(\mathcal{R}, \mathcal{I})$ is an inclusion-maximal set $X \in \mathcal{I}$. Every basis of a matroid has the same cardinality, which is called the *rank* of the matroid. For a *matroid* congestion game Γ , we denote by $\text{rk}(\Gamma)$ the maximal rank of one of the strategy spaces of the players. Examples of matroid congestion games are singleton games and games in which the resources are the edges of a graph and every player has to access a spanning tree. Again, these games can also be seen as an extension of two-sided markets in which each player can propose to

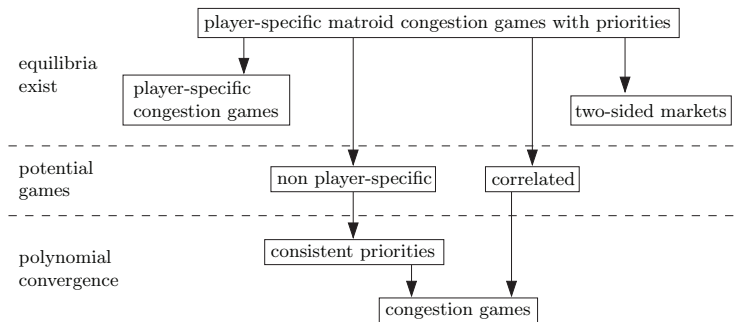


Figure 1. For games on the upper level, equilibria can be computed in polynomial time, games on the mid-level are potential games, and games on the lower level converge in a polynomial number of rounds. An arrow from A to B indicates that B is a special case of A .

a subset of resources instead of only one, so-called *many-to-one markets*, and in which the preference lists of the resources can have ties.

Figure 1 shows a summary of our results and the models we consider.

3. Singleton Congestion Games with Priorities

In this section, we consider singleton congestion games with priorities but without player-specific delay functions. For games with consistent priorities, we show that the better-response dynamics reaches a Nash equilibrium after a polynomial number of *rounds*. We use the term “round” to denote a sequence of activations of players in which every player gets at least once the chance to improve. For example, our result implies that a polynomial (expected) number of better responses suffices if players are activated in a round-robin fashion or uniformly at random. We also prove that games in which different resources can assign different priorities to the players are potential games. We leave open the question whether they converge in a polynomial number of rounds.

Theorem 3.1. *In singleton congestion games with consistent priorities, the better-response dynamics reaches a Nash equilibrium after a polynomial number of rounds.*

Proof. It is proved in [Jeong et al. 05] that in singleton congestion games, every sequence of better responses terminates in a Nash equilibrium after a polynomial number of steps. Since the players with the highest priority are not affected by the other players, the result of Jeong et al. shows that after a polynomial number of rounds, none of them any longer has an incentive to change her strategy. From that point on, the strategies of these players are fixed, and we can again apply the result by Jeong et al. to the players with the second-highest priority. After a polynomial number of rounds, none of them has any longer an incentive to change her strategy. After that, the argument can be applied to the players with the third-highest priority, and so on. \square

Next we consider congestion games in which different resources can assign different priorities to the players.

Theorem 3.2. *Singleton congestion games with priorities are potential games.*

Proof. We set $\mathcal{D} = (\mathbb{N} \cup \{\infty\}) \times \mathbb{N}$, and for elements $x = (x_1, x_2) \in \mathcal{D}$ and $y = (y_1, y_2) \in \mathcal{D}$ we denote by $<$ the lexicographic order on \mathcal{D} in which the first

component is to be minimized and the second component is to be maximized, i.e., we define $x < y$ if $x_1 < y_1$ or if $x_1 = y_1$ and $x_2 > y_2$. We construct a potential function $\Phi: \Sigma_1 \times \dots \times \Sigma_n \rightarrow \mathcal{D}^n$ that maps every state $S = (r_1, \dots, r_n)$ to a vector of values from \mathcal{D} . In state S , every resource $r \in \mathcal{R}$ contributes $n_r(S)$ values to the vector $\Phi(S)$, and $\Phi(S)$ is obtained by sorting all values contributed by the resources in nondecreasing order according to the lexicographic order defined above. Resource r contributes the values $(d_r(1), \text{rk}_r(S)), \dots, (d_r(n_r^*(S)), \text{rk}_r(S))$ to the vector $\Phi(S)$ and $n_r(S) - n_r^*(S)$ times the value $(\infty, 0)$. We claim that if state S' is obtained from S by letting one player play a better response, then $\Phi(S')$ is lexicographically smaller than $\Phi(S)$, i.e., there is k with $\Phi_j(S) = \Phi_j(S')$ for all $j < k$ and $\Phi_k(S') < \Phi_k(S)$.

Assume that in state S , player i plays a better response by changing her allocation from resource r_i to resource r'_i . We compare the two vectors $\Phi(S)$ and $\Phi(S')$, and we show that the smallest element added to the potential vector is smaller than the smallest element removed from the potential vector, showing that the potential decreases lexicographically. Due to the strategy change of player i , either the value $(d_{r_i}(n_{r_i}^*(S)), \text{rk}_{r_i}(S))$ or the value $(\infty, 0)$ is replaced by the value $(d_{r'_i}(n_{r'_i}^*(S')), \text{rk}_{r'_i}(S'))$. Since player i plays a better response, $d_{r'_i}(n_{r'_i}^*(S')) < d_{r_i}(n_{r_i}^*(S))$ or $d_{r'_i}(n_{r'_i}^*(S')) < \infty$, respectively, and hence the term added to the potential is smaller than the term removed from the potential. In the following we show that all values that are contained in $\Phi(S)$ but not in $\Phi(S')$ are larger than $(d_{r'_i}(n_{r'_i}^*(S')), \text{rk}_{r'_i}(S'))$. Clearly, only terms for the resources r_i and r'_i change, and we can restrict our considerations to these two resources.

Let us consider resource r_i first. If the rank of r_i does not decrease by the strategy change of player i or if no player accesses resource r_i in state S' , then only the term $(d_{r_i}(n_{r_i}^*(S)), \text{rk}_{r_i}(S))$ or $(\infty, 0)$ is no longer contained in the vector $\Phi(S')$. All other terms contributed by resource r_i do not change. If the rank of resource r_i is decreased by the strategy change of player i , then additionally some terms $(\infty, 0)$ in the potential are replaced by other terms. Obviously, the removed terms $(\infty, 0)$ are larger than $(d_{r'_i}(n_{r'_i}^*(S')), \text{rk}_{r'_i}(S'))$.

Now we consider resource r'_i . If the rank of r'_i does not increase by the strategy change of player i or if no player accesses r'_i in state S , then only the term $(d_{r'_i}(n_{r'_i}^*(S')), \text{rk}_{r'_i}(S'))$ is added to the potential. All other terms contributed by r'_i do not change. If the rank of r'_i is increased by the strategy change of player i , then additionally the terms $(d_{r'_i}(1), \text{rk}_{r'_i}(S)), \dots, (d_{r'_i}(n_{r'_i}^*(S)), \text{rk}_{r'_i}(S))$ are replaced by $n_{r'_i}^*(S)$ terms $(\infty, 0)$. In this case, $n_{r'_i}^*(S') = 1$, and because $\text{rk}_{r'_i}(S') > \text{rk}_{r'_i}(S)$, the smallest removed term, $(d_{r'_i}(1), \text{rk}_{r'_i}(S))$, is larger than $(d_{r'_i}(1), \text{rk}_{r'_i}(S')) = (d_{r'_i}(n_{r'_i}^*(S')), \text{rk}_{r'_i}(S'))$. \square

4. Player-Specific Singleton Congestion Games with Priorities

In this section we consider singleton congestion games with priorities and player-specific delay functions, and we show that these games always possess Nash equilibria. Our proof also yields an efficient algorithm for finding an equilibrium.

Theorem 4.1. *Every player-specific singleton congestion game with priorities possesses a pure Nash equilibrium that can be computed in polynomial time with $O(m \cdot n^3)$ strategy changes.*

Proof. In order to compute an equilibrium, we compute a sequence of states S^0, \dots, S^k such that S^0 is the state in which no player accesses a resource and S^k is a state in which every player accesses a resource. Keep in mind that in Section 2.3, we distinguish between accessing a resource and being assigned to it. Our construction ensures the invariant that in each state S^a in this sequence, every player who accesses a resource has no incentive to change her strategy. Clearly, this invariant is true for S^0 , and it implies that S^k is a pure Nash equilibrium.

In state S^a we pick an arbitrary player i who is accessing no resource and we let her play her best response. If in state S^a there is no resource to which i can be assigned, then i can access an arbitrary resource without affecting the players who are already accessing a resource and hence without affecting the invariant. It remains to consider the case that after her best response, player i is assigned to a resource r . If we leave the strategies of the other players unchanged, then the invariant may no longer be true after the strategy change of player i . The invariant can, however, be false only for players who are assigned to resource r in state S^a . We distinguish two cases in order to describe how the strategies of these players are modified in order to maintain the invariant.

First we consider the case that the rank of resource r does not change by the strategy change of player i . If there is a player j who is assigned to resource r in S^a and who can improve her strategy after i is also assigned to r , then we change the strategy of j to the empty set, i.e., in state S^{a+1} player j belongs to the set of players who do not access any resource. Besides this, no further modifications of the strategies are necessary because no other player is affected by the replacement of j by i on resource r . In the case that the rank of resource r increases by the strategy change of player i , all players who are assigned to resource r in state S^a are set to their empty strategy in S^{a+1} .

It remains only to show that the described process terminates after a polynomial number of strategy changes in a stable state. We prove this by a potential function that is the lexicographic order of two components. The most important

component is the sum of the ranks of the resources, i.e., $\sum_{r \in \mathcal{R}} \text{rk}_r(S^a)$, which is to be maximized. Observe that this sum does not decrease in any of the two aforementioned cases, and that it increases strictly in the second case. Thus we need to show that after a polynomial number of consecutive occurrences of the first case, the second case must occur. Therefore, we need a second and less-important component in our potential function.

In order to define this component, we associate with every pair $(i, r) \in \mathcal{N} \times \mathcal{R}$ for which i is assigned to r in state S^a a *tolerance* $\text{tol}_a(i, r)$ that describes how many players (including i) can be assigned to r without changing the property that r is an optimal strategy for i , i.e.,

$$\min\{\max\{b \mid \text{in } S^a, r \text{ is best response for } i \text{ if } i \text{ shares } r \text{ with } b - 1 \text{ players}\}, n\}.$$

The second component of the potential function is the sum of the tolerances of the assigned pairs in S^a , which is to be maximized. We denote the set of assignments in state S^a by $E^a \subseteq \mathcal{N} \times \mathcal{R}$ and define the potential function as

$$\Phi(S^a) = \left(\sum_{r \in \mathcal{R}} \text{rk}_r(S^a), \sum_{(i,r) \in E^a} \text{tol}_a(i, r) \right).$$

In every occurrence of the first case, the second component increases by at least 1. Since the values of the components are bounded from above by $m \cdot n$ and n^2 and bounded below away from 0, the potential function implies that there can be at most $m \cdot n^3$ strategy changes before an equilibrium is reached. This does not include the last strategy change of players who are not assigned to any resource in the final state. In their last strategy change, these players access an arbitrary resource, which does not affect the potential. However, there are fewer than n such strategy changes. \square

Let us remark that the potential function does not imply that the considered games are potential games, because it increases only if the strategy changes are made according to the above-described policy.

5. Correlated Two-Sided Markets with Ties

In this section we analyze the better-response dynamics for correlated two-sided markets with ties, and we show that these games are potential games.

Theorem 5.1. *Correlated two-sided markets with ties are potential games.*

Proof. We define a potential function $\Phi: \Sigma_1 \times \dots \times \Sigma_n \rightarrow \mathbb{N}^n$ that is similar to the one used in the proof of Theorem 3.2, and we show that it increases strictly with

every better response that is played. Again each resource r contributes $n_r(S)$ values to the potential, namely the values $q_r(p_r(S), 1), \dots, q_r(p_r(S), n_r^*(S))$ and $n_r(S) - n_r^*(S)$ times the value 0. In the potential vector $\Phi(S)$, all these values are sorted in nonincreasing order. A state S' has a higher potential than a state S if $\Phi(S')$ is lexicographically larger than $\Phi(S)$, i.e., if there exists an index k such that $\Phi_j(S) = \Phi_j(S')$ for all $j < k$ and $\Phi_k(S) < \Phi_k(S')$.

Let S denote the current state and assume that there exists one player $i \in \mathcal{N}$ who plays a better response, leading to state S' . We show that $\Phi(S')$ is lexicographically larger than $\Phi(S)$. Assume that i changes her proposal from r_i to r'_i . Since i plays a better response, she must be matched to r'_i in state S' . That is, the value $q_{r'_i}(p_{i,r'_i}, n_{r'_i}^*(S'))$ is added to the potential. We show that only smaller values are removed from the potential, implying that the potential must lexicographically increase. If i is matched to r_i in state S , then only the value $q_{r_i}(p_{r_i}(S), n_{r_i}^*(S))$ is removed from the vector, and perhaps, if $n_{r_i}^*(S) = 1$, some 0 values are replaced by larger values. Since player i plays a better response, $q_{r_i}(p_{r_i}(S), n_{r_i}^*(S)) < q_{r'_i}(p_{i,r'_i}, n_{r'_i}^*(S'))$. If $n_{r'_i}^*(S') = 1$ and there are players assigned to r'_i in state S , then also the values $q_{r'_i}(p_{r'_i}(S), 1), \dots, q_{r'_i}(p_{r'_i}(S), n_{r'_i}^*(S))$ are removed from the potential vector. In this case, player i displaces the previously assigned players from resource r'_i , which implies $q_{r'_i}(p_{i,r'_i}, n_{r'_i}^*(S')) = q_{r'_i}(p_{i,r'_i}, 1) > q_{r'_i}(p_{r'_i}(S), 1)$, as desired. \square

6. Extensions to Matroid Strategy Spaces

In this section we study player-specific congestion games with priorities in which each strategy space Σ_i consists of the bases of a matroid over the resources. For this setting, we generalize the results that we obtained for the singleton case.

Theorem 6.1. *In matroid congestion games with consistent priorities, the best-response dynamics reaches a Nash equilibrium after a polynomial number of rounds.*

For matroid congestion games, it is known that every sequence of best responses reaches a Nash equilibrium after a polynomial number of steps [Ackermann et al. 08]. Using this result yields the theorem analogously to the proof of Theorem 3.1.

Given a state S of a matroid congestion game with priorities, we call a better response of a player $i \in \mathcal{N}$ from S_i to S'_i *lazy* if it can be decomposed into a sequence of strategies $S_i = S_i^0, S_i^1, \dots, S_i^k = S'_i$ such that $|S_i^{j+1} \setminus S_i^j| = 1$ and the delay of player i in state S_i^{j+1} is strictly smaller than her delay in state S_i^j for

all $j \in \{0, \dots, k - 1\}$. That is, a lazy better response can be decomposed into a sequence of exchanges of single resources such that each step strictly decreases the delay of the corresponding player. In [Ackermann et al. 09], it is observed that for matroid strategy spaces, there always exists a best response that is lazy. In particular, the best response that exchanges the least number of resources is lazy, and in singleton games every better response is lazy.

Theorem 6.2. *Matroid congestion games with priorities are potential games with respect to lazy better responses.*

Since lazy better responses can be decomposed into exchanges of single resources, the same potential function as in the proof of Theorem 3.2 also works for the matroid case. The restriction to lazy better responses in Theorem 6.2 is necessary, as shown by the following result.

Theorem 6.3. *The best-response dynamics in matroid congestion games with priorities can cycle.*

Proof. Let $\mathcal{N} = \{1, 2\}$ denote the players and let $\mathcal{R} = \{a, b, c, d\}$ denote the resources. The set of strategies of player 1 is

$$\Sigma_1 = \{\{a\}, \{d\}\},$$

and the set of strategies of player 2 is

$$\Sigma_2 = \{\{a, b\}, \{b, c\}, \{c, d\}, \{a, d\}\}.$$

Resource a assigns a higher priority to player 2, and resource d assigns the same priority to both players. The delay functions are chosen as follows:

$$d_a(1) = 1, \quad d_b(1) = 3, \quad d_c(1) = 1, \quad d_d(1) = 2, \quad d_d(2) = 4.$$

Given these delays, the following sequence of states is a cycle in the best-response dynamics:

$$(\{d\}, \{a, d\}) \rightarrow (\{d\}, \{b, c\}) \rightarrow (\{a\}, \{b, c\}) \rightarrow (\{a\}, \{a, d\}) \rightarrow (\{d\}, \{a, d\}).$$

This completes the proof. \square

Similar arguments as for Theorem 4.1 yield the following generalization.

Theorem 6.4. *Every player-specific matroid congestion game Γ with priorities possesses a pure Nash equilibrium that can be computed in polynomial time with $O(m \cdot n^3 \cdot \text{rk}(\Gamma))$ strategy changes.*

Proof. For $i \in \mathcal{N}$, we denote by \mathcal{I}_i the set $\{X \mid X \subseteq Y \in \Sigma_i\}$, and we assume that the set system $\mathcal{M}_i = (\mathcal{R}, \mathcal{I}_i)$ is a matroid. We use the same arguments as in the proof of Theorem 4.1, that is, we compute a sequence of states S^0, \dots, S^k such that S^0 is the state in which every player accesses the empty set and S^k is a Nash equilibrium. In contrast to the definition of matroid congestion games, where each player i is required to access a basis from Σ_i , we also allow *partial strategies* from \mathcal{I}_i in states S^a with $a < k$. To be precise, in states S^a with $a < k$ it can happen that the set of resources that a player accesses is a strict subset of a basis. For a player $i \in \mathcal{N}$, let $\mathcal{R}_i^a \subseteq \mathcal{R}$ denote the set of resources to which she can be assigned in state S^a , i.e., \mathcal{R}_i^a contains exactly those resources that are in state S^a not assigned to a player that they strictly prefer to i . Let $\mathcal{M}_i^a = (\mathcal{R}_i^a, \Sigma_i^a)$ denote the matroid that is obtained from \mathcal{M}_i by deleting all resources in $\mathcal{R} \setminus \mathcal{R}_i^a$. The following invariant will be true for all states S^a .

Invariant 6.5. *For every player $i \in \mathcal{N}$, there exists a basis $B_i^a \in \Sigma_i^a$ of the matroid \mathcal{M}_i^a with $S_i^a \subseteq B_i^a$ that has minimum delay given the partial strategies of the other players in S^a .*

That is, if the other players do not change their strategies, no player is forced to leave resources that she currently accesses in order to obtain a basis with minimum delay. If the basis B_i^a of the matroid \mathcal{M}_i^a is not a basis of the matroid \mathcal{M}_i , then player i has no strategy with finite delay given the partial strategies of the other players in S^a .

Now we describe how state S^{a+1} is obtained from state S^a . If in state S^a every player i accesses a basis of the matroid \mathcal{M}_i^a , then due to the invariant, an equilibrium S^{a+1} is obtained from S^a by letting each player i access an arbitrary basis B_i^{a+1} of \mathcal{M}_i with $S_i^a \subseteq B_i^{a+1}$. Assume that there exists a player $i \in \mathcal{N}$ who is not accessing a basis of \mathcal{M}_i^a . In order to obtain S^{a+1} , we choose an arbitrary resource $r \in B_i^a \setminus S_i^a$ and let player i access r , i.e., we set $S_i^{a+1} = S_i^a \cup \{r\}$. Obviously, after this strategy change, player i is assigned to resource r , since $r \in \mathcal{R}_i^a$. If we leave all other strategies unchanged, then the invariant may no longer be true.

We distinguish three cases in order to determine the strategies of the other players in state S^{a+1} :

1. If no player accesses r in S^a , then $S_j^{a+1} = S_j^a$ for all $j \in \mathcal{N} \setminus \{i\}$.
2. If i is ranked higher in r 's preference list than the players assigned to r in S^a , then resource r is removed from the strategies of all players assigned to r in S^a , i.e., for all these players j we set $S_j^{a+1} = S_j^a \setminus \{r\}$. The strategies of all other players remain as in S^a .

3. If i is tied in r 's preference list with the players assigned to r in state S^a , then we check whether the invariant remains true if additionally i is assigned to r . If this is not the case, then we remove one player k from r for whom the invariant becomes false, i.e., we set $S_k^{a+1} = S_k^a \setminus \{r\}$ and $S_j^{a+1} = S_j^a$ for all $j \in \mathcal{N} \setminus \{i, k\}$.

First we show that the invariant remains true in all three cases. This is based on the following property of matroids, which is proven in [Ackermann et al. 09].

Lemma 6.6. *Let $(\mathcal{R}, \mathcal{I})$ be a matroid with weights $w: \mathcal{R} \rightarrow \mathbb{N}$ and let B be a basis of minimum weight. If the weight of a single resource $r \in B$ is increased such that B is no longer of minimum weight, then in order to obtain a basis of minimum weight, it suffices to exchange r with a resource $r' \in \mathcal{R}$ of minimum weight such that $B \cup \{r'\} \setminus \{r\}$ is a basis.*

Consider the first case and assume that the invariant is true in state S^a . Since no player is assigned to resource r in state S^a , there is no player whose current delay is increased by assigning i to r , but there can be players $j \in \mathcal{N}$ with $r \in B_j^a$. For these players, either B_j^a is still a basis of minimum delay or, due to Lemma 6.6, they can choose a basis B_j^{a+1} with $S_j^a \subseteq B_j^{a+1}$ of minimum delay given that i is assigned to r . Since players $j \in \mathcal{N}$ with $r \notin B_j^a$ are not affected by the strategy change of player i , the invariant is also true in state S^{a+1} . In the second case, the invariant remains true for all players who are assigned to r in state S^a , because they just need to exchange r with another resource to obtain a basis with minimum delay again due to Lemma 6.6.

The invariant remains true for all other players j with $r \in B_j^a$ due to Lemma 6.6, and again players $j \in \mathcal{N}$ with $r \notin B_j^a$ are not affected by the strategy change of player i . In the third case, for all players $j \in \mathcal{N} \setminus \{i, k\}$ the effects of the strategy changes of i and k cancel each other out, and hence these players are not affected by the strategy changes of i and k . The invariant remains true for k due to Lemma 6.6.

It remains only to show that the described process terminates after a polynomial number of strategy changes in an equilibrium. This follows by the same potential function as in the proof of Theorem 4.1. The upper bound on the second component of the potential function increases by a factor of $\text{rk}(\Gamma)$, which accounts for the increased number of strategy changes. \square

The previous proof relies on the fact that players who cannot be assigned to a complete basis do not leave the game completely but still access a basis. In fact, it is important only that every player i accesses a basis of the matroid obtained from \mathcal{M}_i by deleting the resources to which she cannot be assigned. If

we redefine the model that players who cannot be assigned to a complete basis do not access any resource, then in general, Nash equilibria do not exist, as the following example shows.

Consider a game with three resources $\mathcal{R} = \{r_1, r_2, r_3\}$ and two players who want to access two resources each, i.e., their strategy spaces are the bases of a uniform matroid of rank 2. Assume that the resources r_1 and r_2 prefer player 1 to player 2 and the players are tied in the preference list of r_3 . If the player-specific delay functions are defined as follows, then no pure Nash equilibrium exists:

	r_1		r_2		r_3	
congestion	1	2	1	2	1	2
player 1	1	-	2	-	1	3
player 2	1	-	1	-	1	2

A state in which the first player accesses the resources r_1 and r_2 cannot be a pure Nash equilibrium, because in such a state the second player cannot be assigned to a basis, and hence the resource r_3 would be empty. In that case, the first player prefers $\{r_1, r_3\}$ to her current strategy $\{r_1, r_2\}$. A state in which the first player accesses the resources r_i and r_3 for either $i = 1$ or $i = 2$ cannot be a Nash equilibrium either, because in such a state, the second player would access the resources r_{3-i} and r_3 , which in turn implies that the first player would prefer $\{r_1, r_2\}$ to her current choice $\{r_{3-i}, r_3\}$.

Next, we consider correlated two-sided matroid markets with ties. Since lazy better responses can be decomposed into exchanges of single resources, the potential function defined in the proof of Theorem 5.1 also works for matroid strategy sets if players play only lazy better responses.

Theorem 6.7. *Correlated two-sided matroid markets with ties are potential games with respect to lazy better responses.*

The restriction in Theorem 6.7 to lazy better responses is necessary.

Theorem 6.8. *The best-response dynamics in correlated two-sided matroid markets with ties can cycle.*

Proof. We choose \mathcal{N} , \mathcal{R} , and the strategy sets as in Theorem 6.3. The payoffs associated with the possible pairs in $\mathcal{N} \times \mathcal{R}$ are defined as follows:

$$p_{1,a} = 5, \quad p_{1,d} = 3, \quad p_{2,a} = 7, \quad p_{2,b} = 1, \quad p_{2,c} = 7, \quad p_{2,d} = 2.$$

The cycle in the best-response dynamics in Remark 6.3 is also a cycle in this example. \square

7. Conclusions and Open Problems

We have considered a model of player-specific congestion games with priorities. We have shown that pure Nash equilibria exist in these games and that the special cases of non-player-specific and correlated games are potential games. We leave open the question whether the better-response dynamics reaches a Nash equilibrium after a polynomial number of rounds in these special cases. This has been shown only for the special case of non-player-specific congestion games with consistent priorities.

In our model, players displace other players with lower priorities. As we mentioned in the introduction, this is reasonable if players control streams of jobs rather than single ones. It would be interesting to find and analyze different models in which jobs are only slowed down by jobs with higher priorities, i.e., models in which they incur a large but finite delay.

Acknowledgements. We thank Fuhito Kojima for pointing out related work and the reviewers of the conference version for their helpful comments.

This work was supported by DFG grant VO 889/2, EPSRC Grant GR/T07343/02, and by the Ultra High-Speed Mobile Information and Communication Research Cluster (UMIC) established under the excellence initiative of the German government. An extended abstract appeared in *Proc. of the 3rd Int. Workshop on Internet and Network Economics (WINE 2007)*.

References

- [Ackermann et al. 08] H. Ackermann, H. Röglin, and B. Vöcking. “On the Impact of Combinatorial Structure on Congestion Games.” *Journal of the ACM* 55:6 (2008), Article No. 25.
- [Ackermann et al. 09] H. Ackermann, H. Röglin, and B. Vöcking. “Pure Nash Equilibria in Player-Specific and Weighted Congestion Games.” *Theoretical Computer Science* 410:17 (2009), 1552–1563.
- [Anshelevich et al. 04] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. “The Price of Stability for Network Design with Fair Cost Allocation.” In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*, pp. 295–304. Los Alamitos, CA: IEEE Press, 2004.
- [Even-Dar et al. 03] E. Even-Dar, A. Kesselman, and Y. Mansour. “Convergence Time to Nash Equilibria.” In *Automata, Languages and Programming: 30th International Colloquium, ICALP 2003, Eindhoven, The Netherlands, June 30–July 4, 2003*,

- Proceedings*, Lecture Notes in Computer Science 2719, pp. 502–513. Berlin: Springer, 2003.
- [Fabrikant et al. 04] A. Fabrikant, C. Papadimitriou, and K. Talwar. “The Complexity of Pure Nash Equilibria.” In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pp. 604–612. New York: ACM Press, 2004.
- [Fischer et al. 07] S. Fischer, M. Petrova, P. Mähönen, and B. Vöcking. “Distributed Load Balancing Algorithm for Adaptive Channel Allocation for Cognitive Radios.” In *Proceedings of the 2nd International Conference on Cognitive Radio Oriented Wireless Networks and Communications, CrownCom 2007*, pp. 508–513. Los Alamitos, CA: IEEE Press, 2007.
- [Fleiner 03] T. Fleiner. “A Fixed-Point Approach to Stable Matchings and Some Applications.” *Mathematics of Operations Research* 28:1 (2003), 103–126.
- [Fleischer et al. 06] L. Fleischer, M. Goemans, V. S. Mirrokni, and M. Sviridenko. “Tight Approximation Algorithms for Maximum General Assignment Problems.” In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms*, pp. 611–620. New York: ACM Press, 2006.
- [Fotakis et al. 02] D. Fotakis, S. C. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. G. Spirakis. “The Structure and Complexity of Nash Equilibria for a Selfish Routing Game.” In *Automata, Languages and Programming: 29th International Colloquium, ICALP 2002, Malaga, Spain, July 8–13, 2002, Proceedings*, Lecture Notes in Computer Science 2380, pp. 123–134. Berlin: Springer, 2002.
- [Gairing et al. 04] M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. “Computing Nash Equilibria for Scheduling on Restricted Parallel Links.” In *Proceedings of the Thirty-Sixth Annual ACM Symposium on Theory of Computing*, pp. 613–622. New York: ACM Press, 2004.
- [Gale and Shapley 62] D. Gale and L. S. Shapley. “College Admissions and the Stability of Marriage.” *American Mathematical Monthly* 69 (1962), 9–15.
- [Goemans et al. 04] M. Goemans, L. Li, V. S. Mirrokni, and M. Thottan. “Market Sharing Games Applied to Content Distribution in Ad-Hoc Networks.” In *Proceedings of the 5th ACM International Symposium on Mobile Ad Hoc Networking and Computing*, pp. 55–66. New York: ACM Press, 2004.
- [Gusfield and Irving 89] D. Gusfield and R. Irving. *The Stable Marriage Problem: Structure and Algorithms*. Cambridge, MA: MIT Press, 1989.
- [Hoffman et al. 06] C. Hoffman, A. Holroyd, and Y. Peres. “A Stable Marriage of Poisson and Lebesgue.” *Annals of Probability* 34:4 (2006), 1241–1272.
- [Jeong et al. 05] S. Jeong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. “Fast and Compact: A Simple Class of Congestion Games.” In *Proceedings of the 20th National Conference on Artificial Intelligence*, Vol. 2, pp. 489–494. Menlo Park, CA: AAAI Press, 2005.
- [Iwama et al. 99] K. Iwama, D. Manlove, S. Miyazaki, and Y. Morita. “Stable Marriage with Incomplete Lists and Ties.” In *Automata, Languages and Programming: 26th International Colloquium, ICALP’99, Prague, Czech Republic, July 11–15, 1999, Proceedings*, Lecture Notes in Computer Science 1644, pp. 443–452. Berlin: Springer, 1999.

- [Kelso and Crawford 82] A. Kelso and V. Crawford. “Job Matchings, Coalition Formation, and Gross Substitute.” *Econometrica* 50 (1982), 1483–1504.
- [Knuth 76] D. Knuth. *Marriages Stables et leurs relations avec d’autres problèmes Combinatoires*. Montreal: Les Presses de l’Université de Montréal, 1976.
- [Kojima and Ünver 08] F. Kojima and M. U. Ünver. “Random Paths to Pairwise Stability in Many-to-Many Matching Problems: A Study on Market Equilibration.” *Int. Journal of Game Theory* 36:3 (2008), 473–488.
- [Milchtaich 96] I. Milchtaich. “Congestion Games with Player-Specific Payoff Functions.” *Games and Economic Behavior* 13:1 (1996), 111–124.
- [Mirrokni 05] V. S. Mirrokni. “Approximation Algorithms for Distributed and Selfish Agents.” PhD thesis, Massachusetts Institute of Technology, 2005.
- [Rosenthal 73] R. W. Rosenthal. “A Class of Games Possessing Pure-Strategy Nash Equilibria.” *Int. Journal of Game Theory* 2 (1973), 65–67.
- [Roth 84] A. E. Roth. “The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory.” *Journal of Political Economy* 92 (1984), 991–1016.
- [Roth and Sotomayor 90] A. E. Roth and M. A. O. Sotomayor. *Two-Sided Matching: A study in Game-Theoretic Modeling and Analysis*. Cambridge, UK: Cambridge University Press, 1990.

Heiner Ackermann, Fraunhofer Institute for Industrial Mathematics, Fraunhofer-Platz 1, 67663 Kaiserslautern, Germany (heiner.ackermann@itwm.fraunhofer.de)

Paul W. Goldberg, Department of Computer Science, University of Liverpool, Ashton Building, Ashton Street, Liverpool, L69 3BX, UK (P.W.Goldberg@liverpool.ac.uk)

Vahab S. Mirrokni, Google Research, 76 9th Ave., 4th floor, New York, NY 10011 (mirrokni@gmail.com)

Heiko Röglin, Department of Quantitative Economics, Maastricht University, P.O. Box 616, 6200 MD Maastricht, The Netherlands (heiko@roeglin.org)

Berthold Vöcking, Department of Computer Science, RWTH Aachen University, Ahornstr. 55, 52074 Aachen, Germany (voecking@cs.rwth-aachen.de)

Received March 27, 2008; accepted January 27, 2009.