A Unified Approach to Congestion Games and Two-Sided Markets*

Heiner Ackermann[1], Paul W. Goldberg[2],
Vahab S. Mirrokni[3], Heiko Röglin[1], and Berthold Vöcking[1]

[1] Department of Computer Science, RWTH Aachen, Germany
{ackermann,roeglin,voecking}@cs.rwth-aachen.de
[2] Department of Computer Science, University of Liverpool, U. K.
P.W.Goldberg@liverpool.ac.uk
[3] Microsoft Research, Redmond, WA, USA
mirrokni@microsoft.com

Abstract. Congestion games are a well-studied model for resource sharing among uncoordinated selfish agents. Usually, one assumes that the resources in a congestion game do not have any preferences over the players that can allocate them. In typical load balancing applications, however, different jobs can have different priorities, and jobs with higher priorities get, for example, larger shares of the processor time. We introduce a model in which each resource can assign priorities to the players and players with higher priorities can displace players with lower priorities. Our model does not only extend standard congestion games, but it can also be seen as a model of two-sided markets with ties. We prove that singleton congestion games with priorities are potential games, and 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 space of each player consists of the bases of a matroid 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 allocate. The delay of a resource depends on the number of players allocating that resource, and every player is interested in allocating 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, network design, and load balancing [3–5]. One appealing property of congestion games is that they are *potential games* [21]. In particular, this implies that every congestion game possesses a *pure Nash equilibrium* and that myopic player eventually reach a Nash equilibrium by iteratively playing better responses.

* This work was supported by DFG grant VO 889/2, EPSRC Grant GR/T07343/02, and by the EU within the 6th Framework Programme under contract 001907 (DELIS).

One drawback of the standard model of congestion games is that resources do not have any preferences over the players. In typical load balancing applications, however, different jobs 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 only allocated 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 allocate a resource, only players with the highest priority incur a delay depending 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, we mainly consider congestion games in which each player has to choose exactly one resource to allocate, 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 [4, 8, 9]. We show that also singleton congestion games with priorities are potential games, implying that uncoordinated players who iteratively play better responses eventually reach a pure Nash equilibrium. If all resources agree on the priorities, then we even obtain polynomial-time convergence to a Nash equilibrium. Milchtaich [19] introduces *player-specific congestion games* 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 not potential games anymore 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 that can be computed efficiently.

Interestingly, our model of player-specific congestion games with priorities does not only extend congestion games but also the well-known model of *two-sided markets*. This model was introduced by Gale and Shapley [10] to model markets on which different kinds of agents are matched to another, for example men and women, students and colleges [10], interns and hospitals [22], 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). As opposed to congestion games, each resource can only be matched to one player. With each pair of player and resource a payoff is associated, and players are interested in maximizing their payoffs. Hence, the payoffs implicitly define a preference list over the resources for each player. Additionally, each resource has a preference list over 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. 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 current proposals of the other players. Gale and Shapley [10] show 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, the book by Knuth [17], the book by Gusfield and Irving [12], or the book by Roth and Sotomayor [23].

In the same way as it is in many situations not realistic to assume that in congestion games the resources have no preferences over the players, it is in two-sided markets often unrealistic 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 that resource. If the most preferred player is not unique, several players share the payoff of the market. 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 for congestion games with priorities. Two-sided markets with ties have been extensively studied in the literature [12, 15]. 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 are 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 [7, 11, 20]. These markets have also been considered for discovering stable geometric configurations with applications in VLSI design [13].

Additionally, we consider player-specific congestion games with priorities in which the strategy space of each player consists of the bases of a matroid over the resources. For this case, we show that pure Nash equilibria exist that can be computed in polynomial time, extending a result for player-specific congestion games without priorities [2]. 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 [6, 16, 18]. Kelso and Crawford [16] show that if the preference list of every player satisfies a certain *substitutability property*, then stable matchings exist. Kojima and Ünver [18] prove that in this case, from every matching there exists a polynomially long better response sequence to a stable matching. 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 notations used throughout the paper.
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 $d_r \colon \mathbb{N} \to \mathbb{N}$ 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* where 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 *Nash equilibrium* if, given the strategies of the others players, no player can decrease her delay by changing her strategy. Rosenthal [21] shows that every congestion game possesses at least one pure Nash equilibrium by considering the potential function $\phi\colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{N}$ with $\phi(S) = \sum_{r\in\mathcal{R}} \sum_{i=1}^{n_r(S)} d_r(i)$. A congestion game is called *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 allocates resource r_i.
Player-Specific Congestion Games. Player-specific congestion games are congestion games in which every player i has her own delay function $d_r^i \colon \mathbb{N} \to \mathbb{N}$ for each resource r. The delay of player i is then computed with respect to the functions d_r^i.
Player-Specific Congestion Games with Priorities. We define this model to be a generalization of player-specific congestion games in which each resource r assigns a *priority* or *rank* $\mathrm{rk}_r(i)$ to every player i. For a state S, let $\mathrm{rk}_r(S) = \max_{i:r\in S_i} \mathrm{rk}_r(i)$. We say that player i *allocates* resource r if $r \in S_i$, and we say that player i is *assigned* to resource r if $r \in S_i$ and $\mathrm{rk}_r(i) = \mathrm{rk}_r(S)$. 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 $\mathrm{rk}_r(i) = \mathrm{rk}_r(S)$. The delay that an assigned player i incurs on r is $d_r^i(n_r^*(S))$. Players who allocate a resource r but are not assigned to it incur an infinite delay on resource r. Congestion games with priorities but without 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.

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 over the resource. Additionally, we assume that every resource has a strict preference list over 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 over i.

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 non-increasing payoff function $p^i_r \colon \mathbb{N} \to \mathbb{N}$ for every resource r. A player i who is matched to resource r receives a payoff of $p^i_r(n^*_r(S))$. 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.

Correlated Two-sided Markets with Ties. In *correlated two-sided markets with ties*, the preferences of players and resources are correlated. We assume that also 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 over 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 non-increasing in the number of players matched to r.

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 $(\mathcal{R}, \mathcal{I})$

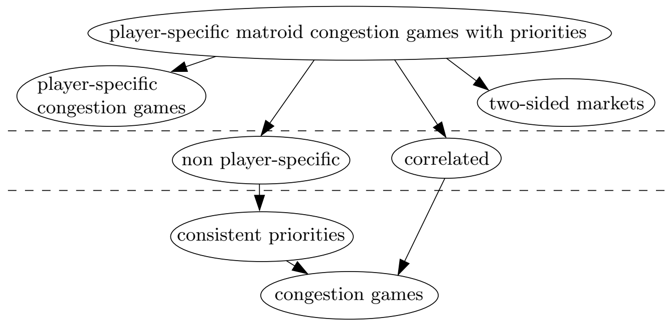

Fig. 1. For games on the upper level, equilibria can be computed in polynomial time, games on the mid-level are potential games, and games on the lower level converge in a polynomial number of rounds.

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 a set $X \in \mathcal{I}$ with maximum cardinality. 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 $\mathrm{rk}(\Gamma)$ 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 allocate a spanning tree. Again, these games can also be seen as an extension of two-sided markets in which each player can propose to 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 1. *In singleton congestion games with consistent priorities, the better response dynamics reaches a Nash equilibrium after a polynomial number of rounds.*

Proof. Ieong et al. [14] prove 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 by Ieong et al. shows that after a polynomial number of rounds, none of them has an incentive to change her strategy anymore. From that point on, the strategies of these players are fixed and we can again apply the result by Ieong et al. to the players with the second highest priority. After a polynomially number of rounds, also none of them has an incentive to change her strategy anymore. After that, the argument can be applied to the players with the third highest priority and so on. □

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

Theorem 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 and only if $x_1 < y_1$ or if $x_1 = y_1$ and $x_2 > y_2$. We construct a potential function $\Phi\colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathcal{D}^n$ that maps every state $S = (r_1, \ldots, r_n)$ 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 non-decreasing order according to the lexicographic order defined above. Resource r contributes the values $(d_r(1), \mathrm{rk}_r(S)), \ldots, (d_r(n_r^*(S)), \mathrm{rk}_r(S))$ 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 a 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 $(d_{r_i}(n_{r_i}^*(S)), \mathrm{rk}_{r_i}(S))$ or the value $(\infty, 0)$ is replaced by the value $(d_{r_i'}(n_{r_i'}^*(S')), \mathrm{rk}_{r_i'}(S'))$. 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 $(d_{r_i'}(n_{r_i'}^*(S')), \mathrm{rk}_{r_i'}(S'))$. 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 allocates resource r_i in state S', then only the term $(d_{r_i}(n_{r_i}^*(S)), \mathrm{rk}_{r_i}(S))$ or $(\infty, 0)$ is not contained in the vector $\Phi(S')$ anymore. 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 $(d_{r'_i}(n^*_{r'_i}(S')), \mathrm{rk}_{r'_i}(S'))$.

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 allocates r'_i in state S, then only the term $(d_{r'_i}(n^*_{r'_i}(S')), \mathrm{rk}_{r'_i}(S'))$ 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 $(d_{r'_i}(1), \mathrm{rk}_{r'_i}(S)), \ldots, (d_{r'_i}(n^*_{r'_i}(S)), \mathrm{rk}_{r'_i}(S))$ are replaced by $n^*_{r'_i}(S)$ terms $(\infty, 0)$. In this case, $n^*_{r'_i}(S') = 1$ and the smallest removed term, $(d_{r'_i}(1), \mathrm{rk}_{r'_i}(S))$, is larger than $(d_{r'_i}(1), \mathrm{rk}_{r'_i}(S')) = (d_{r'_i}(n^*_{r'_i}(S')), \mathrm{rk}_{r'_i}(S'))$ because $\mathrm{rk}_{r'_i}(S') > \mathrm{rk}_{r'_i}(S)$. □

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 3. *Every player-specific singleton congestion game with priorities possesses a pure Nash equilibrium that can be computed in polynomial time by $O(m^2 \cdot n^3)$ strategy changes.*

Proof. In order to compute an equilibrium, we compute a sequence of states $S^0, \ldots, S^k$ such that S^0 is the state in which no player allocates a resource and S^k is a state in which every player allocates a resource. Remember that we distinguish between allocating a resource and being assigned to it. Our construction ensures the invariant that in each state S^a in this sequence, every player who allocates 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 allocating 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 allocate an arbitrary resource without affecting the players who are already allocating 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 not be true anymore after the strategy change of player i. The invariant can, however, only be false for players who are assigned to resource r in state S^a. We distinguish between 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 allocate any resource. Besides this, no further modifications of the strategies are necessary because all other players are not 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 only remains 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., $\sum_{r\in\mathcal{R}} \mathrm{rk}_r(S^a)$, 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 S^a a *tolerance* $\mathrm{tol}_a(i, r)$ 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.,

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

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

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

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 $m \cdot n^2$ and bounded below from 0, the potential function implies that there can be at most $m^2 \cdot n^3$ strategy changes before an equilibrium is reached.

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

Proof. We define a potential function $\Phi\colon \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{N}^n$ that is similar to the one used in the proof of Theorem 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 $q_r(p_r(S), 1), \ldots, q_r(p_r(S), n_r^*(S))$ and $n_r(S) - n_r^*(S)$ times the value 0. In the potential vector $\Phi(S)$, all these values are sorted in non-increasing 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 assigned 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 assigned 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 maybe, 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 $q_{r'_i}(p_{r'_i}(S), 1), \ldots, q_{r'_i}(p_{r'_i}(S), n^*_{r'_i}(S))$ 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. □

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. Due to space limitations, the proofs are omitted.

Theorem 5. *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 [1]. Using this result yields the theorem analogously to the proof of Theorem 1.

Theorem 6. *Matroid congestion games with priorities are potential games with respect to lazy better responses.*

Given a state S, we denote 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 $S_i = S_i^0, S_i^1, \ldots, S_i^k = S'_i$ 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 $j \in \{0, \ldots, k-1\}$. 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 [2], it is observed that for matroid strategy spaces, there does always exist 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. Since lazy best responses can be decomposed into exchanges of single resources, the same potential function as in the proof of Theorem 2 also works for the matroid case. The restriction to lazy better responses in Theorem 6 is necessary.

Remark 7. *The best response dynamics in matroid congestion games with priorities can cycle.*

Similar arguments as for Theorem 3 yield the following generalization.

Theorem 8. *Every player-specific matroid congestion game Γ with priorities possesses a pure Nash equilibrium that can be computed in polynomial time by $O(m^2 \cdot n^3 \cdot \mathrm{rk}(\Gamma))$ strategy changes.*

Since lazy better responses can be decomposed into exchanges of single resources, the potential function defined in the proof of Theorem 4 also works for matroid strategy sets if players play only lazy better responses.

Theorem 9. *Correlated two-sided matroid markets with ties are potential games with respect to lazy better responses.*

The restriction in Theorem 9 to lazy better responses is necessary.

Remark 10. *The best response dynamics in correlated two-sided matroid markets with ties can cycle.*

7 Conclusions and Open Problems

We consider a model of player-specific congestion games with priorities. We show that pure Nash equilibria exist in these games and we show 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 is only shown 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 for their helpful comments.

References

1. H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In *Proc. of the 47th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 613–622, 2006.
2. H. Ackermann, H. Röglin, and B. Vöcking. Pure Nash equilibria in player-specific and weighted congestion games. In *Proc. of the 2nd Int. Workshop on Internet and Network Economics (WINE)*, pages 50–61, 2006.
3. E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. of the 45th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.

4. E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In *Proc. of the 30th Int. Coll. on Automata, Languages and Programming (ICALP)*, pages 502–513, 2003.
5. A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. of the 36th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 604–612, 2004.
6. T. Fleiner. A fixed-point approach to stable matchings and some applications. *Mathematics of Operations Research*, 28(1):103–126, 2003.
7. L. Fleischer, M. Goemans, V. S. Mirrokni, and M. Sviridenko. Tight approximation algorithms for maximum general assignment problems. In *Proc. of the 16th Ann. ACM–SIAM Symp. on Discrete Algorithms (SODA)*, pages 611–620, 2006.
8. 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 *Proc. of the 29th Int. Coll. on Automata, Languages and Programming (ICALP)*, pages 123–134, 2002.
9. M. Gairing, T. Lücking, M. Mavronicolas, and B. Monien. Computing Nash equilibria for scheduling on restricted parallel links. In *Proc. of the 36th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 613–622, 2004.
10. D. Gale and L. S. Shapley. College admissions and the stability of marriage. *American Mathematical Monthly*, 69:9–15, 1962.
11. M. Goemans, L. Li, V. S. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *Proc. of the 5th ACM Int. Symp. on Mobile Ad Hoc Networking and Computing (MobiHoc)*, pages 1020–1033, 2004.
12. D. Gusfield and R. Irving. *The stable Marriage Problem: Structure and Algorithms*. MIT Press, 1989.
13. C. Hoffman, A. Holroyd, and Y. Peres. A stable marriage of poisson and lebesgue. *Annals of Probability*, 34(4):1241–1272, 2006.
14. S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *Proc. of the 20th Nat. Conference on Artificial Intelligence (AAAI)*, pages 489–494, 2005.
15. K. Iwama, D. Manlove, S. Miyazaki, and Y. Morita. Stable marriage with incomplete lists and ties. In *Proc. of the 26th Int. Coll. on Automata, Languages and Programming (ICALP)*, pages 443–452, 1999.
16. A. Kelso and V. Crawford. Job matchings, coalition formation, and gross substitute. *Econometrica*, 50:1483–1504, 1982.
17. D. Knuth. *Marriage Stables et leurs relations avec d'autres problèmes Combinatories.* Les Presses de l'Université de Montréal, 1976.
18. 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*, 2006.
19. I. Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.
20. V. S. Mirrokni. *Approximation Algorithms for Distributed and Selfish Agents.* PhD thesis, Massachusetts Institute of Technology, 2005.
21. R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Int. Journal of Game Theory*, 2:65–67, 1973.
22. A. E. Roth. The evolution of the labor market for medical interns and residents: A case study in game theory. *Journal of Political Economy*, 92:991–1016, 1984.
23. A. E. Roth and M. A. O. Sotomayor. *Two-sided Matching: A study in game-theoretic modeling and analysis.* Cambridge University Press, 1990.

A Proofs from Section 6

Proof (Remark 7). 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:

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

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

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

□

Proof (Theorem 8). Let $\mathcal{M}_i = (\mathcal{R}, \Sigma_i)$ denote the matroid set system that represents the possible strategies of player i. We use the same arguments as in the proof of Theorem 3, that is, we compute a sequence of states $S^0, \ldots, S^k$ such that S^0 is the state in which every player allocates the empty set and S^k is a Nash equilibrium. In contrast to the definition of matroid congestion games, where each player is required to allocate a basis, we also allow *partial strategies* 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 allocates 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 she can be assigned to 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 contracted matroid that is obtained from $\mathcal{M}_i$ by removing all resources $\mathcal{R} \setminus \mathcal{R}_i^a$. The following invariant will be true for all states S^a.

Invariant 11. *For every player $i \in \mathcal{N}$, there exists a basis $\mathcal{B}_i^a \in \Sigma_i^a$ of the contracted matroid $\mathcal{M}_i^a$ with $S_i^a \subseteq \mathcal{B}_i^a$ 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 allocates in order to obtain a basis with minimum delay. If the basis $\mathcal{B}_i^a$ of the contracted 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 allocates a basis of the contracted matroid $\mathcal{M}_i^a$, then an equilibrium S^{a+1} is obtained from S^a by letting each player i allocate an arbitrary basis

$\mathcal{B}_i^{a+1}$ of $\mathcal{M}_i$ with $S_i^a \subseteq \mathcal{B}_i^{a+1}$ due to the invariant. Assume that there exists a player $i \in \mathcal{N}$ who is not allocating a basis of $\mathcal{M}_i^a$. In order to obtain S^{a+1}, we choose an arbitrary resource $r \in \mathcal{B}_i^a \setminus S_i^a$ and let player i allocate r, i. e., we set $S_i^{a+1} = S_i^a \cup \{r\}$. If we leave all other strategies unchanged, then the invariant may not be true anymore.

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

1. If no player allocates 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 stays 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 stays true in all three cases. This is based on the following property of matroids, which is proven in [2].

Lemma 12. *Let $\mathcal{M} = (\mathcal{R}, \mathcal{I})$ be a matroid with weights $w\colon \mathcal{R} \to \mathbb{N}$ and let $\mathcal{B}^*$ be a basis of $\mathcal{M}$ with minimum weight. If the weight of a single resource $r_{opt} \in \mathcal{B}^*$ is increased such that $\mathcal{B}^*$ is no longer of minimum weight, then, in order to obtain a minimum weight basis again, it suffices to exchange r_{opt} with a resource $r^* \in \mathcal{R}$ of minimum weight such that $\mathcal{B}^* \cup \{r^*\} \setminus \{r_{opt}\}$ 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 $r \in \mathcal{B}_j^a$. For these players either $\mathcal{B}_j^a$ is still a basis of minimum delay or, due to Lemma 12, they can choose a basis $\mathcal{B}_j^{a+1}$ with $S_j^a \subseteq \mathcal{B}_j^{a+1}$ of minimum delay given that i is assigned to r. Since players $j \in \mathcal{N}$ with $r \notin \mathcal{B}_j^a$ 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 stays 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 12. It stays true for all other players j with $r \in \mathcal{B}_j^a$ due to Lemma 12, and again players $j \in \mathcal{N}$ with $r \notin \mathcal{B}_j^a$ 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 stays true for k due to Lemma 12.

It only remains 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 3. The upper bound on the second component of the potential function increases by a $\operatorname{rk}(\Gamma)$ factor, which accounts for the increased number of strategy changes. □

Proof (Remark 10). We choose $\mathcal{N}$, $\mathcal{R}$ and the strategy sets as in Remark 7. The payoffs associated with the possible pairs in $\mathcal{N} \times \mathcal{R}$ are defined as follows:

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

The cycle in the best response dynamics in Remark 7 is also a cycle in this example. □