Pure Nash Equilibria in Player-Specific and Weighted Congestion Games ⋆

Heiner Ackermann, Heiko Röglin, Berthold Vöcking

RWTH Aachen University, Lehrstuhl für Informatik I, D-52056 Aachen, Germany

Abstract

Unlike standard congestion games, weighted congestion games and congestion games with player-specific delay functions do not necessarily possess pure Nash equilibria. It is known, however, that there exist pure equilibria for both of these variants in the case of *singleton congestion games*, i. e., if the players' strategy spaces contain only sets of cardinality one. In this paper, we investigate how far such a property on the players' strategy spaces guaranteeing the existence of pure equilibria can be extended. We show that both weighted and player-specific congestion games admit pure equilibria in the case of *matroid congestion games*, i. e., if the strategy space of each player consists of the bases of a matroid on the set of resources. We also show that the matroid property is the maximal property that guarantees pure equilibria without taking into account how the strategy spaces of different players are interweaved.

Additionally, our analysis of player-specific matroid congestion games yields a polynomial time algorithm for computing pure equilibria. We also address questions related to the convergence time of such games. For player-specific matroid congestion games, in which the best response dynamics may cycle, we show that from every state there exists a short sequences of better responses to an equilibrium. For weighted matroid congestion games, we present a superpolynomial lower bound on the convergence time of the best response dynamics showing that players do not even converge in pseudopolynomial time.

Key words: congestion games, Nash equilibria, matroids

⋆ This work was supported in part by the EU within the 6th Framework Programme under contract 001907 (DELIS) and by DFG grant Vo889/2-1. A preliminary version appeared in the Proceedings of the 2nd International Workshop on Internet & Network Economics (WINE 2006).

Email addresses: `ackermann@cs.rwth-aachen.de` (Heiner Ackermann), `roeglin@cs.rwth-aachen.de` (Heiko Röglin), `voecking@cs.rwth-aachen.de` (Berthold Vöcking).

1 Introduction

Congestion games are a natural model for resource allocation in large networks like the Internet. It is assumed that n players share a set $\mathcal{R}$ of m resources. Players are interested in subsets of resources. For example, the resources may correspond to the edges of a graph, and each player may want to allocate a spanning tree of this graph. The delay (cost, negative payoff) of a resource depends on the number of players that allocate the resource, and the delay of a set of allocated resources corresponds to the sum of the delays of the resources in the set. A well known potential function argument of Rosenthal [19] shows that congestion games always possess Nash equilibria [1], i. e., allocations of resources from which no player wants to deviate unilaterally.

The existence of Nash equilibria gives a natural solution concept for congestion games. Unfortunately, this property does not hold anymore if we slightly extend the class of considered games towards congestion games with player-specific delay functions, i. e., games in which different players may have different delay functions, and weighted congestion games, i. e., games in which different players may have different impacts on the delays of the resources they allocate. For both of these classes one can easily construct examples of games that do not possess Nash equilibria. In this paper, we study which conditions on the strategy spaces of individual players guarantee the existence of Nash equilibria. We only consider games with non-decreasing delay functions since otherwise one can construct examples of player-specific or weighted *singleton congestion games*, i. e., games in which the players' strategy spaces contain only sets of cardinality one, that do not possess Nash equilibria.

It is known, however, that there exist Nash equilibria for both of these variants in the case of singleton congestion games with non-decreasing delay functions [7,16]. We extend these results and show that both player-specific and weighted congestion games admit pure equilibria in the case of *matroid congestion games*, i. e., if the strategy space of each player consists of the bases of a matroid on the set of resources. We also show that the matroid property is the maximal condition on the players' strategy spaces that guarantees Nash equilibria without taking into account how the strategy spaces of different players are interweaved. Our negative result shows that for every non-matroid set system there exist a weighted and a player-specific congestion game in which the strategy space of each player is isomorphic to the given set system and that does not possess a Nash equilibrium.

In the case of player-specific matroid congestion games, our analysis also yields a polynomial time algorithm for computing pure equilibria. As the best re-

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

sponse dynamics of such a game may cycle [16], we address the question whether a Nash equilibrium can be found by players iteratively playing better or best responses. Again, we extend results from player-specific singleton congestion games, and show that from every state of such a game there exists a polynomially long sequence of better responses leading to an equilibrium.

For weighted matroid congestion games we do not have an efficient algorithm for computing a Nash equilibrium, but we show that players playing "lazy best responses" reach a Nash equilibrium after a finite number of steps, where a best response is *lazy* if it exchanges the least number of resources compared to the current strategy among all best responses. We also address the convergence time to Nash equilibria in weighted matroid games, and present a superpolynomial lower bound. That is, we present a family of games that possess superpolynomially long best response sequences. Similar results have already been presented by Even-Dar et al. [5]. However, they use players with exponentially large weights, whereas the weights in our construction are polynomially bounded. This implies that players do not converge to a Nash equilibrium in pseudopolynomial time.

1.1 Related Work

Milchtaich [16] considers player-specific singleton congestion games and shows that every such game possesses at least one Nash equilibrium. His existence proof implicitly contains an efficient algorithm for computing an equilibrium. Additionally, he shows that players iteratively playing best responses in such games do not necessarily reach a Nash equilibrium, that is, the best response dynamics may cycle. However, he also shows that from every state of such a game there exists a polynomially long sequence of best responses to a Nash equilibrium. Our work generalizes Milchtaich's analysis from singleton congestion games towards matroid congestion games. Gairing et al. [10] consider the class of player-specific singleton congestion games with linear delay functions without offsets and show that the best response dynamics of such games do not cycle. Milchtaich [17] also observes that player-specific network congestion games do not possess Nash equilibria in general. Ackermann and Skopalik [2] prove that the related decision problem is NP-complete.

Milchtaich [16] also addresses the existence of Nash equilibria in congestion games which are both player-specific and weighted. In this case, a Nash equilibrium does not necessarily exist in singleton congestion games. However, Georgiou et al. [11] and Garing et al. [10] conjecture that these games possess Nash equilibria in the case of linear player-specific delay functions without offsets.

Fotakis et al. [7] consider a selfish routing game in which the players are weighted and their strategy spaces are singleton sets. They show that in this game at least one Nash equilibrium always exists and that players iteratively playing best responses converge to such an equilibrium. Our proof that every weighted matroid congestion game possesses at least one Nash equilibrium reworks the proof in [7]. Even-Dar et al. [5] consider the same game with respect to the convergence time. They distinguish between different types of players' weights and different delay functions, and show that players do not necessarily converge quickly in any of these scenarios.

Fotakis et al. [8] consider weighted network congestion games in which the strategy space of each player corresponds to the set of all paths between possibly different sources and sinks in a network. They show that Nash equilibria do not necessarily exist in these games. On the positive side, they show that every weighted network congestion game possesses a Nash equilibrium if the delay of every resource equals its congestion. Dunkel and Schulz [4] show that it is NP-hard to decide whether a given weighted network congestion game possesses a Nash equilibrium. Milchtaich [17] considers player-specific or weighted network congestion games and tries to characterize which networks possess pure Nash equilibria independent of the number of players, and independent of any assumption on the (player-specific) delay functions except monotonicity.

It is interesting to relate the results about the existence of Nash equilibria in player-specific and weighted matroid congestion games to our recent work on the convergence time of standard congestion games: In [1] we characterize the class of congestion games that admit polynomial time convergence to a Nash equilibrium. Motivated by the fact that in singleton congestion games players converge quickly [14], we show that if the strategy space of each player consists of the bases of a matroid on the set of resources, then players iteratively playing best responses reach a Nash equilibrium in polynomial time. Furthermore, we show that the matroid property is the maximal condition on the players' strategy spaces for guaranteeing polynomial time convergence to a Nash equilibrium if one does not take into account how the players' strategy spaces are interweaved.

1.2 Formal Definition of 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, $\mathcal{R} = \{1, \dots, m\}$ the set of resources, $\Sigma_i \subseteq 2^{\mathcal{R}}$ the strategy space of player i, and $d_r : \mathbb{N} \to \mathbb{N}$ a delay function associated with resource r. We call a congestion game *symmetric* if all players share the same set of strategies, otherwise we call it *asymmetric*. We denote by

$S = (S_1, \ldots, S_n)$ the *state of the game* where player i plays strategy $S_i \in \Sigma_i$. Furthermore, we denote by $S \oplus S'_i$ the state $S' = (S_1, \ldots, S_{i-1}, S'_i, S_{i+1}, \ldots, S_n)$, i. e., the state S except that player i plays strategy S'_i instead of S_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. Players act selfishly and like to play a strategy $S_i \in \Sigma_i$ minimizing their individual delay. The delay $\delta_i(S)$ of player i in state S is given by $\delta_i(S) = \sum_{r \in S_i} d_r(n_r(S))$. Given a state S, we call a strategy S^*_i a *best response* of player i to S if, for all $S'_i \in \Sigma_i$, $\delta_i(S \oplus S^*_i) \leq \delta_i(S \oplus S'_i)$. In the following, we use the term *best response sequence* to denote a sequence of consecutive strategy changes in which each step is a best response which strictly decreases the delay of the corresponding player. Furthermore, we call a state S a *Nash equilibrium* if no player can decrease her delay by changing her strategy, i. e., for all $i \in \mathcal{N}$ and for all $S'_i \in \Sigma_i$, $\delta_i(S) \leq \delta_i(S \oplus S'_i)$. Rosenthal [19] shows that every congestion game possesses at least one Nash equilibrium by considering the potential function $\phi : \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{N}$ with $\phi(S) = \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r(S)} d_r(i)$.

There are two well known extensions of congestion games, namely *player-specific congestion games* and *weighted congestion games*. In a player-specific congestion game every player i has its own delay function $d^i_r : \mathbb{N} \to \mathbb{N}$ for every resource $r \in \mathcal{R}$. Given a state S, the delay of player i is defined as $\delta_i(S) = \sum_{r \in S_i} d^i_r(n_r(S))$. In a weighted congestion game every player $i \in \mathcal{N}$ has a weight $\omega_i \in \mathbb{N}$. Given a state S, we define the congestion on resource r by $n_r(S) = \sum_{i : r \in S_i} \omega_i$, that is, $n_r(S)$ is the total weight of all players sharing resource r in state S.

1.3 Matroids and Matroid Congestion Games

We now introduce *matroid congestion games*. Before we give a formal definition of such games we shortly introduce matroids. For a detailed discussion, we refer the reader to [20].

Definition 1 *A tuple $\mathcal{M} = (\mathcal{R}, \mathcal{I})$ is a matroid if $\mathcal{R} = \{1, \ldots, m\}$ is a finite set of resources and $\mathcal{I}$ is a nonempty family of subsets of $\mathcal{R}$ such that, if $I \in \mathcal{I}$ and $J \subseteq I$, then $J \in \mathcal{I}$, and, if $I, J \in \mathcal{I}$ and $|J| < |I|$, then there exists an $i \in I \setminus J$ with $J \cup \{i\} \in \mathcal{I}$.*

Let $I \subseteq \mathcal{R}$. If $I \in \mathcal{I}$, then we call I an *independent set*, otherwise we call it *dependent*. It is well known that all maximal independent sets of $\mathcal{I}$ have the same cardinality. The *rank* $rk(\mathcal{M})$ of the matroid $\mathcal{M}$ is the cardinality of the maximal independent sets. A maximal independent set B is called a *basis* of $\mathcal{M}$. If additionally a weight function $w : \mathcal{R} \to \mathbb{N}$ is given, $\mathcal{M}$ is called a *weighted matroid* and one is usually interested in finding a basis of minimum weight, where the weight of an independent set I is given by $w(I) = \sum_{r \in I} w(r)$. It is well known that such a basis can be found by a greedy algorithm. Now we are ready to define matroid congestion games.

Definition 2 *We call a congestion game* $\Gamma = (\mathcal{N}, \mathcal{R}, (\Sigma_i)_{i \in \mathcal{N}}, (d_r)_{r \in \mathcal{R}})$ *a* matroid congestion game *if for every player* $i \in \mathcal{N}$, $\mathcal{M}_i := (\mathcal{R}, \mathcal{I}_i)$ *with* $\mathcal{I}_i = \{I \subseteq S \mid S \in \Sigma_i\}$ *is a matroid and* Σ_i *is the set of bases of* M_i. *Additionally, we denote by* $rk(\Gamma) = \max_{i \in \mathcal{N}} rk(\mathcal{M}_i)$ *the* rank of the matroid congestion game Γ.

The obvious application of matroid congestion games are network design problems in which players compete for the edges of a graph in order to build a spanning tree [21]. There are also other interesting applications as even simple matroid structures like uniform matroids, that are rather uninteresting from an optimization point of view, lead to rich combinatorial structures when various players with possibly different strategy spaces are involved. Illustrative examples based on uniform matroids are market sharing games with uniform market costs [12] and scheduling games in which each player has to injectively allocate a given set of tasks (services) to a given set of machines (servers).

Let us remark that, in the case of matroid congestion games, the assumption that all delays are positive is not a restriction. Since all strategies have the same size, one can easily shift all delays by the same value in order to obtain positive delays without changing the better and best response dynamics.

2 Player-Specific Matroid Congestion Games

In this section, we consider player-specific matroid congestion games with non-decreasing player-specific delay functions and prove that every such game possesses at least one Nash equilibrium. The proof we present extends techniques invented for singleton congestion games [16] towards matroid congestion games, and implicitly describes an efficient algorithm to compute an equilibrium of such games.

It is known that the best response dynamics of a player-specific singleton congestion game may cycle, i. e., a sequence of best responses starting in a state S may return to state S. Thus, player-specific singleton congestion games are no potential games. However, if the players play their best responses in a certain order, then they quickly find a Nash equilibrium, that is, from every state of such a game, there exists a sequence of best responses of polynomial length to a Nash equilibrium [16]. Hence, if the players play best responses in a random order, then the expected number of best responses needed to reach a Nash equilibrium is finite. In Section 2.2, we investigate whether a similar property also holds for matroid congestion games. We show that from every state of a player-specific matroid congestion game there exists a sequence of *better* responses of polynomial length leading to a Nash equilibrium. We call such a sequence of better responses an *improvement path*, and we leave it as an open question whether short sequences of *best* responses always exist for player-specific matroid congestion games.

2.1 Existence of Nash Equilibria

Theorem 3 *Every player-specific matroid congestion game* Γ *with non-decreasing delay functions possesses a Nash equilibrium.*

Proof: Recall that since the strategy space of player i corresponds to the set of bases of a matroid $\mathcal{M}_i$, all strategies of player i have the same size $rk(\mathcal{M}_i)$. In the following, we represent a strategy of player i by $rk(\mathcal{M}_i)$ tokens that the player places on the resources she allocates. Suppose that we reduce the number of tokens of some of the players, that is, player i has $k_i \leq rk(\mathcal{M}_i)$ tokens that she places on the resources of an independent set of cardinality k_i. Observe that the independent sets of cardinality k_i form the bases of a matroid $\mathcal{M}'_i$ whose independent sets correspond to those independent sets of $\mathcal{M}_i$ with cardinality at most k_i. The matroid $\mathcal{M}_i$ is also called the k_i-*truncation* of the matroid $\mathcal{M}_i$. Hence, a game in which some of the players have a reduced number of tokens is also a matroid congestion game.

We prove the theorem by induction on the total number of tokens $\tau = \sum_{i\in\mathcal{N}} rk(\mathcal{M}_i)$ that the players are allowed to place, that is, we prove the existence of Nash equilibria for a sequence of games $\Gamma_0, \Gamma_1, \ldots, \Gamma_\tau$, where $\Gamma_{\ell+1}$ is obtained from Γ_ℓ by giving one more token to one of the players. Γ_0 is the game in which each player has only the empty strategy. Obviously, Γ_0 has only one state and this state is a Nash equilibrium.

As induction hypothesis assume that player i has placed $k_i \geq 0$ tokens, for $1 \leq i \leq n$, and this placement corresponds to a Nash equilibrium of the player-specific matroid congestion game $\Gamma_\ell = (\mathcal{N}, \mathcal{R}, (\Sigma_i^{k_i})_{i\in\mathcal{N}}, (d_r^i)_{i\in\mathcal{N}, r\in\mathcal{R}})$ with $\ell = \sum_{i\in\mathcal{N}} k_i$, in which the set of strategies $\Sigma_i^{k_i}$ coincides with the k_i-truncation of $\mathcal{M}_i$.

Now assume that some player i_0 has to place an additional token t_0. We show how to compute a Nash equilibrium for the game $\Gamma_{\ell+1}$ obtained from a Nash equilibrium of Γ_ℓ by changing i_0's strategy space to the set of independent sets of size $k_{i_0} + 1$. Since an optimal basis of a matroid can be computed by a greedy algorithm, there exists a resource r_0 such that placing the token t_0 on r_0 gives an independent set of size $k_{i_0} + 1$ with minimum delay among all independent sets of the same size. Thus, assuming that the tokens of the other players are fixed, an optimal strategy for player i_0 is to place t_0 on r_0 and leave all other tokens unchanged. However, as the congestion on r_0 is increased by one, other players may want to move their tokens from r_0 in order to obtain a better independent set. We now use matroid properties to show that a Nash equilibrium of $\Gamma_{\ell+1}$ can be reached with at most $n \cdot m \cdot rk(\Gamma)$ moves of tokens.

Lemma 4 *Let $\mathcal{M} = (\mathcal{R}, \mathcal{I})$ be a matroid with weights $w : \mathcal{R} \to \mathbb{N}$ and let B_{opt} be a basis of $\mathcal{M}$ with minimum weight. If the weight of a single resource $r_{opt} \in B_{opt}$ is increased such that B_{opt} 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 $B_{opt} \cup \{r^*\} \setminus \{r_{opt}\}$ is a basis.*

Proof: In order to prove the lemma we use the following property of a matroid $\mathcal{M} = (\mathcal{R}, \mathcal{I})$. For a proof of this property, we refer the reader to Lemma 39.12 from [20].

Proposition 5 *(Schrijver [20]). Let $\mathcal{M} = (\mathcal{R}, \mathcal{I})$ be a matroid, and let $I, J \in \mathcal{I}$ with $|I| = |J|$ be independent sets. The bipartite graph $G(I \Delta J) = (V, E)$ with $V = (I \setminus J) \cup (J \setminus I)$ and $E = \{\{i, j\} \mid i \in I \setminus J, j \in J \setminus I, I \cup \{j\} \setminus \{i\} \in \mathcal{I}\}$ contains a perfect matching.*

Let B'_{opt} be a minimum weight basis w. r. t. the increased weight of r_{opt}. Let P be a perfect matching of the graph $G(B_{opt} \Delta B'_{opt})$ and denote by e the edge from P that contains r_{opt}. For every edge $\{r, r'\} \in P \setminus \{e\}$, it holds $w(r) \leq w(r')$ as, otherwise, if $w(r) > w(r')$, the basis $B_{opt} \cup \{r'\} \setminus \{r\}$ would have smaller weight than B_{opt}.

Now denote by r'_{opt} the resource that is matched with r_{opt}, i. e., the resource such that $e = \{r_{opt}, r'_{opt}\} \in P$. As $w(r) \leq w(r')$ for every $\{r, r'\} \in P \setminus \{e\}$, the weight of $B_{opt} \setminus \{r_{opt}\}$ is bounded from above by the weight of $B'_{opt} \setminus \{r'_{opt}\}$. By the definition of the matching P, $B_{opt} \cup \{r'_{opt}\} \setminus \{r_{opt}\}$ is a basis. By our arguments above, the weight of this basis is bounded from above by the weight of B'_{opt}. Hence, this basis is optimal w. r. t. the increased weight of r_{opt}. $\square$

After placing token t_0 of player i_0 on resource r_0, resource r_0 has one additional token in comparison to the initial Nash equilibrium S_ℓ of the game Γ_ℓ. Since we assume non-decreasing delay functions, only the players with a token on r_0 might now have an incentive to change their strategies. Let i_1 be one of these players. It follows from Lemma 4 that i_1 has a best response in which she moves a token t_1 from resource r_0 to another resource that we call r_1. Now r_1 is the only resource with one additional token in comparison to S_ℓ. Suppose we have not yet reached a Nash equilibrium. Only those players with a token on r_1 might have an incentive to change their strategies. Again by applying Lemma 4, we can identify a player i_2 that has a best response in which she moves a token t_2 from r_1 to a resource r_2, which then is the only resource with one additional token.

The token migration process described above can be continued in the same way until it reaches a Nash equilibrium of the game $\Gamma_{\ell+1}$. The correctness of the process is ensured by the following invariant.

Invariant 6 *For every $j \geq 0$, after player i_j moves token t_j onto resource r_j,*

a) only players with a token on r_j may violate the Nash equilibrium condition,
b) the Nash equilibrium condition of all players would be satisfied if one ignores the additional token on r_j, that is, if each player calculates the delay on r_j as if there were one token less on this resource.

The invariant follows by induction on j: For player i_j the invariant is satisfied as this player plays a best response according to Lemma 4. Thus she satisfies the Nash equilibrium condition even without virtually reducing the congestion on r_j. For all other players, the validity of the invariant for j follows directly from the validity of the invariant for $j-1$ as these players do not move their tokens.

Thus, in order to show the existence of a Nash equilibrium for $\Gamma_{\ell+1}$, it suffices to show that the token migration process is finite. Consider an arbitrary token t of player i. For a resource r, let $D_i(r)$ denote the delay of i on r if r has one more token than in the initial state S. Whenever t is moved by the migration process from a resource r to a resource r', it must be $D_i(r) > D_i(r')$. Hence, the token t can visit each resource at most once during the token migration process. As there are at most $n \cdot rk(\Gamma)$ tokens, the migration process terminates after at most $n \cdot m \cdot rk(\Gamma)$ steps in a Nash equilibrium of $\Gamma_{\ell+1}$. □

Observe that the proof of Theorem 3 implicitly describes an efficient algorithm to compute a Nash equilibrium with at most $n^2 \cdot m \cdot rk^2(\Gamma)$ moves of tokens.

Corollary 7 *There exists a polynomial time algorithm to compute a Nash equilibrium of a player-specific matroid congestion game with non-decreasing player-specific delay functions.*

2.2 Existence of Short Improvement Paths

Theorem 8 *Let Γ be a player-specific matroid congestion game with non-decreasing delay functions, and let S be an arbitrary state of Γ. Then there exists a better response sequence of length at most $2 \cdot n^2 \cdot m \cdot rk^2(\Gamma)$ which starts in state S and terminates in a Nash equilibrium.*

Proof: The proof uses similar arguments as the proof of Theorem 3, except that initially every player places all her tokens. After the first placement of the tokens, which corresponds to the given state S, we assume that all tokens are *deactivated*, i. e., players are not allowed to move them in order to decrease their delays. We then consider a sequence of games $\Gamma_0, \ldots, \Gamma_\tau$, where $\Gamma_{\ell+1}$ is obtained from Γ_ℓ by activating one more token. We can achieve that deactivated tokens are not moved by setting the delay of the corresponding player on the corresponding resource to 0. Then activating a token corresponds to restoring the delay to its real value. Thus, each game Γ_ℓ is a player-specific matroid congestion game. Given a Nash equilibrium S_ℓ of Γ_ℓ, we show that there exists a short improvement path in $\Gamma_{\ell+1}$ from the former equilibrium S_ℓ to a Nash equilibrium $S_{\ell+1}$ of $\Gamma_{\ell+1}$. Obviously, by concatenating all these paths we obtain an improvement path from S to a Nash equilibrium of Γ.

As induction hypothesis assume that ℓ tokens have been activated so far and that we are given a Nash equilibrium S_ℓ of Γ_ℓ. Suppose now, that an additional token t_0 of player i_0 is activated, and that i_0 moves t_0 to a resource r_1 in order to decrease her delay. After that, we are in a situation similar to the one in the proof of Theorem 3, that is, the congestion on one resource r_1 is increased by one compared to the Nash equilibrium S_ℓ of Γ_ℓ. In contrast to the situation in the proof of Theorem 3, in which the congestion of the other resources remained unchanged, there exists a resource r_0 whose congestion is decreased by one compared to the congestion in Γ_ℓ. Assume that we place a dummy token on resource r_0 which artificially increases the congestion by one. In this case, we can consider the same token migration process as in the proof of Theorem 3.

In contrast to the previous proof, there are two different ways how this process can terminate. If the process returns to r_0, i. e., if it moves a token onto r_0, we terminate the process and remove the dummy token from r_0. If the process does not return to r_0, then it is not affected by the dummy token and by the same arguments as in the proof of Theorem 3 it follows that it terminates after at most $n \cdot m \cdot rk(\Gamma)$ moves of tokens.

In the first case, if at some time a player moves a token t_j from a resource r_{j-1} to the resource $r_j = r_0$, then after removing the dummy token from r_0 we have reached a Nash equilibrium of $\Gamma_{\ell+1}$ due to Invariant 6. Since the resource r_0 is not involved in the previous moves of tokens, each of these movements reduces the delay of the corresponding player also in the game $\Gamma_{\ell+1}$ without the dummy token. In the last step a player moves a token onto r_0 and improves her delay even if the dummy token is present. Hence, she also decreases her delay in $\Gamma_{\ell+1}$ without the dummy token.

In the second case, we have almost reached a Nash equilibrium. That is, all players were satisfied if we would not remove the dummy token. Suppose now that we remove the dummy token. As the delay functions are non-decreasing, only players who can move tokens onto r_0 may have an incentive to change their strategies. The following lemma, which is a slight variation of Lemma 4, ensures that players who have an incentive to change their strategies with respect to the tokens they are allowed to move only need to move a token onto r_0 in order to obtain a best response.

Lemma 9 *Let $\mathcal{M} = (\mathcal{R}, \mathcal{I})$ be a matroid with weights $w : \mathcal{R} \to \mathbb{N}$ and let B_{opt} be a basis of $\mathcal{M}$ with minimum weight. If the weight of a single resource $r^* \in \mathcal{R} \setminus B_{opt}$ is decreased such that B_{opt} is no longer of minimum weight, then, in order to obtain a minimum weight basis again, it suffices to exchange r^* with a resource $r_{opt} \in B_{opt}$ of maximum weight such that $B_{opt} \cup \{r^*\} \setminus \{r_{opt}\}$ is a basis.*

The proof of Lemma 9 follows the same line of arguments as the proof of Lemma 4 and is therefore omitted. Suppose now that player i'_0 moves a token t'_0 from resource r'_1 to r_0. Afterwards, the congestion on r_0 equals the congestion in the former equilibrium with respect to the dummy token, and the congestion on r'_1 is decreased by one. Again only players who can move a token onto r'_1 have an incentive to change their strategy. We can continue this process obtaining an additional token migration process in which a token t_{j+1} moves to the resource from which token t_j was removed. As before, we have to show that this token migration process is finite and terminates in a Nash equilibrium of $\Gamma_{\ell+1}$. The fact that it terminates in a Nash equilibrium is ensured by the following invariant which is a slight variation of Invariant 6.

Invariant 10 *For every $j \geq 0$, after player i'_j removes token t'_j from resource r'_{j+1},*

a) only players who can move a token onto r'_{j+1} may violate the Nash equilibrium condition,
b) the Nash equilibrium condition of all players would be satisfied if one ignores the missing token on r'_{j+1}, that is, if each player calculates the delay on r'_{j+1} as if there were one additional token on this resource.

Invariant 10 can be proven analogously to Invariant 6. Its proof is therefore omitted. It remains to show that the second token migration process is also finite. Again, the same arguments as in the proof of Theorem 3 show that this is true, and we conclude that the second process terminates after at most $n \cdot m \cdot rk(\Gamma)$ moves of tokens in a Nash equilibrium of $\Gamma_{\ell+1}$.

Altogether, we have shown that there exists an improvement path of length $2 \cdot n \cdot m \cdot rk(\Gamma)$ from S_ℓ to a Nash equilibrium of $\Gamma_{\ell+1}$. As the number of tokens τ is upper bounded by $n \cdot rk(\Gamma)$, the theorem follows. $\square$

3 Weighted Matroid Congestion Games

In this section we consider weighted matroid congestion games with non-decreasing delay functions and show that every such game possesses a Nash equilibrium. Moreover, we show that myopic players do not necessarily reach such an equilibrium if they iteratively play arbitrary best responses. We show, however, that players who are only allowed to play best responses that exchange the least number of resources compared to their current strategies eventually reach a Nash equilibrium. We call such best responses *lazy best responses* and define them formally as follows.

Definition 11 *Given a state S, we call a best response S_i^* of player i* lazy *if it can be decomposed into a sequence of strategies $S_i = S_i^0, S_i^1, \ldots, S_i^k = S_i^*$ with $|S_i^{j+1} \setminus S_i^j| = 1$ and $\delta_i(S \oplus S_i^{j+1}) < \delta_i(S \oplus S_i^j)$, for $0 \leq j < k$.*

From the following proposition one can easily conclude that whenever a player can decrease her delay, then there exists a lazy best response for this player. For a proof we refer the reader to Lemma 39.12 in [20].

Proposition 12 *(Schrijver [20]) Given a matroid $\mathcal{M} = (\mathcal{R}, \mathcal{I})$ with weights $w : \mathcal{R} \to \mathbb{N}$, a basis $B \in \mathcal{I}$ is a minimum weight basis of $\mathcal{M}$ if and only if there exists no basis $B^* \in \mathcal{I}$ with $|B \setminus B^*| = 1$ and $w(B^*) < w(B)$.*

We are now ready to prove that weighted matroid congestion games possess Nash equilibria.

Theorem 13 *Every weighted matroid congestion game Γ with non-decreasing delay functions possesses a Nash equilibrium. Furthermore, players reach an equilibrium after a finite number of lazy best responses.*

Proof: Let S be a state of Γ. With each resource r, we associate a pair $z_r(S) = (d_r(n_r(S)), n_r(S))$ consisting of the delay and the congestion of r in state S. For two resources r and r' and states S and S', let $z_r(S) \geq z_{r'}(S')$ if and only if $d_r(n_r(S)) > d_{r'}(n_{r'}(S'))$ or $d_r(n_r(S)) = d_{r'}(n_{r'}(S'))$ and $n_r(S) \geq n_{r'}(S')$. Let $z_r(S) > z_{r'}(S')$ if and only if $z_r(S) \geq z_{r'}(S')$ and $z_r(S) \neq z_{r'}(S')$. Let $\bar{z}(S)$ denote a vector containing the pairs $z_r(S)$ of all resources $r \in \mathcal{R}$ in non-increasing order, that is, $\bar{z}_j(S) \geq \bar{z}_{j+1}(S)$, where $\bar{z}_j(S)$ denotes the j-th component of $\bar{z}(S)$, for $1 \leq j < |\mathcal{R}|$.

We denote by $<_{\text{lex}}$ the lexicographic order among the vectors $\bar{z}(S)$, i.e., $\bar{z}(S_1) <_{\text{lex}} \bar{z}(S_2)$ if there exists an index l such that $\bar{z}_k(S_1) = \bar{z}_k(S_2)$, for all $k < l$, and $\bar{z}_l(S_1) < \bar{z}_l(S_2)$.

Due to Lemma 12, in every state S which is not a Nash equilibrium there exists at least one player i who can decrease her delay by playing a lazy best response S_i^*. Since S_i^* is a lazy best response, there exists a sequence of strategies $S_i = S_i^0, \ldots, S_i^k = S_i^*$ such that, for every $0 \leq j < k$, $|S_i^{j+1} \setminus S_i^j| = 1$ and

$$\delta_i(S) = \delta_i(S \oplus S_i^0) > \delta_i(S \oplus S_i^1) > \cdots > \delta_i(S \oplus S_i^k) = \delta_i(S \oplus S_i^*) \ .$$

We now claim that $\bar{z}(S \oplus S_i^{j+1}) <_{\text{lex}} \bar{z}(S \oplus S_i^j)$, for every $0 \leq j < k$. Let r_j be the unique resource in S_i^j that is not contained in S_i^{j+1} and let r_j^* be the unique resource that is contained in S_i^{j+1} but not in S_i^j. Since the delay decreases strictly with the exchange, we have

$$d_{r_j}(n_{r_j}(S \oplus S_i^j)) > d_{r_j^*}(n_{r_j^*}(S \oplus S_i^{j+1})) \ .$$

Additionally, since we assume non-decreasing delay functions,

$$d_{r_j}(n_{r_j}(S \oplus S_i^j)) \geq d_{r_j}(n_{r_j}(S \oplus S_i^j) - \omega_i) = d_{r_j}(n_{r_j}(S \oplus S_i^{j+1})) \ .$$

Furthermore, $n_{r_j}(S \oplus S_i^j) > n_{r_j}(S \oplus S_i^{j+1})$. Combining these inequalities implies $z_{r_j}(S \oplus S_i^j) > z_{r_j}(S \oplus S_i^{j+1})$ and $z_{r_j}(S \oplus S_i^j) > z_{r_j^*}(S \oplus S_i^{j+1})$. This yields

$$\max\left\{z_{r_j}(S \oplus S_i^{j+1}), z_{r_j^*}(S \oplus S_i^{j+1})\right\} < \max\left\{z_{r_j}(S \oplus S_i^j), z_{r_j^*}(S \oplus S_i^j)\right\}$$

and hence $\bar{z}(S \oplus S_i^j) >_{\text{lex}} \bar{z}(S \oplus S_i^{j+1})$. That is, the lexicographic order decreases with every exchange and, hence, with every lazy best response. This concludes the proof of the theorem. □

Theorem 13 shows that the number of lazy best responses needed to reach a Nash equilibrium is bounded from above by

$$\min\left\{\left(\sum_{i=1}^{n} \omega_i\right)^m, \binom{m}{rk(\Gamma)}^n\right\} \ .$$

The first term is an upper bound on the maximal number of different vectors $\bar{z}(S)$ and the second one bounds the number of different states of the matroid congestion game Γ. Below we present an example showing that arbitrary best responses do not necessarily lead to a Nash equilibrium. In singleton congestion games, every best response is a lazy best response. Hence, in these games, players playing iteratively best responses always reach a Nash equilibrium. A lower bound on the convergence time in this case is presented in Section 5.

Theorem 14 *The best response dynamics of a weighted matroid congestion game with non-decreasing delay functions can cycle.*

Proof: Consider a weighted matroid congestion game with four resources $\{1,2,3,4\}$ and two players with weights $\omega_1 = 1$ and $\omega_2 = 2$. We define the strategy spaces as follows:

$$\Sigma_1 = \{\{1\},\{3\}\} \qquad \Sigma_2 = \{\{1,2\},\{1,3\},\{2,4\},\{3,4\}\} \ .$$

Observe that both strategy spaces are sets of bases of matroids on subsets of the resources. Additionally, we define non-decreasing delay functions. A dash denotes a value we do not have to care about.

	$n_r = 1$	$n_r = 2$	$n_r = 3$
$\delta_1(n_1)$	2	20	20
$\delta_2(n_2)$	-	9	-
$\delta_3(n_3)$	4	8	10
$\delta_4(n_4)$	-	20	-

Now consider the following cycle of states:

$$(\{3\},\{1,3\}) \rightarrow (\{3\},\{2,4\}) \rightarrow (\{1\},\{2,4\}) \rightarrow (\{1\},\{1,3\}) \rightarrow (\{3\},\{1,3\}) \ .$$

Each strategy change induces a set of inequalities in order to be a best response. One can easily verify that all these inequalities are satisfied by the above defined delay functions. Hence, players playing arbitrary best responses do not necessarily converge to a Nash equilibrium in weighted matroid congestion games. □

The delay functions in the previous example are non-decreasing but not strictly increasing. We leave open the question whether in arbitrary weighted matroid congestion games with strictly increasing delay functions players always converge to an equilibrium.

4 Non-Matroid Strategy Spaces

In the previous two sections, we showed that the matroid property is a sufficient condition on the combinatorial structure of the players' strategy spaces for guaranteeing the existence of Nash equilibria in player-specific or weighted congestion games with non-decreasing delay functions. In this section, we show that the matroid property is also the maximal condition for guaranteeing the existence of Nash equilibria in such games.

Our negative result shows that for every non-matroid set system there exist a weighted and a player-specific congestion game in which the strategy space of each player is isomorphic to the given set system and that does not possess a Nash equilibrium. Our construction assumes that the strategy spaces of different players can be interweaved appropriately. Let us remark that the delay functions are positive and increasing.

If one drops this assumption and considers special classes of congestion games in which the delay functions and/or the way of how the strategy spaces can be interweaved are restricted, then one can identify larger classes of weighted and player-specific congestion games that possess Nash equilibria. For instance, Fotakis et al. [8] prove that every weighted network congestion game possesses a Nash equilibrium if one additionally assumes that the delay on an edge equals the current congestion on that edge. Often there exists a common combinatorial interpretation of the resources in a congestion game; they can, for instance, be the edges of a graph and every player might want to allocate a path in that graph between a given source/sink pair. This restricts the way of how the strategy spaces of different players can be interweaved. For example, Milchtaich [17] shows that every player-specific or weighted network congestion game possesses an equilibrium if the network graph belongs to a certain restricted class of graphs.

Observe that our negative results show that our positive results are tight. In Theorems 3 and 13 we show that every player-specific or weighted congestion game possesses a Nash equilibrium if the strategy space of each player corresponds to the bases of a matroid, regardless of how the strategy spaces of different players are interweaved and for every choice of non-decreasing delay functions. Our negative results show that such a positive result cannot be extended further without placing additional assumptions on the delay functions and/or the relation of the strategy spaces. In addition to that, we also demonstrate that the way of how the strategy spaces are interweaved in our construction is not too restrictive to apply to natural classes of congestion games by showing that our construction can, for instance, easily be embedded into (symmetric) network congestion games.

4.1 A Characterization of Non-Matroid Set Systems

Let Σ be a set system on a set $\mathcal{R}$ of resources. The set system Σ is called an *anti-chain* if for every $X \in \Sigma$, no proper superset $Y \supset X$ belongs to Σ. Moreover, we call Σ a *non-matroid set system* if the tuple $(\mathcal{R}, \{X \subseteq S \mid S \in \Sigma\})$ is not a matroid. In [1] we show that every non-matroid anti-chain possesses the $(1, 2)$-exchange property. Here we need the following variant of this property with positive (instead of non-negative) delays.

Definition 15 (**$(1,2)$-exchange property**) *Let Σ be an anti-chain on a set of resources $\mathcal{R}$. We say that Σ satisfies the* $(1,2)$-exchange property *if we can identify three distinct resources $a, b, c \in \mathcal{R}$ with the property that for every given $k \in \mathbb{N}$ with $k > |\mathcal{R}|$, we can choose a delay $d(r) \in \{1, k + |\mathcal{R}|\}$ for every $r \in \mathcal{R} \setminus \{a, b, c\}$ such that for every choice of the delays of a, b, and c with $|\mathcal{R}| \leq d(a), d(b), d(c) \leq k$, the following property is satisfied: If $d(a) + |\mathcal{R}| \leq d(b) + d(c)$, then for every set $S \in \Sigma$ with minimum delay, $a \in S$ and $b, c \notin S$. If $d(a) \geq d(b) + d(c) + |\mathcal{R}|$, then for every set $S \in \Sigma$ with minimum delay, $a \notin S$ and $b, c \in S$.*

Lemma 16 *Let Σ be an anti-chain on a set of resources $\mathcal{R}$. Furthermore, let $\mathcal{I} = \{X \subseteq S \mid S \in \Sigma\}$, and assume that $(\mathcal{R}, \mathcal{I})$ is not a matroid, i. e., that Σ is not the set of bases of some matroid. Then Σ possesses the $(1,2)$-exchange property.*

Before we prove Lemma 16, we present an additional property of matroids. For a proof of this property, we refer the reader to Theorem 39.6 in [20].

Proposition 17 *(Schrijver [20]) Let Σ be a set system on a finite set $\mathcal{R}$. Then Σ is the set of bases of a matroid if and only if for every pair of sets $S_1, S_2 \in \Sigma$ and every $r_2 \in S_2 \setminus S_1$, there exists an $r_1 \in S_1 \setminus S_2$ such that $S_2 \cup \{r_1\} \setminus \{r_2\} \in \Sigma$.*

Proof: (Lemma 16) Since $(\mathcal{R}, \mathcal{I})$ is not a matroid, there exist due to Proposition 17 two sets $X, Y \in \Sigma$ and a resource $x \in X \setminus Y$ such that for every $y \in Y \setminus X$, the set $X \setminus \{x\} \cup \{y\}$ is not contained in Σ.

Let X and Y be such sets and let $x \in X$ be such a resource. Consider all subsets Y' of the set $X \cup Y \setminus \{x\}$ with $Y' \in \Sigma$. Every such set Y' can be written as $Y' = X \setminus \{x = x_1, \ldots, x_l\} \cup \{y_1, \ldots, y_{l'}\}$ with $x_i \in X \setminus Y$ and $y_i \in Y \setminus X$ and $l + l' > 2$. This is true since $l \geq 1$ holds per definition and $l' \geq 1$ holds because Σ is an anti-chain. Furthermore l and l' cannot both equal 1 as otherwise we obtain a contradiction to the choice of X, Y, and x. Among all these sets Y', let $Y_{\min}$ denote one set for which l' is minimal. Observe that we can replace Y by $Y_{\min}$ without changing the aforementioned properties of X, Y, and x. Hence, in the following, we assume that $Y = Y_{\min}$, that is, we assume that $Y \setminus X = Y' \setminus X$ for all of the aforementioned sets Y'.

We claim that we can always identify resources $a, b, c \in X \cup Y$ such that either $a \in X \setminus Y$ and $b, c \in Y \setminus X$ or $a \in Y \setminus X$ and $b, c \in X \setminus Y$ with the property that for every $Z \subseteq X \cup Y$ with $Z \in \Sigma$, if $a \notin Z$, then $b, c \in Z$. In order to see this, we distinguish between the cases $l' = 1$ and $l' \geq 2$:

(1) Let $Y \setminus X = \{y_1\}$ and hence $X \setminus Y = \{x = x_1, \ldots, x_l\}$ with $l \geq 2$. Then we set $a = y_1$, $b = x_1$, and $c = x_2$. Consider a set $Z \subseteq X \cup Y$ with $Z \in \Sigma$ and $a \notin Z$. Then $Z = X$ since Σ is an anti-chain, and hence $b, c \in Z$.

(2) Let $Y \setminus X = \{y_1, \ldots, y_{l'}\}$ with $l' \geq 2$. Then we set $a = x$, $b = y_1$, and $c = y_2$. Consider a set $Z \subseteq X \cup Y$ with $Z \in \Sigma$ and $a \notin Z$. Since we assumed that $Y = Y_{\min}$, it must be $b, c \in Z$ as otherwise $Z \setminus X \neq Y \setminus X$.

Now we define delays for the resources in $\mathcal{R} \setminus \{a, b, c\}$ such that the properties in Definition 15 are satisfied. Let $k \in \mathbb{N}$ be chosen as in Definition 15, that is, $d(a), d(b), d(c) \in \{|\mathcal{R}|, \ldots, k\}$. We set $d(r) = k + |\mathcal{R}|$ for every resource $r \notin X \cup Y$ and $d(r) = 1$ for every resource $r \in (X \cup Y) \setminus \{a, b, c\}$. First of all, observe that in the first case the delay of Y equals $d(a) + |Y| - 1 < k + |\mathcal{R}|$ and that in the second case the delay of X equals $d(a) + |X| - 1 < k + |\mathcal{R}|$. Hence, a set $Z \in \Sigma$ that contains a resource $r \notin X \cup Y$ can never have minimum delay as its delay is at least $k + |\mathcal{R}|$. Thus, only sets $Z \in \Sigma$ with $Z \subseteq X \cup Y$ can have minimum delay. Since for such sets, $a \notin Z$ implies $b, c \in Z$, we know that every set with minimum delay must contain a or it must contain b and c.

Consider the case $d(a) + |\mathcal{R}| \leq d(b) + d(c)$ and assume for contradiction that there exists an optimal set Z^* with $a \notin Z^*$. Due to the choice of a, b, and c, the set Z^* must then contain b and c. Hence $d(Z^*) \geq d(b) + d(c)$. Furthermore, again due to the choice of a, b, and c, there exists a set $Z' \subseteq X \cup Y$ with $a \in Z'$ and $b, c \notin Z'$. The delay of Z' is $d(Z') = d(a) + |Z'| - 1 < d(a) + |\mathcal{R}| \leq d(b) + d(c) \leq d(Z^*)$, contradicting the assumption that Z^* has minimum delay. Hence every optimal set Z^* must contain a. If Z^* additionally contains b or c, then its delay is at least $d(a) + |\mathcal{R}| > d(Z')$. Hence, in the case $d(a) + |\mathcal{R}| \leq d(b) + d(c)$ every optimal set Z^* contains a but it does not contain b and c.

Consider the case $d(a) \geq d(b) + d(c) + |\mathcal{R}|$ and assume for contradiction that there exists an optimal set Z^* with $b \notin Z^*$ or $c \notin Z^*$. Then Z^* must contain a and hence its delay is at least $d(a)$. Due to the choice of a, b, and c, there exists a set $Z' \subseteq X \cup Y$ with $a \notin Z'$ and $b, c \in Z'$. The delay of Z' is $d(Z') = d(b) + d(c) + |Z'| - 2 < d(b) + d(c) + |\mathcal{R}| \leq d(a) \leq d(Z^*)$, contradicting the assumption that Z^* has minimum delay. Hence every optimal set Z^* must contain b and c. If Z^* additionally contains a, then its delay is at least $d(b) + d(c) + |\mathcal{R}| > d(Z')$. Hence, in the case $d(a) \geq d(b) + d(c) + |\mathcal{R}|$ every optimal set Z^* contains b and c but it does not contain a. □

4.2 The Matroid Property is Maximal for Guaranteeing the Existence of Equilibria

We are now ready to prove that Theorems 3 and 13 cannot be extended further without placing additional assumptions on the delay functions and/or the relation of the strategy spaces. We first consider weighted congestion games.

Theorem 18 *For every non-matroid anti-chain Σ on a set of resources $\mathcal{R}$ there exists a weighted congestion game Γ with two players whose strategy spaces are isomorphic to Σ that does not possess a Nash equilibrium. The delay functions in Γ are positive and increasing.*

Proof: Given a non-matroid anti-chain we describe how to construct a weighted congestion game with the properties stated in the theorem. We first describe how the strategy spaces are defined and then how the delay functions are chosen.

Let Σ_1 and Σ_2 be two set systems on sets of resources $\mathcal{R}_1$ and $\mathcal{R}_2$, respectively. In the following we assume that both sets are isomorphic to Σ and that Σ_i is the strategy space of player i, for $i = 1, 2$. Due to the $(1, 2)$-exchange property we can, for every player i, identify three distinct resources $a_i, b_i, c_i \in \mathcal{R}_i$ with the properties as in Definition 15. Since we have not made any assumption on the global structure of the game, we can arbitrarily decide which resources from $\mathcal{R}_1$ and $\mathcal{R}_2$ coincide. The resources $\mathcal{R}_i \setminus \{a_i, b_i, c_i\}$ are exclusively used by player i. Hence, we can assume that their delays are chosen such that the $(1, 2)$-exchange property is satisfied. Thus, to simplify matters we can assume that

$$\Sigma_1 = \{\underbrace{\{a_1\}}_{S_1^1}, \underbrace{\{b_1, c_1\}}_{S_1^2}\} \text{ and } \Sigma_2 = \{\underbrace{\{a_2\}}_{S_2^1}, \underbrace{\{b_2, c_2\}}_{S_2^2}\} .$$

In the following, we assume that $a_1 = b_2$, $b_1 = a_2$ and $c_1 = c_2$. Thus we can rewrite the strategy spaces as follows: $\Sigma_1 = \{\{x\}, \{y, z\}\}$ and $\Sigma_2 = \{\{y\}, \{x, z\}\}$.

We set $\omega_1 = 2$ and $\omega_2 = 1$ and define the following increasing delay functions for the resources x, y and z, where $m = |\mathcal{R}|$:

	$n_r = 1$	$n_r = 2$	$n_r = 3$
$d_x(n_x)$	m	$20 \cdot m$	$21 \cdot m$
$d_y(n_y)$	$5 \cdot m$	$12 \cdot m$	$15 \cdot m$
$d_z(n_z)$	$3 \cdot m$	$4 \cdot m$	$10 \cdot m$

One can easily verify that $|\delta_i(S \oplus S_i^1) - \delta_i(S \oplus S_i^2)| \geq m$, for $i = 1, 2$, regardless of the choice of the other player. Hence, for every player, one of the inequalities in Definition 15 is always satisfied. This game does not possess a Nash equilibrium since player 1 prefers to play strategy S_1^2 if player 2 plays strategy S_2^1, and S_1^1 if player 2 plays strategy S_2^2. Additionally, player 2 prefers to play strategy S_2^2 if player 1 plays strategy S_1^2, and S_2^1 if player 1 plays strategy S_1^1. $\square$

Theorem 19 *For every non-matroid anti-chain Σ on a set of resources $\mathcal{R}$ there exists a player-specific congestion game Γ with two players whose strategy spaces are isomorphic to Σ that does not possess a Nash equilibrium. The delay functions in Γ are positive and increasing.*

Proof: The proof is similar to the proof of Theorem 18. In particular, the construction of the strategy spaces of the players is identical. The player-specific delay functions are obtained from the delay functions in the proof of Theorem 18 as follows: For the first player $d_r^1(n_r) = d_r(n_r + 1)$, for every resource $r \in \{x, y, z\}$ and every congestion $n_r \in \{1, 2\}$. For the second player $d_r^2(1) = d_r(1)$ and $d_r^2(2) = d_r(3)$, for every resource $r \in \{x, y, z\}$. $\square$

Summarizing, every non-matroid anti-chain can be used to construct a player-specific or weighted congestion game with positive delay functions that does not posses a Nash equilibrium. Observe that this result also holds if the system is not an anti-chain but the *pruned set system*, i. e., the set system obtained after removing all supersets, is not the set of bases of a matroid. This is because supersets cannot occur in a Nash equilibrium in the case of positive delay functions. Correspondingly, our results presented in Theorems 3 and 13 show that a player-specific or weighted congestion game in which all *pruned* strategy spaces correspond to bases of matroids possesses a Nash equilibrium with respect to the pruned and, hence, also with respect to the original strategy spaces because supersets are weakly dominated by subsets in the case of non-negative delay functions. Thus, the matroid property (applied to the pruned strategy spaces) is necessary and sufficient to show the existence of Nash equilibria.

Corollary 20 *The matroid property is the maximal property on the pruned strategy spaces of the individual players that guarantees the existence of Nash equilibria in weighted and player-specific congestion games with non-negative, non-decreasing delay functions.*

4.3 A Comment on Network Congestion Games

Our negative results in Theorems 18 and 19 assume that it is possible to interweave the strategy spaces of the players in a specific manner. A legitimate question is whether our construction can nevertheless be embedded into natural classes of congestion games in which the resources have a common combinatorial interpretation.

Here, we demonstrate that are construction can, for instance, easily be embedded into network congestion games. However, note that we are not the first to present player-specific or weighted network congestion games which do not possess Nash equilibria [8,17].

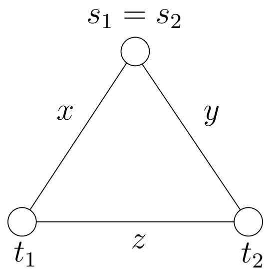

Fig. 1. An example of a network congestion game with the strategy spaces as defined in the proofs of Theorems 18 and 19.

Consider the network depicted in Figure 1. The first player likes to route her traffic from s_1 to t_1, the second player from s_2 to t_2. Observe that the sets of paths of player 1 and 2 coincide with the strategy spaces as defined above. We conclude the following corollary.

Corollary 21 *There exist instances of player-specific and instances of weighted network congestion games with non-decreasing delay functions which do not possess Nash equilibria.*

Observe, that the players are not symmetric, i. e., they like to connect the source to different sinks. However, it is not difficult to make the game symmetric by introducing a common sink t which is connected to t_1 and t_2 and by appropriately defining the delay functions of the edges $\{t_1, t\}$ and $\{t_2, t\}$.

5 Convergence Time in Weighted Matroid Congestion Games

In Section 3 we have shown that players playing lazy best responses eventually reach a Nash equilibrium in every weighted matroid congestion game. In weighted singleton congestion games every best response is a lazy best response, hence, in these games every sequence of best responses leads to a Nash equilibrium. It is an interesting question how many best responses are actually needed to find a Nash equilibrium. This question is addressed by Even-Dar et al. [5] who present a family of weighted singleton congestion games with symmetric players and identical resources with best response sequences of exponential length. However, they use exponentially large weights in their construction. In this section, we present an infinite family of weighted singleton congestion games possessing superpolynomially long best response sequences although every player has either weight one or n and all delays are polynomially bounded in the number of players and resources. This immediately implies that players do not necessarily reach a Nash equilibrium in pseudopolynomial time in a weighted singleton congestion game.

Theorem 22 *There exists a constant $c > 0$ such that for every $n \in \mathbb{N}$, there exists a weighted singleton congestion game Γ with at most cn^2 players and at most cn resources that possesses a best response sequence of length 2^n. The players in Γ have either weight 1 or weight n, and the maximum delay is upper bounded by cn^3.*

From Theorem 22 we can conclude the following corollary.

Corollary 23 *Weighted matroid congestion games do not converge to a Nash equilibrium in pseudopolynomial time.*

Proof: [Theorem 22] A well known technique for constructing instances of local search problems with exponentially long best response sequences is to construct instances that resemble the behavior of a binary counter (see, e. g., [1,3,13,18]). We apply this technique to weighted singleton congestion games.

Let $n \in \mathbb{N}$ be chosen arbitrarily. We construct a weighted singleton congestion game with $O(n^2)$ players and $O(n)$ resources that resembles the behavior of a binary counter counting from 0 to $2^n - 1$. The instance consists of n gadgets $G_0, \ldots, G_{n-1}$ where gadget G_i represents the i-th bit of the counter; G_0 represents the least significant bit, G_{n-1} the most significant bit. For every gadget G_i, we define three main configurations, namely a 0-state, a 1-state and a reset state, with the following properties.

(1) If gadget G_i is in its 0-state and no gadget G_j with $j > i$ is in its reset state, then there exists a best response sequence of gadget G_i such that G_i first changes to its reset state and then to its 1-state.
(2) If gadget G_i is in its 1-state and at least one gadget G_j with $j > i$ is in its reset state, then there exists a best response sequence of gadget G_i such that G_i changes to its 0-state.

One can easily verify that these two properties ensure that there exists a best response sequence of all gadgets that resembles a binary counter counting from 0 to $2^n - 1$: Initially all gadgets are in their 0-state. First gadget G_0 changes to its 1-state, then gadget G_1. However, when gadget G_1 changes to its 1-state it passes its reset state, and therefore resets gadget G_0. Afterwards gadget G_0 may change back to its 1-state. We proceed with gadget G_2 that resets the gadgets G_0 and G_1 by changing to its 1-state. We may continue with gadget G_0 and so on.

Now we describe the gadgets $G_0, \ldots, G_{n-1}$ in detail. Gadget G_i consists of $i+2$ players and 3 resources r_1^i, r_2^i and r_3^i. There are two main players, the *bit player* and the *reset player*, and i additional players, which we call *connection players*. The bit player and the reset player both have weight n, and each connection player has weight 1. Later, we will define delay functions and strategy spaces such that the best responses of the connection players are uniquely determined by the choice of the reset player. The purpose of the connection players is to propagate the decision of the reset player to the gadgets $G_0, \ldots, G_{i-1}$. The delay functions of the resources are defined as follows.

$$d_{r_1^i}(n_{r_1^i}) = \begin{cases} 3(n-i+1)+1 & \text{if } n_{r_1^i} \leq 2n-i-2 \\ 3n^2(i+1)+2 & \text{otherwise} \end{cases}$$

$$d_{r_2^i}(n_{r_2^i}) = \begin{cases} 3(n-i+1)+2 & \text{if } n_{r_2^i} \leq n \\ 3n^2(i+1)+1 & \text{otherwise} \end{cases}$$

$$d_{r_3^i}(n_{r_3^i}) = \begin{cases} 3(n-i+1)+3 & \text{if } n_{r_3^i} \leq i \\ 3n^2(i+1) & \text{otherwise} \end{cases}$$

We denote by Σ^i_{Bit} and Σ^i_{Reset} the strategy spaces of the bit and reset player, respectively, and by $\Sigma^i_{\text{Con}_j}$ the strategy space of the j-th connection player, with $0 \leq j \leq i-1$. Let the strategy spaces be defined as

$$\Sigma^i_{\text{Bit}} = \{\{r_1^i\}, \{r_2^i\}\} \qquad \Sigma^i_{\text{Reset}} = \{\{r_3^i\}, \{r_2^i\}\} \qquad \Sigma^i_{\text{Con}_j} = \{\{r_1^j\}, \{r_3^i\}\} \ .$$

For every player, we call the first strategy according to the above given ordering, her 0-strategy and the second one her 1-strategy. Figure 2 illustrates our construction.

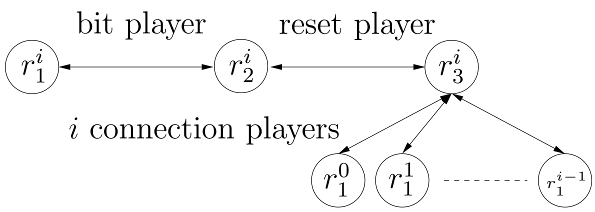

Fig. 2. Illustration of gadget G_i. Nodes represent resources, edges represent players.

In the following, we describe the state of gadget G_i by a pair of bits (x, y), meaning that the bit player plays her x-strategy and that the reset player plays her y-strategy. When describing the state of a gadget by such a pair, we assume that the connection players have played their best responses according to strategy y. We denote by $(0, 0)$ the 0-state of gadget G_i, by $(1, 0)$ the 1-state, and by $(0, 1)$ the reset state. We can then formulate the aforementioned properties of gadget G_i in terms of sequences of states (x, y).

(1) If gadget G_i is in state (0,0) and every gadget G_j with $j > i$ is in state (0,0) or (1,0), then there exists a best response sequence of gadget G_i such that G_i first changes to its reset state (0,1) and then to the state (1,0).

(2) If gadget G_i is in state (1,0) and at least one gadget G_j with $j > i$ is in state (0,1), then there exists a best response sequence of gadget G_i such that G_i changes to state (0,0).

It remains to show that the delay functions are chosen in the right way, that is, all strategy changes are best responses. We first show that the connection players of gadget G_i are solely controlled by the reset player of that gadget. Therefore, consider the following two cases.

(a) If the reset player plays her 0-strategy $\{r_3^i\}$, then the best response for every connection player is her 0-strategy. This is true since in this case the delay on resource r_3^i equals $3n^2(i+1)$ and the maximum delay on any resource r_1^j is at most $3n^2(j+1)+2$ which is less than $3n^2(i+1)$ because $j < i$.

(b) If the reset player plays her 1-strategy $\{r_2^i\}$, then the best response for every connection player is her 1-strategy. This is true since in this case the delay on r_3^i equals $3(n-i+1)+3$, and the minimum delay on any resource r_1^j is at least $3(n-j+1)+1$ which is larger than $3(n-i+1)+3$ because $j < i$.

In the following, we assume that immediately after each strategy change of the reset player, the connection players of the corresponding gadget change their strategies appropriately. Hence, when we say that the reset player of gadget G_i plays her x-strategy, $x \in \{0,1\}$, we implicitly assume that all connection players of that gadget play their x-strategies, too. Now we study the aforementioned best response sequences of the bit and reset players of a gadget G_i in detail.

(1) Gadget G_i is in state $(0,0)$ and all reset players of the gadgets G_j with $j > i$ play their 0-strategy. In this case, the reset player can decrease her delay from $3n^2(i+1)$ to $3(n-i+1)+2$ by changing to her 1-strategy. After that, gadget G_i is in state $(0,1)$, and the bit player can decrease her delay from $3n^2(i+1)+2$ to $3n^2(i+1)+1$. After that gadget G_i is in state $(1,1)$, and the reset player can decrease her delay from $3n^2(i+1)+1$ to $3n^2(i+1)$ by changing to her 0-strategy. After that the gadget is in state $(1,0)$ and as long as no reset player of a gadget G_j with $j > i$ plays her 1-strategy it stays in this state.

(2) Gadget G_i is in state $(1,0)$ and at least one reset player of a gadget G_j with $j > i$ plays her 1-strategy. In this case, the cumulative weight of all players allocating resource r_1^i is at most $n-i-2$. Hence, the bit player can decrease her delay from $(3n-i+1)+2$ to $(3n-i+1)+1$ by changing to her 0-strategy. After that the gadget is in state $(0,0)$.

Altogether this shows that the aforementioned sequence of strategy changes is a best response sequence and results in counting from 0 to 2^n-1. □

Let us briefly mention that our construction can even be implemented with players who have only weights 1 or 2. In order to achieve this, one has to introduce additional players that propagate the decision of the reset players to the connections players. Based on the observation that a player with weight 2 can displace two players of weight 1 from a resource, these players can be arranged in a binary tree with i leaves that propagate the decision to the connection players. As this construction is rather technical and does not give new insights, we do not present the details.

6 Open Problems

In contrast to player-specific congestion games, the proof of Theorem 13 does not yield an efficient algorithm for computing Nash equilibria in weighted matroid congestion games. To the best of our knowledge, efficient algorithms are only known in the case of weighted singleton congestion games with identical resources, i. e., all resources have identical, non-decreasing delay functions. If the players are symmetric, Fotakis et al. [7] show that it suffices to assign the players in non-increasing order of their weights to resources with minimum delay given the choices of the previous players. In the case, of asymmetric players Gairing et al. [9] show how to compute a Nash equilibrium in polynomial time.

We leave it as an open problem whether Nash equilibria in weighted matroid congestion games can be computed efficiently. Since the lexicographic order defined in the proof of Theorem 13 is a potential function with respect to the lazy best response dynamics, the problem of computing an equilibrium belongs to PLS. This has already been observed by Fabrikant et al. [6] for a generalization of weighted singleton congestion games. This implies that the problem of finding an equilibrium cannot be NP-hard, unless NP=co-NP [15]. But it is still open whether it is PLS-complete to find an equilibrium in such games.

References

[1] Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. On the impact of combinatorial structure on congestion games. In *Proceedings of the 47th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 613–622, 2006.

[2] Heiner Ackermann and Alexander Skopalik. On the complexity of pure nash equilibria in player-specific network congestion games. In *Proceedings of the 3rd International Workshop on Internet and Network Economics (WINE)*, pages 419–430, 2007.

[3] Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. The price of stability for network design with fair cost allocation. In *Proceedings of the 45th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.

[4] Juliane Dunkel and Andreas S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. In *Proc. of the 2nd Int. Workshop on Internet and Network Economics (WINE)*, pages 62–73, 2006.

[5] Eyal Even-Dar, Alexander Kesselman, and Yishay Mansour. Convergence time to Nash equilibria. In *Proceedings of the 30th Int. Coll. on Automata, Languages and Programming (ICALP)*, pages 502–513, 2003.

[6] Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. The complexity of pure Nash equilibria. In *Proceedings of the 36th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 604–612, 2004.

[7] Dimitris Fotakis, Spyros C. Kontogiannis, Elias Koutsoupias, Marios Mavronicolas, and Paul G. Spirakis. The structure and complexity of Nash equilibria for a selfish routing game. In *Proceedings of the 29th Int. Coll. on Automata, Languages and Programming (ICALP)*, pages 123–134, 2002.

[8] Dimitris Fotakis, Spyros C. Kontogiannis, and Paul G. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2-3):226–239, 2005.

[9] Martin Gairing, Thomas Lücking, Marios Mavronicolas, and Burkhard Monien. Computing Nash equilibria for scheduling on restricted parallel links. In *Proceedings of the 36th Ann. ACM Symp. on Theory of Computing (STOC)*, pages 613–622, 2004.

[10] Martin Gairing, Burkhard Monien, and Karsten Tiemann. Routing (un-)splittable flow in games with player-specific linear latency functions. In *Proceedings of the 33rd Int. Coll. on Automata, Languages and Programming (ICALP)*, pages 501–512, 2006.

[11] Chryssis Georgiou, Theophanis Pavlides, and Anna Philippou. Uncertainty in selfish routing. In *Proc. of the 20th IEEE Int. Parallel and Distributed Processing Symp. (IPDPS)*, 2006.

[12] Michel X. Goemans, Erran L. Li, Vahab S. Mirrokni, and Marina 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 55–66, 2004.

[13] Armin Haken. Connectionist networks that need exponential time to stabilize. Unpublished manuscript. Department of Computer Science, University of Toronto, 1989.

[14] Samuel Ieong, Robert McGrew, Eugene Nudelman, Yoav Shoham, and Qixiang 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] David S. Johnson, Christos H. Papadimitriou, and Mihalis Yannakakis. How easy is local search? *Journal of Computer and System Sciences*, 37(1):79–100, 1988.

[16] Igal Milchtaich. Congestion games with player-specific payoff functions. *Games and Economic Behavior*, 13(1):111–124, 1996.

[17] Igal Milchtaich. The equilibrium existence problem in finite network congestion games. In *Proc. of the 2nd Int. Workshop on Internet and Network Economics (WINE)*, pages 87–98, 2006.

[18] Pekka Orponen. Computing with truly asynchronous threshold logic networks. *Theoretical Computer Science*, 174(1-2):123–136, 1997.

[19] Robert W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2:65–67, 1973.

[20] Alexander Schrijver. *Combinatorial Optimization: Polyhedra and Efficiency.* Volume B, Matroids, Trees, Stable Sets, Chapters 39-69. Springer, 2003.

[21] Renato Werneck, João Setubal, and Arlindo da Conceicão. Finding minimum congestion spanning trees. *Journal of Experimental Algorithmics*, 5, 2000.