Congestion Games: Optimization in Competition

Berthold Vöcking, RWTH Aachen

ABSTRACT. In a congestion game, several players simultaneously aim at allocating sets of resources, e.g., each player aims at allocating a shortest path between a source/destination pair in a given network or, to give another example, each player aims at allocating a minimum weight spanning tree in a given graph. The cost (length, delay, weight) of a resource (edge) is a function of the congestion, i.e., the number of players allocating the resource. In this paper, we survey recent results about the complexity of computing Nash equilibria for congestion games and the convergence time towards Nash equilibria.

1 Introduction

Congestion games lie in the intersection between optimization and game theory. In optimization, one seeks for an optimal solution among a set of feasible solutions usually specified in a concise way, e.g., the set of feasible solutions is the set of paths between a specified source/destination pair in a given network and the objective is to find a path of minimum length. In game theory, several players simultaneously seek at maximizing their individual payoff. Each player can choose among different strategies and its payoff depends on the choices of all players. A *Nash equilibrium* is a selection of strategies for all players such that none of the players can unilaterally improve its payoff. In a Nash equilibrium, each player has thus chosen an optimal strategy in the form of a best response to the choices of the other players. Suppose several players simultaneously want to allocate a shortest path in a given network and the length of an edge depends on the number of players using the edge. This kind of game is a so-called *network congestion game*, a particular variant of congestion games.

More generally, a *congestion game* Γ is a tuple $(\mathcal{N}, \mathcal{R}, (\Sigma_i)_{i \in \mathcal{N}}, (d_r)_{r \in \mathcal{R}})$ where $\mathcal{N} = \{1, \ldots, n\}$ denotes the set of players, $\mathcal{R} = \{1, \ldots, m\}$ the set of resources, $\Sigma_i \subseteq 2^{\mathcal{R}}$ the strategy space of player i, and $d_r : \mathbb{N} \rightarrow \mathbb{Z}$ a cost function associated with resource r. $S = (S_1, \ldots, S_n)$ is a *state of the game* in which player i chooses strategy $S_i \in \Sigma_i$. For a state S, we define the *congestion* $n_r(S)$ on resource r by $n_r(S) = |\{i \mid r \in S_i\}|$. We assume that players act selfishly and aim at choosing strategies $S_i \in \Sigma_i$ minimizing their individual cost, where the cost $\delta_i(S)$ of player i is given by $\delta_i(S) = \sum_{r \in S_i} d_r(n_r(S))$.

Given any state $S = (S_1, \ldots, S_n)$, an *improvement step* of player i is a change of its strategy from S_i to S'_i such that the cost of player i decreases, i.e., $\delta_i(S \oplus S'_i) < \delta_i(S)$, where $S \oplus S'_i = (S_1, \ldots, S_{i-1}, S'_i, , S_{i+1}, \ldots, S_n)$. A classical result of Rosenthal [6] shows that sequences of improvement steps by possibly different players do not run into cycles but reach a Nash equilibrium after a finite number of steps.

PROPOSITION 1 (Finite Improvement Property). *For every congestion game, every sequence of improvement steps is finite.*

Proof. The proposition is shown by a potential function argument. *Rosenthal's potential function* $\phi : \Sigma_1 \times \cdots \times \Sigma_n \to \mathbb{Z}$ is defined by

$$\phi(S) = \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r(S)} d_r(i) \ .$$

Let $n_r^{(i)}(S)$ denote the number of those players using resource r that have an index from $\{1, \ldots, i\}$, and let $\delta'_i(S) = \sum_{r \in S_i} d_r(n_r^{(i)}(S))$, that is, $\delta'_i(S)$ is a virtual cost that player i would have if the players with index $i+1, \ldots, n$ would not exist. We exchange the two sums in the definition of the potential, and obtain

$$\phi(S) \ = \ \sum_{i=1}^{n} \sum_{r \in S_i} d_r(n_r^{(i)}(S)) \ = \ \sum_{i=1}^{n} \delta'_i(S) \ .$$

This formulation of the potential has the following intuitive interpretation. Players together with their strategy are "inserted" one after the other into the game, and the potential accounts for the cost of each player's strategy at the insertion time of the player.

Now suppose player n can decrease its cost by switching to another strategy. For this player, the virtual cost $\delta'_n(S)$ is the same as the true cost $\delta_n(S)$. Hence, if this player decreases its cost by an unilateral move then the potential is decreased by the same amount. Of course, this is not a special property of player n as the potential can be computed by inserting the players in any order so that any player can be assumed to be the last player.

Consequently, as the potential of any state in a fixed congestion game is upper- and lower-bounded by some finite quantity and every improvement step decreases the potential at least by one, the length of any sequence of improvement steps must be finite as well. ■

Nash equilibria are the only fixed points of the dynamics defined by improvement steps. Hence, the finite improvement property immediately implies the existence of pure Nash equilibria in congestion games, where *pure* means that these equilibria do not involve randomized (mixed) strategies. In the following, we drop the attribute *pure* and the term *Nash equilibrium* always refers to a *pure Nash equilibrium*.

Games that admit an *exact potential function*, i.e., a potential function with the property that an improvement of an individual player decreases the potential by exactly the same amount as the player's cost, are called *potential games*. Rosenthal's potential function is exact and, hence, congestion games are potential games. Monderer and Shapley [2] have shown that every potential game can be represented in form of a congestion game. Thus, congestion games are essentially the only class of games for which one can show the existence of pure equilibria with an exact potential function.

2 The Connection to Local Search

Rosenthal's potential function allows to interprete congestion games as local search problems. In general, a local search problem Π is given by its set of instances $\mathcal{I}_\Pi$. For every instance $I \in \mathcal{I}_\Pi$, we are given a finite set of feasible solutions $\mathcal{F}(I)$, an objective function $c : \mathcal{F}(I) \to \mathbb{Z}$, and for every feasible solution $S \in \mathcal{F}(I)$, a neighborhood $\mathcal{N}(S, I) \subseteq \mathcal{F}(I)$. The so-called *transition graph* contains a node $v(S)$ for every feasible solution $S \in \mathcal{F}(I)$ and a directed edge from a node $v(S_1)$ to a node $v(S_2)$ if S_2 is in the neighborhood of S_1 and if the objective value $c(S_2)$ is strictly better than the objective value $c(S_1)$. The sinks of this graph are the *local optima*. In case of a congestion game, the set of feasible solutions corresponds to the set of states, the objective function is defined by Rosenthal's potential function and the neighborhood of a state S consists of those states that deviate from S only in one player's strategy. Sequences of improvement steps correspond to paths in the transition graph and the sinks of this graph are the Nash equilibria of the game.

Given a congestion game, how difficult is it to compute a Nash equilibrium? – One way to find a Nash equilibrium is to use a heuristic following the *local search paradigm*, i.e., start at any state and perform improvement steps until a Nash equilibrium is reached. Observe that the number of states is exponential in the number of players, even if there are only two alternative strategies for each player. Thus, there might be instances for which a heuristic following this paradigm takes an exponential number of steps. Notice that, even if there are states for which all improvement paths have exponential length, there might be other, indirect algorithmic approaches for computing Nash equilibria efficiently. Fabrikant et al. [3] had the striking idea to relate the complexity of finding Nash equilibria for congestion games to the complexity of computing local optima for other local search problems. In particular, they considered problems from the class PLS (Polynomially Local Search). A local search problem Π belongs to PLS if the following polynomial time algorithms exist:

1. an algorithm A which computes for every instance I of Π an initial feasible solution $S \in \mathcal{F}(I)$,
2. an algorithm B which computes for every instance I of Π and every feasible solution $S \in \mathcal{F}(I)$ the objective value $c(S)$,
3. an algorithm C which determines for every instance I of Π and every feasible solution $S \in \mathcal{F}(I)$ whether S is locally optimal or not and finds a better solution in the neighborhood of S in the latter case.

Let us give some illustrative examples. In the MAX-SAT (maximum satisfiability) problem, we are given a Boolean formula in conjunctive normal form with a positive integer weight for each clause. A solution is an assignment of the value 0 or 1 to all variables. Its weight, to be maximized, is the sum of the weights of all satisfied clauses. The restriction to instances with at most k literals in each clause is called MAX-kSAT. A natural neighborhood for this problem is the *Flip*-neighborhood, where two assignments are neighbors if one can be obtained from the other by flipping the value of a single variable. In a variation of this problem, called POS-NAE-MAX-kSAT, one assumes that clauses contain only positive literals and the objective is to maximize the weighted sum of those clauses whose literals do not all have the same value. Another illustrative example is the MaxCut problem. An instance of this problem consists of a simple undirected graph $G = (V, E)$ with non-negative edge weights. A feasible solution is a partition of V into two sets A and B. Two solutions are neighboring with respect to the *Flip*-neighborhood if one can be obtained from the other by moving a single vertex from one set to the other. The objective is to maximize the weight of the cut (A, B), i.e., the weight of the edges between the two sets A and B. Remarkably, MaxCut is equivalent to POS-NAE-MAX-2SAT.

Johnson et al. [5] introduced the notion of a PLS-*reduction*. A problem Π_1 from PLS is PLS-reducible to a problem Π_2 from PLS if there are polynomial-time computable functions f and g such that:

1. f maps instances I of Π_1 to instances $f(I)$ of Π_2,
2. g maps pairs (S_2, I) with S_2 denoting a solution of $f(I)$ to solutions S_1 of I,
3. for all instances I of Π_1, if S_2 is a local optimum of instance $f(I)$ then $g(S_2, I)$ is a local optimum of I.

A local search problem Π from PLS is PLS-*complete* if every problem in PLS is PLS-reducible to Π. Schäffer and Yannakakis [7] have shown that MAX-kSAT as well as POS-NAE-MAX-kSAT are PLS-complete, for every $k \geq 2$. As MAXCUT is equivalent to POS-NAE-MAX-2SAT it is PLS-complete, too. For further examples and a comprehensive overview on the complexity of local search we refer the reader to [8].

Fabrikant et al. [3] show that computing a Nash equilibrium in congestion games is PLS-complete, even if the class of congestion games is restricted to symmetric games. They use a reduction from POS-NAE-MAX-3SAT. We present a similar reduction from MAXCUT.

THEOREM 2. *Computing a Nash equilibrium in (symmetric) congestion games is* PLS*-complete.*

Proof. Let us first ignore the symmetry requirement. Given an instance of MAXCUT, we construct a congestion game as follows. For each edge e of weight w, we have two resources $r_e^{(A)}$ and $r_e^{(B)}$, with cost 0 if used by only one player and cost w if used by more players. The players correspond to the nodes. Player v has two strategies: one strategy contains all $r_e^{(A)}$'s for edges e incident to v, and another that contains all $r_e^{(B)}$'s for the same edges. The first strategy corresponds to assigning v to the set A and the latter strategy corresponds to assigning v to B. This one-to-one correspondence between the assignments of the nodes in the MAXCUT-instance and the strategies of the players in the congestion game has the property that the local optima of the MAXCUT-instance coincide with the Nash equilibria of the congestion game. Hence, our construction is a PLS-reduction from MAXCUT to (possibly asymmetric) congestion games.

It remains to present a reduction from the general to the symmetric case. Suppose we are given a general congestion game with strategy spaces $\Sigma_1, \ldots, \Sigma_n \subseteq \mathcal{R}$. We extent $\mathcal{R}$ by additional resources $r_1, \ldots, r_n$ with cost 0 if used by one player and cost M, otherwise, where M is a large number. For $i \in \mathcal{N}$, let $\Sigma'_i = \{S \cup \{r_i\} | S \in \Sigma_i\}$. The symmetric game has the common strategy space $\Sigma = \Sigma'_1 \cup \ldots \cup \Sigma'_n$. If M is chosen sufficiently large then any equilibrium of this game has one player using a strategy from Σ'_i. This property yields an obvious correspondence between the Nash equilibria of the symmetric and the asymmetric game, and, hence, gives a PLS-reduction. ■

Hence, designing an efficient algorithm computing Nash equilibria for (symmetric) congestion games requires to discover a general approach for local optimization. It might sound unlikely but one cannot rule out a priori that there exists an algorithm that efficiently computes local optima for every PLS-problem. Let us point out, however, that such an algorithm cannot just use the *local search paradigm*, i.e., start at any state and perform improvement steps until a local optimum is found. This is because there are instances of PLS-problems, e.g., instances of MaxCut, whose transition graph contains nodes (solutions) that have an exponential distance to any sink (local optimum). In fact, the PLS-reduction presented above preserves this property as the transition graph of the congestion game is isomorphic to the transition graph of the MaxCut-instance. Consequently, there are congestion games with states such that all improvement paths that lead from these states to a Nash equilibrium have a length that is exponential in the number of players.

3 Matroid Congestion Games

Now we study the impact of combinatorial structure on congestion games. In particular, we investigate for which classes of congestion games improvement steps converge quickly towards a Nash equilibrium. Suppose players iteratively use *best responses* until they reach a Nash equilibrium, that is, we consider sequences of improvement steps in which the player changing its strategy always switches to an alternative strategy of minimal cost. As best responses are just a special variant of improvement steps, our analysis above shows that there exist congestion games with best response sequences of exponential length. The combinatorial structure underlying these bad examples is based on the MaxCut problem. In the following, we study which kind of combinatorial structure yields fast convergence.

We take a local perspective, that is, we study which property of the strategy spaces of individual players can guarantee that best responses reach a Nash equilibrium in a number of steps that is polynomial in the number of players and resources. *Singleton games* are characterized by the property that the strategy space of each player consist only of strategies that contain a single resource. For this class of games, Ieong et al. [4] show that any sequence of improvement steps reaches a Nash equilibrium after at most $n^2 m$ best responses. Ackermann et al. [1] generalize their analysis towards *matroid congestion games*, i.e., to congestion games in which the strategy space of each player corresponds to the basis of a matroid. An illustrative example are *spanning tree congestion games*. In such a game, each of the players wants to build a minimum cost spanning tree of a given graph and the cost of an edge depends in some arbitrary way on the number of players that use the edge in their spanning tree. Matroid games are not restricted to spanning trees and need not be symmetric. 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. The rank $rk(\Gamma)$ of a *matroid congestion game* Γ is defined to be the maximum rank over all players' matroids.

THEOREM 3. *Let Γ be a matroid congestion game. Then players reach a Nash equilibrium after at most $n^2\, m\, rk(\Gamma) \leq n^2\, m^2$ best response improvement steps.*

Proof. Let L be a sorted list of all cost values that can occur on single resources. Suppose L lists the values $d_r(i)$, for $r \in \mathcal{R}$ and $1 \leq i \leq n$, in non-decreasing order. For each resource r, define an alternative cost function $\tilde{d}_r : \mathbb{N} \to \mathbb{N}$ where, for each possible congestion i, $\tilde{d}_r(i)$ equals the rank of the cost value $d_r(i)$ in L. Equal cost values receive the same rank.

LEMMA 4. *Consider any state S. Let $S_i^* \in \Sigma_i$ be a best response to S of any player $i \in \mathcal{N}$. Then S_i^* decreases the cost of player i with respect to the alternative cost functions $\tilde{d}$.*

Proof. We make use of a special property of matroids: The best response S_i^* can be decomposed in a sequence of *(1,1)-exchange steps* in each of which the player exchanges only one resource in its strategy and does not increase the cost with respect to the original cost functions d. Suppose resource r is exchanged in one of these exchange steps against a resource r' and the original cost of r' is smaller than the original cost of r. Then also the alternative cost of r' is smaller than the alternative cost of r since the original cost value of r' occurs in L before the original cost value of r. Furthermore, if both resources have the same original cost then they also have the same alternative cost. Hence, over all exchange steps the alternative cost decreases because there must be at least one exchange step in which the original cost decreases. ∎

Now we consider Rosenthal's potential function with respect to the alternative cost. Lemma 4 yields that the potential decreases whenever a player makes a best response. Since there are at most $n \cdot m$ different cost values, $\tilde{d}_r(i) \leq n \cdot m$ for all resources $r \in \mathcal{R}$ and for all possible congestion values i. As a consequence,

$$\tilde{\phi}(S) \;=\; \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r(S)} \tilde{d}_r(i) \leq \sum_{r \in \mathcal{R}} \sum_{i=1}^{n_r(S)} n\, m \;\leq\; n^2\, m\, rk(\Gamma) \;,$$

where the latter inequality holds as each of the n players occupies at most $rk(\Gamma)$ resources. Now the theorem follows since each best response decreases the potential by at least one and the potential cannot drop below zero. ∎

It is remarkable that the positive result above holds regardless of the global structure of the game and for any kind of cost functions. Ackermann et al. [1] also show that this analysis is essentially tight in the sense that the matroid property is the maximal property that yields such a result. Their analysis identifies a sub-structure, called *(1,2)-exchange*, that can be found in any inclusion-free non-matroid set system.

LEMMA 5. *Consider any instance $\mathcal{I}$ of a minimization problem over an inclusion-free non-matroid set system Σ over a set of resources $\mathcal{R}$. There exist three resources $a, b, c \in \mathcal{R}$ with the property that, if the weights of the other resources are set appropriately, an optimal solution of $\mathcal{I}$ contains the resource a but not the resources b and c if $w_a < w_b + w_c$, and it contains b and c but not a if $w_a > w_b + w_c$.*

This substructure allows to construct exponentially long improvement sequences of best responses, which implies the following result.

THEOREM 6. *Let Σ be an inclusion-free non-matroid set system over a set of resources $\mathcal{R}$. For every $n \in \mathbb{N}$, there exists a congestion game Γ with $4n$ players each of which having a strategy space isomorphic to Σ, and $O(n \cdot |\mathcal{R}|)$ resources with non-negative and non-decreasing cost functions such that there exists a best response sequence of length 2^n.*

Observe that the assumption that Σ is inclusion-free is natural when all cost values are non-negative as, in this case, supersets are dominated by subsets. Hence, we can conclude that the matroid property is the maximal property on the individual players' strategy spaces that guarantees polynomial convergence of best responses in congestion games.

4 Network Congestion Games

In a *network congestion game*, we are given a directed graph and, for each player, a source and a destination node. Every player seeks for a minimum delay path connecting its source with its destination. The delay of an edge depends on the number of players using that edge. Typically, it is assumed that the delay functions are non-decreasing. In the symmetric variant of the game, all players have the same source and the same destination. In this special case, one can compute a Nash equilibrium with the help of a min-cost flow algorithm [3].

THEOREM 7. *There is a polynomial time algorithm for computing Nash equilibria in symmetric network congestion games with non-decreasing delay functions.*

Proof. The congestion game is reduced to a min-cost flow problem as follows. Each edge e is replaced by n parallel edges $e_1, \ldots, e_n$ between the same nodes. Edge e_i is assigned cost $d_e(i)$, for $1 \leq i \leq n$. All edges have capacity 1. Observe, if a min-cost flow solution uses some of the edges $e_1, \ldots, e_n$, then it sends an integral amount of flow along these edges. If it sends k units of flow along these edges, then it uses the k cheapest edges. W.l.o.g., these are the edges $e_1, \ldots, e_k$ as the delay functions are non-decreasing. Thus, the cost for sending the flow along these edges is $d_e(1) + \ldots + d_e(k)$, which corresponds to the potential that Rosenthal's potential function assigns to edge e if k players use this edge. Consequently, we can translate the optimal solution of the min-cost flow problem into a state of the congestion game whose potential corresponds to the cost of the flow. Hence, the min-cost flow solution corresponds to a Nash equilibrium that globally minimizes Rosenthal's potential function. ■

Let us remark that the result above does not imply that one can reach a Nash equilibrium in symmetric network congestion games with best responses efficiently. In fact, there exist instances of symmetric network congestion games with non-decreasing delay functions that have states with an exponential distance to any Nash equilibrium in the transition graph [1]. This is a remarkable result as it yields an example of a local search problem of rich combinatorial structure for which one cannot find a local optimum efficiently with a direct approach following the local search paradigm but there is an indirect method based on min-cost flow that finds a local optimum efficiently.

One might wonder whether there is such an indirect method also for the general (asymmetric) case. Unfortunately, this is probably not the case as Fabrikant et al. [3] have shown that the problem of finding Nash equilibria in general network congestion games is PLS-complete. Their analysis is quite complicated. Ackermann et al. [1] present a shorter analysis that additionally shows that the problem is PLS-complete even if the delay functions are linear. The analysis leading to this result makes a detour first showing PLS-completeness for another variant of congestion games.

Threshold games are a special class of congestion games. Assume that the set of resources $\mathcal{R}$ is divided into two disjoint subsets $\mathcal{R}_{\text{in}}$ and $\mathcal{R}_{\text{out}}$. The set $\mathcal{R}_{\text{out}}$ contains a resource r_i for every $i \in \mathcal{N}$. This resource has a fixed cost T_i called the *threshold* of player i. Each player i has only two strategies, namely a strategy $S_i^{\text{out}} = \{r_i\}$ with $r_i \in \mathcal{R}_{\text{out}}$, and a strategy $S_i^{\text{in}} \subseteq \mathcal{R}_{\text{in}}$. The preferences of player i can be described in a simple and intuitive way: Player i prefers to choose strategy S_i^{in} against strategy S_i^{out} if the cost of S_i^{in} is smaller than the threshold T_i. *Quadratic threshold games* are a quite restrictive subclass of threshold games. In this variant, the set $\mathcal{R}_{\text{out}}$ contains exactly one resource r_{ij} for every unordered pair of players $\{i, j\} \subseteq \mathcal{N}$. For every player $i \in \mathcal{N}$ of a quadratic threshold game, $S_i^{\text{in}} = \{r_{ij} \mid j \in \mathcal{N} \setminus \{i\}\}$.

THEOREM 8. *Computing a Nash equilibrium of a quadratic threshold game is* PLS*-complete.*

Proof. We prove the theorem by a PLS-reduction from MaxCut. Consider an instance of MaxCut that, w.l.o.g., consists of a complete weighted graph $G = (V, E)$ with non-negative edge weights w_{ij}. The MaxCut problem can be described in terms of a game, the so-called *party affiliation game* in which players correspond to nodes that can choose whether they belong to partition A or B. Edges reflect some symmetric kind of anti-sympathy, that is, a node seeks to choose one of the two sets such that the weighted number of edges leading to the other set is maximized. The Nash equilibria of the party affiliation game coincide with the local optima of the MaxCut instance with respect to the *flip*-neighborhood.

The preferences of the players in the party affiliation game can be described in the following way that points out what could be a suitable threshold for a player. For player i, let W_i denote the sum of the weights of all of its incident edges and $W_i^{(B)}$ the sum of the weights of the edges that connect i with nodes in partition B. Player i prefers strategy A against strategy B if $W_i^{(B)} > \frac{1}{2}W_i$, it prefers strategy B against strategy A if $W_i^{(B)} < \frac{1}{2}W_i$, and it is indifferent if $W_i^{(B)} = \frac{1}{2}W_i$.

Now we show how to represent the party affiliation game in the form of a quadratic threshold game. Both games involve the same number of players. We identify the players in the two games. With each edge $e = \{i, j\}$, we associate the resource $r_{ij} \in \mathcal{R}_{\text{in}}$. The delay of this resource is 0 if the resource is used by only one player and its delay is w_{ij} if it is used by two players. We identify strategy B of player i in the party affiliation game with strategy S_i^{in} in the congestion game. Player i's strategy A in the party affiliation game corresponds to strategy S_i^{out} in the threshold game, and the delay of this strategy is $T_i = \frac{1}{2}W_i$. Observe that the players' preferences in both games are identical so that we have described a PLS-reduction. ■

Now we are ready to show the PLS-completeness of asymmetric network congestion games by a reduction from quadratic threshold games.

THEOREM 9. *Computing a Nash equilibrium for a general network congestion game with non-decreasing, linear delay functions is* PLS*-complete.*

Proof. Let Γ be a quadratic threshold game. We map Γ to an asymmetric network routing game. The network consists of the lower-left triangle of an $n \times n$ grid (including the vertices on the diagonal) in which the column edges are directed downwards and the row edges are directed from left to right. For every player i in Γ, we introduce a player i in the network congestion game whose source node s_i is the i-th node in the first column and whose target node t_i is the i-th node in the last row. For every player $i \in \mathcal{N}$, we add an edge from s_i to t_i, called *threshold edge*. Due to the directions of the grid edges, the threshold edge of player i can only be used by player i.

Our first goal is to define delay functions in such a way that there are only two relevant strategies for player i: the shortcut edge (s_i, t_i) or the *row-column path* from s_i to t_i, i.e., the path from s_i along the edges of row i until column i and then along the edges of column i to t_i. All other paths shall have such high delays that they are dominated by these two paths, regardless of the other players choices. We achieve this goal by assigning the constant delay function 0 to all column edges and the constant delay function $D \cdot i$ to all row edges in row i, where D denotes a large integer. Furthermore, for the time being, we assume that the shortcut edge (s_i, t_i) has the constant delay $D \cdot i \cdot (i-1)$. This way, each player i has only two undominated strategies: its shortcut edge or its row-column path. The delays of these two alternative routes are so far identical.

Now we define additional delay functions for the nodes, that is, we view also the nodes as resources. (It is easy to see how the nodes can be replaced by gadgets such that all resources are edges.) For $1 \leq i < j \leq n$, the node in column i and row j is identified with the resource $r_{ij} \in \mathcal{R}_{\text{in}}$ from the quadratic threshold game. In particular, we assume that the node has the same cost function as the corresponding resource from the threshold game. This way, the row-column path of player i corresponds to the strategy S_i^{in} of the threshold game. Furthermore, we increase the delay on the shortcut edge of player i from $D \cdot i \cdot (i-1)$ to $D \cdot i \cdot (i-1) + T_i$, where T_i is the cost of resource $r_i \in \mathcal{R}_{\text{in}}$ from the threshold game. This way, the shortcut edge of player i corresponds to the strategy S_i^{out} of the threshold game.

If D is chosen sufficiently large then all strategies except for the row-column paths and the shortcut edges are dominated and, hence, can be ignored. The remaining strategy spaces of the players and the corresponding delay functions are isomorphic to the strategies and cost functions of the threshold game. In particular, also the Nash equilibria of the two games coincide. Thus, our construction is a PLS-reduction. Finally, observe that all delay functions can be described in terms of linear functions as each of the resources is used by at most two of the players. ■

5 Conclusions

Finding a Nash equilibrium in general congestion games is PLS-complete and, hence, as hard as solving any other local optimization problem in this class. Assuming, however, that the strategy spaces of all players correspond to the bases of matroids, Nash equilibria can be found in polynomial time just by repeated best responses. The technical reason behind this result is the $(1,1)$-exchange property of matroids that allows to decompose best responses into a sequence of exchanges of single resources that do not increase the cost. In fact, the matroid property is the maximal condition on the individual players' strategy spaces that ensures polynomial convergence time. In instances of optimization problems over (inclusion-free) non-matroid set systems one can identify a substructure called $(1,2)$-exchange. With this substructure as a basic building block, one can construct congestion games with exponentially long best response paths.

Is there a similar structural characterization of those classes of congestion games that are PLS-complete? – In fact, the strategy space of each player in a quadratic threshold game corresponds to a $(1,k)$-exchange, where k grows linearly with the number of players. Despite their simple structure, threshold games are a natural and interesting class of games. Our main interest, however, stems from the fact that threshold games are a good starting point for PLS-reductions because of their simple structure. We have demonstrated the applicability of this approach by showing reductions from quadratic threshold games to network congestion games. Let us remark that this approach has more applications. Further reductions from quadratic threshold games to so-called *market sharing games* and *overlay network design games* can be found in [1].

BIBLIOGRAPHY

[1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. In *Proc. 47th Ann. IEEE Symp. on Foundations of Computer Science (FOCS)*, 2006.

[2] D. D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.

[3] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. 36th Ann. ACM. Symp. on Theory of Comput. (STOC)*, pages 604–612, 2004.

[4] S. Ieong, R. McGrew, E. Nudelman, Y. Shoham, and Q. Sun. Fast and compact: A simple class of congestion games. In *20th National Conference on Artificial Intelligence (AAAI)*, 2005.

[5] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37(1):79–100, 1988.

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

[7] A. A. Schäffer and M. Yannakakis. Simple local search problems that are hard to solve. *SIAM J. Comput.*, 20(1):56–87, 1991.

[8] M. Yannakakis. Computational complexity. In E. Aarts and J. Lenstra, editors, *Local Search in Combinatorial Optimization*, pages 19–55. Wiley, Chichester, 1997.

Berthold Vöcking, Email: voecking@cs.rwth-aachen.de,

Dept. of Computer Science, RWTH Aachen, 52056 Aachen, Germany