

Resource Competition on Integral Polymatroids

Tobias Harks^{1,*}, Max Klimm^{2,**}, and Britta Peis³

¹ Department of Quantitative Economics, Maastricht University, The Netherlands

`t.harks@maastrichtuniversity.nl`

² Department of Mathematics, Technische Universität Berlin, Germany

`klimm@math.tu-berlin.de`

³ School of Business and Economics, RWTH Aachen University, Germany

`britta.peis@oms.rwth-aachen.de`

Abstract. We study competitive resource allocation problems in which players distribute their demands integrally over a set of resources subject to player-specific submodular capacity constraints. Each player has to pay for each unit of demand a cost that is a non-decreasing and convex function of the total allocation of that resource. This general model of resource allocation generalizes both singleton congestion games with integer-splittable demands and matroid congestion games with player-specific costs. As our main result, we show that in such general resource allocation problems a pure Nash equilibrium is guaranteed to exist by giving a pseudo-polynomial algorithm computing a pure Nash equilibrium.

1 Introduction

In an influential paper, Rosenthal [23] introduced *congestion games*, a class of strategic games, where a finite set of players competes over a finite set of resources. Each player is associated with a set of allowable subsets of resources and a pure strategy of a player consists of an allowable subset. In the context of *network games*, the resources may correspond to edges of a graph and the allowable subsets correspond to the paths connecting a source and a sink. The utility a player receives by using a resource depends only on the number of players choosing the same resource and each player wants to maximize (minimize) the utility (cost) of the sum of the resources contained in the selected subset. Rosenthal proved the existence of a pure Nash equilibrium. Up to day congestion games have been used as reference models for describing decentralized systems involving the selfish allocation of congestible resources (e.g., selfish route choices in traffic networks [4,25,29] and flow control in telecommunication networks [17,18,27]) and for decades they have been a focal point of research in (algorithmic) game theory, operations research and theoretical computer science.

* The research of the first author was supported by the Marie-Curie grant “Protocol Design” (nr. 327546) funded within FP7-PEOPLE-2012-IEF.

** The research of the second author was carried out in the framework of MATHEON supported by the Einstein Foundation Berlin.

In the past, the existence of pure Nash equilibria has been analyzed in many variants of congestion games such as singleton congestion games with player-specific cost functions (cf. [11,16,21,22]), congestion games with weighted players (cf. [1,2,5,10,13]), nonatomic and atomic splittable congestion games (cf. [4,14,18,29]) and congestion games with player- and resource-specific and variable demands (cf. [12]).

Most of these previous works can be classified according to the following two categories: (i) the demand of each player is unsplittable and must be completely assigned to exactly one subset of the allowable subsets; (ii) the demand of a player may be fractionally split over the set of allowable subsets. While these assumptions and the resulting models are obviously important (and also realistic for some applications), they do not allow for the requirement that only *integral fractions* of the demand may be assigned to allowable subsets of resources. This requirement is clearly important in many applications, where the demand represents a collection of indivisible items or tasks that need to be placed on subsets of resources. Examples include the scheduling of integer-splittable tasks in the context of load balancing on server farms (cf. [19]) or in logistics where a player controls a fleet of vehicles and each must be assigned to a single route.

Although Rosenthal proposed congestion games with integer-splittable demands as an important and meaningful model already back in 1973 – in his first work on congestion games [24] even published prior to his more famous work [23] – not much is known regarding existence and computability of pure Nash equilibria. Rosenthal gave an example showing that in general, pure Nash equilibria need not exist. Dunkel and Schulz [7] strengthened this result showing that the existence of a pure Nash equilibrium in integer-splittable congestion games is NP-complete to decide. Meyers [20] proved that in games with linear cost functions, a pure Nash equilibrium is always guaranteed to exist. For singleton strategy spaces and non-negative and convex cost functions, Tran-Thanh et al. [28] showed the existence of pure Nash equilibria. They also showed that pure Nash equilibria need not exist (even for the restricted strategy spaces) if cost functions are semi-convex.

Our Results. We introduce congestion games on *integral polymatroids*, where each player may fractionally assign the demand in integral units among the allowable subsets of resources subject to player-specific submodular capacity constraints. This way, the resulting strategy space for each player forms an integral polymatroid base polyhedron (truncated at the player-specific demand). As our main result, we devise an algorithm that computes a pure Nash equilibrium for congestion games on integral polymatroids with player-specific non-negative, non-decreasing and *strongly semi-convex* cost functions. The class of strongly semi-convex functions strictly includes convex functions but is included in the class of semi-convex functions (see Section 2.2 for a formal definition). The runtime of our algorithm is bounded by $n^{\delta+1} \cdot m^{\delta} \delta^{\delta+1}$, where n is the number of players, m the number of resources, and δ is an upper bound on the maximum demand. Thus, for constant δ , the algorithm is polynomial.

Our existence result generalizes that of Tran-Thanh et al. [28] for singleton congestion games with integer-splittable demands and convex cost functions and that of Ackermann et al. [1] for matroid congestion games with unit demands and player-specific non-decreasing costs. For the important class of network design games, where players need to allocate bandwidth in integral units across multiple spanning trees of a player-specific communication graph (cf. [3,6,9]), our result shows for the first time the existence of pure Nash equilibria provided that the cost on each edge is a strongly semi-convex function of total bandwidth allocated.

Techniques. Our algorithm for computing pure Nash equilibria maintains data structures for *preliminary demands*, *strategy spaces*, and *strategies* of the players that all are set to zero initially. Then, it iteratively increases the demand of a player by one unit and recomputes a preliminary pure Nash equilibrium (with respect to the current demands) by following a sequence of best response moves of players. The key insight to prove the correctness of the algorithm is based on two invariants that are fulfilled during the course of the algorithm. As a first invariant, we show that, whenever the demand of a player is increased by one unit, there is a best response that assigns the new unit to some resource *without* changing the allocation of previously assigned demand units. As second invariant, we obtain that, after assigning this new unit to some resource, only those players that use this particular resource with increased load may have an incentive to deviate. Moreover, there is a best response that has the property that at most a single unit of demand is shifted to some other resource. Given the above two invariants, we prove that during the sequence of best response moves a carefully defined vector of *marginal costs* lexicographically decreases, thus, ensuring that the sequence is finite.

The first invariant follows by reducing an integral polymatroid to an ordinary matroid (cf. Helgason [15]) and the fact that for a matroid, a minimum independent set I_d with rank d can be extended to a minimum independent set I_{d+1} with rank $d+1$ by adding a single element to I_d . The second invariant, however, is significantly more complex since a change of the load of one resource results (when using the matroid construction in the spirit of Helgason) in changed element weights for several elements simultaneously. To prove the second invariant we use several exchange and uncrossing arguments that make use of the submodularity of the rank functions and the fact that a non-optimal basis of a matroid can be improved locally. This is the technically most involved part of our paper.

We note that the above invariants have also been used by Tran-Thanh et al. [28] for showing the existence of pure Nash equilibria in singleton integer-splittable congestion games. For singleton games, however, these invariants follow almost directly. The algorithmic idea to incrementally increase the total demand by one unit is similar to the (inductive) existence proof of Milchtaich [21] for singleton congestion games with player-specific cost functions (see also Ackerman et al. [1] for a similar proof for matroid congestion games). The convergence proof for our algorithm and the above mentioned invariants, however, are considerably more involved for general integral polymatroids.

Besides providing new existence results for an important and large class of games, the main contribution of this paper is to propose a unified approach to prove the existence of pure Nash equilibria that connects the seemingly unrelated existence results of Milchtaich [21] and Ackerman et al. [1] on the one hand, and Tran-Thanh et al. [28] on the other hand.

2 Preliminaries

In this section, we introduce polymatroids, strong semi-convexity, and congestion games on integral polymatroids.

2.1 Polymatroids

Let \mathbb{N} denote the set of non-negative integers and let R be a finite and non-empty set of resources. We write \mathbb{N}^R shorthand for $\mathbb{N}^{|R|}$. Throughout this paper, vectors $\mathbf{x} = (x_r)_{r \in R}$ will be denoted with bold face. An integral (set) function $f : 2^R \rightarrow \mathbb{N}$ is *submodular* if $f(U) + f(V) \geq f(U \cup V) + f(U \cap V)$ for all $U, V \in 2^R$. Function f is *monotone* if $U \subseteq V$ implies $f(U) \leq f(V)$, and *normalized* if $f(\emptyset) = 0$. An integral submodular, monotone and normalized function $f : 2^R \rightarrow \mathbb{N}$ is called an *integral polymatroid rank function*. The associated *integral polyhedron* is defined as

$$\mathbb{P}_f = \left\{ \mathbf{x} \in \mathbb{N}^R : \sum_{r \in U} x_r \leq f(U) \text{ for each } U \subseteq R \right\}.$$

Given the integral polyhedron \mathbb{P}_f and some integer $d \in \mathbb{N}$ with $d \leq f(R)$, the *d-truncated integral polymatroid* $\mathbb{P}_f(d)$ is defined as

$$\mathbb{P}_f(d) = \left\{ \mathbf{x} \in \mathbb{N}^R : \sum_{r \in U} x_r \leq f(U) \text{ for each } U \subseteq R, \sum_{r \in R} x_r \leq d \right\}.$$

The corresponding *integral polymatroid base polyhedron* is

$$\mathcal{B}_f(d) = \left\{ \mathbf{x} \in \mathbb{N}^R : \sum_{r \in U} x_r \leq f(U) \text{ for each } U \subseteq R, \sum_{r \in R} x_r = d \right\}.$$

2.2 Strongly Semi-convex Functions

Recall that a function $c : \mathbb{N} \rightarrow \mathbb{N}$ is *convex* if $c(x+1) - c(x) \leq c(x+2) - c(x+1)$ for all $x \in \mathbb{N}$. A function c is called *semi-convex* if the function $x \cdot c(x)$ is convex, i.e.,

$$(x+1)c(x+1) - xc(x) \leq (x+2)c(x+2) - (x+1)c(x+1)$$

for all $x \in \mathbb{N}$.

For the main existence result of this paper, we require a property of each cost function that we call *strong semi-convexity*, which is weaker than convexity

but stronger than semi-convexity. Roughly speaking, it states that the marginal difference $c(a+x)x - c(a+x-1)(x-1)$ does not decrease as a or x increase. We also introduce a slightly weaker notion, termed u -truncated strong semi-convexity, where strong semi-convexity is only required for values of x not larger than u .

Definition 1 (Strong Semi-Convexity). *A function $c : \mathbb{N} \rightarrow \mathbb{N}$ is strongly semi-convex if*

$$c(a+x)x - c(a+x-1)(x-1) \leq c(b+y)y - c(b+y-1)(y-1) \quad (1)$$

for all $x, y \in \mathbb{N}$ with $1 \leq x \leq y$ and all $a, b \in \mathbb{N}$ with $a \leq b$. For an integer $u \geq 1$, c is u -truncated strongly semi-convex, if the above inequality is only required to be satisfied for all $x, y \in \mathbb{N}$ with $1 \leq x \leq y \leq u$ and all $a, b \in \mathbb{N}$ with $a \leq b$.

We note that a similar definition is also given in Tran-Thanh et al. [28]. It is not hard to show that strong semi-convexity is indeed strictly weaker than convexity. In the interest of space, we defer a formal proof to the full version of the paper.

Proposition 1. *Every convex and non-decreasing function $c : \mathbb{N} \rightarrow \mathbb{N}$ is also strongly semi-convex, but not vice versa.*

Finally, we remark that every non-decreasing function is 1-truncated strongly semi-convex.

Remark 1. A function is 1-truncated strongly semi-convex if and only if it is non-decreasing.

2.3 Congestion Games on Integral Polymatroids

In a congestion game on integral polymatroids, there is a non-empty and finite set N of players and a non-empty and finite set R of resources. Each resource is endowed with a player-specific cost function $c_{i,r} : \mathbb{N} \rightarrow \mathbb{N}$, $r \in R, i \in N$ and each player i is associated with a demand $d_i \in \mathbb{N}$, $d_i \geq 1$ and an integral polymatroid rank function $f^{(i)} : 2^R \rightarrow \mathbb{N}$ that together define a d_i -truncated integral polymatroid $\mathbb{P}_{f^{(i)}}(d_i)$ with base polyhedron $\mathcal{B}_{f^{(i)}}(d_i)$ on the set of resources. A strategy of player $i \in N$ is to choose a vector $\mathbf{x}_i = (x_{i,r})_{r \in R} \in \mathcal{B}_{f^{(i)}}(d_i)$, i.e., player i chooses an integral resource consumption $x_{i,r} \in \mathbb{N}$ for each resource r such that the demand d_i is exactly distributed among the resources and for each $U \subseteq R$ not more than $f^{(i)}(U)$ units of demand are distributed to the resources contained in U . Using the notation $\mathbf{x}_i = (x_{i,r})_{r \in R}$, the set X_i of feasible strategies of player i is defined as

$$X_i = \mathcal{B}_{f^{(i)}}(d_i) = \left\{ \mathbf{x}_i \in \mathbb{N}^R : \sum_{r \in U} x_{i,r} \leq f^{(i)}(U) \text{ for each } U \subseteq R, \sum_{r \in R} x_{i,r} = d_i \right\}.$$

The Cartesian product $X = \prod_{i \in N} X_i$ of the players' sets of feasible strategies is the joint strategy space. An element $\mathbf{x} = (\mathbf{x}_i)_{i \in N} \in X$ is a strategy profile.

For a resource r , and a strategy profile $\mathbf{x} \in X$, we write $x_r = \sum_{i \in N} x_{i,r}$. The private cost of player i under strategy profile $\mathbf{x} \in X$ is defined as $\pi_i(\mathbf{x}) = \sum_{r \in R} c_{i,r}(x_r) x_{i,r}$. In the remainder of the paper, we will compactly represent the strategic game by the tuple $G = (N, X, (d_i)_{i \in N}, (c_{i,r})_{i \in N, r \in R})$.

We use standard game theory notation. For a player $i \in N$ and a strategy profile $\mathbf{x} \in X$, we write \mathbf{x} as $(\mathbf{x}_i, \mathbf{x}_{-i})$. A *best response* of player i to \mathbf{x}_{-i} is a strategy $\mathbf{x}_i \in X_i$ with $\pi_i(\mathbf{x}_i, \mathbf{x}_{-i}) \leq \pi_i(\mathbf{y}_i, \mathbf{x}_{-i})$ for all $\mathbf{y}_i \in X_i$. A pure Nash equilibrium is a strategy profile $\mathbf{x} \in X$ such that for each player i the strategy \mathbf{x}_i is a best response to \mathbf{x}_{-i} .

Throughout this paper, we assume that the player-specific cost function $c_{i,r}$ of each player i on each resource r is $u_{i,r}$ -truncated strongly semi-convex, where $u_{i,r} = f^{(i)}(\{r\})$. Note that $u_{i,r}$ is a natural upper bound on the units of demand player i can allocate to resource r in any strategy $\mathbf{x}_i \in X_i$.

Assumption. For all $i \in N, r \in R$, the cost function $c_{i,r} : \mathbb{N} \rightarrow \mathbb{N}$ is non-negative, non-decreasing and $u_{i,r}$ -truncated strongly semi-convex, where $u_{i,r} = f^{(i)}(\{r\})$.

2.4 Examples

We proceed to illustrate that we obtain the well known classes of integer-splittable singleton congestion games and matroid congestion games as special cases of congestion games on integer polymatroids.

Example 1 (Singleton integer-splittable congestion games). For the special case that, for each player i , there is a player-specific subset $R_i \subseteq R$ of resources such that $f^{(i)}(\{r\}) = d_i$, if $r \in R_i$, and $f^{(i)}(\{r\}) = 0$, otherwise, we obtain integer-splittable singleton congestion games previously studied by Tran-Thanh et al. [28]. While they consider the special case of convex and *player-independent* cost functions, our general existence result implies existence of a pure Nash equilibrium even for *player-specific* and strongly semi-convex cost functions.

Example 2 (Matroid congestion games with player-specific costs). For the special case, that for each player i , $f^{(i)}$ is the rank function of a player-specific matroid defined on R , and $d_i = f^{(i)}(R)$, we obtain ordinary matroid congestion games with player-specific costs and unit demands studied by Ackermann et al. [1] as a special case.

Note that the rank function $\text{rk} : 2^R \rightarrow \mathbb{N}$ of a matroid is always subcardinal, i.e., $\text{rk}(U) \leq |U|$ for all $U \subseteq R$. Thus, we obtain in particular that $\text{rk}(\{r\}) \leq 1$ for all $r \in R$. This implies that our existence result continues to hold if we only require that the player-specific cost functions are 1-truncated strongly semi-convex, which is equivalent to requiring that cost functions are non-decreasing as in [1]. Like this, we obtain the existence result of [1] as a special case of our existence result for congestion games on integer polymatroids. As a strict generalization, our model includes the case in which players have a demand $d_i \in \mathbb{N}$ that can be distributed in integer units over bases (or even arbitrary independent sets) of a given player-specific matroid. A prominent application arises in network design (cf. [3,6,9]), where a player needs to allocate bandwidth

along several spanning trees and the cost function for installing enough capacity on an edge is a convex function of the total bandwidth allocated.

3 Equilibrium Existence

In this section, we give an algorithm that computes a pure Nash equilibrium for congestion games on integral polymatroids. Our algorithm relies on two key sensitivity properties of optimal solutions minimizing a linear function over an integral polymatroid base polyhedron (see Lemma 1 and Lemma 2 below). After reducing the usual reduction of integer polymatroids to ordinary matroids (cf. [15]), Lemma 1 follows more or less directly from the respective property for matroids. The proof of Lemma 2 is considerably more involved and relies heavily on uncrossing arguments. The two lemmata will be proven formally in Section 4.

3.1 Key Sensitivity Results

For two vectors $\mathbf{x}_i, \mathbf{y}_i \in \mathbb{N}^R$, we denote their Hamming distance by $H(\mathbf{x}_i, \mathbf{y}_i) = \sum_{r \in R} |x_{i,r} - y_{i,r}|$. Lemma 1 shows that, whenever a strategy \mathbf{x}_i minimizes the cost of player i over the base polyhedron $\mathcal{B}_{f^{(i)}}(d_i)$, then we only need to increase $x_{i,r}$ for some $r \in R$ by one unit to obtain a strategy \mathbf{y}_i minimizing the player's cost over the base polyhedron $\mathcal{B}_{f^{(i)}}(d_i + 1)$.

Lemma 1 (Demand Increase). *Let $\mathbf{x}_i \in \mathcal{B}_{f^{(i)}}(d_i)$ be a best response of player i to $\mathbf{x}_{-i} \in X_{-i}$. Then there exists a best response $\mathbf{y}_i \in \mathcal{B}_{f^{(i)}}(d_i + 1)$ to \mathbf{x}_{-i} such that $H(\mathbf{x}_i, \mathbf{y}_i) = 1$.*

The second result shows that, when some other player $j \neq i$ increases her demand for some resource r that is also used by player i with at least one unit, then player i can simply shift one unit from resource r to some resource $s \in R$ in order to retain minimal costs.

Lemma 2 (Load Increase). *Let $\mathbf{x}_i \in \mathcal{B}_{f^{(i)}}(d_i)$ be a best response of player i to $\mathbf{x}_{-i} \in X_{-i}$ and for each resource r let $a_r = \sum_{j \neq i} x_{j,r}$ be the induced allocation. If for a resource r , the value a_r is increased by 1, then there exists a best response $\mathbf{y}_i \in \mathcal{B}_{f^{(i)}}(d_i)$ towards the new profile with $H(\mathbf{x}_i, \mathbf{y}_i) \in \{0, 2\}$.*

3.2 The Algorithm

Both sensitivity results are used as the main building blocks for Algorithm 1 that computes a pure Nash equilibrium for congestion games on integral polymatroids. Algorithm 1 maintains *preliminary demands, strategy spaces, and strategies* of the players denoted by $\bar{d}_i \leq d_i$, $\bar{X}_i = X_i(\bar{d}_i)$, and $\mathbf{x}_i \in \bar{X}_i$, respectively. Initially, the preliminary demand \bar{d}_i of each player i is set to zero. Trivially, for this game the strategy profile where the strategy of each player equals the zero vector is a pure Nash equilibrium.

Then, in each round, for some player i , the demand is increased from \bar{d}_i to $\bar{d}_i + 1$, and a best response $\mathbf{y}_i \in X(\bar{d}_i + 1)$ with $H(\mathbf{x}_i, \mathbf{y}_i) = 1$ is computed, see Line 5 in Algorithm 1. By Lemma 1, such a best response always exists. In effect, the load on exactly one resource r increases and only those players j with $x_{j,r} > 0$ on this resource can potentially decrease their private cost by a unilateral deviation. By Lemma 2, it is without loss of generality to assume that a best response of such players consists of moving a single unit from this resource to another resource, see Line 8 of Algorithm 1. As a consequence, during the while-loop (Lines 7-10), only one additional unit (compared to the previous iteration) is moved preserving the invariant that only players using a resource to which this additional unit is assigned may have an incentive to profitably deviate. Thus, if the while-loop is left, the current strategy profile \mathbf{x} is a pure Nash equilibrium for the reduced game $\bar{G} = (N, \bar{X}, \bar{d}, (c_{i,r})_{i \in N, r \in R})$. Now we are ready to prove the main existence result.

ALGORITHM 1. Compute PNE

Input: $G = (N, X, (d_i)_{i \in N}, (c_{i,r})_{i \in N, r \in R})$
Output: pure Nash equilibrium \mathbf{x}

- 1 $\bar{d}_i \leftarrow 0, \bar{X}_i \leftarrow X_i(0)$ and $\mathbf{x}_i \leftarrow \mathbf{0}$ for all $i \in N$;
- 2 **for** $k = 1, \dots, \sum_{i \in N} d_i$ **do**
- 3 Choose $i \in N$ with $\bar{d}_i < d_i$;
- 4 $\bar{d}_i \leftarrow \bar{d}_i + 1; \bar{X}_i \leftarrow X_i(\bar{d}_i)$;
- 5 Choose a best response $\mathbf{y}_i \in \bar{X}_i$ with $H(\mathbf{y}_i, \mathbf{x}_i) = 1$;
- 6 $\mathbf{x}_i \leftarrow \mathbf{y}_i$;
- 7 **while** $\exists i \in N$ who can improve in $\bar{G} = (N, \bar{X}, \bar{d}, (c_{i,r})_{i \in N, r \in R})$ **do**
- 8 Compute a best response $\mathbf{y}_i \in \bar{X}_i$ with $H(\mathbf{y}_i, \mathbf{x}_i) = 2$;
- 9 $\mathbf{x}_i \leftarrow \mathbf{y}_i$;
- 10 **end**
- 11 **end**
- 12 Return \mathbf{x} ;

Theorem 1. *Congestion games on integral polymatroids with player-specific non-negative, non-decreasing, and strongly semi-convex cost functions possess a pure Nash equilibrium.*

Proof. We prove by induction on the total demand $d = \sum_{i \in N} d_i$ of the input game $G = (N, X, (d_i)_{i \in N}, (c_{i,r})_{i \in N, r \in R})$ that Algorithm 1 computes a pure Nash equilibrium of G .

For $d = 0$, this is trivial. Suppose that the algorithm works correctly for games with total demand $d - 1$ for some $d \geq 1$ and consider a game G with total demand d . Let us assume that in Line 3, the algorithm always chooses a player with minimum index. Consider the game $G' = (N, X, (d'_i)_{i \in N}, (c_{i,r})_{i \in N, r \in R})$ that differs from G only in the fact that the demand of the last player n is reduced by one, i.e. $d'_i = d_i$ for all $i < n$ and $d'_n = d_n - 1$. Then, when running the algorithm with G' as input, the $d - 1$ iterations (of the for-loop) are equal to the

first $d - 1$ iterations when running the algorithm with G as input. Thus, with G as input, we may assume that after the first $d - 1$ iterations, the preliminary strategy profile that we denote by \mathbf{x}' is a pure Nash equilibrium of G' .

We analyze the final iteration $k = d$ of the algorithm in which the demand of player n is increased by 1 (see Line 4). In Line 5, a best reply \mathbf{y}_n with $H(\mathbf{x}_n, \mathbf{y}_n) = 1$ is computed which exists by Lemma 1. Then, as long as there is a player i that can improve unilaterally, in Line 8, a best response \mathbf{y}_i with $H(\mathbf{y}_i, \mathbf{x}_i) = 2$ is computed which exists by Lemma 2.

It remains to show that the while-loop in Lines 7–10 terminates. To prove this, we give each unit of demand of each player $i \in N$ an identity denoted by $i_j, j = 1, \dots, d_i$. For a strategy profile \mathbf{x} , we define $r(i_j, \mathbf{x}) \in R$ to be the resource to which unit i_j is assigned in strategy profile \mathbf{x} . Let \mathbf{x}^l be the strategy profile after Line 8 of the algorithm has been executed the l -th time, where we use the convention that \mathbf{x}^0 denotes the preliminary strategy profile when entering the while-loop. As we chose in Line 5 a strategy of player n with Hamming distance one, there is a unique resource r_0 such that $x_{r_0}^0 = x'_{r_0} + 1$ and $x_r^0 = x'_r$ for all $r \in R \setminus \{r_0\}$. Furthermore, because we choose in Line 8 a best response with Hamming distance two, a simple inductive claim shows that after each iteration l of the while-loop, there is a unique resource $r_l \in R$ such that $x_{r_l}^l = x'_{r_l} + 1$ and $x_r^l = x'_r$ for all $r \in R \setminus \{r_l\}$.

For any \mathbf{x}^l during the course of the algorithm, we define the *marginal cost* of unit i_j under strategy profile \mathbf{x}^l as

$$\Delta_{i_j}(\mathbf{x}^l) = \begin{cases} c_{i,r}(x_r^l) x_{i,r}^l - c_{i,r}(x_r^l - 1)(x_{i,r}^l - 1), & \text{if } r = r(i_j, \mathbf{x}) = r_l \\ c_{i,r}(x_r^l + 1) x_{i,r}^l - c_{i,r}(x_r^l)(x_{i,r}^l - 1), & \text{if } r = r(i_j, \mathbf{x}) \neq r_l. \end{cases} \quad (2)$$

Intuitively, if $r(i_j, \mathbf{x}) = r_l$, the value $\Delta_{i_j}(\mathbf{x})$ measures the *cost saving* on resource $r(i_j, \mathbf{x})$ if i_j (or any other unit of player i on resource $r(i_j, \mathbf{x})$) is removed from $r(i_j, \mathbf{x})$. If $r(i_j, \mathbf{x}) \neq r_l$, the value $\Delta_{i_j}(\mathbf{x})$ measures the cost saving if i_j is removed from $r(i_j, \mathbf{x})$ after the total allocation has been increased by one unit by some other player. For a strategy profile \mathbf{x} we define $\Delta(\mathbf{x}) = (\Delta_{i_j}(\mathbf{x}))_{i=1, \dots, n, j=1, \dots, d_i}$ to be the vector of marginal costs and let $\bar{\Delta}(\mathbf{x})$ be the vector of marginal costs sorted in non-increasing order. We claim that $\bar{\Delta}(\mathbf{x})$ decreases lexicographically during the while-loop. To see this, consider an iteration l in which some unit i_j of player i is moved from resource r_{l-1} to resource r_l .

For proving $\bar{\Delta}(\mathbf{x}^l) <_{\text{lex}} \bar{\Delta}(\mathbf{x}^{l-1})$, we first observe that we only have to care for Δ -values that correspond to units i_j of the deviating player i , because for all players $h \neq i$ we obtain $\Delta_{h_j}(\mathbf{x}^{l-1}) = \Delta_{h_j}(\mathbf{x}^l)$ for all $j = 1, \dots, d_h$. This follows immediately if h_j is neither assigned to r_{l-1} nor to r_l . If h_j is assigned to r_{l-1} or r_l , then we switch the case in (2), and the claimed equality still holds. It remains to consider the Δ -values corresponding to the units of the deviating player i . Recall that the deviation of player i consists of moving unit i_j from resource r_{l-1} to resource r_l . We obtain

$$\begin{aligned} \Delta_{i_j}(\mathbf{x}^{l-1}) &= c_{i,r_{l-1}}(x_{r_{l-1}}^l) x_{i,r_{l-1}}^l - c_{i,r_{l-1}}(x_{r_{l-1}}^l - 1)(x_{i,r_{l-1}}^l - 1) \\ &> c_{i,r_l}(x_{r_l}^l + 1)(x_{i,r_l}^l + 1) - c_{i,r_l}(x_{r_l}^l) x_{i,r_l}^l = \Delta_{i_j}(\mathbf{x}^l), \end{aligned}$$

where the inequality follows since player i strictly improves. For every unit i_m of player i that is assigned to resource r_l as well, i.e. $r(i_m, \mathbf{x}^l) = r(i_j, \mathbf{x}^l) = r_l$, we have $\Delta_{i_j}(\mathbf{x}^l) = \Delta_{i_m}(\mathbf{x}^l)$ since the Δ -value is the same for all units of a single player assigned to the same resource. The Δ -values of such units i_m might have increased, but only to the Δ -value of unit i_j .

Next, consider the Δ -values of a unit i_m assigned to resource r_{l-1} , i.e., $r(i_m, \mathbf{x}^l) = r(i_j, \mathbf{x}^l) = r_{l-1}$. We obtain

$$\begin{aligned} \Delta_{i_m}(\mathbf{x}^l) &= c_{i,r_l}(x_{r_{l-1}}^l) (x_{i,r_{l-1}}^l - 1) - c_{i,r_{l-1}}(x_{r_{l-1}}^l - 1) (x_{i,r_{l-1}}^l - 2) \\ &\leq c_{i,r_{l-1}}(x_{r_{l-1}}^l) x_{i,r_{l-1}}^l - c_{i,r_{l-1}}(x_{r_{l-1}}^l - 1) (x_{i,r_{l-1}}^l - 1) = \Delta_{i_m}(\mathbf{x}^{l-1}), \end{aligned}$$

where for the inequality we used that $c_{i,r}(x_{r_{l-1}}) \geq c_{i,r}(x_{r_{l-1}} - 1)$ as $c_{i,r}$ is non-decreasing.

Altogether, the Δ -values of all units of all players $h \neq i$ have not changed, for player i , the Δ -values of remaining units assigned to resource r_{l-1} decreased, and the Δ -values assigned to resource r_l increased exactly to $\Delta_{i_j}(\mathbf{x}^l)$ which is strictly smaller than $\Delta_{i_j}(\mathbf{x}^{l-1})$. Thus, $\bar{\Delta}(\mathbf{x}^l) <_{\text{lex}} \bar{\Delta}(\mathbf{x}^{l-1})$ follows. \square

The following corollary states an upper bound on the number of iterations of the algorithm in terms of $\delta = \max_{i \in N} d_i$.

Corollary 1. *The number of iterations is at most $n^{\delta+1} m^\delta \delta^{\delta+1}$, which yields a polynomial algorithm computing a pure Nash equilibrium for constant δ .*

Proof. We analyze the worst-case runtime of Algorithm 1. To this end, let us fix an iteration of the for-loop. In the proof of Theorem 1, we showed that during this iteration, for each player, the sorted vector of marginal costs as defined in (2) decreases lexicographically during the while-loop. Moreover, the marginal cost of a particular unit of demand i_j of player i assigned to a resource r does not depend on the aggregated demand $\sum_{j \in N} x_{j,r}$ of all players for resource r , but only on the number of units of demand $x_{i,r}$ assigned to r by player i . We derive that for each player i and each resource r at most d_i different marginal cost values can occur. This observation bounds the number of different marginal cost vectors of player i by $(m \cdot d_i)^{d_i}$, where $m = |R|$. Since the marginal cost vectors lexicographically decrease, the total number of iterations of the while-loop for each iteration of the for-loop is bounded by $\sum_{i \in N} (m \cdot d_i)^{d_i}$. Setting $\delta = \max_{i \in N} d_i$, this expression is bounded by $(n \cdot m \cdot \delta)^\delta$, where $n = |N|$. Using that there are $\sum_{i \in N} d_i \leq n \cdot \delta$ iterations of the for-loop, one for each unit of demand in the game, we obtain the following corollary. \square

4 Sensitivity Analysis for Integral Polymatroids

It remains to show the key sensitivity results of Lemma 1 and Lemma 2. For ease of notation, let us drop the index i from the statements of the lemmata and let us consider a fixed integral polymatroid base polyhedron

$$\mathcal{B}_f(d) = \left\{ \mathbf{x} \in \mathbb{N}^R : \sum_{r \in U} x_r \leq f(U) \text{ for each } U \subseteq R, \sum_{r \in R} x_r = d \right\}.$$

w.r.t. some submodular, monotone, and normalized function $f : 2^R \rightarrow \mathbb{N}$, and some demand value $d \in \mathbb{N}$.

We identify the points in $\mathcal{B}_f(d)$ with a set family $\mathcal{F}(d)$ on a largely extended ground set E as follows: For each resource $r \in R$, let $u_r = f(\{r\})$ and let $K_r = \{r_1 \prec \dots \prec r_{u_r}\}$ be a totally ordered set (chain) with $|u_r|$ distinct elements r_1, \dots, r_{u_r} . Let further $E = \bigcup_{r \in R} K_r$ be the disjoint union of these chains. Then, $P = (E, \preceq)$ is a partially ordered set (poset) where two elements e, e' are comparable if and only if they are contained in the same chain K_r for some $r \in R$. Furthermore, let $\mathcal{D}(P)$ denote the set of *ideals* of P , i.e., $\mathcal{D}(P)$ consists of all subsets $I \subseteq E$ such that for each $e \in I$ all elements $g \prec e$ also belong to I . Note that there is a one-to-one correspondence between the sets in $\mathcal{D}(P)$ and the integral points in $\{\mathbf{x} \in \mathbb{N}^R : x_r \leq f(\{r\}) \ \forall r \in R\}$. As a consequence, the feasible points in the integral polymatroid \mathbb{P}_f can be identified with the set family

$$\mathcal{F} = \left\{ F \in \mathcal{D}(P) : \left| \bigcup_{r \in U} K_r \cap F \right| \leq f(U) \text{ for each } U \subseteq R \right\}. \quad (3)$$

Accordingly, the vectors contained in the polymatroid base polyhedron $\mathcal{B}_f(d)$ for $d \in \mathbb{N}$ can be identified with the set family

$$\mathcal{F}(d) = \left\{ F \in \mathcal{F} : |F| = d \right\}. \quad (4)$$

In fact, it is known (see, e.g., [26] and [15]) that any integral polymatroid \mathbb{P}_f can be reduced to an ordinary matroid $\mathcal{M} = (E, r)$ on ground set E with rank function $r : 2^E \rightarrow \mathbb{N}$ defined via

$$r(U) = \min_{T \subseteq R} \left(\left| U \setminus \bigcup_{r \in T} K_r \right| + f(T) \right)$$

for all $U \subseteq E$. It turns out that the independent sets in \mathcal{M} of cardinality d are exactly the ideals in $\mathcal{F}(d)$ as defined above. Applying this kind of transformation for each player i , we can identify the strategy set $\mathcal{B}_{f^{(i)}}(d_i)$ of each player i with the set family $\mathcal{F}_i(d_i)$, and this set family, in turn, with the matroid $\mathcal{M}_i = (E_i, r_i)$.

With this notation, let us now return to the problem of finding a best response \mathbf{x}_i of player i towards a strategy profile $\mathbf{a} = \mathbf{x}_{-i} \in \mathbb{N}^R$ of the remaining players. Note that, for $\mathbf{a} \in \mathbb{N}^R$, the player-specific strongly semi-convex cost functions $c_{i,r} : \mathbb{N} \rightarrow \mathbb{N}$ induce weight functions $w_i^{\mathbf{a}} : E \rightarrow \mathbb{N}$ on the ground set E constructed above via

$$w_i^{\mathbf{a}}(r_t) = t c_{i,r}(a_r + t) - (t - 1) c_{i,r}(a_r + t - 1) \quad \text{for all } r \in R, t \in \{1, \dots, d_i\}.$$

Hence, finding a best response $\mathbf{x}_i \in \mathcal{B}_{f^{(i)}}(d_i)$ reduces to the problem of minimizing a linear function over the independent sets of cardinality d_i of the matroid $\mathcal{M}_i = (E_i, r_i)$ associated with the submodular function $f^{(i)}$.

However, the ground set E_i can be of exponential size, so that it is not a priori clear whether the matroid greedy algorithm minimizes a linear weight function

over the base polyhedron $\mathcal{B}_{\mathcal{F}(d)}(d_i)$ in strongly polynomial time. Still, since we assume that the cost functions $c_{i,r} : \mathbb{N} \rightarrow \mathbb{N}$ are strongly semi-convex, it follows that the induced weight functions $w_i^{\mathfrak{a}} : E \rightarrow \mathbb{N}$ are *admissible* in the sense that $e \prec g$ implies $w_i(e) \leq w_i(g)$.

Given an ideal $F \in \mathcal{D}(P)$, we denote by F^+ the set of \preceq -maximal elements in F , and by $(E \setminus F)^-$ the set of \preceq -minimal elements in $E \setminus F$. For $\mathcal{F}(d)$ as defined in (4), and any admissible weight functions $w : E \rightarrow \mathbb{N}$, Faigle [8] showed that the following *ordered greedy algorithm* determines an ideal of minimal weight in $\mathcal{F}(d)$ (provided $\mathcal{F}(d) \neq \emptyset$):

ALGORITHM 2. Ordered Greedy Algorithm

```

1  $F \leftarrow \emptyset$ ;
2 for  $k = 1, \dots, d$  do
3   | Let  $e_k \leftarrow \arg \min \{w(e) : e \in (E \setminus F)^- \text{ and } F + e \in \mathcal{F}\}$ ;
4   |  $F \leftarrow F + e_k$ ;
5 end
6 Return  $\mathcal{F}$ ;

```

In fact, the greedy algorithm determines in each iteration $k \leq d$ an ideal of minimal weight in $\mathcal{F}(k)$. The following proposition arises as a consequence of the discussion above and implies Lemma 1.

Proposition 2. *Let $\mathcal{F} \subseteq \mathcal{D}(P)$ as defined in (3), $k \in \mathbb{N}$, and $w : E \rightarrow \mathbb{R}$ admissible. Suppose F is of minimal w -weight in $\mathcal{F}(k)$. Then there exists $e \in E \setminus F$ such that $F + e$ is of minimal w -weight in $\mathcal{F}(k + 1)$*

Due to the reduction of integral polymatroids to ordinary matroids, most of the structural properties of matroids carry over to integral polymatroids. For example, the following proposition follows as a consequence of the well-known fact, that for any basis B of an ordinary matroid \mathcal{M} which is not of minimal weight, there exists a local improvement step towards a basis $B - e + f$ of smaller weight.

Proposition 3. *Suppose $F \in \mathcal{F}(k)$ is not of minimal w -weight for some admissible function $w : E \rightarrow \mathbb{R}$. Then there exists some local improvement step $F \rightarrow F - e + g \in \mathcal{F}(k)$ such that $w(e) > w(g)$.*

The possibility to improve a non-optimal basis by local steps, as well as the possibility to uncross tight constraints due to the submodularity of the rank functions, are the main ingredients of the proof of the following theorem, which implies Lemma 2. We defer the proof to the full version of the paper.

Theorem 2. *Let $F \in \mathcal{F}(k)$ be of minimal weight w.r.t. the admissible weight function w . If the weight function \bar{w} differs from w only on chain K_{r^*} such that*

$$\bar{w}(r_k) = \begin{cases} w(r_k), & \text{if } r \neq r^*, \\ w(r_{k+1}), & \text{else,} \end{cases}$$

then there exists $F' = F - e + g \in \mathcal{F}(k)$ of minimal weight w.r.t. \bar{w} .

References

1. Ackermann, H., Röglin, H., Vöcking, B.: Pure Nash equilibria in player-specific and weighted congestion games. *Theoret. Comput. Sci.* 410(17), 1552–1563 (2009)
2. Anshelevich, E., Dasgupta, A., Kleinberg, J., Tardos, É., Wexler, T., Roughgarden, T.: The price of stability for network design with fair cost allocation. *SIAM J. Comput.* 38(4), 1602–1623 (2008)
3. Antonakopoulos, S., Chekuri, C., Shepherd, F.B., Zhang, L.: Buy-at-bulk network design with protection. *Math. Oper. Res.* 36(1), 71–87 (2011)
4. Beckmann, M., McGuire, C., Winsten, C.: *Studies in the Economics and Transportation*. Yale University Press, New Haven (1956)
5. Chen, H., Roughgarden, T.: Network design with weighted players. *Theory Comput. Syst.* 45(2), 302–324 (2009)
6. Chen, H.L., Roughgarden, T., Valiant, G.: Designing network protocols for good equilibria. *SIAM J. Comput.* 39(5), 1799–1832 (2010)
7. Dunkel, J., Schulz, A.: On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Math. Oper. Res.* 33(4), 851–868 (2008)
8. Faigle, U.: The greedy algorithm for partially ordered sets. *Discrete Math.* 28(2), 153–159 (1979)
9. von Falkenhausen, P., Harks, T.: Optimal cost sharing for resource selection games. *Math. Oper. Res.* 38(1), 184–208 (2013)
10. Fotakis, D., Kontogiannis, S., Spirakis, P.: Selfish unsplittable flows. *Theoret. Comput. Sci.* 348(2-3), 226–239 (2005)
11. Gairing, M., Monien, B., Tiemann, K.: Routing (un-)splittable flow in games with player-specific linear latency functions. *ACM Trans. Algorithms* 7(3), 1–31 (2011)
12. Harks, T., Klimm, M.: Congestion games with variable demands. In: Apt, K. (ed.) *Proc. 13th Conf. Theoret. Aspects of Rationality and Knowledge*, pp. 111–120 (2011)
13. Harks, T., Klimm, M.: On the existence of pure Nash equilibria in weighted congestion games. *Math. Oper. Res.* 37(3), 419–436 (2012)
14. Haurie, A., Marcotte, P.: On the relationship between Nash-Cournot and Wardrop equilibria. *Networks* 15, 295–308 (1985)
15. Helgason, T.: Aspects of the theory of hypermatroids. In: *Hypergraph Seminar*, pp. 191–213. Springer (1974)
16. Ieong, S., McGrew, R., Nudelman, E., Shoham, Y., Sun, Q.: Fast and compact: A simple class of congestion games. In: *Proc. 20th Natl. Conf. Artificial Intelligence and the 17th Innovative Appl. Artificial Intelligence Conf.*, pp. 489–494 (2005)
17. Johari, R., Tsitsiklis, J.N.: A scalable network resource allocation mechanism with bounded efficiency loss. *IEEE J. Sel. Area Commun.* 24(5), 992–999 (2006)
18. Kelly, F., Maulloo, A., Tan, D.: Rate control in communication networks: Shadow prices, proportional fairness, and stability. *J. Oper. Res. Soc.* 49, 237–252 (1998)
19. Krysta, P., Sanders, P., Vöcking, B.: Scheduling and traffic allocation for tasks with bounded splittability. In: Rován, B., Vojtáš, P. (eds.) *MFCS 2003*. LNCS, vol. 2747, pp. 500–510. Springer, Heidelberg (2003)
20. Meyers, C.: *Network Flow Problems and Congestion Games: Complexity and Approximation Results*. Ph.D. thesis, MIT, Operations Research Center (2006)
21. Milchtaich, I.: Congestion games with player-specific payoff functions. *Games Econom. Behav.* 13(1), 111–124 (1996)
22. Milchtaich, I.: The equilibrium existence problem in finite network congestion games. In: Spirakis, P., Mavronicolas, M., Kontogiannis, S. (eds.) *WINE 2006*. LNCS, vol. 4286, pp. 87–98. Springer, Heidelberg (2006)

23. Rosenthal, R.: A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory* 2(1), 65–67 (1973)
24. Rosenthal, R.: The network equilibrium problem in integers. *Networks* 3, 53–59 (1973)
25. Roughgarden, T.: *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge (2005)
26. Schrijver, A.: *Combinatorial optimization: Polyhedra and efficiency*, vol. 24. Springer (2003)
27. Srikant, R.: *The Mathematics of Internet Congestion Control*. Birkhäuser, Basel (2003)
28. Tran-Thanh, L., Polukarov, M., Chapman, A., Rogers, A., Jennings, N.R.: On the existence of pure strategy Nash equilibria in integer-splittable weighted congestion games. In: Persiano, G. (ed.) *SAGT 2011. LNCS*, vol. 6982, pp. 236–253. Springer, Heidelberg (2011)
29. Wardrop, J.: Some theoretical aspects of road traffic research. *Proc. Inst. Civil Engineers* 1(Part II), 325–378 (1952)