Potential Function Minimizers of Combinatorial Congestion Games: Efficiency and Computation

PIETER KLEER, CWI, Amsterdam, The Netherlands
GUIDO SCHÄFER, CWI and Vrije Universiteit Amsterdam, The Netherlands

We study the inefficiency and computation of pure Nash equilibria in unweighted congestion games, where the strategies of each player i are given implicitly by the binary vectors of a polytope P_i. Given these polytopes, a strategy profile naturally corresponds to an integral vector in the *aggregation polytope* $P_N = \sum_i P_i$. We identify two general properties of the aggregation polytope P_N that are sufficient for our results to go through, namely the *integer decomposition property (IDP)* and the *box-totally dual integrality property (box-TDI).* Intuitively, the IDP is needed to decompose a load profile in P_N into a respective strategy profile of the players, and box-TDI ensures that the intersection of a polytope with an arbitrary integer box is an integral polytope. Examples of polytopal congestion games which satisfy IDP and box-TDI include common source network congestion games, symmetric totally unimodular congestion games, non-symmetric matroid congestion games and symmetric matroid intersection congestion games (in particular, r-arborescences and strongly base-orderable matroids).

Our main contributions for polytopal congestion games satisfying IDP and box-TDI are as follows:

(1) We derive tight bounds on the price of stability for these games. This extends a result of Fotakis (2010) on the price of stability for symmetric network congestion games to the larger class of polytopal congestion games. Our bounds improve upon the ones for general polynomial congestion games obtained by Christodoulou and Gairing (2016).
(2) We show that pure Nash equilibria can be computed in strongly polynomial time for these games. To this aim, we generalize a recent aggregation/decomposition framework by Del Pia et al. (2017) for symmetric totally unimodular and non-symmetric matroid congestion games, both being a special case of our polytopal congestion games.
(3) Finally, we generalize and extend results on the computation of strong equilibria in bottleneck congestion games studied by Harks, Hoefer, Klimm and Skopalik (2013). In particular, we show that strong equilibria can be computed efficiently for symmetric totally unimodular bottleneck congestion games.

In general, our results reveal that the combination of IDP and box-TDI gives rise to an efficient approach to compute a pure Nash equilibrium whose inefficiency is better than in general congestion games.

1 INTRODUCTION

Background and Motivation. Congestion games constitute an important class of non-cooperative games which have been studied intensively since their introduction by Rosenthal (1973). In a *congestion game*, a (finite) set of players compete over a (finite) set of resources. Each resource is associated with a non-negative and non-decreasing cost (or delay) function which specifies its cost depending on the total number of players using it. Every player chooses a subset of resources from a set of available resource subsets (corresponding to the player's strategies) and experiences a cost equal to the sum of the costs of the chosen resources. The goal of each player is to minimize

EC'17, June 26–30, 2017, Cambridge, Massachusetts, USA. © 2017 ACM ISBN 978-1-4503-4527-9/17/06…$15.00.
DOI: http://dx.doi.org/10.1145/3033274.3085149

her individual cost. Congestion games are both theoretically appealing and practically relevant. For example, they find their applications in network routing, resource allocation and scheduling problems.

In a seminal paper, Rosenthal (1973) establishes the existence of pure Nash equilibria in congestion games. He proves this result through the use of an *exact potential function* which assigns a value to each strategy profile such that the difference in potential value of any two strategy profiles corresponding to a unilateral deviation of a player is equal to the difference in cost experienced by that player. Rosenthal proves that every congestion game admits an exact potential function, also known as *Rosenthal's potential*. As a consequence, every best response sequence must converge to a pure Nash equilibrium because the game is finite. Further, this shows that the set of pure Nash equilibria corresponds to the set of local minima of Rosenthal's potential. Especially this correspondence has helped to shed light on several important aspects of congestion games in recent years.

Rosenthal's Potential: Local vs. Global Minimizers. One of the most predominant aspects that has been studied intensively in recent years is the computational complexity of finding a pure Nash equilibrium. In a seminal paper, Fabrikant, Papadimitriou, and Talwar (2004) show that the problem of finding a pure Nash equilibrium is PLS-complete, both for symmetric congestion games and non-symmetric network congestion games. In particular, this suggests that a polynomial time algorithm for finding a pure Nash equilibrium is unlikely to exist for these games. In their proof they construct instances of non-symmetric network congestion games where any best response sequence has exponential length. Ackermann, Röglin, and Vöcking (2008) strengthen this result by exhibiting instances of symmetric network congestion games for which every best response sequence (from certain initial configurations) has exponential length. On the positive side, they prove that best response dynamics converge in polynomial time for non-symmetric matroid congestion games, where the available resource subsets of the players correspond to bases of a given matroid (see below for formal definitions). The authors also show that basically this is the only class of congestion games for which this property holds true.

Most previous works in this context focus on the analysis of decentralized dynamics to reach a pure Nash equilibrium (see, e.g., Ackermann et al. (2008); Chien and Sinclair (2007); Christodoulou et al. (2012); Even-Dar et al. (2007); Fabrikant et al. (2004); Fotakis (2010); Ieong et al. (2005)); said differently, these works focus on finding a local minimum of Rosenthal's potential. Much less is known about the problem of computing a pure Nash equilibrium that corresponds to a *global* minimum. Fabrikant et al. (2004) use this idea to show that a pure Nash equilibrium can be computed in polynomial time for symmetric network congestion games. The authors observe that in this case a global minimum of Rosenthal's potential can be computed by a reduction to a min-cost flow problem (if all cost functions are non-decreasing). Note that this is in stark contrast with the fact that best response dynamics might need exponential time in this case (Ackermann et al. 2008).

Only very recently, Del Pia, Ferris, and Michini (2017) make further progress along these lines. The authors consider congestion games where the strategy sets of the players are given implicitly by a *polyhedral description* (see also (Chan and Jiang 2016)). More precisely, for each player i the incidence vectors of the strategies are defined as the binary vectors in a polytope $P_i = \{x : A_i x \leq b_i\}$, where A_i is an integral matrix and b_i is an integral vector. They (mostly) focus on the case where the matrix A_i is *totally unimodular* (see below for formal definitions) and thus the describing polytope P_i is integral (i.e., all its extreme points are integral); they term these games *totally unimodular (TU) congestion games*. For symmetric TU congestion games (when all A_i, b_i are identical), they devise an aggregation/decomposition framework that reduces the problem of finding a global minimum of Rosenthal's potential to an integer linear programming problem. Using this framework, they show that pure Nash equilibria can be computed efficiently for symmetric TU congestion games. The authors also show that this problem is PLS-complete for non-symmetric TU congestion games. Further, they show that their framework can be adapted to the case of non-symmetric matroid congestion games.

Another important aspect that has has been the subject of intensive research in recent years is the inefficiency of pure Nash equilibria in congestion games (see, e.g., Aland et al. (2011); Awerbuch et al. (2005); Caragiannis et al. (2006); Christodoulou and Gairing (2016); Christodoulou and Koutsoupias (2005); de Jong et al. (2016); Feldman et al. (2016); Fotakis (2010); Gairing et al. (2008); Koutsoupias and Papadimitriou (1999); Lücking et al. (2008); Roughgarden (2015)). Here the goal is to assess the *social cost* (defined as the sum of the costs of the players) of a pure Nash equilibrium relative to an optimal outcome. Koutsoupias and Papadimitriou (1999) introduced the *price of anarchy* as the ratio between the worst social cost of a Nash equilibrium and the social cost of an optimum. Anshelevich et al. (2004) defined the *price of stability* as the ratio between the best social cost of a Nash equilibrium and the social cost of an optimum.

Fotakis (2010) reveals an intriguing connection between the price of stability of network congestion games and the price of anarchy of their *non-atomic* counterparts. More specifically, he shows that for symmetric network congestion games the ratio between the social cost of a global minimum of Rosenthal's potential and the social cost of a social optimum is at most $\rho(\mathcal{D})$, where $\rho(\mathcal{D})$ is a tight bound on the price of anarchy for *non-atomic* network congestion games with latency functions in class $\mathcal{D}$ introduced by Correa, Schulz, and Stier-Moses (2004). In particular, this implies that the price of stability of symmetric network congestion games with cost functions in $\mathcal{D}$ is at most $\rho(\mathcal{D})$. For example, this parameter equals $4/3$ for the class of affine functions and $(27 + 6\sqrt{3})/23 \approx 1.63$ for quadratic functions. These type of bounds fall within Roughgarden's *smoothness framework* (Roughgarden 2015).

Our Contributions. In light of the discussion above a natural question that arises, and which we settle in this paper, is:

> *Which structural properties of the strategy sets of the players are sufficient to*
> *(A) efficiently compute a global minimum of Rosenthal's potential, and*
> *(B) bound the inefficiency of the resulting pure Nash equilibrium?*

In order to tackle this question, we use a polyhedral approach similar to the ones by Chan and Jiang (2016) and Del Pia et al. (2017). But in contrast to these works, we do not restrict our attention to polyhedral descriptions arising from totally unimodular matrices only. Instead, we identify more general polyhedral properties of the describing systems that are sufficient to achieve (A) and (B). Our main contribution in this paper is to unify and extend the results in (Del Pia et al. 2017; Fotakis 2010) to a much larger class of polytopal congestion games.

More specifically, we consider *polytopal congestion games* in which the incidence vectors of the strategies of player i are given by the binary vectors in a polytope $P_i = \{x : Ax \leq b_i\}$, where A is an integral matrix and b_i is an integral vector. Given the polytopes of all players, a strategy profile naturally corresponds to an integral vector in the *aggregation polytope* $P_N = \sum_i P_i$. We identify two general properties of the aggregation polytope P_N which are sufficient for our results to go through, namely the *integer decomposition property (IDP)* and the *box-totally dual integrality property (box-TDI)* (formal definitions are given below). The integer decomposition property is needed to decompose a load profile in P_N to a respective strategy profile of the players. Intuitively, the box-TDI property ensures that the intersection of a polytope with an arbitrary integer box is an integral polytope.

Our main contributions for polytopal congestion games are as follows:

(1) We generalize the upper bound of $\rho(\mathcal{D})$ on the price of stability for symmetric network congestion games by Fotakis (2010) to the much larger class of polytopal congestion games satisfying IDP and box-TDI (Section 4). To this aim, we introduce a novel property, which we term the *symmetric difference decomposition property*, and show that it is satisfied by our polytopal congestion games. By exploiting this property, we can generalize the proof by Fotakis (2010) to these games. We also prove that our bounds are tight for these games.
(2) We provide a framework to derive an efficient algorithm for computing a feasible load profile minimizing Rosenthal's potential for polytopal congestion games satisfying IDP and box-TDI (Section 5). The time complexity of this algorithm is polynomial in the number of players and resources, the largest entry in $\sum_i b_i$ and the complexity of a separation oracle for the aggregation polytope. This generalizes the framework of Del Pia et al. (2017) for symmetric TU congestion games and non-symmetric matroid congestion games, both being special cases of our polytopal congestion games.
(3) We give a series of examples of combinatorial polytopal congestion games satisfying IDP and box-TDI (Section 6). These examples include the symmetric TU congestion games by Del Pia et al. (2017), common source network congestion games, symmetric matroid intersection congestion games (in particular, r-arborescences and strongly base-orderable matroids) and non-symmetric matroid congestion games.
(4) We show that the integer decomposition property and box-total dual integrality can be used to generalize and extend some results on the computation of strong equilibria in bottleneck congestion games studied by Harks, Hoefer, Klimm, and Skopalik (2013) (Section 7). In particular, we obtain the new result that strong equilibria can be computed in strongly polynomial time for symmetric totally unimodular bottleneck congestion games.

Further Implications and Significance. To the best of our knowledge, all previous works addressed either (A) or (B), but not both. Note that the combination of (A) and (B) provides an efficient algorithm for the computation of a pure Nash equilibrium that comes with a provable inefficiency guarantee. Said differently, our contributions (1) and (2) can be seen as an efficient equilibrium selection procedure to find a pure Nash equilibrium whose social cost is at most $\rho(\mathcal{D})$ times the optimal social cost. By exploiting contribution (1), we obtain new bounds on the price of stability which improve upon the ones for general polynomial congestion games: Christodoulou and Gairing (2016) derive tight bounds on the price of stability of congestion games with polynomial cost functions of maximum degree d, which grow like $d + 1$. However, for the class $\mathcal{D}_d$ of polynomial functions of maximum degree d, we have $\rho(\mathcal{D}_d) \approx d/\log(d)$ for large d (see, e.g., (Feldman et al. 2016)), which is a significant improvement over the general case. Also, $\rho(\mathcal{D})$ is well-understood for various classes of delay functions $\mathcal{D}$; for example, a closed form expression is known for $\rho(\mathcal{D}_d)$ (see preliminaries).

Our upper bound of $\rho(\mathcal{D})$ on the price of stability is (asymptotically) tight, even for symmetric singleton congestion games with the class of delay functions $\mathcal{D}$ containing all constant functions and being closed under dilations (see below for formal definitions). Note that singleton congestion games constitute a special case of all polytopal congestion games mentioned in (3). In particular, our results settle the exact price of stability for these applications.

Our results also reveal that the price of stability for matroid congestion games is much more well-behaved than the price of anarchy. de Jong et al. (2016) show that for symmetric k-uniform matroid congestion games with affine cost functions, the price of anarchy lies strictly between $4/3$ and $28/13 \approx 2.15$. Obtaining a tight bound for this case seems (highly) non-trivial. In contrast, for the price of stability we provide tight bounds for arbitrary non-symmetric matroid congestion games with arbitrary cost functions.

Our framework in (2) unifies and extends the aggregation/decomposition framework of Del Pia et al. (2017). In particular, the symmetric TU congestion games and non-symmetric matroid congestion games (considered separately in (Del Pia et al. 2017)) fall into our class of polytopal congestion games satisfying IDP and box-TDI. Similarly, all combinatorial TU congestion games (i.e., network, matching, edge cover, vertex cover and stable set congestion games) and their respective extensions to the maximum (or minimum) cardinality versions in (Del Pia et al. 2017) can be handled by our framework.

Contribution (2) can also be regarded as a "black-box" approach for the computation of a pure Nash equilibrium. Given a congestion game that exhibits some combinatorial structure, checking whether our approach applies reduces to the following three tasks: (i) derive a polytopal description P_i for the strategy set of each player i, (ii) verify whether the resulting aggregation polytope P_N satisfies the IDP, (iii) check that the system describing the aggregation polytope P_N is box-TDI. In particular, if the integer decomposition of P_N can be done in polynomial time, then this approach provides an efficient algorithm to compute a pure Nash equilibrium. By exploiting this idea, we derive strongly polynomial time algorithms for the computation of Rosenthal's potential minimum for all applications mentioned in (3).

It is interesting to note that the IDP seems to be the limiting property for our approach to apply. For example, non-symmetric network congestion games can naturally be modeled as polytopal congestion games satisfying box-TDI. But it is easy to see that the IDP does not hold. In fact, it is unlikely that an efficient algorithm to find a pure Nash equilibrium exists because this problem is PLS-complete (Fabrikant et al. 2004).

2 PRELIMINARIES

Congestion Games and Rosenthal's Potential. A congestion game Γ is given by a tuple $(N, E, (\mathcal{S}_i)_{i\in N}, (c_e)_{e\in E})$, where $N = [n]$ is a finite set of players, $E = [m]$ is a finite set of resources (or facilities), $\mathcal{S}_i \subseteq 2^E$ is a set of strategies of player $i \in N$, and $c_e : \mathbb{R}_{\geq 0} \to \mathbb{R}$ is a cost function of resource $e \in E$. Unless stated otherwise, the cost functions are assumed to be non-negative and non-decreasing. For a strategy profile $s = (s_1, \dots, s_n) \in \times_i \mathcal{S}_i$, we define x_e as the number of players using resource e, i.e., $x_e = x_e(s) = |\{i \in N : e \in s_i\}|$. We call x the *load profile* corresponding to strategy profile s. The cost of player $i \in N$ under a strategy profile $s = (s_1, \dots, s_n) \in \times_i \mathcal{S}_i$ is given by $C_i(s) = \sum_{e\in s_i} c_e(x_e)$. If $\mathcal{S}_i = \mathcal{S}_j$ for all $i, j \in N$, the game is called *symmetric*. The *social cost* $C(s)$ of a strategy profile refers to the sum of the players' individual costs, i.e., $C(s) = \sum_i C_i(s)$.

We say that $\Phi : \times_i \mathcal{S}_i \to \mathbb{R}$ is an *exact potential function* for a congestion game Γ if for every strategy profile $s \in \times_i \mathcal{S}_i$, for every player $i \in N$ and every unilateral deviation $s_i' \in \mathcal{S}_i$ of i it holds: $\Phi(s) - \Phi(s_{-i}, s_i') = C_i(s) - C_i(s_{-i}, s_i')$. Rosenthal (1973) shows that $\Phi(s) = \sum_{e\in E} \sum_{k=1}^{x_e} c_e(k)$ is an exact potential. Subsequently, we refer to this potential function simply as *Rosenthal's potential*. Further, a strategy profile minimizing Rosenthal's potential is said to be a *Rosenthal minimizer*.

Inefficiency of equilibria. A strategy profile s is a *Nash equilibrium* if for every player $i \in N$ it holds that $C_i(s) \leq C_i(s_i', s_{-i})$ for all $s_i' \in \mathcal{S}_i$. Further, a strategy profile s is a *strong equilibrium* if for every group of players $I \subseteq N$ and every deviation $s_I' \in \times_{i\in I} \mathcal{S}_i$ of the players in I, it holds that $C_i(s) \leq C_i(s_I', s_{-I})$ for some $i \in I$.

The price of anarchy (POA) and the price of stability (POS) of a game Γ are defined as

$$\mathrm{POA}(\Gamma) = \frac{\max_{s \in \mathrm{NE}} C(s)}{\min_{s^* \in \times_i \mathcal{S}_i} C(s^*)} \quad \text{and} \quad \mathrm{POS}(\Gamma) = \frac{\min_{s \in \mathrm{NE}} C(s)}{\min_{s^* \in \times_i \mathcal{S}_i} C(s^*)},$$

where NE denotes the set of all Nash equilibria of Γ. For a collection of games $\mathcal{H}$ we define $\mathrm{POA}(\mathcal{H}) = \sup_{\Gamma \in \mathcal{H}} \mathrm{POA}(\Gamma)$ and $\mathrm{POS}(\mathcal{H}) = \sup_{\Gamma \in \mathcal{H}} \mathrm{POS}(\Gamma)$. These notions naturally generalize to the solution concept of strong equilibria.

Smoothness parameter. Correa et al. (2004) show that for *non-atomic* network congestion games with latency functions in class $\mathcal{D}$ the price of anarchy of an instance is at most

$$\rho(\mathcal{D}) := (1 - \beta(\mathcal{D}))^{-1}, \quad \text{where} \quad \beta(\mathcal{D}) = \sup_{d \in \mathcal{D}} \sup_{x, y \in \mathbb{R}: x \geq y > 0} \frac{y(d(x) - d(y))}{x d(x)}. \tag{1}$$

The value of $\rho(\mathcal{D})$ is well-understood for many important classes of latency functions. For example, let $\mathcal{D}_d = \{g : \mathbb{R}_{\geq 0} \to \mathbb{R}_{\geq 0} : g(\mu x) \geq \mu^d g(x)\ \forall \mu \in [0, 1]\}$. In particular, $\mathcal{D}_d$ contains all polynomial latency functions with non-negative coefficients and maximum degree d. We have $\rho(\mathcal{D}_d) = (1 - \frac{d}{(d+1)^{(d+1)/d}})^{-1}$. The parameter $\rho(\mathcal{D})$ plays a crucial role in bounding the price of stability of our congestion games.

Integral polytopes. We review some basic definitions and results from polyhedral combinatorics which are used in this paper. A polytope $P \subseteq \mathbb{R}^m$ is the convex hull of a finite set of vectors in $\mathbb{R}^m$. We say that P is *integral* if all its extreme points are integral vectors. P is said to be *box-integral* if the intersection of P with any integral box, i.e., $P \cap \{x : c \leq x \leq d\}$ for arbitrary integral c and d, yields an integral polytope.

A matrix $A \in \mathbb{Q}^{r \times m}$ is *totally unimodular (TUM)* if the determinant of each square submatrix of A is in $\{0, \pm 1\}$. In particular, each entry of a totally unimodular matrix is in $\{0, \pm 1\}$. If A is totally unimodular and $b \in \mathbb{Z}^m$ is an integer vector, then the polyhedron $P = \{x : Ax \leq b\}$ is integral (Schrijver 1986, Theorem 19.1).

Edmonds and Giles (1977) introduced the powerful notion of total dual integrality. A rational system $Ax \leq b$ with $A \in \mathbb{Q}^{r \times m}$ and $b \in \mathbb{Q}^r$ is *totally dual integral (TDI)* if for every integral $c \in \mathbb{Z}^m$, the dual of minimizing $c^\top x$ over $Ax \leq b$, i.e., $\max\{y^\top b : y \geq 0,\ y^\top A = c^\top\}$, has an integer optimum solution y, if it is finite. If $Ax \leq b$ is a TDI-system and b is integral, then the polyhedron $P = \{x : Ax \leq b\}$ is integral (Schrijver 1986, Corollary 22.1c). Note that TDI is a weaker sufficient condition for the integrality of P than TUM.

The system $Ax \leq b$ is *box-totally dual integral (box-TDI)* if the system $Ax \leq b,\ l \leq x \leq u$ is TDI for all rational vectors l and u. We say that a polytope P is *box-TDI*, if it can be described by a box-TDI system. If P has some box-TDI describing system, then every TDI-system describing P is also box-TDI (Schrijver 1986, Theorem 22.8). We will use the following properties of box-TDI descriptions:

Proposition 2.1. *(Schrijver 1986, Section 22.5) The following statements are equivalent:*

(i) The system $Ax \leq b,\ x \geq 0$ is box-TDI.
(ii) The system $Ax + \mu = b,\ \mu \geq 0,\ x \geq 0$ is box-TDI.
(iii) The system $Ax \leq \alpha b,\ x \geq 0$ is box-TDI for all $\alpha \geq 0$.
(iv) The system $a\zeta_0 + Ax \leq b$ is box-TDI, where a is a column of A and ζ_0 is a new variable.

Moreover, if a polytope P is box-integral, then every edge of P is in the direction of a $\{0, \pm 1\}$-vector.

For our computational results we need the notion of a separation oracle. Let $P \in \mathbb{R}^m$ be a polytope defined by a finite set of rational inequalities. Given a vector $y \in \mathbb{Q}^n$, we assume that there is a *separation oracle* that decides whether $y \in P$ or not, and in the latter case it returns a vector $a \in \mathbb{Q}^n$ such that $a^\top x < a^\top y$ for all $x \in P$. All applications that we consider in this paper are known to have efficient separation oracles. Finally, for a vector $y \in \mathbb{Q}^r$ we define the *size* of y as $\mathrm{size}(y) = \max\{\log(y_i) + 1 : i = 1, \dots, r\}$.

Matroids. We introduce some general terminology and facts for matroids (an extensive treaty of matroids can be found, e.g., in (Schrijver 2003)). Let E be a finite set of elements and $\mathcal{I} \subseteq 2^E$ be a collection of subsets of E (called *independent sets*). The pair $\mathcal{M} = (E, \mathcal{I})$ is a *matroid* if the following three properties hold: (1) $\emptyset \in \mathcal{I}$, (2) if $A \in \mathcal{I}$ and $B \subseteq A$, then $B \in \mathcal{I}$, (3) if $A, B \in \mathcal{I}$ and $|A| > |B|$, then there exists an $a \in A \setminus B$ such that $B + a \in \mathcal{I}$. An independent set $B \in \mathcal{I}$ of maximum size is called a *basis*. We use $\mathcal{B}$ to denote the set of all bases of $\mathcal{M}$. The matroid $\mathcal{M}$ also has a rank function $r : 2^E \to \{1, \dots, |E|\}$ which maps every subset $A \subseteq E$ to the cardinality of the largest independent set contained in A.

The *base matroid polytope* is given by

$$P_{\mathcal{M}} = \{x : x(A) \le r(A)\ \forall A \subset E,\ x(E) = r(E),\ x \ge 0\},$$

where $x(A) = \sum_{a \in A} x_a$ for all $A \subseteq E$. It is the convex hull of the incidence vectors of the bases in $\mathcal{B}$ (Schrijver 2003). If in the description above the equality $x(E) = r(E)$ is replaced by $x(E) \le r(E)$, we obtain the *independent set polytope* which is the convex hull of the incidence vectors of the independent sets.

We assume that the matroid is given by an *independence oracle* that takes as input a subset $A \subseteq 2^E$ and returns whether or not $A \in \mathcal{I}$. Given an independence oracle, we can determine in time polynomial in $|E|$ and the complexity of the oracle, whether a set is a basis and what the rank of a set is. Further, there exists a separation oracle for $P_{\mathcal{M}}$ that runs in time polynomial in $|E|$ and the complexity of an independence oracle. This follows from the fact that the *most violated inequality problem* can be solved in time polynomial in $|E|$ and the complexity of an independence oracle. The most violated inequality problem takes as input a vector $x \in \mathbb{Q}^m$ and returns whether or not $x \in P$, and if not, it returns a subset A for which $r(A) - x(A)$ is minimized, see, e.g., (Schrijver 2003, Section 40.3).

Given two matroids $\mathcal{M}_1$ and $\mathcal{M}_2$ on a common ground set E, the polytope

$$P_{\mathcal{M}_1, \mathcal{M}_2} = \{x : x(A) \le r_i(A)\ \forall A \subset E,\ x(E) = r_i(E) \text{ for } i = 1, 2,\ x \ge 0\} \tag{2}$$

is the convex hull of the common bases of matroids $\mathcal{M}_1$ and $\mathcal{M}_2$, see, e.g., (Schrijver 2003, Corollary 41.12d). It follows directly that $P_{\mathcal{M}_1, \mathcal{M}_2}$ also has a separation oracle which runs in time polynomial in $|E|$ and the complexity of the independence oracles for $\mathcal{M}_1$ and $\mathcal{M}_2$.

3 POLYTOPAL CONGESTION GAMES

We consider *polytopal congestion games* $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ with $N = [n]$ and $E = [m]$, where the set of strategies $\mathcal{S}_i$ of each player $i \in N$ is given implicitly by a polytopal representation. More precisely, let $\mathcal{X}_i$ be the finite set of all incidence vectors of the strategies of player i, i.e., $\mathcal{X}_i = \{\chi_i \in \{0, 1\}^m : \chi_{ie} = 1 \text{ iff } e \in s_i \text{ for } s_i \in \mathcal{S}_i\}$. The polytope P_i representing the strategies of player i is defined as the convex hull of $\mathcal{X}_i$, i.e., $P_i = \mathrm{conv}(\mathcal{X}_i) \subseteq [0, 1]^m$. We assume that P_i is given by $P_i = \{x : Ax \le b_i\} \subseteq [0, 1]^m$, where $A \in \mathbb{Z}^{r \times m}$ is an integral $r \times m$-matrix and $b_i \in \mathbb{Z}^r$ is an integral vector. Note that $\mathcal{X}_i = P_i \cap \{0, 1\}^m$. For notational convenience, we subsequently use $\mathcal{S}_i$ also to refer to the set of incidence vectors $\mathcal{X}_i$; no confusion shall arise.

We say that $(N, E, (\mathcal{S}_i)_{i \in N})$ is the *polytopal tuple* given by $P_1, \dots, P_n$, where $\mathcal{S}_i = P_i \cap \{0, 1\}^m$. Moreover, we define $P_N = \{y : Ay \le \sum_i b_i\} \subseteq \mathbb{R}^m$ to be the *aggregation polytope* of the tuple. If $b_i = b_j = b$ for all $i, j \in N$, the tuple is called *symmetric* and denoted by $(N, E, \mathcal{S})$ where

$\mathcal{S} = P \cap \{0,1\}^m$, with $P = \{x : Ax \leq b\}$. If additionally we equip the tuple with cost functions $(c_e)_{e \in E}$, we call $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ the *polytopal congestion game* given by $P_1, \dots, P_n$.

3.1 Aggregation Polytope: IDP and box-TDI

We identify two crucial properties that the aggregation polytope P_N has to satisfy for our results to go through:

1. P_N satisfies the integer decomposition property (IDP).
2. The system $Ay \leq \sum_i b_i$ describing P_N is box-totally dual integral (box-TDI).

The aggregation polytope $P_N = \{y : Ay \leq \sum_i b_i\} \subseteq \mathbb{R}^m$ has the *integer decomposition property (IDP)* if every integral $z \in P_N$ can be written as $z = \sum_{i=1}^n z^i$, where $z^i \in P_i \cap \mathbb{Z}^m$ for all $i = 1, \dots, n$.[1]

We next introduce the notion of a feasible load profile. Given a tuple $(N, E, (\mathcal{S}_i)_{i \in N})$, a load profile $y \in \mathbb{N}^m$ is said to be *feasible* for $(N, E, (\mathcal{S}_i)_{i \in N})$ if there exists a strategy profile $s = (s_1, \dots, s_n) \in \times_i \mathcal{S}_i$ such that y is the load profile corresponding to s, i.e., $y_e = x_e(s)$ for all $e \in E$. We omit the explicit reference to the tuple if it is clear from the context.

The IDP is crucial to establish a correspondence between feasible load profiles for $(N, E, (\mathcal{S}_i)_{i \in N})$ and the integral vectors in P_N.

Proposition 3.1. *If the aggregation polytope P_N of a polytopal tuple $(N, E, (\mathcal{S}_i)_{i \in N})$ has the IDP, then the feasible load profiles of the tuple correspond precisely to the integral vectors in P_N.*

Proof. Let $s = (s_1, \dots, s_n) \in \times_i \mathcal{S}_i$ be a strategy profile and let x be the load profile corresponding to s. It follows directly that $x \in P_N$ by definition of P_N. Moreover, any integral vector z in P_N can be decomposed as $z = \sum_{i=1}^n z^i$ where $z^i \in P_i \cap \mathbb{Z}^m$ for all $i = 1, \dots, n$. This implies that for every i the vector z^i is the incidence vector of some strategy of player i and thus z is a feasible load profile. □

The main reason as to why box-TDI is useful, is that it serves as a sufficient condition to show that the polytope it describes is box-integral.

Proposition 3.2. *If the system $Ax \leq b$ describing a polytope P is box-TDI and b is integral, then P is box-integral.*

Proof. By assumption, the describing system $Ax \leq b$ of P is box-TDI. Thus the system $Ax \leq b,\ l \leq y \leq u$ is TDI for all rational vectors l and u. In particular, $Ax \leq b,\ c \leq y \leq d$ is TDI for arbitrary integral vectors c and d. Because b, c and d are integral, we can conclude that the polytope $P \cap \{y : c \leq y \leq d\}$ is integral (see, e.g., (Schrijver 1986, Corollary 22.1c)). □

It seems that most 0/1-polytopes for which the integer decomposition property is known in the literature, also have a box-TDI describing system. We do not know if this is true in general, but it would imply that the box-TDI assumption is redundant in all our statements below.

4 PRICE OF STABILITY

We first introduce the *symmetric difference decomposition property*. This property is crucial to derive the bound on the price of stability.

[1] Note that in the symmetric case ($b_i = b_j$ for all $i, j \in N$) this definition reduces to the integer decomposition property for a polytope $P_N = nP$ as introduced by Baum and Trotter (1978).

4.1 Symmetric Difference Decomposition Property

Definition 4.1 (Symmetric difference decomposition). A tuple $(N, E, (\mathcal{S}_i)_{i\in N})$ satisfies the *symmetric difference decomposition property (SDD)* if for all feasible load profiles f and g, there exist vectors $a^1, \ldots, a^q \in \{0, \pm 1\}^m$ such that $g - f = \sum_{k=1}^{q} a^k$, and, for all $k = 1, \ldots, q$, the load profile $f + a^k$ is feasible and a^k satisfies

$$a_e^k = -1 \;\Rightarrow\; f_e - g_e > 0 \qquad \text{and} \qquad a_e^k = 1 \;\Rightarrow\; f_e - g_e < 0. \tag{3}$$

As an example, let us consider symmetric network congestion games, where the common strategy set of all players is the set of all directed simple s,t-paths in some directed graph $G = (V, A)$ with $s, t \in V$. Here each feasible load profile corresponds to an integral feasible s,t-flows of value $n = |N|$. The symmetric difference of two flows f and g can be written as the sum of unit circuit flows on cycles.[2] The incidence vectors of these unit circuit flows correspond to the vectors a^k in Definition 4.1.

The following theorem establishes the symmetric difference decomposition property.

Theorem 4.2. *Let $(N, E, (\mathcal{S}_i)_{i\in N})$ be a polytopal tuple given by $P_1, \ldots, P_n$ and let P_N be the corresponding aggregation polytope. If P_N satisfies the IDP and has a box-TDI description, then the tuple has the symmetric difference decomposition property.*

Proof. We start by adding slack-variables to the system $Ay \leq \sum_{i\in N} b_i$ describing P_N. Note that by Proposition 2.1 (ii) box-TDI is preserved under adding slack variables. As a result, we obtain the polytope

$$Q_N = \{(y, \mu) : Ay + \mu = \textstyle\sum_i b_i,\; \mu \geq 0,\; y \geq 0\}$$

for which its describing system is box-TDI. Also, Q_N is integral.

Let f and g be two feasible load profiles with $f \neq g$. By Proposition 3.1, we have $f, g \in P_N$. Therefore, there are non-negative integral slack vectors τ, σ such that $(f, \tau), (g, \sigma) \in Q_N$. Let $f' = (f, \tau)$ and $g' = (g, \sigma)$. Observe that τ and σ are integral because of the integrality of A, $\sum_i b_i$ and f and g, respectively.

Note that the pairs $f' = (f, \tau)$ and $g' = (g, \sigma)$ are vectors in $\mathbb{Z}^{m+r}$ since A is an $r \times m$-matrix. Let $c, d \in \mathbb{Z}^{m+r}$ be vectors defined by $c_j = \min\{f'_j, g'_j\}$ and $d_j = \max\{f'_j, g'_j\}$ for $j = 1, \ldots, r + m$, and let B be the integral box defined by $B = \{z : c \leq z \leq d\} \subseteq \mathbb{R}^{m+r}$.

Lemma 4.3. *The polytope $Q_N \cap B$ is integral and every edge of $Q_N \cap B$ is in the direction of a $\{0, \pm 1\}$-vector.*

Proof. The integrality follows from box-TDI of the integral system Q_N. For the second part of the claim, we first show that $Q_N \cap B$ is box-integral. Note that Q_N is box-integral by Proposition 3.2. Let $B' = \{x : \gamma \leq x \leq \delta\} \in \mathbb{Z}^{m+r}$ be an arbitrary integral box. Note that $(Q_N \cap B) \cap B' = Q_N \cap (B \cap B')$ and that $B \cap B'$ is again an integral box, since B is integral as well (because f' and g' are integral). It follows that $Q_N \cap (B \cap B')$ is an integral polytope. Thus, $(Q_N \cap B) \cap B'$ is integral which proves that $Q_N \cap B$ is box-integral. The claim now follows from Proposition 2.1. □

Note that $f', g' \in Q_N \cap B$. Further, both f' and g' are extreme points of this polytope because they are extreme points of the box B. We now fix some edge of $Q_N \cap B$ containing f'. Such an edge must exist because $Q_N \cap B$ contains at least two elements (since $f' \neq g'$). Let $(a^1)' = (a^1, \mu^1)$ be the non-zero $\{0, \pm 1\}$-vector describing the direction of the edge.[3] Since $Q_N \cap B$ is an integral polytope we can show that $f' + (a^1)' \in Q_N \cap B$. To see this, let $h(\lambda) = f' + \lambda \cdot (a^1)'$ be a parametrization of

[2] A unit circuit flow is a $\{0, \pm 1\}$-flow that satisfies flow-conservation at every node, including s and t.

[3] Here, μ^1 corresponds to the slack variables, and a^1 to the original variables.

the edge for some range $0 \leq \lambda \leq \lambda^*$, where $h' = h(\lambda^*)$ is the other extreme point of the edge $(a^1)'$. Since f' is integral and $(a^1)'$ a $\{0, \pm 1\}$-vector, it must be that λ^* is a strictly positive integer.

We have shown that $f' + (a^1)' \in Q_N \cap B$. It follows that $A(f + a^1) + (\tau + \mu^1) = \sum_{i \in N} b_i$. Thus, $Aa^1 + \mu^1 = 0$ because $Af + \tau = \sum_{i \in N} b_i$. Moreover, by construction of the box B it follows that

$$(a^1)'_j = -1 \;\Rightarrow\; f'_j - g'_j > 0 \qquad \text{and} \qquad (a^1)'_j = 1 \;\Rightarrow\; g'_j - f'_j > 0 \tag{4}$$

for $j = 1, \dots, r + m$. Using the fact that $Aa^1 + \mu_1 = 0$, it now also follows that $g' - (a^1)'_j \in Q_N \cap B$. To see this, note that $A(g - a^1) + (\sigma - \mu^1) = Ag + \sigma - (Aa^1 + \mu^1) = \sum_{i \in N} b_i + 0$. Moreover, we also have $g' - (a^1)' \geq 0$ by construction, since if $(a^1)'_j = 1$ for some j then $g'_j > f'_j \geq 0$, so in particular $g'_j - 1 \geq 0$ (because of the integrality of g'_j).

We can now apply the same argument to the vectors f' and $g' - (a^1)'$ in order to obtain a vector $(a^2)'$ satisfying (4) and for which $f' + (a^2)'$, $g' - (a^1)' - (a^2)' \in Q_N$. Repeating this procedure we find vectors $(a^1)', \dots, (a^q)'$ satisfying (4), and such that $g' - f' = \sum_{k=1}^{q} (a^k)'$ with $f' + (a^k)' \in Q_N$ for $k = 1, \dots, q$.[4]

We argue that this process terminates. For the K-th step of this procedure, we have by construction of the $(a^k)'$,

$$T(K) = \left\| \left(g' - \sum_{k=1}^{K} (a^k)' \right) - f' \right\|_1 < \left\| \left(g' - \sum_{k=1}^{K-1} (a^k)' \right) - f' \right\|_1 = T(K-1)$$

where $||\cdot||_1$ is the L_1-norm. Since f', g' and the a^k are all integral this guarantees that the expression $T(K)$ decreases by at least one in every step.

We conclude the proof by showing that f and g can be decomposed according to Definition 4.1. We have $(a^k)' = (a^k, \mu^k)$ as defined before. It then follows that $a^1, \dots, a^q$ are vectors satisfying (3) such that $g - f = \sum_{k=1}^{q} a^k$ with $f + a^k \in P_N$ for $k = 1, \dots, q$. Note that a^k might be the zero-vector, if $(a^k)'$ only contained non-zero elements in the part of the vector corresponding to slack variables. These a^k can be left out. It remains to show that $f + a^k$ corresponds to a feasible strategy profile for $k = 1, \dots, q$. This follows directly from the fact that P_N has the IDP. The decomposition yields the strategies of the players. □

For symmetric polytopal tuples with common polytope P, we have $P_N = nP = \{y : y/n \in P\}$. The latter polytope has a box-TDI description if and only if P has a box-TDI description, which follows from Proposition 2.1 (iii). This yields the following corollary.

Corollary 4.4. *Let $(N, E, \mathcal{S})$ be a symmetric polytopal tuple given by P. If P satisfies the IDP and has a box-TDI description, then the tuple has the symmetric difference decomposition property.*

4.2 Upper bound for price of stability

The following is the main result of this section.

Theorem 4.5 (price of stability). *Let $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ be a congestion game with cost functions in class $\mathcal{D}$. Suppose that the tuple $(N, E, (\mathcal{S}_i)_{i \in N})$ satisfies the symmetric difference decomposition property. Then $POS(\Gamma) \leq \rho(\mathcal{D})$. Further, this bound is (asymptotically) tight, even for symmetric singleton congestion games.*

Recall that $\rho(\mathcal{D})$ is defined as in (1) and refers to the price of anarchy of *non-atomic* network congestion games with latency function in class $\mathcal{D}$. We need the following lemma to prove Theorem 4.5. Its proof relies on the symmetric difference decomposition property.

[4]This construction is essentially a *conformal circuit decomposition*, see, e.g., Onn, Rothblum, and Tangir (2005).

Lemma 4.6. *Let $(N, E, (\mathcal{S}_i)_{i \in N})$ satisfy the symmetric difference decomposition property and let $(c_e)_{e \in E}$ be arbitrary cost functions. Let f be a feasible load profile that minimizes Rosenthal's potential $\Phi(\cdot)$. Then for every feasible load profile g*

$$\Delta(f, g) := \sum_{e: f_e > g_e} (f_e - g_e) c_e(f_e) - \sum_{e: f_e < g_e} (g_e - f_e) c_e(f_e + 1) \leq 0.$$

Proof. Let f be a global minimizer of Rosenthal's potential and let g be an arbitrary feasible load profile. Then by the SDD property, there exist vectors $a^1, \ldots, a^q$ such that $g - f = \sum_{k=1}^{q} a^k$ for some q. Moreover,

$$\Phi(f) - \Phi(f + a^k) = \sum_{e:\ a_e^k = -1} c_e(f_e) - \sum_{e:\ a_e^k = 1} c_e(f_e + 1) \leq 0$$

for all $k = 1, \ldots, q$, where the inequality holds because f minimizes Rosenthal's potential Φ. By adding up these inequalities for all $k = 1, \ldots, q$, we obtain that $\Delta(f, g) \leq 0$. To see this, note that if $e \in E$ with $f_e > g_e$ then there are precisely $f_e - g_e$ vectors a^k with $a_e^k = -1$; similarly, if $e \in E$ with $g_e > f_e$ then there are precisely $g_e - f_e$ vectors a^k with $a_e^k = 1$. □

Proof of Theorem 4.5. The upper bound proof follows the same line of arguments as in (Fotakis 2010, Lemma 3). We sketch the main ideas here only and refer to (Fotakis 2010) for more details. Let f be a minimizer of Rosenthal's potential and g an arbitrary feasible load profile. Note that f is a pure Nash equilibrium. It can be shown that

$$C(f) \leq C(g) + \beta(\mathcal{D}) C(f) + \sum_{e: f_e > g_e} (f_e - g_e) c_e(f_e) - \sum_{e: f_e < g_e} (g_e - f_e) c_e(f_e + 1).$$

By Lemma 4.6, the sum difference $\Delta(f, g)$ is non-negative. By rearranging terms, we obtain $C(f)/C(g) \leq (1 - \beta(\mathcal{D}))^{-1} = \rho(\mathcal{D})$, which establishes the upper bound. The proof that this bound is tight will be given in the full version of the paper. □

5 MINIMIZING ROSENTHAL'S POTENTIAL

Del Pia, Ferris, and Michini (2017) introduce an *aggregation/decomposition method* for computing a global minimum of Rosenthal's potential. It consists of two phases: In the *aggregation phase*, we find a feasible load profile f^* minimizing Rosenthal's potential. In the *decomposition phase*, f^* is then decomposed into a feasible strategy profile. The authors provide an aggregation approach (detailed below) that works for totally unimodular matrices A and one common vector $b = b_i$ for all $i \in N$. Here we extend this result to aggregation polytopes P_N that have a box-TDI description.

Recall from Proposition 3.1 that if the aggregation polytope P_N of a polytopal congestion game has the integer decomposition property, then the feasible load profiles correspond to the integral vectors of P_N. As a result, the problem we need to solve in the aggregation phase is equivalent to

$$(Z) \qquad \min \sum_{e \in E} \sum_{k=1}^{f_e} c_j(k) \quad \text{s.t.} \quad f \in P_N \cap \mathbb{Z}^m$$

Note that this formulation is not a linear program in the variables $(f_e)_{e \in E}$. As in the approach of Del Pia et al. (2017), this problem can be resolved by introducing binary variables $h_j^k \in \{0, 1\}$ for $k = 1, \ldots, n$ and $j = 1, \ldots, m$. The interpretation is that $h_j^k = 1$ if at least k players are using resource $j \in E$, and $h_j^k = 0$ otherwise. In particular, if the cost functions $(c_e)_{e \in E}$ are non-decreasing, then the aggregation problem (Z) is equivalent to the problem (R) stated below:

$$(R) \qquad \min \sum_{j=1}^{m} \sum_{k=1}^{n} c_j(k) h_j^k$$

$$\text{s.t.} \quad [A, A, \ldots, A](h_1^1, \ldots, h_m^1, h_1^2, \ldots, h_m^2, \ldots, h_1^n, \ldots, h_m^n)^\top \leq \textstyle\sum_{i \in N} b_i \tag{5}$$

$$h_j^k \in \{0,1\} \quad \forall k = 1, \ldots, n,\ j = 1, \ldots, m \tag{6}$$

The equivalence of (Z) and (R) follows from the following observations: If $f = (f_1, \ldots, f_m) \in P_N \cap Z^m$ is optimal for (Z), we define $h_j^k = 1$ for $k = 1, \ldots, f_j$ and $h_j^k = 0$ for $k = f_j + 1, \ldots, n$. The resulting solution $h = (h_j^k)$ is feasible for (R). Similarly, if $h = (h_j^k)$ is an optimal solution for (R), then the vector f defined by $f_j = \sum_{k=1}^n h_j^k$ is feasible for (Z). Note that here we implicitly exploit that the cost functions are non-decreasing.

We show that the integer program (R) can be solved efficiently for box-TDI aggregation polytopes.

Lemma 5.1. *If P_N has a box-TDI description and A is a $\{0, \pm 1\}$-matrix, then (R) can be solved in time polynomial in n, m, $\operatorname{size}(\sum_i b_i)$ and the complexity of a separation oracle of P_N.*

Proof. Define $A' = [A, A, \ldots, A] \in \mathbb{Z}^{r \times mn}$ and $h = (h_j^k) \in \mathbb{Q}^{mn}$. The relaxation of the integral system (5) and (6) can then be written as the system $A'h \leq \sum_i b_i,\ 0 \leq h \leq 1$. Let $Q_N = \{h : A'h \leq \sum_i b_i,\ 0 \leq h \leq 1\}$ be the polytope described by this system.

We first show that Q_N is integral. By assumption the description of $P_N = \{f : Af \leq \sum_{i \in N} b_i\}$ is box-TDI. In particular, by applying Proposition 2.1(iv) repeatedly, we obtain that the system

$$[A, A, \ldots, A](h_1^1, \ldots, h_m^1, h_1^2, \ldots, h_m^2, \ldots, h_1^n, \ldots, h_m^n)^\top \leq \textstyle\sum_{i \in N} b_i$$

is box-TDI as well. In particular, this implies that the system $A'h \leq \sum_i b_i,\ 0 \leq h \leq 1$ is TDI because the intersection of a box-TDI system with an arbitrary box yields a TDI system. Because $\sum_i b_i$ and the restrictions on h are integral vectors, we conclude that Q_N is indeed integral.

We now show how to construct a separation oracle for Q_N from a separation oracle for P_N. For $h = (h_1^1, \ldots, h_m^1, h_1^2, \ldots, h_m^2, \ldots, h_1^n, \ldots, h_m^n) \in \mathbb{Q}^{mn}$, let the aggregated vector $f \in \mathbb{Q}^m$ be defined as $f_j = \sum_{k=1}^n h_j^k$ for $j = 1, \ldots, m$. Then $h \in Q_N$ if and only if $f \in P_N$. We now give a separation oracle for Q_N. Let $y = (y_j^k) \in \mathbb{Q}^{mn}$ be an arbitrary rational vector and let f be defined as above. We use the separation oracle of P_N to check if $f \in P_N$ or not. If $f \in P_N$, then also $h \in Q_N$ and we are done. Otherwise if $f \notin P_N$ the oracle returns a vector $a \in \mathbb{Q}^m$ such that $a^\top x < a^\top f$ for all $x \in P_N$. In particular this means that $(a^\top, a^\top, \ldots, a^\top)z < (a^\top, a^\top, \ldots, a^\top)y$ for all $z = (z_j^k) \in Q_N$. Thus, we obtain a separation oracle for Q_N.

We now conclude the proof with a running-time analysis. From the claims above, it follows that we can efficiently solve (R) with the ellipsoid method. In particular, by (Schrijver 1986, Corollary 14.1a) this can be done in time polynomial in n, m, $\operatorname{size}(\sum_i b_i)$, $\operatorname{size}(c_e)$ for $e \in E$, and the running time of the separation oracle for P_N. Frank and Tardos (1987) show that for every linear program $\max\{c^\top x : Ax \leq b\}$ with a $\{0, \pm 1\}$-matrix A, the objective function c can be replaced by an objective function c', which is polynomially bounded in m, that yields the same set of optimal solutions. The function c' can be computed in strongly polynomial time. This concludes the proof. □

We obtain the following main result from the discussion above.

Theorem 5.2 (Aggregation). *Let $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ be a polytopal congestion game with aggregation polytope P_N that satisfies the IDP and has a box-TDI description. Then a feasible load profile minimizing Rosenthal's potential can be computed in time polynomial in n, m, $\operatorname{size}(\sum_i b_i)$ and the complexity of a separation oracle of P_N.*

For symmetric polytopal congestion games we obtain the following corollary.

Corollary 5.3. *Let $\Gamma = (N, E, \mathcal{S}, (c_e)_{e\in E})$ be a symmetric polytopal congestion game given by P. If P satisfies the IDP and has a box-TDI description, then a feasible load profile minimizing Rosenthal's potential can be computed in time polynomial in n, m, $\mathrm{size}(\sum_i b_i)$, and the complexity of a separation oracle of P.*

To the best of our knowledge, there is no universal algorithm that can perform integer decomposition of an arbitrary polytope satisfying the IDP in time polynomial in n, m, $\mathrm{size}(\sum_i b_i)$ and the complexity of a separation oracle. However, under a slightly stronger integer decomposition property such a decomposition can be done as explained below. Here we focus on symmetric congestion games for clarity; but these arguments can be extended to the non-symmetric case as well (details will be given in the full version of the paper).

We say that a polytope P satisfies the *middle integral decomposition property* (McDiarmid 1983) if for $n \in \mathbb{N}$ and $w \in \mathbb{Z}^m$, the polytope $P \cap (w - (n-1)P)$ is integral. If this property is satisfied, the decomposition algorithm of Baum and Trotter (1978) can then be used to perform the integer decomposition; details are given in the proof of Theorem 5.4.

Theorem 5.4 (Aggregation/Decomposition). *Let $\Gamma = (N, E, \mathcal{S}, (c_e)_{e\in E})$ be a symmetric polytopal congestion game given by P. If P satisfies the middle integral decomposition property and has a box-TDI description, then a feasible strategy profile minimizing Rosenthal's potential can be computed in time polynomial in n, m, $\mathrm{size}(\sum_i b_i)$ and the complexity of a separation oracle of P.*

We remark that all results in this section also hold for computing a social optimum of congestion games with *weakly convex* cost functions, since this problem can be reduced to computing a global optimum of Rosenthal's potential (see (Del Pia et al. 2017)).

6 APPLICATIONS

We now give examples of polytopal tuples $(N, E, (\mathcal{S}_i)_{i\in N})$ for which the aggregation polytope has the IDP (or middle integral decomposition property), a box-TDI description and an efficient separation oracle. As a consequence, our results on the price of stability (Theorem 4.5) and the computation of Rosenthal's potential minimizer (Theorem 5.2 and Corollary 5.3) apply.

6.1 Symmetric totally unimodular congestion games

Totally unimodular congestion games (Del Pia et al. 2017) capture a wide range of combinatorial congestion games. Here the common strategy set of the players is described by a polytope $P = \{x : Ax \le b\}$ with a totally unimodular $r \times m$-matrix A and an integral vector b. In particular, such a system satisfies the IDP and is box-TDI. The integer decomposition property was shown by Baum and Trotter (1978). We argue that the system is box-TDI. The constraint matrix describing the intersection of P with $\{x : c \le x \le d\}$ for $c, d \in \mathbb{Q}^m$ is again totally unimodular (Veinott and Dantzig 1968). Any totally unimodular system is TDI (see, e.g., (Schrijver 1986, Section 22.1)), and therefore the system $Ax \le b, c \le x \le d$ is TDI. We conclude that the system $Ax \le b$ is box-TDI. If, as in (Del Pia et al. 2017), the parameter r is considered as part of the input size as well, then there is a trivial (strongly) polynomial separation oracle that simply checks all inequalities of the system $Ax \le b$. For all combinatorial applications given in (Del Pia et al. 2017), the parameter r is actually polynomially bounded in n and m, so then this assumption is justified.

6.2 Common source network congestion games

In a *common source network congestion game* we are given a directed graph $G = (V, A)$ and a source $s \in V$. The strategy set of player $i \in N$ is the set of all directed s, t_i-paths for some $t_i \in V$.

Ackermann et al. (2008) already showed that one can compute a global optimum of Rosenthal's potential function for these games. We outline how this case can be cast in our framework. The strategies of player i can be described by a polytope $P_i = \{x : Ax = b_i, 0 \le x \le 1\}$, where A is the arc-incidence matrix of the network G, and b is the vector with $(b_i)_s = 1$, $(b_i)_{t_i} = -1$ and zero otherwise.[5] The aggregation polytope is then $P_N = \{y : Ay = \sum_{i \in N} b_i, 0 \le y \le n\}$. Any feasible load profile minimizing Rosenthal's potential can be decomposed into a feasible strategy profile, using a similar argument as in (Ackermann et al. 2008). Further, the describing system of P_N is totally unimodular and thus box-TDI.

6.3 Symmetric matroid intersection congestion games

In *symmetric matroid intersection congestion games* the (symmetric) strategy set of all players is given by the common bases of two matroids $\mathcal{M}_1 = (E, \mathcal{I}_1)$ and $\mathcal{M}_2 = (E, \mathcal{I}_2)$ over a common element set E. The polytope P of the players corresponds to the common base polytope $P_{\mathcal{M}_1, \mathcal{M}_2}$ as defined in (2). The describing system of P is box-TDI (see, e.g., (Schrijver 2003, Corollary 41.12e)). Further, as noted in the preliminaries there is a separation oracle for P (and thus P_N) which runs in time polynomial in $|E|$ and the complexity of the independence oracles for $\mathcal{M}_1$ and $\mathcal{M}_2$. However, it is not precisely known for which cases of matroid intersection the integer decomposition property holds.

Example 6.1 (r-Arborescences). Let $D = (V, A)$ be a directed graph. An *r-arboresence* in D is a directed spanning tree rooted in $r \in V$. The set of all r-arboresences can be seen as the set of common bases of two matroids. The first matroid $\mathcal{M}_1$ is the graphic matroid on the undirected graph $D' = (V, A')$, where A' is the set formed by replacing every directed arc in A with its undirected version, i.e., $A' = \{\{u, v\} : (u, v) \in A\}$. The second matroid $\mathcal{M}_2$ is the partition matroid in which independent sets are given by sets of arcs for which there is at most one incoming arc at every node $v \neq r$ (we assume there are no incoming arcs at r). Thus, the common base polytope $P_{\mathcal{M}_1, \mathcal{M}_2}$ describes the arborescences of D and we let $P = P_{\mathcal{M}_1, \mathcal{M}_2}$.

We argue that there is a strongly polynomial time algorithm for computing a minimum of Rosenthal's potential. First note that the describing system of $P_{\mathcal{M}_1, \mathcal{M}_2}$ is box-TDI (see (Schrijver 2003, Corollary 41.12e)). Also, $P_{\mathcal{M}_1, \mathcal{M}_2}$ satisfies the integer decomposition property, which follows from Edmonds' Disjoint Arborescences Theorem (Edmonds 2003). By Corollary 5.3, we can compute a minimum of Rosenthal's potential in time polynomial in n, m, size$(\sum_i b_i)$ and the complexity of a separation oracle for $P_{\mathcal{M}_1, \mathcal{M}_2}$. The elements of the vector b are bounded by $|E|$, by the definition of the rank functions. Moreover, it is not hard to see that there exist independence oracles for both $\mathcal{M}_1$ and $\mathcal{M}_2$ that run in time polynomial in m. These oracles can be used for separation oracles as described in the preliminaries. It is not hard to see that if both base matroid polytopes have a polynomial time separation oracle, then the intersection of these polytopes has one too. This shows that there is an algorithm for computing an optimal feasible load profile in time polynomial in n and m. Integer decomposition can also be done in time polynomial in n and m (see, e.g., (Harks et al. 2013, Theorem 5)).

Example 6.2 (Intersection of strongly base-orderable matroids). A matroid $\mathcal{M} = (E, \mathcal{I})$ is strongly base-orderable if for every pair of bases $B_1, B_2 \in \mathcal{B}$ there exists a bijection $\tau : B_1 \to B_2$ such that for every $X \subseteq B_1$, we have $B_1 - X + \tau(X) \in \mathcal{B}$. As in the previous example, a box-TDI description follows from (Schrijver 2003, Corollary 41.12e). It is also known that the independent set polytope

[5]Technically, this polytope can also contain paths with a finite number of disjoint cycles, but these can always be removed in the end.

of the intersection of strongly base-orderable matroids has the integer decomposition property (McDiarmid 1983, Theorem 5.1).[6]

6.4 Non-symmetric matroid congestion games

In a non-symmetric matroid congestion game $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E}$, the strategy set of player i is given by the bases $\mathcal{B}_i$ of a matroid $\mathcal{M}_i = (E, \mathcal{I}_i)$ for $i \in N$.[7] The incidence vectors of the bases of $\mathcal{B}_i$ can be described by the base matroid polytope

$$P_i = \{x : x(A) \leq r_i(A), A \subseteq E,\ x(E) = r_i(E),\ x \geq 0\}$$

as introduced in the preliminaries. That is, for every player we have a polytope of the form $P_i = \{x : Ax \leq b_i, x \geq 0\}$ where b_i is the rank function r_i of the matroid $\mathcal{M}_i$. In particular, it follows that the aggregation polytope is given by

$$P_N = \{y : y(A) \leq \textstyle\sum_i r_i(A), A \subseteq E,\ y(E) = \sum_i r_i(E),\ y \geq 0\}.$$

The polytope P_N has a box-TDI description, which follows from (Schrijver 2003, Theorem 46.2).[8] The integer decomposition property is also satisfied (see, e.g., (Schrijver 2003, Corollary 46.2c)). Using similar arguments as for r-arborescences, we can thus derive a strongly polynomial time algorithm to compute a minimum of Rosenthal's potential.

We also derive a result that is of independent interest: We can give a combinatorial approach for computing the symmetric difference decomposition (which is of a specific form) of non-symmetric matroid congestion games. Our analysis also gives rise to a strongly polynomial time local search algorithm that computes a global optimum of Rosenthal's potential function. This local search algorithm can be seen as a natural generalization of best-response dynamics. The details regarding these results will be given in the full version of the paper.

7 BOTTLENECK CONGESTION GAMES.

A *bottleneck congestion game* $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ is defined similar as a congestion game, with the only difference that the objective of a player is to minimize the *maximum* (rather than the aggregated) congestion over all resources that she occupies. Formally, the cost of player $i \in N$ under strategy profile $s = (s_1, \dots, s_n)$ is given by $C_i(s) = \max_{e \in s_i} c_e(x_e)$.

Harks, Hoefer, Klimm, and Skopalik (2013) give a *dual greedy algorithm* to compute a strong equilibrium, which uses a *strategy packing oracle* as a subroutine. They give efficient packing oracles for symmetric network congestion games, non-symmetric matroid congestion games, and (a slight generalization of) r-arborescences. In particular, this leads to polynomial time algorithms for computing a strong equilibrium in these cases.

We adapt the algorithm in (Harks et al. 2013) to compute a load profile of a strong equilibrium for bottleneck polytopal congestion games satisfying the IDP and box-TDI property.

Theorem 7.1 (Aggregation). *Let $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ be a polytopal bottleneck congestion game with aggregation polytope P_N that satisfies the IDP and has a box-TDI description. Then there is an algorithm for computing a load profile of a strong equilibrium in time polynomial in n, m,* size$(\sum_i b_i)$ *and the complexity of a separation oracle of P_N.*

[6]This also implies that the common base polytope has the integer decomposition property, since the integer decomposition property is preserved if we restrict ourselves to a face of a polytope with the integer decomposition property.

[7]Our framework also captures the *independent set congestion games* studied by Del Pia et al. (2017). However, we mainly focus on non-negative cost functions here (because of the inefficiency measures) and then these games are trivial.

[8]To see this, we use the fact that the rank function is submodular and that the sum of submodular functions is again submodular. We can then apply Theorem 46.2 in (Schrijver 2003).

ALGORITHM 1: Load profile-dual greedy algorithm.

Input : Bottleneck congestion game $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$, load profile oracle $\mathfrak{O}$
Output : Load profile of strong equilibrium of Γ.
set $N' = N$, $u_e = n$ for all $e \in E$, $T = \emptyset$, $L = E$ **and** $a = \mathfrak{O}(T \cup L, (u_e)_{e \in E})$;
while $e \in L : u_e > 0 \neq \emptyset$ **do**
  choose $e' \in \text{argmax}_{e \in L: u_e > 0}\{c_e(u_e)\}$
  $u_{e'} := u_{e'} - 1$;
  if $\mathfrak{O}(T \cup L, (u_e)_{e \in E}) = NO$ **then**
    $u_{e'} := u_{e'} + 1$;
    $L = L \setminus e',\ T = T \cup \{e'\}$
  end
  $a = \mathfrak{O}(T \cup L, (u_e)_{e \in E})$
end
return $(u_e)_{e \in E}$;

We adapt the definition of the strategy packing oracle of Harks et al. (2013) to load profiles:

Load profile oracle

Input: finite set of resources $E = T \cup L$ with upper bounds $(u_e)_{e \in E}$ and n collections $\mathcal{S}_1, \ldots, \mathcal{S}_n \subseteq 2^E$ (given implicitly by a certain combinatorial property)

Output: yes, if there exists a feasible load profile f such that $f_e = u_e$ for all $e \in T$ and $f_e \leq u_e$ for all $e \in L$; no otherwise.

Our adaptation of the dual greedy algorithm in (Harks et al. 2013) is given in Algorithm 1. Although the ideas are similar to the ones in (Harks et al. 2013), our algorithm only works with load profiles; in particular, we do not have to explicitly compute decompositions of feasible load profiles in intermediate steps of the algorithm (which significantly improves the running time).

Our algorithm works roughly as follows. We start with capacities of n on every resource. In every step we pick a resource $e' \in L$ with maximum cost among all resources that are called *loose*, and check whether there is a feasible load profile if we reduce the capacity on e' by one. If this is not possible, we remove e' from L and add e' to the set T of so-called *tight* resources. Note that after the algorithm has terminated, all resources are in the set T.

Proof of Theorem 7.1. It can be shown that Algorithm 1 computes a load profile of a strong equilibrium (details will be given in the full version of the paper). It is clear that Algorithm 1 can be executed in time polynomial in n, m and the complexity of a load profile oracle. We now give an efficient load profile oracle, based on a separation oracle of P_N. From the fact that P_N has a box-TDI description, it follows that the polytope

$$\{y : Ay \leq \textstyle\sum_{i \in N} b_i\} \cap \{y_e = u_e : e \in T\} \cap \{0 \leq y_e \leq u_e : e \in L\}$$

is integral. We can then use a separation oracle for P_N to find an integral vector in this polytope in time polynomial in n, m, size$(\sum_i b_i)$ and the complexity of the separation oracle. This concludes the proof. □

Once we have obtained the feasible load profile, we can use an integer decomposition algorithm to find the corresponding strategies of the players. If the integer decomposition can be done within the same time bounds as in Theorem 7.1, we obtain a (strongly) polynomial algorithm for computing a strong equilibrium in a polytopal bottleneck congestion game. In particular, this applies to all applications mentioned in Section 6.

For matroid bottleneck congestion games, we can also derive an upper bound on the strong price of stability (SPOS). The proof of the following result essentially relies on the fact that Algorithm 1 actually also calculates a global optimum of Rosenthal's potential in the case of matroid bottleneck congestion games.

Corollary 7.2 (Strong price of stability). *Let* $\Gamma = (N, E, (\mathcal{S}_i)_{i \in N}, (c_e)_{e \in E})$ *be a matroid bottleneck congestion game with cost functions in class* $\mathcal{D}$. *Then SPOS*$(\Gamma) \leq \rho(\mathcal{D})$.

Acknowledgements. We thank Carla Groenland, Bart de Keijzer and Daniel Dadush for helpful discussions.

REFERENCES

Heiner Ackermann, Heiko Röglin, and Berthold Vöcking. 2008. On the Impact of Combinatorial Structure on Congestion Games. *J. ACM* 55, 6, Article 25 (2008), 22 pages.

Sebastian Aland, Dominic Dumrauf, Martin Gairing, Burkhard Monien, and Florian Schoppmann. 2011. Exact Price of Anarchy for Polynomial Congestion Games. *SIAM J. Comput.* 40, 5 (2011), 1211–1233.

Elliot Anshelevich, Anirban Dasgupta, Jon Kleinberg, Eva Tardos, Tom Wexler, and Tim Roughgarden. 2004. The Price of Stability for Network Design with Fair Cost Allocation. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science*. 295–304.

Baruch Awerbuch, Yossi Azar, and Amir Epstein. 2005. The Price of Routing Unsplittable Flow. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*. 57–66.

Steven P. Baum and Leslie. E. Trotter. 1978. Integer Rounding and Polyhedral Decomposition for Totally Unimodular Systems. In *Optimization and Operations Research: Proceedings of a Workshop Held at the University of Bonn, October 2–8, 1977*, Rudolf Henn, Bernhard Korte, and Werner Oettli (Eds.). 15–23.

Ioannis Caragiannis, Michele Flammini, Christos Kaklamanis, Panagiotis Kanellopoulos, and Luca Moscardelli. 2006. Tight Bounds for Selfish and Greedy Load Balancing. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming*. 311–322.

Hau Chan and Albert Xin Jiang. 2016. Congestion Games with Polytopal Strategy Spaces. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence*. 165–171.

Steve Chien and Alistair Sinclair. 2007. Convergence to Approximate Nash Equilibria in Congestion Games. In *Proceedings of the Eighteenth Annual ACM-SIAM Symposium on Discrete Algorithms*. 169–178.

George Christodoulou and Martin Gairing. 2016. Price of Stability in Polynomial Congestion Games. *ACM Transactions on Economics and Computation* 4, 2, Article 10 (2016), 17 pages.

George Christodoulou and Elias Koutsoupias. 2005. The Price of Anarchy of Finite Congestion Games. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*. 67–73.

George Christodoulou, Vahab S. Mirrokni, and Anastasios Sidiropoulos. 2012. Convergence and approximation in potential games. *Theoretical Computer Science* 438 (2012), 13–27.

José R. Correa, Andreas S. Schulz, and Nicolás E. Stier-Moses. 2004. Selfish Routing in Capacitated Networks. *Mathematics of Operations Research* 29, 4 (2004), 961–976.

Jasper de Jong, Max Klimm, and Marc Uetz. 2016. Efficiency of Equilibria in Uniform Matroid Congestion Games. In *Proceedings of the 9th International Symposium on Algorithmic Game Theory*. 105–116.

Alberto Del Pia, Michael Ferris, and Carla Michini. 2017. Totally Unimodular Congestion Games. In *Proceedings of the Twenty-Eighth Annual ACM-SIAM Symposium on Discrete Algorithms*. 577–588.

Jack Edmonds. 2003. *Combinatorial Optimization — Eureka, You Shrink! Papers Dedicated to Jack Edmonds.* Chapter Submodular Functions, Matroids, and Certain Polyhedra.

Jack Edmonds and Rick Giles. 1977. A Min-Max Relation for Submodular Functions on Graphs. In *Studies in Integer Programming*. Annals of Discrete Mathematics, Vol. 1. Elsevier, 185–204.

Eyal Even-Dar, Alex Kesselman, and Yishay Mansour. 2007. Convergence Time to Nash Equilibrium in Load Balancing. *ACM Transactions on Algorithms* 3, 3 (2007), 1–32.

Alex Fabrikant, Christos Papadimitriou, and Kunal Talwar. 2004. The Complexity of Pure Nash Equilibria. In *Proceedings of the Thirty-sixth Annual ACM Symposium on Theory of Computing*. 604–612.

Michal Feldman, Nicole Immorlica, Brendan Lucier, Tim Roughgarden, and Vasilis Syrgkanis. 2016. The Price of Anarchy in Large Games. In *Proceedings of the Forty-eighth Annual ACM Symposium on Theory of Computing*. 963–976.

Dimitris Fotakis. 2010. Congestion Games with Linearly Independent Paths: Convergence Time and Price of Anarchy. *Theory of Computing Systems* 47, 1 (2010), 113–136.

András Frank and Éva Tardos. 1987. An application of simultaneous diophantine approximation in combinatorial optimization. *Combinatorica* 7, 1 (1987), 49–65.

Martin Gairing, Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. 2008. Nash equilibria in discrete routing games with convex latency functions. *J. Comput. System Sci.* 74, 7 (2008), 1199–1225.

Tobias Harks, Martin Hoefer, Max Klimm, and Alexander Skopalik. 2013. Computing pure Nash and strong equilibria in bottleneck congestion games. *Mathematical Programming* 141, 1 (2013), 193–215.

Samuel Ieong, Robert McGrew, Eugene Nudelman, Yoav Shoham, and Qixiang Sun. 2005. Fast and Compact: A Simple Class of Congestion Games. In *Proceedings of the 20th National Conference on Artificial Intelligence*. 489–494.

Elias Koutsoupias and Christos Papadimitriou. 1999. Worst-case Equilibria. In *Proceedings of the 16th Annual Conference on Theoretical Aspects of Computer Science*. 404–413.

Thomas Lücking, Marios Mavronicolas, Burkhard Monien, and Manuel Rode. 2008. A new model for selfish routing. *Theoretical Computer Science* 406, 3 (2008), 187–206.

Colin McDiarmid. 1983. Integral decomposition in polyhedra. *Mathematical Programming* 25, 2 (1983), 183–198.

Shmuel Onn, Uriel G. Rothblum, and Yoav Tangir. 2005. Edge-Directions of Standard Polyhedra with Applications to Network Flows. *Journal of Global Optimization* 33, 1 (2005), 109–122.

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

Tim Roughgarden. 2015. Intrinsic Robustness of the Price of Anarchy. *J. ACM* 62, 5 (2015), 32.

Alexander Schrijver. 1986. *Theory of Linear and Integer Programming*. John Wiley & Sons, Inc., New York, NY, USA.

Alexander Schrijver. 2003. *Combinatorial optimization: polyhedra and efficiency. Vol. B, Matroids, trees, stable sets. Chapters 39–69*. Springer-Verlag.

Arthur F. Veinott and George B. Dantzig. 1968. *Integral Extreme Points*. Defense Technical Information Center.