

# Complexity and Parametric Computation of Equilibria in Atomic Splittable Congestion Games via Weighted Block Laplacians\*

Max Klimm<sup>†</sup>

Philipp Warode<sup>†</sup>

## Abstract

We settle the complexity of computing an equilibrium in atomic splittable congestion games with player-specific affine cost functions  $l_{e,i}(x) = a_{e,i}x + b_{e,i}$  showing that it is PPAD-complete. To prove that the problem is contained in PPAD, we develop a homotopy method that traces an equilibrium for varying flow demands of the players. A key technique is to describe the evolution of the equilibrium locally by a novel block Laplacian matrix. This leads to a path following formulation where states correspond to supports that are feasible for some demands and neighboring supports are feasible for increased or decreased flow demands. A closer investigation of the block Laplacian system allows to orient the states giving rise to unique predecessor and successor states thus putting the problem into PPAD. For the PPAD-hardness, we reduce from computing an approximate equilibrium of a bimatrix win-lose game. As a byproduct of our reduction we further show that computing a multi-class Wardrop equilibrium with class-dependent affine cost functions is PPAD-complete as well. As a byproduct of our PPAD-completeness proof, we obtain an algorithm that computes all equilibria parametrized by the players' flow demands. For player-specific costs, this computation may require several increases and decreases of the demands leading to an algorithm that runs in polynomial space but exponential time. For player-independent costs only demand increases are necessary. If the coefficients  $b_{e,i}$  are in general position, this yields an algorithm computing all equilibria as a function of the flow demand running in time polynomial in the output.

## 1 Introduction

Congestion games are a central topic in algorithmic game theory with applications in traffic [4, 13, 49, 45], telecommunication [3, 41, 43], and logistics [11]. We are given a graph  $G = (V, E)$  with a finite set of  $k$  commodities, each specified by a triplet  $(s_i, t_i, r_i)$

consisting of a source node  $s_i \in V$ , a target node  $t_i \in V$ , and a fixed demand rate  $r_i \in \mathbb{R}_{\geq 0}$ . Each edge  $e \in E$  is endowed with a flow-dependent (and possibly commodity-specific) cost function  $l_{e,i} : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$  that maps its flow  $\bar{x}_e = \sum_{i=1}^k x_e^i$  to a cost  $l_{e,i}(\bar{x}_e)$  experienced by all flow particles of commodity  $i$  using that edge.

In Wardrop's basic model [49], each commodity corresponds to a continuum of users. Each user acts selfishly in minimizing their total cost, i.e., the sum of the costs of the used edges. The corresponding equilibrium concept of a *Wardrop equilibrium* is defined as a multi-commodity flow with the additional property that each commodity only uses paths of minimum cost. Wardrop equilibria exist under mild assumptions on the cost functions [4, 48]. When the cost functions of an edge are equal for all commodities, an equilibrium can be computed in polynomial time by convex programming techniques [4]. For the case of commodity-dependent affine cost functions, Meunier and Pradeau [38] very recently showed that the problem to compute a Wardrop equilibrium lies in PPAD, but left it as an open problem whether the problem is actually PPAD-complete.

In the light of the rise of navigation systems such as Waze and TomTom and ride sharing platforms such as Lyft and Uber, and in view of the anticipated market penetration of autonomous cars, it is sensible to assume that in the near future several competing companies will control significant portions of the road traffic. Similarly, the ongoing discussion on Net Neutrality Rules for the Internet are fueled by the fact that few companies constitute and control large portions of the internet traffic, e.g., these days Netflix and YouTube each constitute about 15% of total downstream traffic worldwide [47]. In these scenarios some players may be willing to sacrifice the cost experienced by some of their traffic in order to improve the overall cost of their flow, see also the discussion in Catoni and Pallottino [8]. Atomic splittable congestion are a much more compelling model in these situations. In such a game, each commodity corresponds to a single player who controls a splittable flow of traffic in the network. The goal of the player is to minimize the total cost of their flow defined as  $C(\mathbf{x}) = \sum_{e \in E} x_e^i l_{e,i}(\bar{x}_e)$ . A multi-

\*Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) under Germany's Excellence Strategy – The Berlin Mathematics Research Center MATH+ (EXC-2046/1, project ID: 390685689).

<sup>†</sup>Humboldt-Universität zu Berlin, School of Business and Economics, Spandauer Str. 1, 10178 Berlin, Germany.

commodity flow  $\mathbf{x} = (\mathbf{x}^1, \dots, \mathbf{x}^k)$  is a Nash equilibrium if  $C^i(\mathbf{x}) \leq C^i(\tilde{\mathbf{x}}^i, \mathbf{x}^i)$  for all players  $i \in \{1, \dots, k\}$  and all  $s_i$ - $t_i$ -flows of rate  $r_i$ , where  $(\tilde{\mathbf{x}}^i, \mathbf{x}^i)$  denotes the multi-commodity flow where all players  $j \neq i$  send the flow  $\mathbf{x}^j$  and player  $i$  sends the flow  $\tilde{\mathbf{x}}^i$ .

Despite considerable progress regarding the computational complexity of equilibria in general bimatrix games [9, 14], Wardrop equilibria [4, 38], and atomic-unsplittable congestion games [1, 17], much less is known regarding the computation of equilibria in atomic splittable congestion games. For affine player-independent cost functions, Cominetti et al. [11] showed that an equilibrium can be found by computing the minimum of a convex potential function, see also Huang [29] for a combinatorial algorithm for special graph topologies. Bhaskar and Lolakapuri [7] proposed two algorithms with exponential worst-case complexity that compute  $\epsilon$ -approximate Nash equilibria in singleton games with convex costs. Harks and Timmermans [23] developed a polynomial time algorithm that computes an equilibrium in singleton games with player-specific affine cost functions.

Besides leaving the complexity of computing an equilibrium wide open, the approaches above also yield only a *single* equilibrium for a *fixed* vector of player demands. Moreover, the algorithms of Bhaskar and Lolakapuri [5] and Harks and Timmermans [23] work only for singleton games played on a network with two nodes. In actual traffic scenarios, the assumption that the players' demand vector is fully known and fixed is unrealistic since demands often fluctuate. In this paper, we will work towards understanding how the equilibria in atomic splittable games change as a function of the players' demand vectors. These insights will lead to a proof that the computation of equilibria is PPAD-complete.

**1.1 Our results and techniques** We settle the complexity of computing an equilibrium in an atomic splittable congestion game with player-specific affine costs showing that it is PPAD-complete. The complexity class PPAD ("polynomial parity argument on directed graphs") captures the complexity of search problems that can be solved by directed path-following algorithms [42]. In a nutshell, a problem is in PPAD if there is an exponential set of states  $\mathfrak{S}$  and polynomially computable functions  $\text{start}$ ,  $\text{pred}(\cdot)$ , and  $\text{succ}(\cdot)$  computing a start state, and well-defined predecessor and successor states for a given state. A state  $\mathfrak{s}$  is a source if  $\text{pred}(\mathfrak{s}) = \mathfrak{s}$  and a sink if  $\text{succ}(\mathfrak{s}) = \mathfrak{s}$ . We are guaranteed that  $\text{start}$  is a source and the goal is to compute a sink or another source. Most notably the problem to compute an equilibrium of a bimatrix game is PPAD-complete, even

in special cases [9, 14, 37]. In this paper, we show the PPAD-completeness of the following problem:

**NASH-ATOMIC-SPLITTABLE**

INPUT: game  $(G, K, l)$  graph  $G = (V, E)$ ,  
 commodities  $K = ((s_i, t_i, r_i))_{i \in [k]}$ ,  
 and cost functions  $l_{e,i}(x) = a_{e,i}x + b_{e,i}$   
 for some  $a_{e,i} \in \mathbb{R}_{>0}$ ,  $b_{e,i} \in \mathbb{R}_{\geq 0}$ .  
 OUTPUT: Nash equilibrium  $\mathbf{x}$  of  $(G, K, l)$ .

**RESULT 1.** (CF. THEOREMS 4.1 AND 5.2) *NASH-ATOMIC-SPLITTABLE is PPAD-complete.*

To show PPAD-hardness, we reduce from the problem to compute an  $n^{-\beta}$ -approximate Nash equilibrium for a  $n \times n$  win-lose bimatrix game that is PPAD-complete for all  $\beta > 0$  [10]. Our reduction requires only two players and produces an atomic splittable congestion game on a planar graph implying hardness even for this special case. This is an interesting contrast to the PLS-completeness results for computing Nash equilibria in *unsplittable* congestion games using a non-constant number of players and highly non-planar graphs [1, 17]. As a byproduct, we also obtain PPAD-completeness for a related problem settling an open question from Meunier and Pradeau [38].

**RESULT 2.** (CF. THEOREM 5.3) *Computing a multi-class Wardrop equilibrium is PPAD-complete.*

The more challenging part of the proof is to show that NASH-ATOMIC-SPLITTABLE is contained in PPAD. To this end, we develop a path following algorithm that pivots over player supports in a similar fashion to the Lemke-Howson-algorithm for bimatrix games [35]. Our algorithm follows a continuous path of Nash equilibria  $\mathbf{x}(\lambda)$  for demand rates  $\lambda r_1, \dots, \lambda r_k$  with  $\lambda \in [0, 1]$ . During the course of the algorithm,  $\lambda$  is changed in a continuous but non-monotonic matter. To describe the evolution of  $\lambda$  consider an arbitrary value  $\lambda^* \in [0, 1]$  and a corresponding equilibrium  $\mathbf{x}(\lambda^*)$  and let us fix the supports of  $(\mathbf{x}^1, \dots, \mathbf{x}^k)$ . By the Karush-Kuhn-Tucker conditions for each player, we derive that in  $\mathbf{x}(\lambda^*)$  every player only uses paths that minimize the marginal total cost. This implies that for every player  $i$ , there is a vector of vertex potentials  $\pi^i$  such that that player  $i$  uses edge  $e = (u, v)$  if and only if the difference in vertex potentials  $\pi_v^i - \pi_u^i$  is at least  $b_{e,i}$ . Going further, we can reformulate the Nash equilibrium conditions as a system of linear equations of the form  $\mathbf{y} = \mathbf{L}\pi - \mathbf{d}$  where  $\mathbf{y} = ((\mathbf{y}^1)^\top, \dots, (\mathbf{y}^k)^\top)^\top$  is a block excess vector containing the excess  $y_v^i$  for each player  $i$  and each vertex  $v$ ,  $\pi = ((\pi^1)^\top, \dots, (\pi^k)^\top)^\top$  is the block

potential vector, and  $\mathbf{d} \in \mathbb{R}^{nk}$  is an appropriate offset. The matrix  $\mathbf{L} \in \mathbb{R}^{nk \times nk}$  is a block matrix of the form

$$\mathbf{L} := \begin{pmatrix} \mathbf{L}^{11} & -\mathbf{L}^{12} & \dots & -\mathbf{L}^{1k} \\ -\mathbf{L}^{21} & \mathbf{L}^{22} & \dots & -\mathbf{L}^{2k} \\ \vdots & \vdots & \ddots & \vdots \\ -\mathbf{L}^{k1} & -\mathbf{L}^{k2} & \dots & \mathbf{L}^{kk} \end{pmatrix},$$

where the diagonal matrices  $\mathbf{L}^{ii}$  are weighted Laplacian matrices for the graph containing only the edges in the support of player  $i$ . The off-diagonal matrices  $-\mathbf{L}^{ij}$  and  $-\mathbf{L}^{ji}$  with  $i \neq j$  are the negative of weighted Laplacian matrices for the graph containing only the edges in the support of both player  $i$  and  $j$ . The weights of the matrices  $\mathbf{L}^{ij}$  depend on the coefficients  $a_{e,j}$  while the weights of the matrices  $\mathbf{L}^{ji}$  depend on the coefficients  $a_{e,i}$  so that  $\mathbf{L}$  is non-symmetric. We call  $\mathbf{L}$  the *block Laplacian matrix* of the graph. We show that  $\text{rank}(\mathbf{L}) = k(n-1)$ , except for a degenerated case that will be discussed later. This implies that there is a bijection between excess vectors and potentials (after fixing without loss of generality the potentials  $\pi_{s_i}^i = 0$  for all players  $i$ ). In particular, we obtain that for a given support the set of equilibria  $\{\mathbf{x}(\lambda) : \lambda \in [0, 1]\}$  has dimension at most one. When the dimension is exactly one, the one-dimensional linear space  $\{\mathbf{x}(\lambda) : \lambda \in [0, 1]\}$  hits the boundaries of the support at exactly two points, uniquely defining two neighboring support sets. These support sets will form *succ* and *pred* of the current support set. To obtain a unique orientation, we show that it suffices to consider the determinant of the block Laplacian after erasing the row and columns corresponding to  $s_i$  for each player  $i$ .

As a byproduct, we obtain an algorithm to solve the following problem of computing *all* Nash equilibria of an atomic splittable congestion game as a function of the players' demand rates:

PARAMETRIC-NASH-ATOMIC-SPLITTABLE

INPUT: game  $(G, K, l)$  with graph  $G = (V, E)$ , commodities  $K = ((s_i, t_i, r_i))_{i \in [k]}$ , and cost functions  $l_{e,i}(x) = a_{e,i}x + b_{e,i}$  for some  $a_{e,i} \in \mathbb{R}_{>0}$ ,  $b_{e,i} \in \mathbb{R}_{\geq 0}$ .

OUTPUT: piece-wise affine function  $\mathbf{f}: [0, 1] \rightarrow \mathbb{R}^{mk}$  such that  $\mathbf{f}(\lambda)$  is a Nash equilibrium for demand rates  $\lambda r_1, \dots, \lambda r_k \forall \lambda \in [0, 1]$ .

In order to solve PARAMETRIC-NASH-ATOMIC-SPLITTABLE, we simply run the trivial PPAD-algorithm that starts in the Nash equilibrium given by *start*. Then, it consecutively applies *succ* on the current state. During the course of the algorithm, it may be necessary to

decrease  $\lambda$ , which leads to operations that are not reflected in the output of the algorithm. In these cases, we recall the maximal lambda  $\bar{\lambda}$  such that  $\mathbf{f}$  has been outputted for  $[0, \bar{\lambda}]$  and only proceed to output Nash equilibria once  $\bar{\lambda}$  is reached again. Not counting the space of the output (which may be exponential in the input), this algorithm can clearly be implemented in polynomial space.

RESULT 3. (CF. THEOREM 7.1) PARAMETRIC-NASH-ATOMIC-SPLITTABLE can be solved in polynomial space.

For the special case of player-independent cost functions  $l_{e,i} = l_{e,j}$  for all  $e \in E$  and  $i, j \in [k]$ , we can show that in our algorithm no decrease of  $\lambda$  is necessary. We then obtain the following.

RESULT 4. (CF. THEOREM 7.2) PARAMETRIC-NASH-ATOMIC-SPLITTABLE can be solved in output-polynomial time for non-b-degenerate atomic splittable congestion games with player-independent cost functions. In particular, the runtime is in  $\mathcal{O}((kn)^{2.4} + \tau(kn)^2)$ , where  $\tau$  is the number of break-points of the piecewise affine function  $\mathbf{f}$  returned by the algorithm.

Our analysis further allows to obtain the following results regarding the multiplicity of equilibria.

RESULT 5. (CF. COROLLARY 4.1) When the constants  $b_{e,i}$  are in general position, then the games have an odd number of Nash equilibria, except for a nullset of demands.

We also exhibit an example of a game with an infinite number of equilibria (cf. Example 8.1). Due to space constraints, most proofs are deferred to the full version of the paper [33].

**1.2 Related work** Atomic splittable congestion games can be seen as a coalitional version of Wardrop equilibria [49] where a finite set of player each controls a non-negligible amount of flow [25, 36]. The existence of pure Nash equilibria follows from standard fixed point arguments [31, 44]. Games with player-specific cost functions were studied by Orda et al. [41] who showed that Nash equilibria are unique in networks of parallel edges. Richman and Shimkin [43] characterized the set of two-terminal network topologies that are necessary and sufficient for uniqueness showing that equilibria are unique if and only if the networks are nearly parallel. The latter class of networks has been introduced by Milchtaich [40] to characterize the uniqueness of multi-class Wardrop equilibria. Harks and Timmermans [24] characterized the uniqueness of equilibria in atomic splittable congestion games in terms of

the combinatorial structure of the strategy set showing that equilibria are unique when the strategy space of each player is a bidirectional flow matroid. Bhaskar et al. [5] showed that edge flows in an atomic splittable game need not be unique even when the cost functions are player-independent. For games with cost functions that are monomials of degree at most three, edge flows are known to be unique [3]. The price of anarchy of atomic splittable congestion games has been studied by Cominetti et al. [11], Harks [21], and Roughgarden and Schoppmann [46]. Catoni and Pallottino [8] provided a paradox of a non-atomic game where replacing the non-atomic players of one commodity by an atomic player with the same demand decreases the overall performance of that commodity. Hayrapetyan et al. [26] and Bhaskar et al. [6] studied this effect in more detail.

Cominetti et al. [11] showed that for games with player-independent affine cost functions an equilibrium can be computed efficiently by solving a quadratic program. For special network topologies (including series-parallel graphs), Huang [29] gave a combinatorial algorithm. For the case of parallel links, Harks and Timmermans [23] gave a polynomial algorithm that computes the equilibrium of an atomic splittable congestion game with player-specific affine costs. Bhaskar and Lolakapuri [7] provided an exponential algorithm that computes approximate equilibria in games with player-independent convex costs. They also showed that some decision problems involving equilibria are NP-complete.

Further related are *unsplittable* congestion games [45] where each commodity chooses a single path of the network. For results on the existence of equilibria, see [1, 15, 18, 19, 22, 39, 45]. Computing a pure Nash equilibrium in a congestion game with unweighted players and player-independent cost functions reduces to finding the local minimum of a potential function and is, thus, contained in the complexity class PLS, the class of all local search problems with polynomially searchable neighborhoods as defined by [30]. Fabrikant et al. [17] and Ackermann et al. [1] showed that computing a pure Nash equilibrium is in fact PLS-complete.

The problem of computing a Nash equilibrium in an atomic splittable congestion game with player-specific affine costs may also be formulated as a linear complementarity problem (LCP). Harks and Timmermans [23] showed this explicitly for the singleton case, but it is not hard to convince ourselves that such a formulation is also possible in the general case by using the vertex potentials. However, it is not clear whether the resulting LCP belongs to any of the classes for which it is known that Lemke's algorithm terminates [2, 12, 16]. In addition, Lemke's algorithm introduces an *artificial*

variable that is traced along the course of the algorithm. Our algorithm, in contrast, traces the equilibria along a meaningful increase of demand rates. In previous work [34], we developed an algorithm that computes all Wardrop equilibria parametrized by the flow demand. While relying on a similar idea, Wardrop equilibria are much more easy to handle since they can be computed in polynomial time and are essentially unique. For further homotopy methods for computing equilibria, see [20, 27, 28, 32].

## 2 Preliminaries

An *atomic splittable congestion game* is a tuple  $(G, K, l)$  where  $G = (V, E)$  is a directed, weakly connected graph with  $n$  vertices  $V$  and  $m$  edges  $E$ , the family  $K = ((s_1, t_1, r_1), \dots, (s_k, t_k, r_k))$  contains  $k$  triples each of which consisting of a source node  $s_i \in V$ , a sink node  $t_i \in V$ , and a demand rate  $r_i \in \mathbb{R}_{\geq 0}$  for each of the  $k$  players, and  $l$  is a family of player-specific, strictly increasing affine linear cost functions  $l = (l_{e,i})_{e \in E, i \in [k]}$  with  $l_{e,i}(x) = a_{e,i}x + b_{e,i}$  for some  $a_{e,i} \in \mathbb{R}_{> 0}$  and  $b_{e,i} \in \mathbb{R}_{\geq 0}$ .

A feasible strategy for every player  $i \in [k]$  is to route their demand  $r_i$  between their terminal vertices  $s_i$  and  $t_i$ . Thus, a strategy for player  $i$  is a  $s_i$ - $t_i$ -flow of rate  $r_i$ , i.e., a non-negative vector  $\mathbf{x}^i = (x_e^i)_{e \in E}$  in  $\mathbb{R}_{\geq 0}^m$  satisfying the flow conservation constraints

$$\sum_{e \in \delta^+(v)} x_e^i - \sum_{e \in \delta^-(v)} x_e^i = \begin{cases} r_i & \text{if } v = s_i, \\ -r_i & \text{if } v = t_i, \\ 0 & \text{otherwise} \end{cases}$$

for every vertex  $v \in V$ . A strategy profile is a vector  $\mathbf{x} = ((\mathbf{x}^1)^\top, \dots, (\mathbf{x}^k)^\top)^\top \in \mathbb{R}_{\geq 0}^{km}$  containing the flow vectors of all players. We use the notation  $(\tilde{\mathbf{x}}^i, \mathbf{x}^{-i})$  for the strategy profile where player  $i$  uses the flow  $\tilde{\mathbf{x}}^i$  and all other players use their flow as in the strategy profile  $\mathbf{x}$ . The cost  $l_{e,i}$  experienced by the flow of the player  $i$  on some edge  $e$  depends on the *total flow*  $\bar{\mathbf{x}} = (\bar{x}_e)_{e \in E}$  where  $\bar{x}_e = \sum_{i=1}^k x_e^i$ . Every player wants to minimize the *total cost*  $C_i(\mathbf{x})$  experienced by the flow sent by this player, i.e.,  $C_i(\mathbf{x}) = \sum_{e \in E} x_e^i l_{e,i}(\bar{x}_e)$ . We say that  $\mathbf{x}$  is a *Nash equilibrium* if for every player  $i \in [k]$  there is no profitable deviation from  $\mathbf{x}$ , i.e.,  $C_i(\mathbf{x}) \leq C_i(\tilde{\mathbf{x}}^i, \mathbf{x}^{-i})$  for all  $s_i$ - $t_i$ -flows  $\tilde{\mathbf{x}}^i$  of rate  $r_i$ . The *marginal total cost* of player  $i$  on edge  $e$  given the flow  $\mathbf{x}$  is given by  $\mu_e^i(\mathbf{x}) := \frac{\partial}{\partial x_e^i} x_e^i l_{e,i}(\bar{x}_e) = a_{e,i} \bar{x}_e + b_{e,i} + a_{e,i} x_e^i$ . We obtain the following characterization of Nash equilibria, see, e.g., Bhaskar et al. [5] for a reference.

**LEMMA 2.1.** *The strategy profile  $\mathbf{x}$  is a Nash equilibrium flow if and only if, for every player  $i$ ,  $\mathbf{x}^i$  is a  $s_i$ - $t_i$ -flow and  $\sum_{e \in P} \mu_e^i(\mathbf{x}) \leq \sum_{e \in Q} \mu_e^i(\mathbf{x})$  for all  $s_i$ - $t_i$ -paths  $P, Q$  with  $x_e^i > 0$  for all  $e \in P$ .*

Lemma 2.1 states that  $\mathbf{x}$  is a Nash equilibrium if and only if all path used by player  $i$  are also shortest path for that player with respect to the marginal costs. This enables us to give another characterization based on (shortest path) potentials.

LEMMA 2.2. *The flow  $\mathbf{x}$  is a Nash equilibrium if and only if for all  $i \in [k]$  there is a potential vector  $\boldsymbol{\pi}^i = (\pi_v^i)_{v \in V}$  with  $\pi_w^i - \pi_v^i = \mu_e^i(\mathbf{x})$  if  $x_e^i > 0$ , and  $\pi_w^i - \pi_v^i \leq \mu_e^i(\mathbf{x})$  if  $x_e^i = 0$  for all  $e = (v, w) \in E$ .*

We denote the block flow vector with  $\mathbf{x} = ((\mathbf{x}^1)^\top, \dots, (\mathbf{x}^k)^\top)^\top$  and the block potential vector with  $\boldsymbol{\pi} = ((\boldsymbol{\pi}^1)^\top, \dots, (\boldsymbol{\pi}^k)^\top)^\top$ . For every edge  $e \in E$ , let  $S_e \subseteq [k]$  be some subset of the players. Then we call the family  $\mathcal{S} := (S_e)_{e \in E}$  of these sets a *support*. We say an edge  $e$  is *active* for player  $i$  if  $i \in S_e$ , and  $e$  is *inactive* for player  $i$  otherwise. We say a Nash equilibrium  $\mathbf{x}$  has the support  $\mathcal{S}$  if for every edge  $e = (v, w) \in E$ ,  $i \in S_e$  implies that  $\pi_w^i - \pi_v^i = \mu_e^i(\mathbf{x})$ . Thus,  $\mathcal{S}$  is a support of the Nash equilibrium  $\mathbf{x}$  if for every active edge the inequality  $\pi_w^i - \pi_v^i \leq \mu_e^i(\mathbf{x})$  is satisfied with equality. In general, there are multiple supports for the same equilibrium, since for every player  $i$ , any edge with  $x_e^i = 0$  and  $\pi_w^i - \pi_v^i = \mu_e^i(\mathbf{x})$  can be considered to be an active or an inactive edge. Given a support of a Nash equilibrium, the equilibrium flow can be obtained by solving the system of equations  $\pi_w^i - \pi_v^i = \mu_e^i(\mathbf{x})$ ,  $i \in S_e$ ,  $e = (v, w) \in E$  with the unknowns  $\boldsymbol{\pi}$  and  $\mathbf{x}$ . For affine cost functions this reduces to a system of linear equations that can be solved explicitly. This will be the key element of our analysis and we discuss this in further detail in the next section.

### 3 Equilibrium structure

In this section we develop a characterization of equilibria in atomic splittable congestion games for different demand rates parametrized by a factor  $\lambda$ . To this end, we use the conditions of Lemma 2.2 to characterize all flows  $\mathbf{x}$  and potentials  $\boldsymbol{\pi}$  that describe Nash equilibria. Given a fixed support  $\mathcal{S}$ , we call a flow  $\mathbf{x}$  *induced by a potential  $\boldsymbol{\pi}$  wrt. the support  $\mathcal{S}$*  if for every edge  $e \in E$  we have  $\pi_w^i - \pi_v^i = \mu_e^i(\mathbf{x})$  for all  $i \in S_e$  and  $x_e^i = 0$  for all  $i \notin S_e$ . As we will prove in Lemma 3.1, for a fixed support  $\mathcal{S}$ , the flow  $\mathbf{x}$  induced by a potential  $\boldsymbol{\pi}$  is unique. That is, for every support  $\mathcal{S}$ , there is a well-defined function  $\boldsymbol{\nu}_{\mathcal{S}} : \mathbb{R}^{nk} \rightarrow \mathbb{R}^{mk}$  such that  $\mathbf{x} = \boldsymbol{\nu}_{\mathcal{S}}(\boldsymbol{\pi})$  is a flow satisfying

$$(3.1) \quad \begin{aligned} \mu_e^i(\mathbf{x}) &= \pi_w^i - \pi_v^i & \text{if } i \in S_e, \\ x_e^i &= 0 & \text{if } i \notin S_e \end{aligned}$$

for all edges  $e = (v, w)$ . Note that  $\mathbf{x}$  need not be a flow satisfying any given demands and may even be negative.

In the following, we will work towards finding those potential vectors  $\boldsymbol{\pi}$  for which the induced flow  $\boldsymbol{\nu}_{\mathcal{S}}(\boldsymbol{\pi})$  are feasible for some demand vector  $\lambda \mathbf{r}$  and some support  $\mathcal{S}$ . To this end, we call a potential vector  $\boldsymbol{\pi} \in \mathbb{R}^{nk}$  a  $\lambda$ -*potential* for the support  $\mathcal{S}$  if  $\boldsymbol{\nu}_{\mathcal{S}}(\boldsymbol{\pi})$  is feasible for the demand  $\lambda \mathbf{r}$  and satisfies the conditions of Lemma 2.2.

The remainder of the section is structured as follows. In Section 3.1 we show that the induced flow functions  $\boldsymbol{\nu}_{\mathcal{S}}$  are well-defined and express necessary and sufficient conditions on  $\lambda$ -potentials in terms of a block Laplacian matrix of the underlying graph. In Section 3.2, we use these characterizations in order to show that for a fixed support  $\mathcal{S}$ , the space of  $\lambda$ -potentials for some  $\lambda \in [0, 1]$  is a zero or one dimensional subset of  $\mathbb{R}^{nk}$  under a degeneracy assumption. Finally, in Section 3.3, we show that for a given support that admits a Nash equilibrium there are two well-defined *neighboring supports* that admitting Nash equilibria for slightly higher or slightly lower demand rates. The latter will be used in Section 4 in order to define the functions `pred` and `succ` that put the equilibrium computation problem into PPAD. This also yields an algorithm for parametric equilibrium computation as a byproduct that we will study in more detail in Section 7.

**3.1 Weighted block Laplacians** In this subsection, we assume that we are given a fixed support  $\mathcal{S}$ . For ease of exposition, we omit the dependence on  $\mathcal{S}$  for  $\boldsymbol{\nu} = \boldsymbol{\nu}_{\mathcal{S}}$  and related notations in this chapter. In order to find a closed form representation of the solution  $\mathbf{x}$  to the system (3.1), we need to introduce some notation. For any two players  $i, j \in [k]$  and any edge  $e \in E$  we define the coefficient  $\omega_e^{ij}$  as follows:  $\omega_e^{ij} := 1$  if  $\{i, j\} \subseteq S_e$  and  $\omega_e^{ij} := 0$  otherwise, i.e., the coefficient  $\omega_e^{ij}$  indicates whether the players  $i$  and  $j$  both use the edge  $e$ . We also write  $\omega_e^i$  as a shorthand for  $\omega_e^{ii}$ . For further reference, we define the diagonal matrices  $\boldsymbol{\Omega}^i := \text{diag}(\omega_{e_1}^i, \dots, \omega_{e_m}^i)$  and the block-diagonal matrix  $\boldsymbol{\Omega}$  that contains all matrices  $\boldsymbol{\Omega}^i$  as block-diagonal elements. For any edge  $e \in E$ , let  $\kappa_e := |S_e|$  be the number of players using edge  $e$ . We denote by  $\bar{\mathbf{K}} := \text{diag}(\frac{1}{\kappa_{e_1}+1}, \dots, \frac{1}{\kappa_{e_m}+1}) \in \mathbb{R}^{m \times m}$  the diagonal matrix containing all  $\kappa$  values and by  $\mathbf{K} \in \mathbb{R}^{mk \times mk}$  the block matrix that contains  $k \times k$ -times the block  $\bar{\mathbf{K}}$ . We further define diagonal matrix  $\tilde{\mathbf{C}}^{ij} \in \mathbb{R}^{m \times m}$  containing weights for every edge by

$$\tilde{\mathbf{C}}^{ii} := \text{diag}\left(\frac{\kappa_{e_1}+1-\omega_{e_1}^i}{(\kappa_{e_1}+1)a_{e_1,i}}\omega_{e_1}^i, \dots, \frac{\kappa_{e_m}+1-\omega_{e_m}^i}{(\kappa_{e_m}+1)a_{e_m,i}}\omega_{e_m}^i\right)$$

for all  $i \in [k]$  and

$$\tilde{\mathbf{C}}^{ij} := \text{diag}\left(\frac{1}{(\kappa_{e_1}+1)a_{e_1,j}}\omega_{e_1}^j, \dots, \frac{1}{(\kappa_{e_m}+1)a_{e_m,j}}\omega_{e_m}^j\right)$$

for all  $i, j \in [k]$  with  $i \neq j$ . We define the  $mk \times mk$  block-matrix  $\tilde{\mathbf{C}}$  with the blocks  $\tilde{\mathbf{C}}^{ii}$  on the block-diagonal and

the blocks  $-\mathbf{C}^{ij}$  as off-diagonal blocks, i.e.,

$$\tilde{\mathbf{C}} = \begin{pmatrix} \tilde{\mathbf{C}}^{11} & \cdots & -\tilde{\mathbf{C}}^{1k} \\ \vdots & \ddots & \vdots \\ -\tilde{\mathbf{C}}^{k1} & \cdots & \tilde{\mathbf{C}}^{kk} \end{pmatrix}.$$

Finally, we define the matrix  $\mathbf{C} := \Omega \tilde{\mathbf{C}}$ . In particular,  $\mathbf{C}$  has the same block structure as  $\tilde{\mathbf{C}}$ , the diagonal blocks have the form  $\mathbf{C}^{ii} = \tilde{\mathbf{C}}^{ii}$  and the off diagonal blocks have the form

$$\mathbf{C}^{ij} := \text{diag} \left( \frac{1}{(\kappa_{e_1} + 1)a_{e_1, j}} \omega_{e_1}^{ij}, \dots, \frac{1}{(\kappa_{e_m} + 1)a_{e_m, j}} \omega_{e_m}^{ij} \right).$$

The matrices  $\Omega^i$  encode the activity status of all edges for every player  $i$ , hence,  $\Omega$  encodes the support. The matrices  $\bar{\mathbf{K}}$  and  $\mathbf{K}$  contain normalization factors depending on the number of players. At last, the matrices  $\tilde{\mathbf{C}}$  and  $\mathbf{C}$  encode the linear relationship between potentials and flows in the conditions of Lemma 2.2—we will prove this in Lemma 3.1 and Theorem 3.1 below. Overall, the above definitions imply

$$\tilde{\mathbf{C}} = (\mathbf{I}_{km} - \mathbf{K}\Omega)\mathbf{A}^{-1} \text{ and } \mathbf{C} = \Omega(\mathbf{I}_{km} - \mathbf{K}\Omega)\mathbf{A}^{-1},$$

where  $\mathbf{I}_{km}$  is the  $km \times km$  identity matrix and  $\mathbf{A}^{-1} = \text{diag} \left( \frac{1}{a_{e_1, 1}}, \dots, \frac{1}{a_{e_m, k}} \right)$  is the diagonal matrix containing all reciprocals of the slopes of all cost functions.

We denote by  $\Gamma$  the vertex-edge-incidence matrix of the graph  $G$  and let  $\mathbf{G} \in \mathbb{R}^{nk \times mk}$  be the block-diagonal matrix containing  $k$  copies of  $\Gamma$  on the diagonal. Then, in particular, the vector  $\mathbf{G}^\top \boldsymbol{\pi}$  contains all potential differences  $\pi_w^i - \pi_v^i$  for all edges  $e = (v, w)$  and players  $i$ , and the vector  $\mathbf{G}\mathbf{x}$  contains the excess flow of every player at every vertex. Using all these definitions, we obtain the following representation of the induced flow function  $\boldsymbol{\nu}$ .

LEMMA 3.1. *Given a fixed support  $\mathcal{S}$  and a potential vector  $\boldsymbol{\pi} \in \mathbb{R}^{nk}$  the system (3.1) has the unique solution*

$$\mathbf{x} = \mathbf{C}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b})$$

where  $\mathbf{b} = (b_{e,i})_{e \in E, i \in [k]}$  is the vector of all offsets of the cost functions. Thus, the induced flow function  $\boldsymbol{\nu}$  is well-defined. Furthermore, the total flow wrt.  $\mathbf{x}$  on edge  $e$  can be computed as  $\bar{x}_e = \mathbf{u}_{e,i}^\top \mathbf{K}\Omega\mathbf{A}^{-1}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b})$  for all  $i \in [k]$ , where  $\mathbf{u}_{e,i}$  is the unit vector corresponding to player  $i$  and edge  $e$ .

Lemma 3.1 allows to compute for every potential vector  $\boldsymbol{\pi}$  the induced flow  $\mathbf{x} = \boldsymbol{\nu}(\boldsymbol{\pi}) = \mathbf{C}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b})$ . Then the vector  $\mathbf{y} := \mathbf{G}\mathbf{x}$  contains the excess flows for every player at every vertex, i.e., the values  $y_v^i = \sum_{e \in \delta^+(v)} x_e^i - \sum_{e \in \delta^-(v)} x_e^i$ .

DEFINITION 3.1. *For every support, we refer to the matrix  $\mathbf{L} := \mathbf{G}\mathbf{C}\mathbf{G}^\top$  as the block Laplacian matrix of the support  $\mathcal{S}$ .*

By the definition of the matrices  $\mathbf{C}$  and  $\mathbf{G}$ , we observe that  $\mathbf{L}$  has the form

$$\mathbf{L} = \begin{pmatrix} \mathbf{L}^{11} & -\mathbf{L}^{12} & \cdots & -\mathbf{L}^{1k} \\ -\mathbf{L}^{21} & \mathbf{L}^{22} & \cdots & -\mathbf{L}^{2k} \\ \vdots & \vdots & \ddots & \vdots \\ -\mathbf{L}^{k1} & -\mathbf{L}^{k2} & \cdots & \mathbf{L}^{kk} \end{pmatrix}$$

where every block  $\mathbf{L}^{ij} = \Gamma\mathbf{C}^{ij}\Gamma^\top$  is a weighted Laplacian matrix of  $G$  with weights  $\mathbf{C}^{ij}$  justifying the name *block Laplacian matrix*. By definition, the excess vector of the flow  $\mathbf{x}$  induced by some potential  $\boldsymbol{\pi}$  wrt. the support  $\mathcal{S}$  can be expressed as

$$(3.2) \quad \mathbf{y} = \mathbf{L}\boldsymbol{\pi} - \mathbf{d},$$

where  $\mathbf{d} := \mathbf{G}\mathbf{C}\mathbf{b}$ . Overall, we can compute for every potential vector  $\boldsymbol{\pi} \in \mathbb{R}^{nk}$  an induced flow vector  $\mathbf{x} = \boldsymbol{\nu}(\boldsymbol{\pi})$  such that  $\boldsymbol{\pi}$  and  $\mathbf{x}$  satisfy (3.1). To ensure that, furthermore, the induced flow  $\mathbf{x} = \boldsymbol{\nu}(\boldsymbol{\pi})$  is a Nash equilibrium, we additionally need that  $\mathbf{x}$  is non-negative and satisfies the flow conservation as well as the inequality conditions of Lemma 2.2. To this end, we define the vector  $\Delta\mathbf{y} := ((\Delta\mathbf{y}^1)^\top, \dots, (\Delta\mathbf{y}^k)^\top)^\top$  with  $\Delta y_v^i = -r_i$  if  $v = s_i$ ,  $\Delta y_v^i = r_i$  if  $v = t_i$  and  $\Delta y_v^i = 0$  otherwise. We refer to  $\Delta\mathbf{y}$  as the *excess direction*.

We introduce a new coefficient  $\sigma_e^i := 1$  if  $i \in S_e$  and  $\sigma_e^i = -1$  if  $i \notin S_e$ . We then define the diagonal matrices  $\boldsymbol{\Sigma}^i := \text{diag}(\sigma_{e_1}^i, \dots, \sigma_{e_m}^i)$  for every player  $i$  and the block diagonal matrix  $\boldsymbol{\Sigma}$  containing the matrices  $\boldsymbol{\Sigma}^i$  as block entries, similar to the matrices  $\Omega^i$  and  $\Omega$ . Then we obtain the following result characterizing  $\lambda$ -potentials, i.e., potentials that induce a Nash equilibrium flow for demands  $\lambda\mathbf{r}$  for some  $\lambda \in [0, 1]$ .

THEOREM 3.1. *The potential vector  $\boldsymbol{\pi} \in \mathbb{R}^{nk}$  is a  $\lambda$ -potential with  $\lambda \in [0, 1]$  if and only if*

$$(3.3a) \quad \mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda\Delta\mathbf{y}$$

$$(3.3b) \quad \mathbf{W}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b}) \geq \mathbf{0},$$

where  $\mathbf{W} := \boldsymbol{\Sigma}\tilde{\mathbf{C}}$ .

Informally, Theorem 3.1 can be read as follows. If  $\boldsymbol{\pi}$  is a solution to (3.3) then Equation (3.3a) ensures that the induced flow  $\boldsymbol{\nu}(\boldsymbol{\pi})$  is a flow satisfying the demands  $\lambda\mathbf{r}$ . The induced flow  $\boldsymbol{\nu}(\boldsymbol{\pi})$  and  $\boldsymbol{\pi}$  satisfy system (3.1) by definition. Finally, Equation (3.3b) ensures that the flow is non-negative on active edges and that the potential inequalities from Lemma 2.2 are satisfied for inactive edges.

**3.2 Potential space and line segments** Theorem 3.1 gives a characterization of  $\lambda$ -potentials, i.e., of potentials that induce to Nash equilibrium flows. The aim of this section is to characterize the set of *all*  $\lambda$ -potentials of some fixed support, if such potentials exist. Before we proceed, we have to address some natural sources of ambiguity in our setting. It is obvious from Lemma 2.2 as well as from Lemma 3.1 that the induced flow does only depend on potential differences on edges rather than absolute potential values. In fact, the vertex-edge-incidence matrix  $\mathbf{\Gamma}$  has rank  $n - 1$  (assuming that  $G$  is weakly connected) and its left nullspace is spanned by the all-ones-vector  $\mathbf{1} \in \mathbb{R}^n$ . Hence, the left nullspace of the matrix  $\mathbf{G}$  is

$$\mathcal{N} := \{(\alpha_1 \mathbf{1}^\top, \dots, \alpha_k \mathbf{1}^\top)^\top \in \mathbb{R}^{nk} \mid \alpha_1, \dots, \alpha_k \in \mathbb{R}\}.$$

To avoid ambiguity in the potential space, we restrict ourselves to potentials that are normalized to zero at the source vertices of the players. That is, we consider the vector space

$$\mathcal{P} := \{\boldsymbol{\pi} \in \mathbb{R}^{nk} \mid \pi_{s_i}^i = 0 \text{ for all players } i \in [k]\}$$

that we refer to as the *potential space*. The potential space is isomorphic to  $\mathbb{R}^{nk}/\mathcal{N}$  and has the advantage that the potentials  $\pi_v^i$  can be interpreted as the marginal costs of any  $s_i$ - $v$ -path used by player  $i$ . Another technical problem with ambiguity may arise if not all vertices of  $G$  are connected by active edges for some player. In this case, the potential of that player can be shifted by a constant on each of the connected components without changing the flow. To avoid this problem, we say a support  $\mathcal{S}$  is a *total support* if, for every player  $i$ , the subgraph of  $G$  containing only edges  $e$  with  $i \in S_e$  is connected. It is not hard to show, that for every Nash equilibrium  $\mathbf{x}$ , there is a corresponding total support  $\mathcal{S}$  such that it is no restriction to consider total supports only.

Denote by  $\mathcal{P}_{\mathcal{S}}$  the set of all  $\lambda$ -potentials for the support  $\mathcal{S}$ . With the characterization from Theorem 3.1 we obtain

$$\mathcal{P}_{\mathcal{S}} = \{\boldsymbol{\pi} \in \mathcal{P} \mid \exists \lambda \in [0, 1] : \mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda \Delta \mathbf{y}, \\ \mathbf{W}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b}) \geq \mathbf{0}\}.$$

We say, a support is *feasible* if it is total and  $\mathcal{P}_{\mathcal{S}} \neq \emptyset$ , i.e., if there is some  $\lambda$  such that there is a Nash equilibrium for the demand vector  $\lambda \mathbf{r}$ . Our next step is to find solutions to the equation  $\mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda \Delta \mathbf{y}$ . First, we observe the following properties of the matrix  $\mathbf{C}$ .

LEMMA 3.2. *If  $\mathcal{S}$  is total, then  $\ker(\mathbf{C}\mathbf{G}^\top) = \mathcal{N}$ .*

Lemma 3.2 shows that the induced flow function  $\boldsymbol{\nu}(\boldsymbol{\pi}) = \mathbf{C}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b})$  is injective on  $\mathcal{P}$ . Furthermore, we see that

the nullspace of the block Laplacian matrix  $\mathbf{L}$  is a subset of  $\mathcal{N}$ . If  $\ker(\mathbf{L}) = \mathcal{N}$ , we can solve  $\mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda \Delta \mathbf{y}$  uniquely for  $\boldsymbol{\pi}$  in  $\mathcal{P}$ . In this case, there is a unique equilibrium potential (and flow) with the support  $\mathcal{S}$  for every excess vector, and thus for every demand rate  $\lambda \mathbf{r}$ . If, on the other hand,  $\dim(\ker(\mathbf{L})) > \dim(\mathcal{N})$ , there are actually infinitely many equilibrium potentials, and thus also flows, for the same demand. If this situation occurs for some support  $\mathcal{S}$ , we call  $\mathcal{S}$  *a-degenerate*.

DEFINITION 3.2. *A feasible support  $\mathcal{S}$  is called a-degenerate if  $\dim(\ker(\mathbf{L})) > \dim(\mathcal{N})$ . We say a game is a-degenerate, if there is a feasible support  $\mathcal{S}$  that is a-degenerate.*

See Example 8.1 in Section 8.1 for a concrete example with a degenerate support  $\mathcal{S}$  and a continuum of equilibria for fixed demands.

For a matrix  $\mathbf{A}$ , we denote by  $\mathbf{A}^+$  a generalized inverse of  $\mathbf{A}$ , that is a matrix  $\mathbf{A}^+$  satisfying  $\mathbf{A}\mathbf{A}^+\mathbf{A} = \mathbf{A}$ . If the linear system  $\mathbf{A}\mathbf{x} = \mathbf{b}$  is consistent, then  $\mathbf{x} = \mathbf{A}^+\mathbf{b}$  is a solution to the system for every generalized inverse  $\mathbf{A}^+$ . Using generalized inverses, we can express a solution to (3.3a) as  $\boldsymbol{\pi} = \mathbf{L}^+(\lambda \Delta \mathbf{y} + \mathbf{d})$ . For non-a-degenerate supports, there is a unique solution in  $\mathcal{P}$  to this system. Thus there is a unique generalized inverse that we denote by  $\mathbf{L}^*$  such that  $\boldsymbol{\pi} = \mathbf{L}^*(\lambda \Delta \mathbf{y} + \mathbf{d})$  is the unique solution to (3.3a).

COROLLARY 3.1. *Every non-a-degenerate game has a finite number of Nash equilibria for every demand  $\mathbf{r}$ .*

*Proof.* If the game is non-a-degenerate, then for a fixed  $\lambda = 1$  there is at most one  $\boldsymbol{\pi} \in \mathcal{P}$  satisfying  $\mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda \Delta \mathbf{y}$ . Hence, there is at most one equilibrium flow for every support  $\mathcal{S}$ . Since there are only finitely many supports, the claim follows.  $\square$


We assume that all supports  $\mathcal{S}$  are non-a-degenerate and discuss how to deal with a-degeneracy in Section 8.1. Let us denote with  $\Delta \boldsymbol{\pi} := \mathbf{L}^* \Delta \mathbf{y}$  the (*unoriented*) *potential direction of the support  $\mathcal{S}$*  and by  $\bar{\mathbf{d}} := \mathbf{L}^* \mathbf{d}$  the *offset potential of support  $\mathcal{S}$* . Then the unique solution to (3.3a) can be expressed as  $\boldsymbol{\pi} = \lambda \Delta \boldsymbol{\pi} + \bar{\mathbf{d}}$  and we see that all  $\lambda$ -potentials for a given support lie on a line in the potential space. The following lemma gives an alternative representation of  $\mathcal{P}_{\mathcal{S}}$ .

LEMMA 3.3. *For every feasible support  $\mathcal{S}$ , there are numbers  $0 \leq \lambda_{\mathcal{S}}^{\min} \leq \lambda_{\mathcal{S}}^{\max} \leq 1$  such that*

$$\mathcal{P}_{\mathcal{S}} = \{\lambda \Delta \boldsymbol{\pi} + \bar{\mathbf{d}} \mid \lambda_{\mathcal{S}}^{\min} \leq \lambda \leq \lambda_{\mathcal{S}}^{\max}\}.$$

*Proof.* Using that  $\lambda \Delta \boldsymbol{\pi} + \bar{\mathbf{d}}$  is the unique solution in  $\mathcal{P}$  of  $\mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda \Delta \mathbf{y}$ , we get

$$\mathcal{P}_{\mathcal{S}} = \{\boldsymbol{\pi} \in \mathcal{P} \mid \exists \lambda \in [0, 1] : \mathbf{W}(\mathbf{G}^\top(\lambda \Delta \boldsymbol{\pi} + \bar{\mathbf{d}}) - \mathbf{b}) \geq \mathbf{0}\}.$$



**Figure 1:** The set of all  $\lambda$ -potentials  $\mathcal{P}_S$  is a segment of the line with direction  $\Delta\pi$  with extreme points  $\pi_S^{\min}$  and  $\pi_S^{\max}$ . All potentials on this line segment induce Nash equilibria for demands  $\lambda r$  with  $\lambda_S^{\min} \leq \lambda \leq \lambda_S^{\max}$ .

Denote by  $\mathbf{u}_{e,i}$  the unit vector corresponding to the edge-player-pair  $(e,i)$ . Then  $\mathbf{w}_{e,i} := \mathbf{W}^\top \mathbf{u}_{e,i}$  is the  $(e,i)$ -th row of the matrix  $\mathbf{W}$ . Then, we express the constraint in  $\mathcal{P}_S$  as

$$\lambda \mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi \geq \mathbf{w}_{e,i}^\top (\mathbf{b} - \mathbf{G}^\top \bar{\mathbf{d}})$$

for every  $e \in E$  and  $i \in [k]$ . Solving for  $\lambda$ , we obtain that the value

$$\lambda_S^{\min} := \max \left\{ \frac{\mathbf{w}_{e,i}^\top (\mathbf{b} - \mathbf{G}^\top \bar{\mathbf{d}})}{\mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi} \mid \mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi < 0 \right\} \cup \{0\}$$

$$\lambda_S^{\max} := \min \left\{ \frac{\mathbf{w}_{e,i}^\top (\mathbf{b} - \mathbf{G}^\top \bar{\mathbf{d}})}{\mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi} \mid \mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi > 0 \right\} \cup \{1\}$$

are the lower and upper bounds on  $\lambda$ . (Note that, additionally,  $\mathbf{w}_{e,i}^\top (\mathbf{b} - \mathbf{G}^\top \bar{\mathbf{d}}) \leq 0$  must be satisfied for all  $(e,i)$  with  $\mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi = 0$ . Since we assumed that  $\mathcal{S}$  is feasible, this is true.)  $\square$

Lemma 3.3 shows that the set of all  $\lambda$ -potentials of a given support  $\mathcal{S}$  is a line segment, potentially reduced to a single point when  $\lambda_S^{\min} = \lambda_S^{\max}$ . In particular, all  $\lambda$ -potentials of the support  $\mathcal{S}$  are convex combinations of the two potentials  $\pi_S^{\min} := \lambda_S^{\min} \Delta\pi + \bar{\mathbf{d}}$  and  $\pi_S^{\max} := \lambda_S^{\max} \Delta\pi + \bar{\mathbf{d}}$ , c.f. Figure 1. We denote by  $\partial\mathcal{P}_S = \{\pi_S^{\min}, \pi_S^{\max}\}$  the set of *boundary potentials of the support  $\mathcal{S}$* .

**3.3 Neighboring supports** So far, we considered a fixed support  $\mathcal{S}$  and characterized the  $\lambda$ -potentials for this demand. In this subsection, we study how different supports are connected in order to find a way to find feasible supports for all  $\lambda \in [0, 1]$ . We begin by introducing the notion of *neighboring supports*.

**DEFINITION 3.3.** We say, two supports  $\mathcal{S}, \mathcal{S}'$  are  $(e,i)$ -neighboring if

- (i)  $S_{\tilde{e}} = S'_{\tilde{e}}$  for all  $\tilde{e} \neq e$  and
- (ii)  $S'_e \setminus S_e = \{i\}$  or  $S_e \setminus S'_e = \{i\}$ .

We denote by  $N(\mathcal{S}, e, i)$  the unique  $(e,i)$ -neighboring support  $\mathcal{S}'$ . By definition, two neighboring supports only differ by the activity status of one edge  $e$  for one particular player  $i$ . Hence, the matrices  $\mathbf{C}_S$  and  $\mathbf{C}_{S'}$  as well as the block Laplacians  $\mathbf{L}_S$  and  $\mathbf{L}_{S'}$  are closely related. In fact, the following theorem shows that the respective matrices can be obtained from each other by a rank-1-update. To this end, we introduce for every support  $\mathcal{S}$  and any pair of edge  $e \in E$  and player  $i \in [k]$  two new vectors  $\mathbf{q}_{S,e,i} := \mathbf{W}_S^\top \mathbf{G}^\top \mathbf{u}_{e,i}$  and  $\bar{\mathbf{q}}_{S,e,i} := (\mathbf{I}_{km} - \mathbf{K}_S \Omega_S) \mathbf{G}^\top \mathbf{u}_{e,i}$ . Geometrically, the vector  $\mathbf{q}_{S,e,i}$  is the normal vector of the hyperplane corresponding to the  $(e,i)$ -th inequality of  $\mathbf{W}(\mathbf{G}^\top \pi - \mathbf{b}) \geq 0$  bounding the polytope  $\mathcal{P}_S$ .

**THEOREM 3.2.** For a pair  $(e,i) \in E \times [k]$ , let  $\mathcal{S}' := N(\mathcal{S}, e, i)$  be the  $(e,i)$ -neighbor of the support  $\mathcal{S}$ . Then,

$$(i) \mathbf{L}_{S'} = \mathbf{L}_S + \bar{\mathbf{q}}_{S',e,i} \mathbf{q}_{S,e,i}^\top,$$

$$(ii) \mathbf{L}_{S'}^+ = \mathbf{L}_S^+ - \frac{1}{1 + \mathbf{q}_{S,e,i}^\top \mathbf{L}_S^+ \bar{\mathbf{q}}_{S',e,i}} \mathbf{L}_S^+ \bar{\mathbf{q}}_{S',e,i} \mathbf{q}_{S,e,i}^\top \mathbf{L}_S^+,$$

$$(iii) \text{sgn}(\mathbf{q}_{S,e,i}^\top \mathbf{G}^\top \Delta\pi_S) = -\frac{\sigma_S}{\sigma_{S'}} \text{sgn}(\bar{\mathbf{q}}_{S',e,i}^\top \mathbf{G}^\top \Delta\pi_{S'}),$$

where  $\sigma_S := \text{sgn}(\det(\hat{\mathbf{L}}_S))$  is the sign of the determinant of the submatrix of  $\mathbf{L}_S$  obtained by deleting the first row and first column.

Theorem 3.2 establishes an easy update formula for the block Laplacians and their generalized inverses as well as a relation between the orientation of the direction vectors  $\Delta\pi_S$  of two neighboring supports.


For a feasible support  $\mathcal{S}$ , we are interested in neighboring feasible supports  $\mathcal{S}'$ . Neighboring supports with  $\mathcal{P}_S \cap \mathcal{P}_{S'} \neq \emptyset$  are of particular interest, as the line segment  $\mathcal{P}_S$  is extended with the line segment  $\mathcal{P}_{S'}$ . We call such pairs of neighboring support *continuative neighbors* as the line segment  $\mathcal{P}_S$  is “continued” by  $\mathcal{P}_{S'}$ . In order to characterize these continuative neighbors, we introduce for every feasible support  $\mathcal{S}$  the (possibly empty) sets

$$\mathfrak{N}_S^{\min} := \left\{ N(\mathcal{S}, e, i) \mid \frac{\mathbf{w}_{e,i}^\top (\mathbf{b} - \mathbf{G}^\top \bar{\mathbf{d}})}{\mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi} = \lambda_S^{\min} \right\}$$

$$\mathfrak{N}_S^{\max} := \left\{ N(\mathcal{S}, e, i) \mid \frac{\mathbf{w}_{e,i}^\top (\mathbf{b} - \mathbf{G}^\top \bar{\mathbf{d}})}{\mathbf{w}_{e,i}^\top \mathbf{G}^\top \Delta\pi} = \lambda_S^{\max} \right\}$$

of  $(e,i)$ -neighboring supports of  $\mathcal{S}$  that induce the values  $\lambda_S^{\min}$  and  $\lambda_S^{\max}$ . We also refer to these sets as *upper and lower continuative neighbors of  $\mathcal{S}$* . The following lemma shows that these sets contain all continuative neighbors.





**Figure 2:** The polytope of  $\lambda$ -potentials  $\mathcal{P}_{S_1}$  of some support  $S_1$  and the respective polytopes of two continuative neighbors. The boundary potentials  $\pi$  of the neighboring supports coincide. The polytopes are separated by hyperplanes with normalvector  $\mathbf{q}_{S,e,i}$ .

LEMMA 3.4. Denote by  $\mathfrak{N}_S := \{S' \mid \mathcal{P}_S \cap \mathcal{P}_{S'} \neq \emptyset\}$  the set of all continuative neighbors of  $S$ . Then,

$$\mathfrak{N}_S = \mathfrak{N}_S^{\min} \cup \mathfrak{N}_S^{\max}.$$

Further, the sets of  $\lambda$ -potentials do only intersect in exactly one potential, i.e.,

$$\mathcal{P}_S \cap \mathcal{P}_{S'} = \{\pi^*\} = \partial\mathcal{P}_S \cap \partial\mathcal{P}_{S'}$$


whenever  $S$  and  $S'$  are continuative neighbors. Moreover,  $\nu_S(\pi^*) = \nu_{S'}(\pi^*)$  in this case.

Lemma 3.4 can be interpreted as follows. The continuative neighbors of some support are fully described by  $\mathfrak{N}_S^{\min}$  and  $\mathfrak{N}_S^{\max}$  and can thus be easily identified with the computation of  $\lambda_S^{\min}$  and  $\lambda_S^{\max}$ . Moreover, the  $\lambda$ -potentials of two continuatively neighboring supports intersect in exactly one potential that is a boundary potential of both supports. And, finally, this potential in the intersection induces the same flow wrt. both supports, i.e., continuative neighbors do not only “continue” the sets of  $\lambda$ -potentials, but also the sets of Nash equilibrium flows.


#### 4 Membership in PPAD

With the insights of the previous section we now turn to show that NASH-ATOMIC-SPLITTABLE is in PPAD. The high level idea for this is the following: We consider all feasible supports as states of the PPAD-graph. We use the support of the empty flow for  $\lambda = 0$  as the start state. Then, under some non-degeneracy assumption, we can show that every support other than the initial support and the support of the equilibrium for  $\lambda = 1$  have exactly one predecessor and successor support. In order to define these predecessors and successors, we use the previously defined lower and upper continuative neighbors that are unique if the game is non-degenerate.

Before defining the state set of our PPAD-graph, we have to care about one further source of ambiguity. Consider for example the network from Figure 3a. If there is one player with source  $s$  and sink  $t$ , there



(a) Game with two total supports for  $\mathbf{x} = \mathbf{0}$ .



(b) Serial dependent edges.

**Figure 3:** (a) A 1-player game network (leftmost graph) that admits two potentials (values at the vertices) and two supports (active edges in red) for the equilibrium flow  $\mathbf{x} = \mathbf{0}$  (middle and rightmost graph). (b) The thick edges with the same color are serial-dependent. In particular,  $e_1$  is closer to the marked vertex  $v^*$  than  $e_2$ , as well as  $e_3$  is closer to  $v^*$  than  $e_4$  and  $e_5$ .

are, even if we restrict ourselves to total supports, two different potentials that induce the same flow (in this example the flow  $\mathbf{x} = (0, 0, 0)^\top$ ). In particular these two potentials belong to different supports. This ambiguity makes it harder to define predecessor and successor functions in a well-defined way. To avoid this problem, we will restrict ourselves to special potentials—namely potentials that can be interpreted as shortest path potentials of the players. To ensure that all potentials are shortest path potentials, we restrict ourselves to certain supports, which we will characterize with the following definition.

We say in any graph  $G = (V, E)$  two edges  $e, e' \in E$  are *serial-dependent* if there is a subset of vertices  $V' \subseteq V$  such that  $\delta^+(V') = \{e\}$  and  $\delta^-(V') = \{e'\}$  or  $\delta^+(V') = \{e'\}$  and  $\delta^-(V') = \{e\}$ , i.e., the edge-cut induced by  $V'$  contains exactly one incoming and one outgoing edge, and these cut edges are exactly  $e$  and  $e'$ . For some vertex  $v$  and two serial-dependent edges  $e, e'$  we say  $e$  is *closer to  $v$  than  $e'$*  if there is a subset of vertices  $V' \subseteq V$  with  $v \notin V'$  and  $\delta^+(V') = \{e'\}$  and  $\delta^-(V') = \{e\}$ . See Figure 3b for an illustration of these definitions. We say a total support  $S$  is a *shortest-path-support* if the following holds for all players  $i$ : If  $e$  and  $e'$  are serial-dependent in the graph  $G$  and  $e$  is closer to  $s_i$  than  $e'$  then the edge  $e$  is active for player  $i$ , i.e.,  $i \in S_e$ . In fact, it is not hard to show that if  $S$  is a shortest-path-support and  $\pi \in \mathcal{P}_S$ , then  $\pi$  contains the length of shortest paths for all players to all vertices<sup>1</sup> with respect to the edge lengths  $a_e^i \bar{x}_e + b_e^i$ , where  $\mathbf{x} = \nu(\pi)$  is the induced flow of potential  $\pi$ .

<sup>1</sup>For this to be true, we require that the graph  $G$  is strongly connected. If this is not the case, additional edges can be added with latency functions with high offsets.

Regarding the example from Figure 3a, observe that the edges  $(s, v)$  and  $(v, t)$  are serial-dependent, since  $\delta^+(\{v\}) = \{(v, t)\}$  and  $\delta^-(\{v\}) = \{(s, v)\}$ . Thus, only supports where  $(s, v)$  is active are shortest path supports which excludes the second support in the rightmost graph.

Finally, we define the set of PPAD-states as

$$\mathfrak{S} := \{\mathcal{S} : \mathcal{S} \text{ is a feasible, non-}a\text{-degenerate, shortest-path support}\}.$$

Next, we want to define successor and predecessor functions that yield predecessor and successor supports for every support  $\mathcal{S} \in \mathfrak{S}$ . These predecessors and successors will be based on upper and lower continuative neighbors. Although we eliminated many sources of ambiguity there still can be some degeneracy yielding multiple upper and/or lower continuative neighbors.

**DEFINITION 4.1.** *We say a support  $\mathcal{S} \in \mathfrak{S}$  is  $b$ -degenerate if  $|\mathfrak{N}_S^{\min} \cap \mathfrak{S}| > 1$  or  $|\mathfrak{N}_S^{\max} \cap \mathfrak{S}| > 1$ . A game is called  $b$ -degenerate if it has a  $b$ -degenerate support.*

For the remainder of this section, we assume that every game is non- $b$ -degenerate. We discuss  $b$ -degenerate games in Section 8.2.

**LEMMA 4.1.** *Let  $\mathcal{S} \in \mathfrak{S}$ . Then,*

- (i)  $\pi_S^{\min}$  is a 0-potential if  $\mathfrak{N}_S^{\min} \cap \mathfrak{S} = \emptyset$ .
- (ii)  $\pi_S^{\max}$  is a 1-potential if  $\mathfrak{N}_S^{\max} \cap \mathfrak{S} = \emptyset$ .

Lemma 4.1 states, that whenever there is no lower continuative neighbor, then the lower boundary potential  $\pi_S^{\min}$  is a 0-potential, i.e., a potential inducing the zero-flow. Similarly, whenever a support has no upper continuative neighbor, then the support admits the 1-potential  $\pi_S^{\max}$ , i.e., a potential inducing a Nash equilibrium for demand  $\mathbf{r}$  and thus a solution to NASH-ATOMIC-SPLITTABLE.

Finally, we introduce the *orientation*  $\sigma_S$  of the support  $\mathcal{S}$  by defining  $\sigma_S := \text{sgn}(\det(\hat{\mathbf{L}}_S))$ . That is, the orientation of the support  $\mathcal{S}$  is the sign of the determinant of the submatrix  $\hat{\mathbf{L}}_S$  obtained from  $\mathbf{L}_S$  by deleting the first row and first column. Theorem 3.2 implies that the relative orientation of the directions  $\Delta\pi_S$  of two neighboring supports depends on the orientation of the respective supports. This implies that, for example, the intersection  $\mathcal{P}_S \cap \mathcal{P}_{S'}$  contains the upper boundary potential  $\pi_S^{\max}$  of one of the supports and the lower boundary potential  $\pi_{S'}^{\min}$  whenever both supports have the same orientation. If two neighboring, continuative neighbors have opposite orientation the intersection of their  $\lambda$ -potentials  $\mathcal{P}_S \cap \mathcal{P}_{S'}$  contains either both lower

or both upper boundary potentials. For the example in Figure 2, we see that supports  $\mathcal{S}_0$  and  $\mathcal{S}_1$  have the same orientation, while  $\mathcal{S}_1$  and  $\mathcal{S}_2$  have opposite orientation. Using the orientation, we finally define a predecessor function  $\text{pred} : \mathfrak{S} \rightarrow \mathfrak{S} \cup \{\emptyset\}$  and successor function  $\text{succ} : \mathfrak{S} \rightarrow \mathfrak{S} \cup \{\emptyset\}$  as follows:

$$\text{pred}(\mathcal{S}) := \begin{cases} \emptyset & \text{if } \pi_S^{\min} \text{ is a 0-potential,} \\ \mathcal{S}' \in \mathfrak{N}_S^{\min} \cap \mathfrak{S} & \text{if } \sigma_S = 1 \text{ and } \pi_S^{\min} \\ & \text{is not a 0-potential,} \\ \mathcal{S}' \in \mathfrak{N}_S^{\max} \cap \mathfrak{S} & \text{if } \sigma_S = -1 \text{ and } \pi_S^{\min} \\ & \text{is not a 0-potential,} \end{cases}$$

$$\text{succ}(\mathcal{S}) := \begin{cases} \emptyset & \text{if } \pi_S^{\max} \text{ is a 1-potential,} \\ \mathcal{S}' \in \mathfrak{N}_S^{\max} \cap \mathfrak{S} & \text{if } \sigma_S = 1 \text{ and } \pi_S^{\max} \\ & \text{is not a 1-potential,} \\ \mathcal{S}' \in \mathfrak{N}_S^{\min} \cap \mathfrak{S} & \text{if } \sigma_S = -1 \text{ and } \pi_S^{\max} \\ & \text{is not a 1-potential.} \end{cases}$$

Additionally, we define the constant function **start**, that maps to the unique support  $\mathcal{S}^0$  of the zero flow. The following lemma shows that the  $\mathcal{S}^0$  as well as the functions **pred** and **succ** are well-defined, computable in polynomial time and compliant.

**LEMMA 4.2.** *The functions **start**, **pred**, **succ** are well-defined and computable in polynomial time. Furthermore,*

- (i) *The support  $\mathcal{S}^0 := \text{start}$  has positive orientation.*
- (ii) *If  $\text{pred}(\mathcal{S}) \neq \emptyset$ , then  $\text{succ}(\text{pred}(\mathcal{S})) = \mathcal{S}$ .*
- (iii) *If  $\text{succ}(\mathcal{S}) \neq \emptyset$ , then  $\text{pred}(\text{succ}(\mathcal{S})) = \mathcal{S}$ .*
- (iv) *If  $\text{succ}(\mathcal{S}) = \emptyset$ , then there is a 1-potential with support  $\mathcal{S}$ .*

Lemma 4.2 immediately proves that NASH-ATOMIC-SPLITTABLE is in PPAD.

**THEOREM 4.1.** *NASH-ATOMIC-SPLITTABLE is in PPAD.*

We obtain the following corollary which mirrors a similar result for non-degenerated bimatrix games. For the proof, we use that support changes appear only for a nullset of demands since there are only finitely many supports. This implies that except for a nullset of demands, the equilibria appear in the relative interior of the polytopes  $\mathcal{P}_S$  for some supports  $\mathcal{S}$ . This further implies that these equilibria are all unique startpoints or endpoints of a PPAD-paths. Subtracting the artificial equilibrium at  $\lambda = 0$ , we have shown that there is an odd number of equilibria.

COROLLARY 4.1. *A non- $b$ -degenerate atomic splittable congestion game has an odd number of Nash equilibria except for a nullset of demands.*

As we will see in Section 8.2, a small perturbation of the additive coefficients  $b_{e,i}$  makes a  $b$ -degenerate game non- $b$ -degenerate, implying that for  $b_{e,i}$  in general position,  $b$ -degeneracy does not occur.

## 5 PPAD-hardness

In this section, we show that it is PPAD-hard to compute an equilibrium in an integer-splittable congestion game with affine player-specific costs.

To prove that computing an equilibrium is PPAD-hard, we reduce from the problem of computing an  $\epsilon$ -approximate Nash equilibrium in Win-Lose-Games.

### APPROXIMATE-NASH-WIN-LOSE-GAME

INPUT: matrices  $\mathbf{U}, \mathbf{V} \in \{0, 1\}^{n \times n}$ ,  $\epsilon > 0$   
 OUTPUT: strategies  $\bar{\mathbf{y}}, \bar{\mathbf{z}} \in \mathbb{R}_{\geq 0}^n$   
 with  $\sum_{j=1}^n \bar{y}_j = \sum_{j=1}^n \bar{z}_j = 1$  such that  
 $\mathbf{y}^\top \mathbf{U} \bar{\mathbf{z}} \leq \bar{\mathbf{y}}^\top \mathbf{U} \bar{\mathbf{z}} + \epsilon$   
 $\bar{\mathbf{y}}^\top \mathbf{V} \bar{\mathbf{z}} \leq \bar{\mathbf{y}}^\top \mathbf{V} \bar{\mathbf{z}} + \epsilon$   
 for all strategies  $\mathbf{y}, \mathbf{z}$ .

The following result is due to Chen et al. [10].

THEOREM 5.1. (CHEN ET AL. [10]) *For any constant  $\beta > 0$ , APPROXIMATE-NASH-WIN-LOSE-GAME for matrices  $\mathbf{U}, \mathbf{V} \in \{0, 1\}^{n \times n}$  and  $\epsilon = n^{-\beta}$  is PPAD-complete.*

We show the following result.

THEOREM 5.2. NASH-ATOMIC-SPLITTABLE is PPAD-hard.

*Proof.* We reduce from APPROXIMATE-NASH-WIN-LOSE-GAME. Let  $(\mathbf{U}, \mathbf{V})$  be a win-lose game for some matrices  $\mathbf{U}, \mathbf{V} \in \{0, 1\}^{n \times n}$ . We proceed to describe the construction of a corresponding atomic splittable congestion games  $(G_{\mathbf{U}, \mathbf{V}}, \{1, 2\}, l)$  with two players with demands  $d_1 = d_2 = 1$ . The macro structure of the underlying graph  $G_{\mathbf{U}, \mathbf{V}}$  is shown in Figure 4. We note that our construction uses constant cost functions (in particular cost functions with constant cost 0) for the benefit of a simpler exposition. It is not hard to see that the same construction is valid for cost functions with non-constant functions with very small slopes  $a_{e,i} = \delta$  for a very small  $\delta > 0$ .

There are  $n^2$  gadgets  $G_{r,c}$  with  $r, c \in [n]$  that are arranged in a grid like fashion. The horizontal edges





Figure 4: Macro structure of the reduction.

of the grid as well as the edges connecting  $s_1$  and  $t_1$  to the grid, shown dashed and blue in Figure 4, have constant cost 0 for player 1 and constant cost  $4n$  for player 2. We call these edges *type-1 auxiliary edges*. Similarly, the vertical edges of the grid as well as the edges connecting  $s_2$  and  $t_2$  to the grid, shown dot-dashed and red in Figure 4 have constant cost  $4n$  for player 1 and constant cost 0 for player 2. These edges are called *type-2 auxiliary edges*. Every gadget  $G_{r,c}$  with  $r, c \in \{1, \dots, n\}$  has four designated verices  $\bar{s}_r, \bar{t}_r, \bar{s}_c, \bar{t}_c$ . In the macro structure, incoming auxiliary type-1 edges to  $G_{r,c}$  from the left are connected to  $\bar{s}_r$ , incoming auxiliary type-2 edges from above are connected to  $\bar{s}_c$ , outgoing auxiliary type-1 edges to the right are connected to  $\bar{t}_r$ , and outgoing auxiliary type-2 edges to below are connected to  $\bar{t}_c$ .

We introduce an additional parameter  $T \in \mathbb{N}$  in our construction of the gadgets  $G_{r,c}$ , parametrizing the number of edges in every gadget. We specify  $T$  in Claim 3 below. For some given  $T \geq 1$ , every gadget  $G_{r,c}$  has  $T$  *type-1 main edges*  $e_{r,c,1}^1, \dots, e_{r,c,T}^1$  and  $T$  *type-2 main edges*  $e_{r,c,1}^2, \dots, e_{r,c,T}^2$  as well as some auxiliary edges, see Figure 5. The dashed blue edges are type-1 auxiliary edges that have constant cost 0 for player 1 and constant cost  $4n$  for player 2 while the dash-dotted red edges are type-2 auxiliary edges that have constant cost  $4n$  for player 1 and constant cost 0 for player 2. The structure of  $G_{r,c}$  is such that every path from  $\bar{s}_r$  to  $\bar{t}_r$  that does not use type-2 auxiliary edges uses exactly one of the type-1 main edges  $e_{r,c,1}^1, \dots, e_{r,c,T}^1$  and all type-2 main edges  $\{e_{r,c,1}^2, \dots, e_{r,c,T}^2\}$ . Similarly, every path from  $\bar{s}_c$  to  $\bar{t}_c$  that does not use any of the type-1 auxiliary edges uses exactly one of the type-2 main



**Figure 5:** Gadget  $G_{r,c}$  used in the macro construction.

edges  $e_{r,c,1}^2, \dots, e_{r,c,T}^2$  as well as all of the type-1 main edges  $\{e_{r,c,1}^1, \dots, e_{r,c,T}^1\}$ .

For the type-1 and type-2 main edges, we define the player-specific affine costs as follows. The player-specific costs are of the form  $c_{e,i}(x) = a_{e,i}x + b_{e,i}$  with

$$a_{e,1} = \begin{cases} 1 - u_{r,c} & \text{if } e = e_{r,c,t}^1 \text{ with } r, c \in [n], t \in [T], \\ 0 & \text{otherwise;} \end{cases}$$

$$a_{e,2} = \begin{cases} 1 - v_{r,c} & \text{if } e = e_{r,c,t}^2 \text{ with } r, c \in [n], t \in [T], \\ 0 & \text{otherwise;} \end{cases}$$

$$b_{e,i} = 0 \quad \text{for all } e_{r,c,t}^1, e_{r,c,t}^2 \text{ with } r, c \in [n], t \in [T].$$

We claim that in every Nash equilibrium, player 1 does not use any of the auxiliary type-2 edges and player 1 does not use any of the auxiliary type-1 edges.

**CLAIM 1.** *Let  $\mathbf{x}$  be a Nash equilibrium of  $(G_{\mathbf{U}, \mathbf{V}}, \{1, 2\}, l)$ . Then,  $x_e^2 = 0$  for every type-2 auxiliary edge  $e$  and  $x_e^1 = 0$  for every type-1 auxiliary edge.*

*Proof of Claim 1.* We show the claim only for player 1 since the argumentation for player 2 is symmetrical. By Lemma 2.1, player 1 only uses paths with minimal marginal costs in  $\mathbf{x}$ . The marginal cost of every path containing a type-2 auxiliary edge  $e$  is at least

$$a_{e,i}\bar{x}_e + b_{e,i} + a_{e,i}x_e^i \geq 4n.$$

On the other hand, every path that does not contain a type-2 auxiliary edge traverses all gadgets  $G_{r,1}, \dots, G_{r,n}$  for some  $r \in [n]$ . In each gadget,  $G_{r,c}$  with  $c \in [n]$ ,

player 1 traverses one of the type-1 main edges  $e_{r,c,t_c}$  with  $t_c \in [T]$  as well as some type-1 auxiliary edges and some type-2 main edges each of which have no cost for player-1. Thus, the marginal cost of the path is determined by the  $n$  type-1 main edges. Let  $E_r = \{e_{r,c,t_c}^1 \mid c \in [n]\}$  be the set of these edges. The marginal cost of the path is then given by

$$\sum_{e \in E_r} a_{e,i}\bar{x}_e + b_{e,i} + a_{e,i}x_e^i \leq \sum_{e \in E_r} \bar{x}_e + x_{e,i} \leq 3n.$$

We conclude that every path of player 1 not containing a type-2 auxiliary edge has lower marginal cost than a path containing a type-2 auxiliary edge and the claim follows.  $\square$

Notice that player 1 has an exponential number of paths from  $s_1$  to  $t_1$  that do not use any of the type-2 edges. Every such path visits one row of gadgets  $G_{r,1}, \dots, G_{r,n}$  for some  $r \in \{1, \dots, n\}$ . Within each gadget  $G_{r,c}$ ,  $c \in \{1, \dots, n\}$  the path uses one of the  $T$  type-1 main edges  $e_{r,c,1}^1, \dots, e_{r,c,T}^1$ . We claim that in every Nash equilibrium, all paths that visit the same row of gadgets have the same flow of player 1.

**CLAIM 2.** *Let  $\mathbf{x}$  be a Nash equilibrium of  $(G_{\mathbf{U}, \mathbf{V}}, \{1, 2\}, l)$ . Then,  $x_{e_{r,c,t}^i}^i = x_{e_{r,c,t'}^i}^i$  for all  $i \in \{1, 2\}$ ,  $r, c \in [n]$ , and  $t, t' \in [T]$ .*

*Proof of Claim 2.* We show the claim only for player 1 since the argumentation for player 2 is symmetrical. Let  $r, c \in [n]$  and  $t, t' \in [T]$  be arbitrary and consider the edges  $e = e_{r,c,t}^1$  and  $e' = e_{r,c,t'}^1$ . Since  $\mathbf{x}$  is a Nash equilibrium, by Claim 1, we may assume that player 1 does not use any type-2 auxiliary edges and player 2 does not use any type-1 auxiliary edges. Since every path of player 2 that does not use any type-1 auxiliary edges and traverses the gadget  $G_{r,c}$  uses all type-1 main edges, this implies in particular that  $x_e^2 = x_{e'}^2$ .

If  $x_e^1 = x_{e'}^1 = 0$ , there is nothing left to show, so it is without loss of generality to assume that  $x_e^1 > 0$ . This implies that there is a path  $P$  from  $s_1$  to  $t_1$  carrying positive flow. Let  $P'$  be the path that uses the same edges as  $P$  except that in gadget  $G_{r,c}$  edge  $e'$  is used instead of edge  $e$ . (This means that also some auxiliary type-1 edges are swapped in  $G_{r,c}$  but since they have no cost for player 1, we may ignore them for the following arguments.) Lemma 2.1 implies

$$(5.4) \quad a_{e,1}\bar{x}_e + b_{e,1} + a_{e,1}x_e^1 \leq a_{e',1}\bar{x}_{e'} + b_{e',1} + a_{e',1}x_{e'}^1$$

which is equivalent to

$$(1 - u_{r,c})(2x_e^1 + x_{e'}^2) \leq (1 - u_{r,c})(2x_{e'}^1 + x_e^2)$$

using the definition of the cost functions. Further using that  $x_e^2 = x_{e'}^2$ , and that  $1 - u_{r,c} \geq 0$  this implies  $x_e^1 \leq x_{e'}^1$

and, thus,  $x_e^1 > 0$ . We conclude that (5.4) is actually satisfied with equality implying that  $x_e^1 = x_e^1$ .  $\square$

For player 1 and  $r \in \{1, \dots, n\}$  let  $x_r^1$  denote the total flow sent along the paths using the  $r$ th gadget row  $G_{r,1}, \dots, G_{r,n}$ . Similarly, for player 2 and  $c \in \{1, \dots, n\}$  let  $x_c^2$  denote the total flow sent along the paths using the  $c$ th gadget column  $G_{1,c}, \dots, G_{n,c}$ .

**CLAIM 3.** *Let  $\mathbf{x}$  be a Nash equilibrium of  $(G_{\mathbf{U}, \mathbf{V}}, \{1, 2\}, l)$  with  $T = 2n^{\beta+1}$  for some  $\beta > 0$ . Then, the strategy profile  $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$  with  $\bar{y}_r = x_r^1$  for all  $r \in [n]$  and  $\bar{z}_c = x_c^2$  for all  $c \in [n]$  is an  $n^{-\beta}$ -approximate Nash equilibrium of  $(\mathbf{U}, \mathbf{V})$ .*

*Proof of Claim 3.* For the sake of a contradiction, suppose that  $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$  is not an  $n^{-\beta}$ -approximate Nash equilibrium of  $(\mathbf{U}, \mathbf{V})$ . Due to symmetry, it is without loss of generality to assume that player 1 has an alternative strategy  $\mathbf{y}$  with  $\mathbf{y}^\top \mathbf{U} \bar{\mathbf{z}} > \bar{\mathbf{y}}^\top \mathbf{U} \bar{\mathbf{z}} + n^{-\beta}$ . This implies in particular that

$$(5.5) \quad \max_{r \in [n]} \sum_{c \in [n]} u_{r,c} \bar{z}_c > \min_{r \in [n]: \bar{y}_r > 0} \sum_{c=1}^n u_{r,c} \bar{z}_c + n^{-\beta}.$$

Let  $r^* \in \arg \max_{r \in [n]} \sum_{c \in [n]} u_{r,c} \bar{z}_c$  and let  $r' \in \operatorname{argmin}_{r \in [n]: \bar{y}_r > 0} \sum_{c=1}^n u_{r,c} \bar{z}_c$  be arbitrary. Since  $\mathbf{x}$  is a Nash equilibrium for  $G_{\mathbf{U}, \mathbf{V}}$ , all used paths have the same marginal total cost, and unused paths have higher marginal total cost. Using Claim 2, this implies in particular that

$$\begin{aligned} & \frac{\partial}{\partial x_{r'}^1} \left( \sum_{c \in [n]} T(1 - u_{r',c}) \left( \frac{x_{r'}^1}{T} + x_c^2 \right) \frac{x_{r'}^1}{T} \right) \\ & \leq \frac{\partial}{\partial x_{r^*}^1} \left( \sum_{c \in [n]} T(1 - u_{r^*,c}) \left( \frac{x_{r^*}^1}{T} + x_c^2 \right) \frac{x_{r^*}^1}{T} \right) \end{aligned}$$

which gives

$$\begin{aligned} & \sum_{c \in [n]} \frac{2(1 - u_{r',c}) x_{r'}^1}{T} + (1 - u_{r',c}) x_c^2 \\ & \leq \sum_{c \in [n]} \frac{2(1 - u_{r^*,c}) x_{r^*}^1}{T} + (1 - u_{r^*,c}) x_c^2. \end{aligned}$$

Using the definition of  $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$  we obtain

$$\begin{aligned} & \sum_{c \in [n]} \frac{2(1 - u_{r',c}) \bar{y}_{r'}}{T} + (1 - u_{r',c}) \bar{z}_c \\ & \leq \sum_{c \in [n]} \frac{2(1 - u_{r^*,c}) \bar{y}_{r^*}}{T} + (1 - u_{r^*,c}) \bar{z}_c. \end{aligned}$$

Finally, we obtain

$$\begin{aligned} & \sum_{c \in [n]} u_{r^*,c} \bar{z}_c \\ & \leq \sum_{c \in [n]} u_{r',c} \bar{z}_c + \sum_{c \in [n]} \frac{2(1 - u_{r^*,c}) \bar{y}_{r^*} - 2(1 - u_{r',c}) \bar{y}_{r'}}{T} \\ & \leq \sum_{c \in [n]} u_{r',c} \bar{z}_c + \sum_{c \in [n]} \frac{2 \bar{y}_{r^*}}{T} \\ & \leq \sum_{c \in [n]} u_{r',c} \bar{z}_c + \frac{2n}{T} = \sum_{c=1}^n u_{r',c} \bar{z}_c + n^{-\beta}, \end{aligned}$$

contradicting (5.5).  $\square$

Using that it is PPAD-complete to compute an  $n^{-\beta}$ -approximate equilibrium of a two-player win-lose game for any  $\beta > 0$ , we conclude that the computation of a Nash equilibrium of an atomic-splittable congestion game is PPAD-hard as well.  $\square$

With similar arguments, we can also show that the computation of a multi-class Wardrop equilibrium is PPAD-complete. A multi-class Wardrop equilibrium is a multi-commodity flow that satisfies the characterization via shortest path potentials of Lemma 2.2 for the original cost functions  $l_{e,i}$  instead for the marginal cost functions  $\mu_{e,i}^i$ , i.e.  $\mathbf{x}$  is a Wardrop equilibrium if and only if for all  $i \in [k]$  there is a potential vector  $\boldsymbol{\pi}^i$  with  $\pi_w^i - \pi_v^i = l_{e,i}(\bar{x}_e)$  if  $x_e^i > 0$ , and  $\pi_w^i - \pi_v^i \leq l_{e,i}(\bar{x}_e)$  if  $x_e^i = 0$  for all  $e = (v, w) \in E$ . We prove the following result settling an open question in [38].

**THEOREM 5.3.** *Computing a multi-class Wardrop equilibrium for commodity-specific affine costs is PPAD-hard.*

*Proof.* We use a similar construction as in the proof of Theorem 5.2 with the only exception that player 1 is replaced by a population of players of size 1 and player 2 is replaced by a population of players of size 1.

**CLAIM 4.** *Let  $\mathbf{x}$  be a Wardrop equilibrium of  $(G_{\mathbf{U}, \mathbf{V}}, \{1, 2\}, l)$ . Then  $x_e^1 = 0$  for every type-2 auxiliary edge  $e$  and  $x_e^2 = 0$  for every type-1 auxiliary edge  $e$ .*

*Proof of Claim 4.* We again show the claim only for population 1 since the argumentation for population 2 is symmetric. In a Wardrop equilibrium only paths with minimum cost are used. The total cost of any path containing a type-2 auxiliary edge  $e$  is at least  $4n$ . On the other hand, every path that does not contain a type-2 edge contains  $n$  type-1 main edges contained in some set  $E_r = \{e_{r,c,t_c} \mid c \in [n]\}$  with  $t_c \in [T]$  for all  $c \in [n]$ .

The total cost of these edges is equal to

$$\sum_{e \in E_r} a_{e,i} \bar{x}_e + b_{e,i} \leq \sum_{e \in E_r} x_e^1 + x_e^2 \leq 2n,$$

so that we conclude that no path containing a type-2 edge is used in a Wardrop equilibrium.  $\square$

We proceed to show that population  $i$  chooses all type- $i$  main edges within a gadget  $G_{r,c}$  with the same flow.

**CLAIM 5.** *Let  $\mathbf{x}$  be a Wardrop equilibrium of  $(G_{\mathbf{U},\mathbf{V}}, \{1, 2\}, l)$ . Then,  $x_{e_i}^i = x_{e_{r,c,t}^i}^i$  for all  $i \in \{1, 2\}$ ,  $r, c \in [n]$  and  $t, t' \in [T]$ .*

*Proof of Claim 5.* We again show the result only for population 1 since the argumentation for population 2 is symmetric. Let  $r, c \in [n]$  and  $t, t' \in [T]$  be arbitrary and consider the type-1 main edges  $e = e_{r,c,t}^1$  and  $e' = e_{r,c,t'}^1$  in gadget  $G_{r,c}$ . Using Claim 4, no population  $i$  uses auxiliary edges that are not of type  $i$ . This implies in particular that the flow of population 2 on edges  $e$  and  $e'$  is equal, i.e.,  $x_e^2 = x_{e'}^2$ .

If  $x_e^1 = x_{e'}^1 = 0$  there is nothing left to show, so it is without loss of generality to assume that  $x_e^1 > 0$ . This implies that there is a path  $P$  from  $s_1$  to  $t_1$  carrying positive flow. Let  $P'$  be the path that uses the same edges as  $P$  except that in gadget  $G_{r,c}$  the type-1 main edge  $e'$  is used instead of edge  $e$ . The equilibrium condition of Wardrop flows implies that

$$(5.6) \quad a_{e,1} \bar{x}_e + b_{e,1} \leq a_{e',1} \bar{x}_{e'} + b_{e',1}.$$

Using the definition of the cost functions, this is equivalent to

$$(1 - u_{r,c})(x_e^1 + x_e^2) \leq (1 - u_{r,c})(x_{e'}^1 + x_{e'}^2).$$

Further using that  $x_e^2 = x_{e'}^2$ , this implies that  $x_e^1 \leq x_{e'}^1$  and, thus,  $x_{e'}^1 > 0$ . We conclude that (5.6) was satisfied with equality implying that  $x_e^1 = x_{e'}^1$ .  $\square$

For population 1 and  $r \in [n]$ , let  $x_r^1$  denote the total flow sent along the paths using the  $r$ th gadget row  $G_{r,1}, \dots, G_{r,n}$ . Similarly, define  $x_c^2$  as the total flow sent by population 2 along the paths in the  $c$ th gadget column  $G_{1,c}, \dots, G_{n,c}$ .

**CLAIM 6.** *Let  $\mathbf{x}$  be a Wardrop equilibrium of  $(G_{\mathbf{U},\mathbf{V}}, \{1, 2\}, l)$  with  $T = n^{\beta+1}$  for some  $\beta > 0$ . Then, the strategy profile  $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$  with  $\bar{y}_r = x_r^1$  for all  $r \in [n]$  and  $\bar{z}_c = x_c^2$  for all  $c \in [n]$  is an  $n^{-\beta}$ -approximate Nash equilibrium of  $(\mathbf{U}, \mathbf{V})$ .*

*Proof of Claim 6.* For the sake of a contradiction, suppose that  $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$  is not an  $n^{-\beta}$ -approximate Nash

equilibrium of  $(\mathbf{U}, \mathbf{V})$ . Due to symmetry, it is without loss of generality to assume that player 1 has an alternative strategy  $\mathbf{y}$  with  $\mathbf{y}^\top \mathbf{U} \bar{\mathbf{z}} > \bar{\mathbf{y}}^\top \mathbf{U} \bar{\mathbf{z}} + n^{-\beta}$ . This implies in particular that

$$(5.7) \quad \max_{r \in [n]} \sum_{c \in [n]} u_{r,c} \bar{z}_c > \min_{r \in [n]: \bar{y}_r > 0} \sum_{c \in [n]} u_{r,c} \bar{z}_c + n^{-\beta}.$$

Let  $r^* \in \arg \max_{r \in [n]} \sum_{c \in [n]} u_{r,c} \bar{z}_c$  and let  $r' \in \arg \min_{r \in [n]: \bar{y}_r > 0} \sum_{c \in [n]} u_{r,c} \bar{z}_c$  be arbitrary. Since  $\mathbf{x}$  is a Wardrop equilibrium for  $G_{\mathbf{U},\mathbf{V}}$ , all used paths have the same cost, and unused paths have higher costs. Using Claim 5, this implies

$$\begin{aligned} & \sum_{c \in [n]} T(1 - u_{r',c}) \left( \frac{x_{r'}^1}{T} + x_c^2 \right) \\ & \leq \sum_{c \in [n]} T(1 - u_{r^*,c}) \left( \frac{x_{r^*}^1}{T} + x_c^2 \right). \end{aligned}$$

Using the definition of  $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ , we obtain

$$\begin{aligned} & \sum_{c \in [n]} T(1 - u_{r',c}) \left( \frac{\bar{y}_{r'}}{T} + \bar{z}_c \right) \\ & \leq \sum_{c \in [n]} T(1 - u_{r^*,c}) \left( \frac{\bar{y}_{r^*}}{T} + \bar{z}_c \right) \end{aligned}$$

which yields

$$\begin{aligned} & \sum_{c \in [n]} T u_{r^*,c} \bar{z}_c \\ & \leq \sum_{c \in [n]} T u_{r',c} \bar{z}_c + \sum_{c \in [n]} (1 - u_{r^*,c}) \bar{y}_{r^*} - (1 - u_{r',c}) \bar{y}_{r'} \end{aligned}$$

and equivalently

$$\begin{aligned} & \sum_{c \in [n]} u_{r^*,c} \bar{z}_c \\ & \leq \sum_{c \in [n]} u_{r',c} \bar{z}_c + \sum_{c \in [n]} \frac{(1 - u_{r^*,c}) \bar{y}_{r^*} - (1 - u_{r',c}) \bar{y}_{r'}}{T} \\ & \leq \sum_{c \in [n]} u_{r',c} \bar{z}_c + \sum_{c \in [n]} \frac{\bar{y}_{r^*}}{T} \leq \sum_{c \in [n]} + n^{-\beta} \end{aligned}$$

contradicting (5.6).  $\square$

Since the computation of an  $\epsilon$ -approximate equilibrium for win-lose games is PPAD-complete, we conclude that also the computation of a multi-class Wardrop equilibrium is PPAD-complete.  $\square$

## 6 Player-independent cost functions

In this section, we briefly discuss the special case of player-independent cost functions, i.e., games where the

cost functions on the edges are  $l_{e,i}(x_e) := l_e(x_e) = a_e x_e + b_e$  and, thus, independent of the player. In this case, we obtain the following result for the block Laplacians of any total support.

**LEMMA 6.1.** *If the coefficients of the cost functions  $l_{e,i}$  are independent of the player  $i$ , then, for any total support  $\mathcal{S}$ ,  $\mathbf{L}_{\mathcal{S}}$  is symmetric and positive semi-definite and  $\ker(\mathbf{L}_{\mathcal{S}}) = \mathcal{N}$ .*

Sylvester's criterions states that a matrix is positive semi-definite if and only if all principal minors of that matrix are positive semi-definite. Thus, by Lemma 6.1, the principal minor  $\sigma_{\mathcal{S}} = \det(\hat{\mathbf{L}}_{\mathcal{S}})$  is non-negative. Additionally, Lemma 6.1(ii) states that all supports are non- $a$ -degenerate implying that  $\sigma_{\mathcal{S}} > 0$ , and we obtain the following result.

**THEOREM 6.1.** *If the coefficients of the cost functions  $l_{e,i}$  are independent of the player  $i$ , then*

- (i) *the game is non- $a$ -degenerate.*
- (ii) *the orientation  $\sigma_{\mathcal{S}}$  of every total support  $\mathcal{S}$  is positive.*
- (iii) *For every demand vector  $\mathbf{r}$ , there is a unique Nash equilibrium  $\mathbf{x}$ .*

## 7 Parametric computation

Using the trivial PPAD-algorithm, beginning with the support given by `start`, we can iterate through (possibly exponentially many) supports using the function `succ`, until we eventually obtain a support  $\mathcal{S}$  for the Nash equilibrium for the demand  $\mathbf{r}$ .

As a byproduct of this procedure, we obtain a sequence of supports  $\mathcal{S}_0, \dots, \mathcal{S}_{\tau}$ , where  $\mathcal{S}_l$  is the support of Nash equilibria for demands  $\lambda \mathbf{r}$  with  $\lambda \in [\lambda_{\mathcal{S}_l}^{\min}, \lambda_{\mathcal{S}_l}^{\max}]$ . By construction of the successor function, we know that  $\bigcup_{l=0}^{\tau} [\lambda_{\mathcal{S}_l}^{\min}, \lambda_{\mathcal{S}_l}^{\max}] = [0, 1]$ . Thus, we can define a function  $\lambda \mapsto \mathcal{S}(\lambda)$  that maps every  $\lambda$  to a support of some Nash equilibrium for demand  $\lambda \mathbf{r}$ . Given the support  $\mathcal{S}(\lambda)$  and  $\lambda$ , we can easily compute a Nash equilibrium for the demand  $\lambda$  as we have discussed in subsection 3.1. Hence, we obtain a function  $\mathbf{f} : [0, 1] \rightarrow \mathbb{R}^{mk}$  such that  $\mathbf{f}(\lambda)$  is a Nash equilibrium for the demand  $\lambda \mathbf{r}$ , i.e., a function solving PARAMETRIC-NASH-ATOMIC-SPLITTABLE.

**THEOREM 7.1.** *PARAMETRIC-NASH-ATOMIC-SPLITTABLE can be solved in polynomial space.*

*Proof.* Since the function  $\lambda \mapsto \pi(\lambda)$  that maps a value  $\lambda$  to the potential of a Nash equilibrium for demand  $\lambda \mathbf{r}$  is linear for every support (see Lemma 3.3), and flows depend linearly on the potentials, the function

$\mathbf{f}$  is piecewise-linear. Thus, it is enough to compute the potentials at the breakpoints  $\lambda_{\mathcal{S}}^{\min}, \lambda_{\mathcal{S}}^{\max}$  explicitly. For any support, this can be clearly done in polynomial time. Since the functions `start` and `succ` can be computed in polynomial time as well, PARAMETRIC-NASH-ATOMIC-SPLITTABLE can be computed in polynomial space.  $\square$

A game has unique equilibria for all demands if and only if, for every feasible support  $\mathcal{S}$ , we have  $\sigma_{\mathcal{S}} > 0$ . In this case, the function  $\mathbf{f}$  computes *all Nash Equilibria* of the game. If, in addition, the game is non- $b$ -degenerate, then every support corresponds to a breakpoint of the piecewise affine function  $\mathbf{f}$ , yielding an output polynomial algorithm. Using Theorem 6.1 this implies in particular an output-polynomial algorithm for the parametric Nash equilibrium problem for player-independent costs.

**THEOREM 7.2.** *PARAMETRIC-NASH-ATOMIC-SPLITTABLE can be solved in output-polynomial time for non- $b$ -degenerate games with  $\sigma_{\mathcal{S}} > 0$  for all feasible supports  $\mathcal{S}$ . In particular, the runtime is in  $\mathcal{O}((kn)^{2.4} + \tau(kn)^2)$ , where  $\tau$  is the number of breakpoints of the piecewise affine function  $\mathbf{f}$  returned by the algorithm.*

## 8 Degeneracy

**8.1  $a$ -Degeneracy** We are now going to show how the functions `start`, `pred`, and `succ` can be modified to still work as intended, even if  $a$ -degenerate supports exist. First, recall that a support  $\mathcal{S}$  is  $a$ -degenerate, if  $\dim(\ker(\mathbf{L}_{\mathcal{S}})) > \dim(\mathcal{N}) = k$ , i.e., if  $\hat{\mathbf{L}}_{\mathcal{S}}$  is singular. This means that equation  $\mathbf{L}_{\mathcal{S}}\boldsymbol{\pi} = \mathbf{y}$  has no longer a unique solution in  $\mathcal{P}$ . In fact, if there is an  $a$ -degenerate support  $\mathcal{S}$  that is feasible, i.e.,  $\mathbf{L}_{\mathcal{S}}\boldsymbol{\pi} = \mathbf{y}$  has a solution, then there must be infinitely many solutions to this equation. Lemma 3.2 shows that the induced flow function  $\nu$  is injective. Thus, if there exists an  $a$ -degenerate support, then there must be also infinitely many Nash equilibrium flows for the same excess vector  $\mathbf{y}$  and, hence, infinitely many Nash equilibria for the same demand. See Example 8.1 below for a concrete example where such support exists.

In general, for an  $a$ -degenerate support, all potentials induce the same excess vector  $\mathbf{y}$ . Thus,  $\lambda_{\mathcal{S}}^{\min} = \lambda_{\mathcal{S}}^{\max}$  (we prove this formally in Lemma 8.2) and the potentials  $\boldsymbol{\pi}_{\mathcal{S}}^{\min}$  and  $\boldsymbol{\pi}_{\mathcal{S}}^{\max}$  are not unique. Therefore, the functions `pred` and `succ` are not well-defined in the case of an  $a$ -degenerate support. We now proceed by adapting these functions for the case of  $a$ -degenerate games. As before, we still (only) use

$$\mathfrak{S} := \{\mathcal{S} : \mathcal{S} \text{ is a feasible, non-}a\text{-degenerate, shortest-path-support}\}$$

as the set of states, i.e., even in an  $a$ -degenerate game, we only consider non- $a$ -degenerate supports as states. We observe, that **start** as defined before yields the support  $\mathcal{S}^0$  which is non- $a$ -degenerate with the same proof as in Lemma 4.2. Hence, we only need to define new successor and predecessor functions.

The first observation that we are going to make is that any (continuative) neighboring support  $\mathcal{S}'$  of some non- $a$ -degenerate support  $\mathcal{S}$  is not “too degenerate”. Formally, we say a support  $\mathcal{S}$  is *weakly- $a$ -degenerate* if  $\text{rank}(\mathbf{L}_{\mathcal{S}}) = k(n-1) - 1$  (or equivalently, if  $\dim(\ker(\mathbf{L}_{\mathcal{S}})) = k+1$ ). With the rank-1-update formula from Theorem 3.2 we directly obtain

LEMMA 8.1. *Let  $\mathcal{S}$  be a non- $a$ -degenerate support. Then every neighboring support  $\mathcal{S}'$  of  $\mathcal{S}$  is non- $a$ -degenerate or weakly- $a$ -degenerate.*

For every weakly- $a$ -degenerate support  $\mathcal{S}$  we know by definition that  $\dim(\ker(\hat{\mathbf{L}}_{\mathcal{S}})) = 1$ , i.e., there is a direction vector  $\Delta\pi_{\mathcal{S}}^N$  such that  $\ker(\hat{\mathbf{L}}_{\mathcal{S}}) = \text{span}(\Delta\pi_{\mathcal{S}}^N)$ . (This direction is of course not unique, but it is unique up to a scalar multiplication.) We refer to this direction as the *nullspace direction* of the weakly- $a$ -degenerate support  $\mathcal{S}$ . As it will turn out, if, for example, the successor  $\mathcal{S}' := \text{succ}(\mathcal{S})$  of some non- $a$ -degenerate support  $\mathcal{S}$  is  $a$ -degenerate, then we can use this direction to find another, non- $a$ -degenerate support  $\mathcal{S}''$  by following the nullspace direction  $\Delta\pi_{\mathcal{S}'}^N$ . We can then use  $\mathcal{S}''$  as successor instead. In order to prove this formally, we first need the following technical observations.

LEMMA 8.2. *Let  $\mathcal{S}$  be a non- $a$ -degenerate support and  $\mathcal{S}'$  be a  $(e, i)$ -neighboring,  $a$ -degenerate support.*

(i) *If  $\mathbf{w}_{\mathcal{S},e,i}^\top \Delta\pi_{\mathcal{S}} \neq 0$  and  $\mathcal{S}'$  is feasible, then the equation  $\mathbf{L}_{\mathcal{S}'} \boldsymbol{\pi} = \lambda \Delta\mathbf{y} + \mathbf{d}_{\mathcal{S}'}$  has a solution if and only if  $\lambda = \lambda_{\mathcal{S}'}^*$  for some fixed  $\lambda_{\mathcal{S}'}^* \in [0, 1]$ .*

(ii) *Some potential vector  $\boldsymbol{\pi} \in \mathcal{P}$  satisfies  $\mathbf{L}_{\mathcal{S}'} \boldsymbol{\pi} = \lambda_{\mathcal{S}'}^* \Delta\mathbf{y} + \mathbf{d}_{\mathcal{S}'}$  if and only if there is some  $\xi \in \mathbb{R}$  such that  $\boldsymbol{\pi} = \mathbf{L}_{\mathcal{S}'}^+(\lambda_{\mathcal{S}'}^* \Delta\mathbf{y} + \mathbf{d}_{\mathcal{S}'}) + \xi \Delta\pi_{\mathcal{S}'}^N$ , for some generalized inverse  $\mathbf{L}_{\mathcal{S}'}^+$ , with  $\mathbf{L}_{\mathcal{S}'}^+(\lambda_{\mathcal{S}'}^* \Delta\mathbf{y} + \mathbf{d}_{\mathcal{S}'}) \in \mathcal{P}$ .*

Note, that the statements of Lemma 8.2 hold for any (fixed) choice of the generalized inverses of  $\mathbf{L}_{\mathcal{S}}$  and  $\mathbf{L}_{\mathcal{S}'}$ . In the non- $a$ -degenerate case, we use the fact that there is a unique generalized inverse  $\mathbf{L}_{\mathcal{S}}^*$  that maps into  $\mathcal{P}$ . For  $a$ -degenerate regions, there is no longer a unique such choice of the generalized inverse. For the remainder of this section we assume we have chosen some fixed generalized inverse  $\mathbf{L}_{\mathcal{S}'}^+$  that maps into  $\mathcal{P}$ .

Assume that the continuative  $(e, i)$ -neighbor  $\mathcal{S}'$  of some non- $a$ -degenerate support  $\mathcal{S}$  is  $a$ -degenerate. Then, we know that  $\mathcal{S}'$  is feasible (since  $\mathcal{S}$  and  $\mathcal{S}'$  share

a boundary potential), thus we can use Lemma 8.2 to describe the polytope of  $\lambda$ -potentials  $\mathcal{P}_{\mathcal{S}'}$  of  $\mathcal{S}'$  as

$$\begin{aligned} \mathcal{P}_{\mathcal{S}'} &= \{ \boldsymbol{\pi} \in \mathcal{P} \mid \mathbf{L}\boldsymbol{\pi} - \mathbf{d} = \lambda_{\mathcal{S}'}^* \Delta\mathbf{y}, \mathbf{W}(\mathbf{G}^\top \boldsymbol{\pi} - \mathbf{b}) \leq \mathbf{0} \} \\ &= \{ \mathbf{L}^+(\lambda_{\mathcal{S}'}^* \Delta\mathbf{y} + \mathbf{d}) + \xi \Delta\pi_{\mathcal{S}'}^N \mid \exists \xi \in \mathbb{R} : \\ &\quad \mathbf{W}(\mathbf{G}^\top(\mathbf{L}^+(\lambda_{\mathcal{S}'}^* \Delta\mathbf{y} + \mathbf{d}) + \xi \Delta\pi_{\mathcal{S}'}^N) - \mathbf{b}) \leq \mathbf{0} \}. \end{aligned}$$

We see that the polytope  $\mathcal{P}_{\mathcal{S}'}$  of an  $a$ -degenerate region is parametrized by  $\xi \in \mathbb{R}$  along the nullspace direction  $\Delta\pi_{\mathcal{S}'}^N$ , in contrast to a non- $a$ -degenerate region  $\mathcal{S}$ , where  $\mathcal{P}_{\mathcal{S}}$  is parametrized by  $\lambda$  and the potential direction  $\Delta\pi_{\mathcal{S}}$ . Similar to Lemma 3.3, we can rewrite the polytope of  $\lambda$ -potentials as

$$\mathcal{P}_{\mathcal{S}'} = \{ \mathbf{L}_{\mathcal{S}'}^+(\lambda_{\mathcal{S}'}^* \Delta\mathbf{y} + \mathbf{d}_{\mathcal{S}'}) + \xi \Delta\pi_{\mathcal{S}'}^N \mid \xi_{\mathcal{S}'}^{\min} \leq \xi \leq \xi_{\mathcal{S}'}^{\max} \}.$$

We note that  $\xi_{\mathcal{S}'}^{\min}, \xi_{\mathcal{S}'}^{\max}$  are in fact finite values by the following argument. The induced flow of the nullspace direction  $\boldsymbol{\nu}(\Delta\pi_{\mathcal{S}'}^N)$  is a flow circulation that satisfies the potential equations  $\Delta\pi_w^{N,i} - \Delta\pi_v^{N,i} = a_e^i \Delta\bar{x}_e$ . Hence, there must be at least one edge with positive total induced flow and one edge with negative total induced flow implying that adding or subtracting  $\boldsymbol{\nu}(\Delta\pi_{\mathcal{S}'}^N)$  violates at some point the non-negativity of the flow for some edge. We therefore know that a non- $a$ -degenerate support  $\mathcal{S}'$  has well-defined boundary potentials  $\pi_{\mathcal{S}'}^{\min}, \pi_{\mathcal{S}'}^{\max}$ . As for non- $a$ -degenerate supports we can find continuative neighbors by looking at the  $(e, i)$ -neighbors that are shortest-path-supports for all pairs  $(e, i)$  that induce  $\xi_{\mathcal{S}'}^{\min}$  and  $\xi_{\mathcal{S}'}^{\max}$ .

It can be shown that for every boundary potential of the degenerate support  $\mathcal{S}$ , we can find a continuative neighbor. This neighbor is also unique, as long as the  $a$ -degenerate support is not also  $b$ -degenerate—we cover this case of  $a$ - $b$ -degeneracy in the next section. The following lemma proves the existence, uniqueness, and the non- $a$ -degeneracy of such neighbor supports.

LEMMA 8.3. *Let  $\mathcal{S}$  be a non- $a$ -degenerate support. Further, let  $\mathcal{S}'$  be a continuative  $(e, i)$ -neighbor of  $\mathcal{S}$  that is  $a$ -degenerate. Then, there is a uniquely defined support  $\mathcal{S}_{\mathcal{S},\mathcal{S}'}^* \neq \mathcal{S}$  that is a neighbor of  $\mathcal{S}'$  and non- $a$ -degenerate.*

We define a new predecessor function

$$\text{pred}^*(\mathcal{S}) := \begin{cases} \text{pred}(\mathcal{S}) & \text{if } \text{pred}(\mathcal{S}) \in \mathfrak{S} \\ \mathcal{S}_{\mathcal{S},\text{pred}(\mathcal{S})}^* & \text{if } \text{pred}(\mathcal{S}) \notin \mathfrak{S} \end{cases}$$

and a new successor  $\text{succ}^*$  function analogously that are based on the the basic functions  $\text{pred}$  and  $\text{succ}$  except for the cases where these function would map to a  $a$ -degenerate support  $\mathcal{S}' \notin \mathfrak{S}$ . Finally, the next Lemma proves, that these functions are still compliant with each other.






Figure 6: Graph from Example 8.1 with costs of player 1.

LEMMA 8.4. The functions  $\text{pred}^*$  and  $\text{succ}^*$  are well-defined and computable in polynomial time. Further,

- (i)  $\text{pred}^*(\mathcal{S}) \neq \emptyset \Rightarrow \text{succ}^*(\text{pred}^*(\mathcal{S})) = \mathcal{S}$
- (ii)  $\text{succ}^*(\mathcal{S}) \neq \emptyset \Rightarrow \text{pred}^*(\text{succ}^*(\mathcal{S})) = \mathcal{S}$ .

EXAMPLE 8.1. We consider a game with 8 players on the graph given in Figure 6. Every player has two adjacent vertices as source and sink vertex and a demand of  $r_i = 2$ . We define the costs in a symmetric way, meaning that for every player the edge connecting the source and the sink vertex (e.g.  $e_8$  for player 1) is equipped with the cost function  $c_1(x) = 9x + 3$  and the edges on the longer path from the source to the sink (e.g.  $e_1, e_3, e_5$  for player 1) are equipped with the cost function  $c_2(x) = x + 6$ . All other edges are assumed to have cost functions with high offsets and slopes such that they are never in the support of the respective players. By definition, the game is player-symmetric, i.e., the cost functions look the same for every player. Yet, not all players have the same cost function on every edge. Because of this player symmetry, there are overall three possible strategy profiles:

- a) All players use the direct path between  $s_i$  and  $t_i$ .
- b) All players use both path between  $s_i$  and  $t_i$ .
- c) All players use the long path between  $s_i$  and  $t_i$ .

Since all strategies and all cost functions are symmetric we give all potentials, flows and directions for the first player routing demand from  $v_1$  to  $v_4$ . All other potential and flow values are given implicitly.

For small demands  $\lambda \mathbf{r}$ ,  $\lambda \leq \frac{1}{2}$ , because of the relatively high offset of the long path, all players will route their demand on the direct path to their source. The potential for the zero flow  $\mathbf{x}^0 = \mathbf{0}$  is the shortest path potential  $\boldsymbol{\pi}^0 = (0, 6, 12, 3)^\top$ . This yields the initial support  $\mathcal{S}^0$  where the edges  $e_1, e_3$  and  $e_8$  are active for the player 1. (Note that we consider the edges  $e_1, e_3$

active although they are not used by the player. They are only part of the shortest path network.) We obtain the potential direction  $\Delta \boldsymbol{\pi}^0 = (0, 2, 4, 36)^\top$  that induces the flow direction  $\Delta \mathbf{x}^0 = (0, 0, 0, 0, 0, 0, 0, 2)^\top$ . For  $\lambda = \frac{1}{2}$ , we obtain the potential  $\boldsymbol{\pi}^1 = (0, 7, 14, 21) = \boldsymbol{\pi}^0 + \frac{1}{2} \Delta \boldsymbol{\pi}^0$  inducing the flow  $\mathbf{x}^1 = (0, 0, 0, 0, 0, 0, 0, 1)^\top$ , i.e., every player routes exactly one flow unit on the direct path from source to sink.

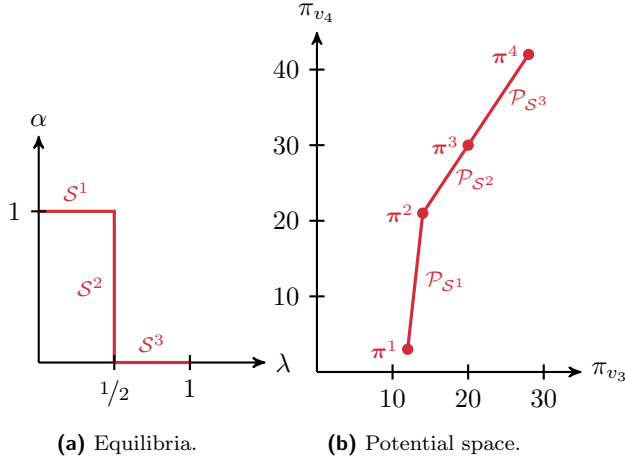
Denote by  $e_{i_1}$  the last edge of the long path for player  $i$  (e.g.  $e_{1_1} = e_5$ ). In the potential  $\boldsymbol{\pi}^1$ , all edges  $e_{i_1}$  become active at once. (We note that this implies that the game is actually  $b$ -degenerate, but we ignore this for the benefit of an easier example.) This defines a new support  $\mathcal{S}^1$  that is  $a$ -degenerate—the Laplacian matrix has rank 23 and  $\ker(\mathbf{L}_{\mathcal{S}^1}) \cap \mathcal{P} = \text{span}((0, 1, 2, 3)^\top)$ . Hence, the nullspace direction is  $\Delta \boldsymbol{\pi}_{\mathcal{S}^1}^1 = (0, 1, 2, 3)^\top$  inducing the circulation  $\Delta \mathbf{x}^1 = (\frac{1}{3}, 0, \frac{1}{3}, 0, \frac{1}{3}, 0, 0, -\frac{1}{3})^\top$ . The potential  $\boldsymbol{\pi}^2 := \boldsymbol{\pi}^1 + 3 \Delta \boldsymbol{\pi}_{\mathcal{S}^1}^1 = (0, 10, 20, 30)^\top$  induces the flow  $\mathbf{x}^2 = (1, 0, 1, 0, 1, 0, 0, 0)^\top$ . Note that this flow is still an equilibrium for the demand  $\frac{1}{2} \mathbf{r}$ . In fact, all flows

$$\mathbf{x}(\alpha) = ((1-\alpha), 0, (1-\alpha), 0, (1-\alpha), 0, 0, \alpha)^\top, 0 \leq \alpha \leq 1$$

are equilibria for the demand  $\frac{1}{2} \mathbf{r}$ , i.e., there are infinitely many equilibria when every player wants to route a demand of 1.

Finally, denote by  $e_{d_i}$  the direct edge connecting the source and sink of player  $i$  (e.g.  $e_{d_1} = e_8$ ). All these edges become inactive in the potential  $\boldsymbol{\pi}^2$  which defines the next support  $\mathcal{S}^2$  which is non- $a$ -degenerate. (Again, we ignore the  $b$ -degeneracy at this point.) With this support, we obtain the direction  $\Delta \boldsymbol{\pi}^2 = (0, 8, 16, 24)^\top$  inducing the flow direction  $\Delta \mathbf{x} = (2, 0, 2, 0, 2, 0, 0, 0)^\top$ . For  $\lambda = 1$  we then get the potential  $\boldsymbol{\pi}^3 := \boldsymbol{\pi}^2 + \frac{1}{2} \Delta \boldsymbol{\pi}^2 = (0, 14, 28, 42)^\top$  inducing the (solution) flow  $\mathbf{x}^3 = (2, 0, 2, 0, 2, 0, 0, 0)^\top$ .

**8.2  $b$ -Degeneracy** In this section, we develop a rule for handling  $b$ -degenerate games. A game is  $b$ -degenerate, if there is a  $b$ -degenerate support. That is, a support with  $|\mathcal{N}_{\mathcal{S}}^{\min} \cap \mathcal{G}| > 1$  or  $|\mathcal{N}_{\mathcal{S}}^{\max} \cap \mathcal{G}| > 1$ , i.e., a support with a boundary potential that is feasible in more than one  $(e, i)$ -neighboring supports. This is the case if and only if there is not a unique minimizer or maximizer in the computation of  $\lambda_{\mathcal{S}}^{\min}$  or  $\lambda_{\mathcal{S}}^{\max}$ . Concretely, this means that when increasing or decreasing the demand, multiple edge-player-pairs change their activity status at once, e.g., if two edges become active for one player or one edge becomes inactive for two players and so on. Roughly speaking,  $b$ -degeneracy occurs if the offsets of the cost functions  $\mathbf{b}$  are such that the hyperplanes induced by the inequalities in the polytopes



**Figure 7:** Strategies and  $\lambda$ -potentials in the game defined in Example 8.1. **(a)** The strategies of the players described by the fraction  $\alpha$  of the demand that is routed on the direct edge between source and sink, parametrized by the demand multiplier  $\lambda$ . **(b)** The polytopes of  $\lambda$ -potentials  $\mathcal{P}_{\mathcal{S}}$  of the three supports with the boundary potentials in a projection of the potential space  $\mathcal{P}$ .

$\mathcal{P}_{\mathcal{S}}$  intersect exactly in some boundary potential of some support. Intuitively, a small random perturbation of the offsets  $\mathbf{b}$  would translate the hyperplanes in a way that, almost surely, they do not intersect in boundary potentials anymore. Further, a small perturbation should not change the equilibrium to much. In fact, we will show that there is an explicit, small perturbation  $\epsilon \in \mathbb{R}^{mk}$  that can be added to the vector  $\mathbf{b}$  such that the game with these perturbed offsets is non- $b$ -degenerate. Further, we will show that all feasible supports in this game are also feasible for the original, unperturbed game and, hence, it will be enough to compute the states  $\mathcal{S}$  of the perturbed game. Additionally, we will see that it is not necessary to carry out the perturbation explicitly. Rather, we can use a lexicographic criterion. We initially assume that every game is non- $a$ -degenerate and discuss the interplay of  $a$ -degeneracy and  $b$ -degeneracy later.


Let  $\epsilon > 0$ . Then we define the *perturbation vector*

$$\epsilon := (\epsilon, \epsilon^2, \dots, \epsilon^{km-1}, \epsilon^{km})^\top \in \mathbb{R}^{mk}$$

and the vector of perturbed offsets  $\mathbf{b}^\epsilon := \mathbf{b} + \epsilon$ . We call the game with this offset vector the  $\epsilon$ -perturbed game. Analogous to Theorem 3.1 and Lemma 3.3 we now obtain the following result for the  $\epsilon$ -perturbed game.

**LEMMA 8.5.** *For  $\epsilon > 0$  and some support  $\mathcal{S}$ , let  $\mathcal{P}_{\mathcal{S}}^\epsilon$  be the set of all  $\lambda$ -potentials in the  $\epsilon$ -perturbed game. Then there are values  $\lambda_{\mathcal{S}}^{\min, \epsilon}$  and  $\lambda_{\mathcal{S}}^{\max, \epsilon}$  such that*

$$\mathcal{P}_{\mathcal{S}}^\epsilon = \{ \lambda \Delta \pi + \bar{\mathbf{d}}^\epsilon \mid \lambda_{\mathcal{S}}^{\min, \epsilon} \leq \lambda \leq \lambda_{\mathcal{S}}^{\max, \epsilon} \},$$



**Figure 8:** The  $\lambda$ -potentials of a game (thick line segments) and its  $\epsilon$ -perturbed variant (dotted line segments). The dashed lines are the hyperplanes induced by the normal vectors  $\mathbf{q}_{\mathcal{S}, e, i}$ . The support  $\mathcal{S}^0$  is non- $b$ -degenerate, thus, both in the perturbed and the unperturbed game, there is the unique continuative neighbor  $\mathcal{S}^1$ . This support is  $b$ -degenerate in the unperturbed game: It has five continuative neighbors—for four of them  $\mathcal{P}_{\mathcal{S}}$  is a single point. In the perturbed game, all supports are non- $b$ -degenerate.

where  $\bar{\mathbf{d}} := \mathbf{L}^* \mathbf{G} \mathbf{C} \mathbf{b}^\epsilon$ .

With the next theorem, we will prove two results. First, we show that every support that is feasible in the perturbed game is also feasible in the original justifying that it is sufficient to consider just these supports. Second, we show that the perturbed game is non- $b$ -degenerate. Figure 8 illustrates the  $\lambda$ -potentials for  $b$ -degenerate and non- $b$ -degenerate supports for the unperturbed and the  $\epsilon$ -perturbed game.

For any support  $\mathcal{S}$ , we say two edges  $e, e'$  are *serial-dependent for player  $i$  in support  $\mathcal{S}$*  if  $e, e'$  are serial-dependent in the subgraph of  $G$  containing only the active edges of player  $i$ .

**THEOREM 8.1.** *Let  $\mathcal{S}$  be a non- $a$ -degenerate, total support. Then there is  $\epsilon_{\mathcal{S}}^* > 0$  such that for all  $0 < \epsilon < \epsilon_{\mathcal{S}}^*$  the following holds.*

- (i) *If  $\mathcal{S}$  is infeasible, then  $\mathcal{P}_{\mathcal{S}}^\epsilon = \emptyset$  as well.*
- (ii) *If  $\mathcal{S}$  is feasible and non- $b$ -degenerate, then  $\mathcal{S}$  is also non- $b$ -degenerate in the  $\epsilon$ -perturbed game. Further, the support  $\mathcal{S}$  has the same unique continuative neighbors in the unperturbed and the  $\epsilon$ -perturbed game.*
- (iii) *If  $\mathcal{S}$  is feasible and  $b$ -degenerate, then  $\mathcal{P}_{\mathcal{S}}^\epsilon = \emptyset$  or  $\mathcal{S}$  is non- $b$ -degenerate in the  $\epsilon$ -perturbed game.*

Denote by

$$\mathfrak{S}^\epsilon := \{ \mathcal{S} : \mathcal{S} \text{ is a feasible, non-}a\text{-degenerate shortest-path-support for offsets } \mathbf{b} + \epsilon \}$$

the support states of the game with perturbed offsets  $\mathbf{b} + \epsilon$ . Then, by Theorem 8.1(i), we have that  $\mathfrak{S}^\epsilon \subseteq \mathfrak{S}$ .

Since the perturbed game is non- $b$ -degenerate, it has a unique initial support for the zero flow as well as well-defined predecessor and successor functions. Since  $\mathfrak{S}^\varepsilon \subseteq \mathfrak{S}$  we can also use these functions for the unperturbed, degenerate game.

Note that Theorem 8.1 also implies that, for every game with offset vector  $\mathbf{b}$  that is in general position, the game is non- $b$ -degenerate.

**Implicit perturbation.** In order to find the continuative neighbors of some support of the  $\varepsilon$ -perturbed game it is not necessary to actually compute the values  $\lambda_S^{\min,\varepsilon}$  and  $\lambda_S^{\max,\varepsilon}$  explicitly for some small  $\varepsilon$ . It can be shown that there are vectors  $\mathbf{m}_{S,e,i}$  such that minimizer and maximizer in the computation of  $\lambda_S^{\min,\varepsilon}$  and  $\lambda_S^{\max,\varepsilon}$  can be determined by a lexicographic comparison of the vectors  $\mathbf{m}_{S,e,i}$ . Hence, it is not necessary to compute these values explicitly for some small  $\varepsilon > 0$ .

**$a$ - $b$ -degenerate games.** Finally, for the case of  $a$ - $b$ -degenerate supports, we note there is an analogous statement of Theorem 8.1 using the values  $\xi_S^{\min}$  and  $\xi_S^{\max}$  implying that  $a$ - $b$ -degeneracy does not occur in a  $\varepsilon$ -perturbed game. For more details, we refer to the full version of this paper [33].

## References

- [1] H. Ackermann, H. Röglin, and B. Vöcking. On the impact of combinatorial structure on congestion games. *J. ACM*, 55(6):25:1–25:22, 2008.
- [2] I. Adler and S. Verma. The linear complementarity problem, Lemke algorithm, perturbation, and the complexity class PPAD. [http://www.optimization-online.org/DB\\_HTML/2011/03/2948.html](http://www.optimization-online.org/DB_HTML/2011/03/2948.html), 2011.
- [3] E. Altman, T. Başar, T. Jiménez, and N. Shimkin. Competitive routing in networks with polynomial costs. *IEEE Trans. Automat. Contr.*, 47:92–96, 2002.
- [4] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, New Haven, CT, 1956.
- [5] U. Bhaskar, L. Fleischer, D. Hoy, and C.-C. Huang. On the uniqueness of equilibrium in atomic splittable routing games. *Math. Oper. Res.*, 40:634–654, 2015.
- [6] U. Bhaskar, L. Fleischer, and C.-C. Huang. The price of collusion in series-parallel networks. In F. Eisenbrand and F. Shepherd, editors, *Proc. 14th Internat. Conf. Integer Program. and Comb. Optim. (IPCO)*, pages 313–326, 2010.
- [7] U. Bhaskar and P. R. Lolakapuri. Equilibrium computation in atomic splittable routing games. In Y. Azar, H. Bast, and G. Herman, editors, *Proc. 26th Annu. Europ. Symp. on Algorithms (ESA)*, volume 112 of *LIPICs*, pages 58:1–58:14, 2018.
- [8] S. Catoni and S. Pallottino. Technical note – Traffic equilibrium paradoxes. *Transp. Sci.*, 25:240–244, 1991.
- [9] X. Chen, X. Deng, and S.-H. Teng. Settling the complexity of computing two-player Nash equilibria. *J. ACM*, 56(3):14:1–14:57, 2009.
- [10] X. Chen, S.-H. Teng, and P. Valiant. The approximation complexity of win-lose games. In *Proc. 18th ACM-SIAM Sympos. on Discrete Algorithms (SODA)*, pages 159–168, 2007.
- [11] R. Cominetti, J. R. Correa, and N. E. Stier-Moses. The impact of oligopolistic competition in networks. *Oper. Res.*, 57:1421–1437, 2009.
- [12] R. Cottle, J.-S. Pang, and R. Stone. *The Linear Complementarity Problem*. Academic-Press, 1992.
- [13] S. Dafermos. The traffic assignment problem for multiclass-user transportation networks. *Transp. Sci.*, 6:73–87, 1972.
- [14] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a Nash equilibrium. *SIAM J. Comput.*, 39(1):195–259, 2009.
- [15] J. Dunkel and A. S. Schulz. On the complexity of pure-strategy Nash equilibria in congestion and local-effect games. *Math. Oper. Res.*, 33:851–868, 2008.
- [16] B. C. Eaves. The linear complementarity problem. *Management Sci.*, 17:612–634, 1971.
- [17] A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure Nash equilibria. In *Proc. 36th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 604–612, 2004.
- [18] M. Gairing and M. Klimm. Congestion games with player-specific costs revisited. In B. Vöcking, editor, *Proc. 6th Proc. Internat. Sympos. Algorithmic Game Theory (SAGT)*, volume 8146 of *Lecture Notes in Computer Science*, 2013.
- [19] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proc. 46th Annu. IEEE Sympos. Foundations Comput. Sci. (FOCS)*, pages 142–151, 2005.
- [20] P. W. Goldberg, C. H. Papadimitriou, and R. Savani. The complexity of the homotopy method, equilibrium selection, and Lemke-Howson solutions. *ACM Trans. Econ. Comput.*, 1, 2013.
- [21] T. Harks. Stackelberg strategies and collusion in network games with splittable flow. *Theory Comput. Syst.*, 48:781–802, 2011.
- [22] T. Harks and M. Klimm. On the existence of pure Nash equilibria in weighted congestion games. *Math. Oper. Res.*, 37(3):419–436, 2012.
- [23] T. Harks and V. Timmermans. Equilibrium computation in atomic splittable singleton congestion games. In F. Eisenbrand and J. Koenemann, editors, *Proc. 19th Internat. Conf. Integer Program. and Comb. Optim. (IPCO)*, volume 10328 of *LNCS*, pages 442–454, 2017.
- [24] T. Harks and V. Timmermans. Uniqueness of equilibria in atomic splittable polymatroid congestion games. *J. Comb. Optim.*, 36:812–830, 2018.
- [25] A. Haurie and P. Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15:295–308, 1985.
- [26] A. Hayrapetyan, É. Tardos, and T. Wexler. The effect

- of collusion in congestion games. In *Proc. 38th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 89–98, 2006.
- [27] P. J.-J. Herings and R. Peeters. Homotopy methods to compute equilibria in game theory. *Econ. Theory*, 42:119–156, 2010.
- [28] P. J.-J. Herings and A. H. van den Elzen. Computation of the Nash equilibrium selected by the tracing procedure in  $n$ -person games. *Games Econ. Behav.*, 38:89–117, 2002.
- [29] C.-C. Huang. Collusion in atomic splittable routing games. *Theory Comput. Syst.*, 52:763–801, 2013.
- [30] D. S. Johnson, C. H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Comput. Syst. Sci.*, 37:79–100, 1988.
- [31] S. Kakutani. A generalization of Brouwer’s fixed point theorem. *Duke Mathematical Journal*, 8:457–459, 1941.
- [32] J. Katzenelson. An algorithm for solving nonlinear resistor networks. *The Bell System Technical Journal*, 44:1605–1620, 1965.
- [33] M. Klimm and P. Warode. Complexity and parametric computation of equilibria in atomic splittable congestion games via weighted block Laplacians. 2019. arXiv:1811.08354.
- [34] M. Klimm and P. Warode. Computing all Wardrop equilibria parametrized by the flow demand. In *Proc. 30th ACM-SIAM Sympos. on Discrete Algorithms (SODA)*, pages 917–934, 2019.
- [35] C. E. Lemke and J. T. Howson, Jr. Equilibrium points of bimatrix games. *SIAM J. Appl. Math.*, 12:413–423, 1964.
- [36] P. Marcotte. Algorithms for the network oligopoly problem. *J. Oper. Res. Soc.*, 38:1051–1065, 1987.
- [37] R. Mehta. Constant rank bimatrix games are PPA-hard. In *Proc. 46th Annu. ACM Sympos. Theory Comput. (STOC)*, pages 545–554, 2014.
- [38] F. Meunier and T. Pradeau. Computing solutions of the multiclass network equilibrium problem with affine cost functions. *Ann. Oper. Res.*, 274:447–469, 2019.
- [39] I. Milchtaich. Congestion games with player-specific payoff functions. *Games Econ. Behav.*, 13:111–124, 1996.
- [40] I. Milchtaich. Topological conditions for uniqueness of equilibrium in networks. *Math. Oper. Res.*, 30:225–244, 2005.
- [41] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Trans. Netw.*, 1:510–521, 1993.
- [42] C. Papadimitriou. On the complexity of the parity argument and other inefficient proofs of existence. *J. Comput. Syst. Sci.*, 48:498–532, 1994.
- [43] O. Richman and N. Shimkin. Topological uniqueness of the Nash equilibrium for selfish routing with atomic users. *Math. Oper. Res.*, 32:215–232, 2007.
- [44] J. B. Rosen. Existence and uniqueness of equilibrium points for concave  $n$ -person games. *Econometrica*, 33:520–534, 1965.
- [45] R. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Internat. J. Game Theory*, 2:65–67, 1973.
- [46] T. Roughgarden and F. Schoppmann. Local smoothness and the price of anarchy in splittable congestion games. *J. Econ. Theory*, 156:317–342, 2015.
- [47] Sandvine. The global Internet phenomena report. <https://www.sandvine.com/hubfs/downloads/phenomena/2018-phenomena-report.pdf>, October 2018.
- [48] D. Schmeidler. Equilibrium points on nonatomic games. *J. Statist. Phys.*, 7:295–300, 1970.
- [49] J. G. Wardrop. Some theoretical aspects of road traffic research. *Proceedings of the Institution of Civil Engineers Part II*, 1:325–378, 1952.