Uniqueness and Computation of Equilibria in Resource Allocation Games

Veerle Timmermans

UNIQUENESS AND COMPUTATION OF EQUILIBRIA IN RESOURCE ALLOCATION GAMES

VEERLE TIMMERMANS

Cover design by Stijn Timmermans.

This book was typeset by the author using LaTeX and the classicthesis package.

Published by Universitaire Pers Maastricht
ISBN: 978 94 6380 173 7
Printed in The Netherlands by Datawyse Maastricht

UNIQUENESS AND COMPUTATION OF EQUILIBRIA IN RESOURCE ALLOCATION GAMES

DISSERTATION

to obtain the degree of Doctor at Maastricht University,
on the authority of the Rector Magnificus,
Prof. dr. Rianne M. Letschert,
in accordance with the decision of the Board of Deans,
to be defended in public
on Friday, 1st of February 2019, at 10:00 hours

by

Veerle Marloes Timmermans

First Promotor:
Prof. dr. T. Harks (Augsburg University)

Second Promotor:
Prof. dr. ir. C.P.M. van Hoesel (Maastricht University)

Assessment Committee:
Prof. dr. A.J. Vermeulen (chair, Maastricht University)
Prof. dr. R.J. Müller (Maastricht University)
Prof. dr. B. Peis (RWTH Aachen)
Prof. dr. G. Schäfer (CWI Amsterdam)

This work is part of the research programme *Optimal Coordination Mechanisms for Distributed Resource Allocation* with project number 617.001.302, which is (partly) financed by the Netherlands Organisation for Scientific Research (NWO). This research is also financially supported by the Graduate School of Business and Economics (GSBE).

To my husband and my family, for their unlimited love and support.

ACKNOWLEDGEMENTS

This thesis is the result of a journey that started in 2014, at Maastricht University. There have been many people that have walked alongside me during the last four years and I'd like to take this opportunity to express my gratitude to some of them in particular.

First, and foremost, I would like to express my appriciation for my supervisor Tobias Harks. Thank you for offering me this PhD-position. With my background in theoretical mathematics I was fairly new to the field of algorithmic game theory. You were patient, and gave me the books and papers I needed to expand my expertise. You shared your knowledge, your skills and your network with me, and, though I still have a lot to learn, you handed me all the tools I needed to become an independent researcher. It is you who encouraged me to pursue a career in academia. And, I am convinced that all the papers I wrote under your supervision are just the start of many more collaborations.

Then, I would like to thank my promotor Stan van Hoesel for having me as his PhD student, and of course for reading and approving the thesis. I would also like to thank all the members of my assessment committee, Rudolf Müller, Britta Peis, Dries Vermeulen and Guido Schäfer, for taking the time to read my dissertation over the summer holiday.

I'd also like to take a moment to thank my paranymphs, Tim en Vincent. Tim, I remember that on my first day, you immediately invited me for a dinner and game-night at your place. You made me feel welcome, and we became very good friends almost instantly. Your kindness and hard-working attitude have been an inspiration. Vincent, we shared our office, and our supervisor, which kind of makes us research-siblings. Your enthusiasm for teaching was contagious, and a good example for us all.

Also a big thank you to the whole QE department, with Karin and Yolanda at its heart. To both of you, you have been incredibly kind and helpful during these last four years. Tjark, you were the supervisor of my master thesis, which we rewrote and published at WAOA. After that, you helped me to find a place within Maastricht University, for which I am incredibly grateful. And, especially after Tobias moved to Augsburg, I always felt as if you were looking out for me. Thank you. Dries, we taught the same course for four years and a row. Though the material itself was easy,

it was the diversity of the students and that made the course a bit more challenging to tutor. Alex, I enjoyed our project together, your enthusiasm for problems the and the fact that your door is always open. Thank you János, for the time we spend in the cinema, and all the dinners and drinks.

Then, there are the many PhD students that made work even better. I'd like to thank a couple of them in particular. First of all, Anna, I loved having you around. We always agreed that balance is the key of a happy PhD life. For us that meant going to the gym, and making oreo muffins after. Martijn, thank you for all the tea breaks, cinema visits and Zelda fights. Julian, I am happy we got to share some beers in Rotterdam, and work on a more practical project together. Thank you Aditya, Benoit and Roland for all the time we spend in the cinema, and Aïda, Catherina, Dewi, Hanno, Li, Shash, Verena and Yuliya for the tea breaks, games and lunches we shared.

In my last year as a PhD, I also regularly visited RWTH Aachen, where I felt at home immediately. A special thanks to Britta Peis. Britta, I am very grateful for all advice you gave me in the past, and opportunity to work at RWTH Aachen in the future. I enjoyed being your co-author and I am looking forward to collaborating with you in the coming years. Thank you Björn and Laura, for introducing me to the interesting research field of flows over time. And also a big thanks to Daniël, Dina, Marc, Martin and Vipin for making me feel welcome and appreciated.

During my PhD, I had the pleasure to visit many conferences, workshops and talks. I went to Canada, the USA, Greece, Belgium, Germany, Switzerland, Italy and though traveling on your own doesn't sound ideal, there were plenty of people that made sure I never felt lonely. Thank you Anja Huber en Manuel Surek, for showing me around in Augsburg and sharing (too) many cocktails. And lets not forget Loe Schlicher, Martijn van Ee, Miriam Schöter, Bart de Keijzer, Pieter Kleer and Leon Sering for the many games we played during the conferences.

Then, there are also the people who supported me not at work, but kept me sane during these past four years. Thank you, Lenie, you are one of my best and oldest friends. I love our hikes, cooking adventures, laughter, games and honest conversations. Anneloes, my LinkdIn page would have been empty without you. Your braveness inspired me to speak to people I would have been too shy to talk to otherwise (gotta catch'em all). Paul and Saskia, thank you for being worthy opponents in the many, many games we played and all the weekends, dinners and trips we made. Patrique and Pascal, thank you for all the fridays we spend baking, playing games and making foolish word puns (the more the butter). Bart, Floortje, Corine, Job,

Rianne and Killian, thank you for the dinners, parties, phone calls and teaching jokes.

This brings me to the closest and most important group of people: my family. Willeke, Noud, Peter en Els, thank you for all your love and support, that made me the person who I am today. Stijn, thank you for designing this gorgeous cover and all the philosophical advise. Ilse, thank you for sharing all my childhood memories and grown up dreams, and let's not forget all time you spend on doing my hair and make-up. *Ohana*.

At last, a word of thanks towards my husband, Rupert Tan. You left your job, friends and family behind, to build this new life with me in Maastricht. You let me grow, and as a result, we grew together. You are the most warm, loving and caring person I know, and you always make me feel strong and confident. In your eyes, I am a hero that could solve any problem I encounter. And, throughout this journey, you have been mine.

CONTENTS

1 INTRODUCTION

Game theory is the study of conflict and cooperation between rational decision makers, or *players*, within a competitive environment. It offers a framework in which the strategic interaction between players can be modelled, which is used to understand the behavior of these players in practice. Game theory comprises two branches: *cooperative* game theory and *non-cooperative* game theory. Cooperative game theory analyzes situations where enforceable, binding agreements between the players are possible. In such situations, the main issue is to find a reasonable redistribution of the joint revenues. Non-cooperative games do not allow for these binding contracts, and individual incentives play a prominent role. In this thesis we focus on non-cooperative games. The essential elements of such a game are its players, their available strategies and the pay-offs that assign a value to each player for each combination of strategies.

Definition 1.0.1. *A* non-cooperative strategic game $\mathcal{G} := (N, (X_i)_{i \in N}, (u_i)_{i \in N})$ *consists of a set of players* N*, a set of feasible strategies* X_i *for each player* $i \in N$ *and a pay-off function* $u_i : \times_{i \in N} X_i \to \mathbb{R}$ *for each player* $i \in N$*. The goal of each player is to selfishly maximize her pay-off.*

A collection of strategies that maximizes the *social cost*, i.e., the aggregated pay-off, is called a *social optimum*. One of the most important tools that game theorists have at their disposal is the *Nash equilibrium*, named after John Forbes Nash Jr. [54]. A Nash equilibrium is a set of strategies that players act out, with the property that no player benefits from an unilateral deviation. Intuitively, this means that if a player is told the strategies of all her opponents, she would still choose to retain her original strategy. A version of this equilibrium concept was already used in 1838 by Cournot [20], who developed game models of oligopolistic competition. In Cournot's theory, firms choose how much output they need to produce in order to maximize their own profit.

The modern game-theoretic concept of a Nash equilibrium is broader than Cournot's. In general, it is defined in mixed strategies, where players choose a probability distribution over their available strategies instead of one available strategy. A Nash equilibrium in mixed strategies is called a *mixed Nash equilibrium*, and using Kakutani's fixed-point theorem [40] Nash proved that a mixed Nash equilibrium exists in each finite-player strategic game. An equilibrium in pure strategies is called a *pure Nash equilibrium*, and does not exist in each strategic game.

The three papers discussed in this thesis all belong to the field of *algorithmic game theory*, and address questions regarding the computational complexity of constructing pure Nash equilibria.

1.1 ALGORITHMIC GAME THEORY

This relatively new field lies in the intersection of game theory and computer science, and concerns itself with the design and analysis of algorithms in strategic environments. In doing so, it deals with perhaps the most fundamental discrepancy between computer science and game theory: the latter studies (the dynamics of) competitive behavior, but disregards the computational complexity issues that arise when computing equilibrium states - which would be the foremost concern of any computer scientist. As pointed out by several researchers (e.g. [16, 21]), the computational tractability of a solution concept contributes to its credibility as a plausible prediction of the outcome of competitive environments in practice. Thus, the computation of equilibria has been one of the earliest research goals of algorithmic game theory.

More general, algorithmic game theory concerns itself with issues regarding the computational complexity of algorithms for strategic games. On the one hand, it analyzes existing games: it studies algorithms that compute or approximate equilibria, and determines if these equilibria are unique and efficient. On the other hand, it studies the design of strategic games that are guaranteed to have 'good' equilibria: equilibria that can be computed efficiently and have a social cost that is close to the social cost of a social optimum. It combines ideas from the classical economic mechanism design with algorithm design and computational complexity. This side of algorithmic game theory is called *algorithmic mechanism design*.

In this thesis, we focus on the uniqueness and computation of pure Nash equilibria in *resource allocation games*. In the rest of this thesis we use the term equilibrium as shortcut for pure Nash equilibrium.

1.2 RESOURCE ALLOCATION GAMES

A resource allocation game is a collective name for all strategic games where players compete over a set of resources. In the games we consider, we assume that all players have *full information* about the game, i.e., the demand, strategy space and pay-off function of each player are publicly known. Furthermore, the games we consider are *one-shot simultaneous move* games, as all players simultaneously choose an action from their available strategies. Among others, resource allocation games contain the well-known classes of *congestion games* and *multimarket oligopolies*.

1.2.1 *Congestion Games*

UNSPLITTABLE CONGESTION GAMES Congestion games as introduced in Rosenthal [64] constitute an elegant game-theoretic model describing the distributed allocation of resources among selfish players. Specifically, such a game comprises a finite set of players N, a finite set of resources E and the pure strategies of a player are given by a set of allowable subsets of resources $\mathcal{S}_i \subseteq 2^E$, where 2^E denotes the power set of E. In the context of *network games*, the resources may correspond to edges of a graph and the allowable subsets correspond to paths connecting a source and a sink. Resources have cost functions $c_e : \mathbb{N} \to \mathbb{R}_{\geqslant 0}$ that depend on the number of players currently using the resource. We denote the strategy of player $i \in N$ by x_i and x denotes a collection of pure strategies x_i of the players. For a given strategy profile x, the disutility of each player is just the sum of the resource costs of the chosen subset of resources:

$$\pi_i(x) := \sum_{e \in E} c_e(d_e(x)),$$

where $d_e(x)$ denotes the number of players that use resource e in their chosen strategy x_i. Rosenthal proved in his seminal paper that congestion games always admit a pure Nash equilibrium. Nowadays, we refer to this classical model as an *unsplittable congestion game*.

Example 1.2.1 (Unsplittable Congestion Game)**.** Consider the unsplittable *network* congestion game where two players need to choose an s-t path in graph G_1 depicted in Figure 1.1. The cost of using a specific edge is depicted next to the edge, where x stands for the number of players that pick this edge in their s-t path. In the feasible strategies depicted in Figure 1.2

and Figure 1.3, player 1 has a total cost of $0+1+0=1$ and player 2 a total cost of $2+0=2$. As no player can strictly decrease her cost, these strategies form a pure Nash equilibrium.

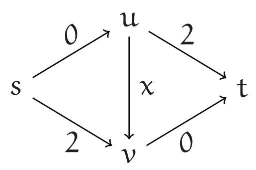

Figure 1.1: Graph G_1.

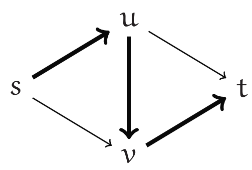

Figure 1.2: Player 1.

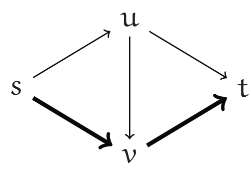

Figure 1.3: Player 2.

In this classical model all players have an equal weight, and such games are known to be *potential games*. A game is said to be a potential game if the incentive of all players to change their strategy can be expressed using a single global function called the *potential function*, where pure Nash equilibria are the local minima of this function. For example, for unweighted, unsplittable congestion games a potential function is:

$$\Phi(x) := \sum_{e \in E} \sum_{k=1}^{d_e(x)} c_e(k).$$

If one would assign a weight to each player, the game is called a *weighted* unsplittable congestion game. For weighted congestion games, the cost of using a resource does not depend on the number of players using this resource, but on the total weight of all players using this resource instead. a weighted unsplittable congestion game is not a potential game, and a pure Nash equilibrium might not exist.

NONATOMIC CONGESTION GAMES Since the initial work of Rosenthal, several works studied related or generalized variants of congestion games, starting with *nonatomic* congestion games, introduced in 1973 by Schmeidler [66]. In these games, there is an infinite number of players that are all infinitesimally small. Hence, no player can individually influence the game. Instead of having a set N of players, N now represents a set of different *player types*. Each type $i \in N$ is then represented by a continuum of $[0, d_i]$ players, and all players of the same type have the same available strategies. The cost of a resource depends on the fraction of players using it.

Example 1.2.2 (Nonatomic Congestion Game)**.** Consider the nonatomic network congestion game where we have players of only one type, and the continuum of players is represented by the interval $[0, 1]$. Similar to Example 1.2.1, each player needs to choose an s-t path in graph G_2 depicted in Figure 1.4. In an equilibrium state all players will take the lower path, resulting in a cost of 1 for each player. In a social optimum, the players will divide equally over both edges, resulting in an average cost of $\frac{3}{4}$. This famous example by Pigou [60] maximizes the ratio between the social cost in a Nash equilibrium and the social cost in a social optimum ($\frac{4}{3}$ here), also known as the *price of anarchy* [43].

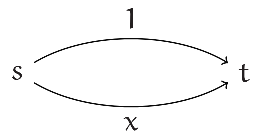

Figure 1.4: Graph G_2.

For non-atomic congestion games the equilibrium strategy profile is also called a Wardrop equilibrium. Also for this setting, a potential function is known:

$$\sum_{e \in E} \int_0^{f_e} c_e(y)dy.$$

The existence of a potential function implies that each non-atomic game posesses at least one pure Nash equilibrium.

ATOMIC SPLITTABLE CONGESTION GAMES *Atomic splittable* congestion games resemble the nonatomic congestion games, but assume that all players in one type cooperate in order to minimize the aggregated cost of all players in this type. For example, assume that multiple oil companies send oil through the same system of pipes. Then, all infinitesimally small oil drops owned by the same company might 'cooperate' in order to reduce their average flow time or cost. In such games, one could also say that N is a set of players, where each player has a positive *demand* d_i that she can fractionally distribute over her available strategies.

Example 1.2.3 (Atomic Splittable Congestion Game)**.** Consider the atomic splittable network congestion game where we have two players with demands $d_1 = d_2 = \frac{1}{2}$. Again, each player needs to fractionally distribute her demand over the s-t paths in graph G_2 depicted in Figure 1.4. In an equilibrium state both players send $\frac{1}{3}$ over the lower path, and the remaining $\frac{1}{6}$ over the upper path, resulting in a cost of $\frac{2}{3} \cdot \frac{1}{3} + 1 \cdot \frac{1}{6} = \frac{7}{18}$ for each player, and a social cost of $\frac{14}{18}$. Each solution where the total load on both edges is equal to $\frac{1}{2}$ is a social optimum, with a social cost of $\frac{3}{4}$.

Haurie and Marcotte [38] showed that classical nonatomic congestion games (cf. Beckmann et al. [7] and Wardrop [75]) can be modeled as atomic splittable congestion games by constructing a sequence of games and taking the limit with respect to the number of players. It follows that atomic splittable congestion games are strictly more general as their nonatomic counterpart.

INTEGRAL SPLITTABLE CONGESTION GAMES Lastly, an *integral splittable* congestion game is similar to an atomic splittable congestion game, but players can only split their demand in multiples of some common 'packet size'.

Example 1.2.4 (Integral Splittable Congestion Game)**.** Consider the integral splittable network congestion game where we have two players with demands $d_1 = d_2 = \frac{1}{2}$, and a packet size of $\frac{1}{10}$. Again, each player needs to divide her demand over the s-t paths in graph G_2 depicted in Figure 1.4, but can only divide her demand in multiples of $\frac{1}{10}$. In an equilibrium state both players send $\frac{3}{10}$ over the lower path, and the remaining $\frac{2}{10}$ over the upper path, resulting in a cost of $\frac{6}{10} \cdot \frac{3}{10} + 1 \cdot \frac{2}{10} = \frac{38}{100}$ for each player, and a social cost of $\frac{76}{100}$. Again, each solution were the total load on both edges is equal to $\frac{1}{2}$ is a social optimum, with a social cost of $\frac{3}{4}$.

There are integral splittable congestion games that do not posess a pure Nash equilibrium. Though, in some special cases equilibria do exist. For example, when the strategy set of each player is the set of bases of a matroid.

In this subsection we introduced the basic setting for four different types of congestion games. In this thesis we mostly study the atomic splittable congestion games. Hence, in the next section we introduce these games formally, and in their most general form.

1.2.2 *Atomic Splittable Congestion Games*

Atomic splittable (network) congestion games have first been proposed by Orda et al. [59] in the context of modelling routing in communication networks. They can be compactly represented as:

$$\mathcal{G} = (N, E, (\mathcal{S}_i)_{i \in N}, (d_i)_{i \in N}, (c_{i,e})_{i \in N, e \in E})\,.$$

Here, E is a finite, non-empty set of resources and N a finite, non-empty set of players. Each player $i \in N$ is associated with a demand $d_i \geqslant 0$ and a collection of allowable subsets of resources $\mathcal{S}_i \subseteq 2^E$. A strategy for player $i \in N$ is then a (possibly fractional) distribution $x_i \in \mathbb{R}^{|\mathcal{S}_i|}_{\geqslant 0}$ of the demand over the allowable subsets $S \in \mathcal{S}_i$. Thus, one can compactly represent the strategy space of every player $i \in N$ by the following polytope:

$$P_i := \left\{ x_i \in \mathbb{R}^{|\mathcal{S}_i|}_{\geqslant 0} \mid \sum_{S \in \mathcal{S}_i} x_{i,S} = d_i \right\}.$$

We denote by $x = (x_i)_{i \in N}$ the overall strategy profile. The induced load under x_i at e is defined as $x_{i,e} := \sum_{S \in \mathcal{S}_i : e \in S} x_{i,S}$ and the total load on e is then given as $x_e := \sum_{i \in N} x_{i,e}$. Resources have player-specific cost functions $c_{i,e} : \mathbb{R}_{\geqslant 0} \to \mathbb{R}_{\geqslant 0}$ which are assumed to be non-negative, increasing, differentiable and convex. The total cost of player i in strategy profile x is defined as:

$$\pi_i(x) = \sum_{e \in E} c_{i,e}(x_e)\, x_{i,e}.$$

Each player aims to minimize her private disutility $\pi_i(x)$. For $i \in N$, we write $\mathcal{S}_{-i}(d_{-i}) = \times_{j \neq i} \mathcal{S}_j(d_j)$ and $x = (x_i, x_{-i})$ meaning that $x_i \in \mathcal{S}_i(d_i)$ and $x_{-i} \in \mathcal{S}_{-i}(d_{-i})$. In this context, a Nash equilibrium is defined as follows.

Definition 1.2.5. *A strategy profile x is a Nash equilibrium if and only if for all $i \in N$ and $y_i \in \mathcal{S}_i(d_i)$:*

$$\pi_i(x) \leqslant \pi_i(y_i, x_{-i}).$$

Using that the strategy space is compact and cost functions are increasing and convex, Kakutani's fixed-point theorem implies the existence of a Nash equilibrium [63, Theorem 1].

1.2.3 *Cournot Competition and Multimarket Oligopolies*

The Cournot Competition model also falls into the class of resource allocation games. In the basic model of Cournot [20] introduced in 1838, firms produce homogeneous goods and sell them in a common market. The selling price of the goods depends on the total quantity of goods that is offered in the market. Each firm aims to maximize its profit, which is equal to the revenue minus the production costs. In a *multimarket oligopoly* (cf. Bulow [13]), firms compete over a *set* of markets and each firm has access to a firm-specific subset of the markets. In their most general form, these games are represented by the tuple:

$$\mathcal{M} = (N, E, (E_i)_{i\in N}, (p_{i,e})_{i\in N, e\in E_i}, (C_i)_{i\in N}),$$

where N is a finite, non-empty set of n firms and E a finite, non-empty set of m markets. Each firm i only has access to a subset $E_i \subseteq E$ of the markets. Each market $e \in E$ is endowed with firm-specific, non-increasing price functions $p_{i,e}(t), i \in N$. In a strategy profile, a firm chooses a non-negative production quantity $x_{i,e} \in \mathbb{R}_{\geqslant 0}$ for each market $e \in E_i$. We denote a strategy profile for a firm by $x_i = (x_{i,e})_{e\in E_i}$, and a joint strategy profile by $x = (x_i)_{i\in N}$. The functions $C_i(\cdot)$ represent the production costs of player i, which depend on the total quantity player i produces and they are assumed to be increasing. The goal of each firm $i \in N$ is to maximize its utility, which is given by:

$$u_i(x) = \sum_{e\in E_i} p_{i,e}(x_e)x_{i,e} - C_i\Big(\sum_{e\in E_i} x_{i,e}\Big),$$

where $x_e := \sum_{i\in N} x_{i,e}$. The existence of equilibria in single market Cournot models (beyond quasi-polynomial utility functions) has been studied extensively in the past decades (see Vives [74] for a good survey).

1.3 OUTLINE OF THE THESIS

This section contains a detailed outline of the remaining chapters of this thesis.

1.3.1 *Uniqueness of Equilibria*

In Chapter 2 we study uniqueness of Nash equilibria in atomic splittable congestion games. This property is key to actually predict the outcome of distributed resource allocation: if there are multiple equilibria it is not clear upfront which equilibrium will be selected by the players. This issue has been raised explicitly by Aumann [6]: "...it is by no means clear how the players would arrive at an equilibrium, why they should play equilibrium strategies, and how a specific equilibrium would be chosen from among the set of all equilibria."

As our main result we give a sufficient condition for uniqueness based on the theory of polymatroids. We show that if the strategy space of every player is a polymatroid base polytope satisfying a special exchange property – we term this class of polymatroids *bidirectional flow polymatroids* – the equilibria are unique. We demonstrate that bidirectional flow polymatroids are quite general as they contain base-orderable matroids, gammoids, transversal and laminar matroids.

We complement our uniqueness result by showing that multiple equilibria exist when certain assumptions are dropped. We consider a game with at least three players for which the set systems $\mathcal{S}_i$ of all players $i \in N$ are *not* bases of a matroid. Then, there exists a game isomorphic to it which admits multiple equilibria. Here, the term *isomorphic* means that there is no a priori description on how the individual strategy spaces of players interweave in the ground set of resources. Our results leave a gap between general matroids and base orderable matroids, for which we do not know whether or not equilibria are unique.

We also study the uniqueness of equilibria if the set systems $\mathcal{S}_i$ correspond to paths in an undirected graph. The instance used for showing multiplicity of equilibria of non-matroid games can be seen as a 3-player game played on an undirected 3-vertex cycle graph. From this we can derive a new characterization of uniqueness of equilibria in undirected graphs. If we assume at least three players and if we do not specify beforehand which vertices of the graph serve as sources or sinks, an undirected graph induces unique equilibria if and only if the graph has no cycle of length at least 3.

Parts of Chapter 2 are based on [34].

1.3.2 *Equilibrium Computation for Games with Affine Costs*

In Chapter 3, we devise the first polynomial time algorithm computing a pure Nash equilibrium for atomic splittable congestion games with singleton strategies and player-specific affine cost functions. Our algorithm is purely combinatorial and computes the *exact* equilibrium assuming rational input. The idea is to compute a pure Nash equilibrium for an associated integral splittable singleton congestion game. While integral games have been considered in the literature before, no polynomial time algorithm computing an equilibrium was known. Also for this class, we devise the first polynomial time algorithm and use it as a building block for our main algorithm.

Polynomial running time of the algorithm is shown by several structural results on the sensitivity of integral splittable equilibria with respect to different packet sizes. Specifically, when halving the packet size, we derive bounds on the difference of the resulting global load vectors as well as the individual load vectors of players. The main idea of our algorithm is to first compute a Nash equilibrium for a large packet size, and then iteratively halving the packet size and recomputing the equilibrium by a sequence of best responses. We use the sensitivity results to show that best responses converge to an equilibrium within a polynomial number of steps.

Then, we develop a polynomial time computable transformation mapping a multimarket Cournot competition game with firm-specific affine price functions and quadratic costs to an associated atomic splittable congestion game. The transformation preserves equilibria in either games and, thus, leads – via our first algorithm – to a polynomial time algorithm computing Cournot equilibria. Finally, our analysis for integral splittable games implies new bounds on the difference between real and integral Cournot equilibria. The bounds can be seen as a generalization of the recent bounds for single market oligopolies obtained by Todd [69] towards multimarkets.

Chapter 3 is based on [37].

1.3.3 *Equilibrium Computation for Games with Convex Costs*

In Chapter 4 we construct Nash equilibria in atomic splittable congestion games with convex cost functions, where the strategy space of each player is the base polytope of a polymatroid. As equilibria in these games are not guaranteed to be rational, we look for ϵ-approximate equilibria. Here, we say a strategy profile is an ϵ-approximate equilibrium when no player can deviate from her current strategy and decrease her cost by at least ϵ.

Again, the idea is to compute a pure Nash equilibrium for an associated integral splittable congestion game. It is known that one can compute pure Nash equilibria for integral splittable polymatroid congestion games with convex cost functions within a running time that is pseudo-polynomial in the aggregated demand of the players (see Harks et al. [32, 36]). In this paper we compute for each $\epsilon > 0$ a packet size k_ϵ such that the k_ϵ-integral equilibrium is guaranteed to be an ϵ-approximate equilibrium. Thus, by using the algorithm by Harks et al. [32, 36] for packet size k_ϵ, one can compute an ϵ-approximate equilibrium within pseudo-polynomial time.

We then consider multimarket oligopolies with decreasing, concave price functions and quadratic production costs, and prove that there exists a polynomial time transformation from atomic splittable congestion games to multimarket oligopolies. Using our first result, this implies that we are also able to find ϵ-approximate Cournot-Nash equilibria for multimarket oligopolies within pseudo-polynomial time. This result complements our results in Chapter 3 on multimarket oligopolies.

1.4 PUBLICATIONS

The chapters in this dissertation are based upon the following publications.

Published

- T. Harks and V. Timmermans. "Uniqueness of Equilibria in Atomic Splittable Polymatroid Congestion Games". In: *Journal of Combinatorial Optimization* (2017), pp. 1-19.
- T. Harks and V. Timmermans. "Computation of Equilibria in Atomic Splittable Singleton Games". In: *International Conference on Integer Programming and Combinatorial Optimization (IPCO 2017)*. 2017, pp. 442-454.

2 UNIQUENESS OF EQUILIBRIA

2.1 INTRODUCTION

An intriguing question in the field of atomic splittable congestion games is the possible non-uniqueness of equilibria. Let x and y be two equilibria. We say that x and y are different whenever there exists a player i and resource e such that $x_{i,e} \neq y_{i,e}$. A variant on this question is whether or not there exist multiple equilibria such that there exists at least one resource e for which $x_e \neq y_e$. We call this variant "uniqueness up to induced load on the resources".

For non-atomic players and network congestion games on directed graphs, Milchtaich [49] proved that Nash equilibria are not unique when cost functions are player-specific. Uniqueness is only guaranteed if the underlying graph is two terminal s-t-*nearly-parallel*. Richman and Shimkin [62] extended this result to hold for atomic splittable network games. Bhaskar et al. [9] looked at uniqueness up to induced load on the resources. They proved that even when all players experience the same cost on a resource, there can exist multiple equilibria. They further proved that for two players, the Nash equilibrium is unique if and only if the underlying undirected graph is generalized series-parallel. For multiple players of two types (players are of the same type if they have the same weight and share the same origin-destination pair), there is a unique equilibrium if and only if the underlying undirected graph is s-t-series-parallel. For more than two types of players, there is a unique equilibrium if and only if the underlying undirected graph is generalized nearly-parallel.

2.1.1 *Our Results and Techniques*

We consider atomic splittable congestion games as defined in Section 1.2.2 and study the uniqueness of equilibria for general set systems $(\mathcal{S}_i)_{i \in N}$. Interesting combinatorial structures of the $\mathcal{S}_i$'s beyond paths may be trees, forests, Steiner trees or tours all in a directed or undirected graph or bases of matroids.

As our main result we give a sufficient condition for uniqueness based on the theory of polymatroids. We show that if the strategy space of every player is a polymatroid base polytope satisfying a special exchange property – we term this class of polymatroids *bidirectional flow polymatroids* – the equilibria are unique. The formal definition of bidirectional flow polymatroids appears in Definition 2.3.3. We demonstrate that bidirectional flow polymatroids are quite general as they contain *base-orderable matroids, gammoids, transversal and laminar matroids*. For an overview of special cases that follow from our main result, see Figure 2.1.

The uniqueness result is stated in Section 2.4. In Section 2.5 we show that base-orderable matroids are a special case of bidirectional flow polymatroids. Definitions of polymatroid congestion games and bidirectional flow polymatroids are introduced in Section 2.2 and Section 2.3, respectively. In Section 2.6 and Section 2.7 we complement our uniqueness result by showing multiple equilibria exist when certain assumptions are dropped. In Section 2.6 we discuss why it is necessary for cost functions to be differentiable. In Section 2.7 we consider a game with at least three players for which the set systems $\mathcal{S}_i$ of all players $i \in N$ are *not* bases of a matroid. Then there exists a game with strategy spaces $\phi(\mathcal{S}_i)$ isomorphic to $\mathcal{S}_i$ which admits multiple equilibria. Here, the term *isomorphic* means that there is no a priori description on how the individual strategy spaces of players interweave in the ground set of resources. Our results leave a gap between general matroids and base orderable matroids for which we do not know whether or not equilibria are unique.

In Section 2.8 we consider uniqueness of equilibria if the set systems $\mathcal{S}_i$ correspond to paths in an undirected graph. The instance used for showing multiplicity of equilibria of non-matroid games can be seen as a 3-player game played on an undirected 3-vertex cycle graph. From this we can derive a new characterisation of uniqueness of equilibria in undirected graphs. If we assume at least three players and if we do not specify beforehand which vertices of the graph serve as sources or sinks, an undirected graph induces unique equilibria if and only if the graph has no cycle of length at least 3.

2.1.2 *Related Work*

Related work on atomic splittable congestion games in general can be found in Section 1.2.2 . Matroid congestion games were first considered by Ackermann et al. [2, 3]. They showed that (unsplittable) weighted con-

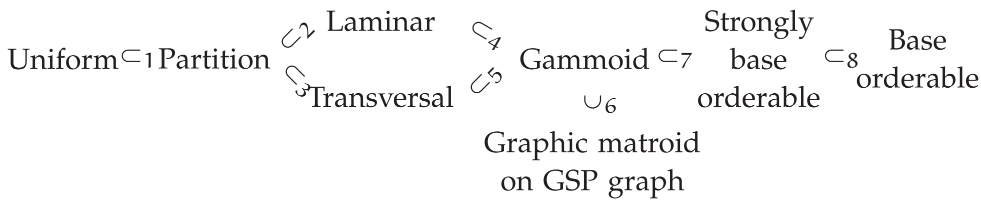

Figure 2.1: Several well-known classes of matroids and the relations between them. Here GSP is short for generalized series-parallel. References and arguments for the inclusions can be found in Section 2.9.

gestion games possess pure Nash equilibria even for player-specific non-decreasing cost functions. They also showed that the matroid property is the maximal property that gives rise to a pure Nash equilibrium, that is, for any strategy space not satisfying the matroid property, there is an instance of a weighted congestion game not having a pure Nash equilibrium. Integral polymatroid congestion games, a generalization of matroid congestion games, were later introduced in Harks, Klimm and Peis [36] (see also [33]). In addition, polymatroid theory was recently used in the context of nonatomic congestion games, where it is shown that matroid set systems are immune to the Braess paradox, see Fujishige et al. [27].

2.2 POLYMATROID CONGESTION GAMES

In polymatroid congestion games we assume that the strategy space for every player corresponds to a polymatroid base polytope.

In order to define polymatroids we first have to introduce submodular functions. A function $\rho : 2^E \to \mathbb{R}$ is called:

- *submodular* if $\rho(U) + \rho(V) \geqslant \rho(U \cup V) + \rho(U \cap V)$ for all $U, V \subseteq E$.
- *monotone* if $\rho(U) \leqslant \rho(V)$ for all $U \subseteq V$.
- *normalised* if $\rho(\emptyset) = 0$.

Given a submodular, monotone and normalised function ρ, the pair (E, ρ) is called a *polymatroid*. The associated *polymatroid base polytope* is defined as:

$$P_\rho := \left\{x \in \mathbb{R}^E_{\geqslant 0} \mid x(U) \leqslant \rho(U)\ \forall U \subseteq E,\ x(E) = \rho(E)\right\},$$

where $x(U) := \sum_{e \in U} x_e$ for all $U \subseteq E$.

In a polymatroid congestion game, we associate with every player i a player-specific polymatroid (E, ρ_i) and assume that the strategy space of player i is defined by the (player-specific) polymatroid base polytope P_{ρ_i}.

$$P_{\rho_i} := \left\{x_i \in \mathbb{R}^E_{\geqslant 0} \mid x_i(U) \leqslant \rho_i(U) \ \forall U \subseteq E,\ x_i(E) = \rho_i(E)\right\}.$$

From now on, when we mention a polymatroid congestion game, we mean an atomic splittable polymatroid congestion game. We give three examples of polymatroid congestion games:

Example 2.2.1 (Queueing Games (cf. [41]))**.** Let $Q = \{q_1, \dots q_m\}$ be a set of $M/M/1$ queues served in a first-come-first-served fashion and let $N = \{1, \dots, n\}$ be a set of companies who independently send packets with arrival rates $d_1, \dots, d_n$. Every queue q has a single server with exponentially distributed service time with mean $1/\mu_q$, where $\mu_q > 0$. Each packet is routed to a single server q out of a set of allowable queues, depending on the company. Given a distribution of packets $x \in \mathbb{R}^m_{\geqslant 0}$, the mean delay of queue q can be computed as $c_q(x_q) = \frac{1}{\mu_q - x_q}$. In this case the sets $\mathcal{S}_i$ are uniform rank-1 matroids, which are also called *singleton games*.

Example 2.2.2 (Transversal games)**.** Consider a finite set E of storing facilities, a finite set A of locations and a finite set N of players. Each player has to store an amount of d_i of divisible goods in each area $j \in A$. Each area j can be served from any storing facility within a given set $S_j \subseteq E$. The sets S_j may overlap, even for the same player i. However, due to reliability reasons, a player cannot store more than d_i goods in one storing facility. The cost $c_{i,e}$ for using a specific storing facility depends on the total amount of goods that have to be stored in storing facility e. The more goods need to be stored, the larger the cost to use it.

This setting can be modelled as a bipartite graph G on vertex sets E and A, where an edge between area $j \in A$ and storage facility $e \in E$ exists if and only if area j can be served from storage facility e. In a feasible strategy a player divides its goods over bases of the transversal matroid of this graph: subsets of storage facilities that are the endpoints of a maximal matching in G. Hence, the strategy space of every player $i \in N$ corresponds to the base polytope $P_{d_i \cdot rk_i}$, where rk_i is the rank function of a *transversal matroid*.

Example 2.2.3 (Matroid Congestion Games)**.** A *matroid* $\mathcal{M}$ is a pair $(E, \mathcal{I})$, where E is a finite set of resources and $\mathcal{I}$ is a family of subsets of E, called the *independent sets*. Set $\mathcal{I}$ has the following three properties:

1. The empty set is an independent set: $\emptyset \in \mathcal{I}$.

2. Set $\mathcal{I}$ is closed under taking subsets: if $I \subseteq J$ and $J \in \mathcal{I}$, then $I \in \mathcal{I}$.

3. Set $\mathcal{I}$ has the *exchange property*: if $I, J \in \mathcal{I}$ and $|I| < |J|$, then there exists an $e \in J$ such that $I \cup \{e\} \in \mathcal{I}$.

A *basis* is an independent set that becomes dependent on adding any element of E. The *base set* $\mathcal{B}$ contains all bases of $(E, \mathcal{I})$.

Consider an *atomic splittable matroid congestion model*, where for every $i \in N$ the allowable subsets are the base set $\mathcal{B}_i$ of a matroid $\mathcal{M}_i = (E, \mathcal{I}_i)$. The rank function $rk_i : 2^E \to \mathbb{R}$ of matroid $\mathcal{M}_i$ is defined as:

$$rk_i(S) := \max\{|U| \mid U \subseteq S \text{ and } U \in \mathcal{I}_i\},$$

for all $S \subseteq E$, and is submodular, monotone and normalised [61]. Moreover, the characteristic vectors of the bases in $\mathcal{B}_i$ are exactly the vertices of the polymatroid base polytope P_{rk_i}. It follows that the polytope P_i, where:

$$P_i := \left\{ x \in \mathbb{R}_{\geqslant 0}^{|\mathcal{B}_i|} \mid \sum_{B \in \mathcal{B}_i} x_B = d_i \right\},$$

corresponds to strategy distributions that lead to load vectors in the following polytope:

$$P_{d_i \cdot rk_i} = \left\{ x_i \in \mathbb{R}_{\geqslant 0}^{E} \mid x_i(U) \leqslant d_i \cdot rk_i(U) \ \forall U \subseteq E, x_i(E) = d_i \cdot rk_i(E) \right\}.$$

Hence matroid congestion models are a special case of polymatroid congestion models. Both the singleton games in Example 2.2.1 and the transversal games in Example 2.2.2 are a special case of matroid congestion games.

2.3 BIDIRECTIONAL FLOW POLYMATROIDS

We provide a sufficient condition for a class of polymatroid congestion games to have a unique Nash equilibrium. We prove that if the strategy space of every player is the base polytope of a *bidirectional flow polymatroid*, Nash equilibria are unique. In order to define the class of bidirectional flow polymatroids we first discuss some basic properties of polymatroids. We start with a generalization of the strong exchange property for matroids. Let $\chi_e \in \mathbb{Z}^{|E|}$ be the characteristic vector with $\chi_e(e) = 1$, and $\chi_e(e') = 0$ for all $e' \neq e$.

Lemma 2.3.1 (Strong exchange property polymatroids (Murota [52])). *Let P_ρ be a polymatroid base polytope defined on (E, ρ). Let $x, y \in P_\rho$ and suppose $x_e > y_e$ for some $e \in E$. Then there exists an $e' \in E$ with $x_{e'} < y_{e'}$ and an $\epsilon > 0$ such that:*

$$x + \epsilon(\chi_{e'} - \chi_e) \in P_\rho \text{ and } y + \epsilon(\chi_e - \chi_{e'}) \in P_\rho.$$

This exchange property will play an important role in the definition of bidirectional flow polymatroids. Given a strategy x in the base polytope of polymatroid (E, ρ), we are interested in the exchanges that can be made between x_e and $x_{e'}$ for some resources in $e, e' \in E$. For that, we define a directed exchange graph $D(x) = (E, V)$, where the set of vertices equals the set of resources E and the set of edges is defined by:

$$V := \left\{(e, e') | \exists\ \epsilon > 0 \text{ s.t. } x + \epsilon(\chi_{e'} - \chi_e) \in P_\rho\right\}.$$

We define exchange capacities $\hat{c}_x(e, e')$ (following notation of Fujishige [26]), which denotes the maximal amount of load that can be exchanged in x between resources e and e'. More formally:

$$\hat{c}_x(e, e') := \max\{\alpha | x + \alpha(\chi_{e'} - \chi_e) \in P_\rho\}.$$

We use Lemma 2.3.1 to prove the following:

Lemma 2.3.2. *Let P_ρ be a polymatroid base polytope defined on (E, ρ). Then, for $x, y \in P_\rho$, there exists a flow in $D(x)$ satisfying all supplies and demands, where a resource e with $x_e > y_e$ has supply of $x_e - y_e$ and e with $x_e < y_e$ has a demand of $y_e - x_e$.*

Proof. Consider Algorithm 1. Note that this algorithm is a slightly changed version of Fujishige [26, Theorem 3.27]. The only difference is that we do not change y to x with exchanges that only can be made on strategy y (which is proven in Fujishige [26, Theorem 3.27]) but with exchanges that can be executed on both x and y. As these exchanges always exists, the results by Fujishige [26, Theorem 3.27] are still valid for our algorithm. Hence, this algorithm transforms y into x with at most $\lfloor |E|^2/4 \rfloor$ elementary transformations described in Lemma 2.3.1, such that each component y_e with $y_e < x_e$ monotonically increases and each component y_e with $y_e > x_e$ monotonically decreases. Therefore f satisfies all supplies and demands as described in the lemma. Flow f also satisfies all capacity constraints, as every pair of resources (e, e') is considered at most once, and all exchanges can be done on x. Hence $f_{(e,e')} \leqslant \hat{c}_x(e, e')$, thus f is a flow in $D(x)$ satisfying all supplies and demands. ∎

Algorithm 1: Computing a flow in $D(x)$.

1. Let f be the *zero flow*, a flow where we send zero flow along all edges in $D(x)$.

2. If $x = y$, then stop and output flow f.

3. Choose any element $e \in E$ such that $x_e > y_e$.

4. Use Lemma 2.3.1 to find $e' \in E$ such that $x_{e'} < y_{e'}$ and $\epsilon > 0$ with:
$$x + \epsilon(\chi_{e'} - \chi_e) \in P_\rho \text{ and } y + \epsilon(\chi_e - \chi_{e'}) \in P_\rho.$$
Put $\alpha = \min\{\hat{c}_x(e, e'), \hat{c}_y(e', e), x_e - y_e, y_{e'} - x_{e'}\}$, define $y \leftarrow y + \alpha(\chi_e - \chi_{e'})$ and add α flow to edge (e, e') in flow f.

5. If $\alpha < x_e - y_e$, then go to step 4. Otherwise ($\alpha = x_e - y_e$), go to step 2.

The flow f mentioned in Lemma 2.3.2 is a flow from the perspective of strategy x and therefore we call this a *directed flow*. In the following we define a *bidirectional flow*. Let P_ρ again be a polymatroid base polytope on set E. For any $x, y \in P_\rho$ define the capacitated graph $D(x, y)$ on vertices E. An edge (e, e') exist if there is an $\epsilon > 0$ such that $x + \epsilon(\chi'_e - \chi_e) \in P_\rho$ and $y + \epsilon(\chi_e - \chi'_e) \in P_\rho$. For edges (e, e') we define capacities $\hat{c}_{x,y}(e, e')$ as follows:

$$\hat{c}_{x,y}(e, e') := \max\{\alpha | x + \alpha(\chi_{e'} - \chi_e) \in P_\rho \text{ and } y + \alpha(\chi_e - \chi'_e) \in P_\rho\}$$

A *bidirectional flow* is a flow in $D(x, y)$ where every resource e with $x_e > y_e$ has supply of $x_e - y_e$ and every resource e with $x_e < y_e$ has a demand of $y_e - x_e$. Such a flow might not exist. In that case we say that x and y are *conflicting strategies*.

We are ready to define the class of *bidirectional flow polymatroids*:

Definition 2.3.3 (Bidirectional flow polymatroid)**.** *A polymatroid* (E, ρ) *is called a* bidirectional flow polymatroid *if for every pair of vectors* x, y *in the base polytope* P_ρ*, there exists a bidirectional flow in* $D(x, y)$*.*

We give a simple example of a bidirectional flow polymatroid.

Example 2.3.4. We consider polymatroid P_ρ defined by the graphic matroid on the graph depicted in Figure 2.2. In this polymatroid, a total load of 1 is divided over the bases of the graphic matroid. Here, for any two strategies x and y there exists a bidirectional flow in $D(x, y)$. In particular, in Figure 2.3 we show the existence for a bidirectional flow for strategy x and y defined in Figure 2.2.

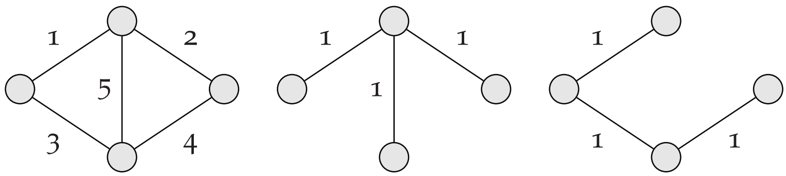

Figure 2.2: Left: the original graph with numbered resources. Middle: Load distribution for strategy x. Right: Load distribution for strategy y.

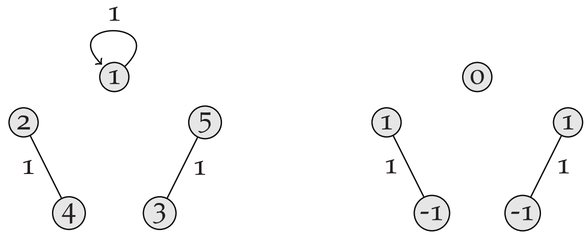

Figure 2.3: Left: Graph $D(x, y)$ with corresponding capacities. Right: the bidirectional flow in $D(x, y)$, including supplies and demands.

2.4 A UNIQUENESS RESULT

In this section we prove that when the strategy space of every player is the base polytope of a bidirectional flow polymatroid, equilibria are unique. We denote the *marginal cost* of player i on resource $e \in E$ by:

$$\mu_{i,e}(x) = c_{i,e}(x_e) + x_{i,e} c'_{i,e}(x_e).$$

An equilibrium condition for polymatroid congestion games, a result that follows from [31, Lemma 1], is as follows:

Lemma 2.4.1. *Let x be a Nash equilibrium in an atomic splittable polymatroid congestion game. If $x_{i,e} > 0$, then for all $e' \in E$ for which there is an $\epsilon > 0$ such that $x_i + \epsilon(\chi_{e'} - \chi_e) \in P_{\rho_i}$, we have $\mu_{i,e}(x) \leqslant \mu_{i,e'}(x)$.*

In the rest of this section we will prove the following theorem:

Theorem 2.4.2. *If for a polymatroid congestion game, the strategy space for every player is the base polytope of a bidirectional flow polymatroid, then the equilibria of this game are unique.*

From now on we assume $x = (x_i)_{i \in N}$ and $y = (y_i)_{i \in N}$ are strategy profiles, where strategies x_i and y_i are taken from the base polytope P_{ρ_i} of a player-specific bidirectional flow polymatroid. Before we prove Theorem 2.4.2, we introduce some new notation. We define $E^+ = \{e \in E | x_e > y_e\}$ and $E^- = \{e \in E | x_e < y_e\}$ as the sets of *globally* overloaded and underloaded resources. We define $E^= = \{e \in E | x_e = y_e\}$ as the set of resources on which the total load does not change. In the same way we define player-specific sets of *locally* underloaded and overloaded resources $E^{i,+} = \{e \in E | x_{i,e} > y_{i,e}\}$ and $E^{i,-} = \{e \in E | x_{i,e} < y_{i,e}\}$. We also introduce four player sets:

$$\begin{aligned}
N^+_> &= \{i \in N | \sum_{e \in E^+} x_{i,e} - y_{i,e} > 0\}, \\
N^-_> &= \{i \in N | \sum_{e \in E^- \cup E^=} x_{i,e} - y_{i,e} > 0\}, \\
N^+_< &= \{i \in N | \sum_{e \in E^+} x_{i,e} - y_{i,e} \leqslant 0\}, \\
N^-_< &= \{i \in N | \sum_{e \in E^- \cup E^=} x_{i,e} - y_{i,e} \leqslant 0\}.
\end{aligned}$$

We distinguish between two cases. Either $E = E^=$, thus $x_e = y_e$ for all resources $e \in E$, or $E \neq E^=$, which implies that E^+ and E^- are non-empty.

Lemma 2.4.3. *If* $E \neq E^=$, *then* $N^+_> \neq \emptyset$.

Proof. Every player distributes the same weight over the resources in x_i and y_i, thus $\sum_{e \in E} x_{i,e} - y_{i,e} = 0$ and $N^+_> = N^-_<$ and $N^+_< = N^-_>$. As $E^+ \neq \emptyset$ we have:

$$0 < \sum_{e \in E^+} x_e - y_e = \sum_{i \in N^+_>} \sum_{e \in E^+} x_{i,e} - y_{i,e} + \sum_{i \in N^+_<} \sum_{e \in E^+} x_{i,e} - y_{i,e}.$$

Note that the first term in the last expression is non-negative and the second one is non-positive. As the whole equation should be positive, we need that this first term is strictly positive and therefore $N^+_> \neq \emptyset$. ∎

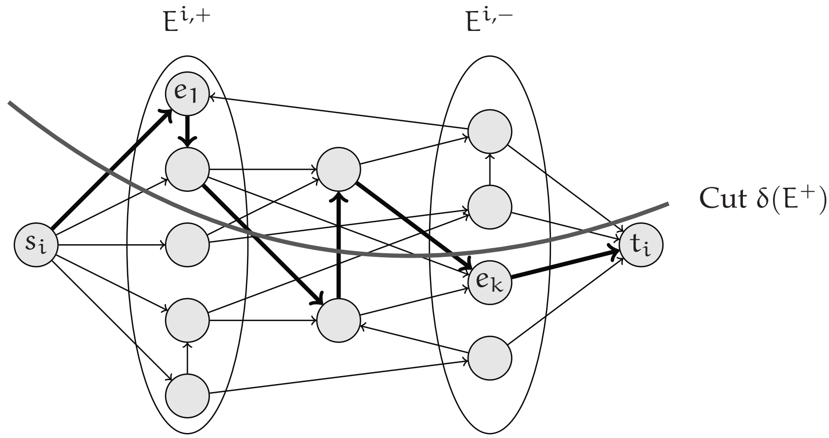

Figure 2.4: Visualization of graph $G(x_i, y_i)$ and cut $\delta(E^+)$ used in the proof of Lemma 2.4.4.

For each player i we create a graph $G(x_i, y_i)$ from graph $D(x_i, y_i)$ by adding a super-source s_i and a super-sink t_i to $D(x_i, y_i)$. We add edges from s_i to $e \in E^{i,+}$ with capacity $x_{i,e} - y_{i,e}$ and edges from $e \in E^{i,-}$ to t_i with capacity $y_{i,e} - x_{i,e}$. Graph $G(x_i, y_i) = (V_G, E_G)$ is visualized in Figure 2.4.

Recall that strategies x_i and y_i are both chosen from the base polytope of a bidirectional flow polymatroid. Therefore there exists a flow f_i in $D(x_i, y_i)$ where every resource $e \in E^{i,+}$ has a supply of $x_{i,e} - y_{i,e}$ and $e \in E^{i,-}$ a demand of $y_{i,e} - x_{i,e}$. Using f_i we define a flow f'_i in $G(x_i, y_i)$ as follows:

$$f'_i(e, e') = \begin{cases} x_{i,e} - y_{i,e}, & \text{if } e = s_i \text{ and } e' \in E^{i,+}, \\ y_{i,e} - x_{i,e}, & \text{if } e \in E^{i,-} \text{ and } e' = t_i, \\ f_i(e, e'), & \text{otherwise.} \end{cases} \tag{2.1}$$

Lemma 2.4.4. *There exists a player* i *and a path* $(s_i, e_1, \ldots, e_k, t_i)$ *in* $G(x_i, y_i)$ *such that* $e_1 \in E^{i,+} \cap (E^+ \cup E^=)$ *and* $e_k \in E^{i,-} \cap (E^- \cup E^=)$.

Proof. If $E \neq E^=$, then using Lemma 2.4.3 we have that $N^+_> \neq \emptyset$, and we pick a player $i \in N^+_>$. Flow f'_i can be decomposed into flow carrying s_i-t_i paths, and we will show that there exists a path in this path decomposition that goes from s_i to a vertex $e_1 \in E^{i,+} \cap E^+$, and, after visiting possibly other vertices, finally goes through a vertex $e_k \in E^{i,-} \cap (E^- \cup E^=)$ to t_i. To see this consider the cut

$$\delta(E^+) := \{(u, v) \in E_G \mid u \in E^+ \text{ and } v \notin E^+, \text{ or } u \notin E^+ \text{ and } v \in E^+)\},$$

as visualized in Figure 2.4. Recall that $i \in N^+_>$, hence, $\sum_{e \in E^+} x_{i,e} - y_{i,e} > 0$. Thus, in f'_i more load enters E^+ from s_i, than leaves E^+ to t_i. This implies that in the flow decomposition of f'_i there must be a path that goes from s_i to a vertex $e_1 \in E^{i,+} \cap E^+$, crosses cut $\delta(E^+)$ an odd number of times to a vertex $e_k \in E^{i,-} \cap (E^- \cup E^=)$ before ending in t_i. As this is a flow-carrying path in f'_i, it exists in $G(x_i, y_i)$.

If $E = E^=$, pick a player i for who there exists a resource e with $x_{i,e} \neq y_{i,e}$ and look at the path decomposition of f'_i. Every path $(s_i, e_1, \ldots, e_k, t_i)$ in this decomposition is a path such that $e_1 \in E^{i,+}$ and $e_k \in E^{i,-}$. As $E = E^=$, it also holds that $e_1 \in E^{i,+} \cap E^=$ and $e_k \in E^{i,-} \cap E^=$. As this is a flow-carrying path in f'_i, it exists in $G(x_i, y_i)$ ∎

Proof of Theorem 2.4.2. Assume that x and y are both Nash equilibria. Using Lemma 2.4.4 we find a path $(s_i, e_1, \ldots, e_k, t_i)$ in $G(x_i, y_i)$ such that both $e_1 \in E^{i,+} \cap (E^+ \cup E^=)$ and $e_k \in E^{i,-} \cap (E^- \cup E^=)$. Since every edge (e_j, e_{j+1}) exists in $G(x_i, y_i)$, for all $j \in \{1, \ldots, k-1\}$ we get for sufficiently small $\epsilon > 0$:

$$x_i + \epsilon(\chi_{e_{j+1}} - \chi_{e_j}) \in P_{\rho_i} \text{ and } y_i + \epsilon(\chi_{e_j} - \chi_{e_{j+1}}) \in P_{\rho_i}.$$

Using Lemma 2.4.1 we obtain for x:

$$\mu_{i,e_1}(x) \leqslant \mu_{i,e_2}(x) \leqslant \cdots \leqslant \mu_{i,e_k}(x), \tag{2.2}$$

and similarly for y:

$$\mu_{i,e_k}(y) \leqslant \mu_{i,e_{k-1}}(y) \leqslant \cdots \leqslant \mu_{i,e_1}(y). \tag{2.3}$$

Recall that $\mu_{i,e}(x) = c_{i,e}(x_e) + x_{i,e}c'_{i,e}(x_e)$. As $e_1 \in E^{i,+}$, we have that $x_{i,e_1} > y_{i,e_1}$. Because c_{i,e_1} is strictly increasing and $e_1 \in (E^+ \cup E^=)$ we get $c_{i,e_1}(x_{e_1}) \geqslant c_{i,e_1}(y_{e_1})$ and $c'_{i,e_1}(x_{e_1}) > 0$ using $x_{e_1} \geqslant x_{i,e_1} > 0$. Moreover, since c_{i,e_1} is convex, the slope of c_{i,e_1} is non-decreasing and, hence, $c'_{i,e_1}(x_{e_1}) \geqslant c'_{i,e_1}(y_{e_1})$. Putting things together, we get

$$\mu_{i,e_1}(y) < \mu_{i,e_1}(x). \tag{2.4}$$

Similarly, as $e_k \in E^{i,-} \cap (E^- \cup E^=)$, we have:

$$\mu_{i,e_k}(x) \leqslant \mu_{i,e_k}(y). \tag{2.5}$$

Combining (2.2), (2.3), (2.4) and (2.5), we have:

$$\mu_{i,e_k}(x) \leqslant \mu_{i,e_k}(y) \leqslant \mu_{i,e_1}(y) < \mu_{i,e_1}(x) \leqslant \mu_{i,e_k}(x).$$

This is a contradiction and therefore either strategy x_i or y_i is not a Nash equilibrium for player i. ■

2.5 APPLICATIONS

In this section we demonstrate that bidirectional flow polymatroids are general enough to allow for meaningful applications. As described in Example 2.2.3, matroid congestion games belong to polymatroid congestion games. A subclass of matroids are *base orderable* matroids introduced by Brualdi [11] and Brualdi and Scrimger [12].

Definition 2.5.1 (Base orderable matroid). *A matroid $\mathcal{M} = (E, \mathcal{I})$ is called base orderable if for every pair of bases (B, B') there exists a bijective function $g_{B,B'} : B \to B'$ such that both $B - e + g_{B,B'}(e) \in \mathcal{B}$ and $B' + e - g_{B,B'}(e) \in \mathcal{B}$ for all $e \in E$.*

We prove that polymatroids defined by the rank function of a base orderable matroid belong to the class of bidirectional flow polymatroids. Therefore, all matroid congestion games for which the player-specific matroids are base orderable have unique equilibria.

Theorem 2.5.2. *Let rk be the rank function of base orderable matroid $M = (E, \mathrm{rk})$. Then, for any $d \geqslant 0$, the polymatroid $(E, d \cdot \mathrm{rk})$ is a bidirectional flow polymatroid.*

Proof. Similar as in Example 2.2.3, polytope P describes how weight d can be divided over bases in $\mathcal{B}$ to obtain a feasible strategy $x \in P_{d \cdot \mathrm{rk}}$. We call vector $x' \in P$ a *base decomposition* of x if it satisfies $x_e = \sum_{B \in \mathcal{B}; e \in B} x'_B$ for all $e \in E$. Note that a base composition of $x \in P_{d \cdot \mathrm{rk}}$ always exists, as $P_{d \cdot \mathrm{rk}}$ is the convex hull of all characteristic vectors (multiplied by d) of all the bases of matroid M (see [26, Corollary 3.25]). Given two vectors $x, y \in P_{d \cdot \mathrm{rk}}$, we look at the differences between two base decompositions $x', y' \in P$. We introduce sets $\mathcal{B}^+, \mathcal{B}^- \subset \mathcal{B}$ that will contain respectively the *overloaded* and *underloaded bases*: $\mathcal{B}^+ = \{B \in \mathcal{B} | x'_B > y'_B\}$ and $\mathcal{B}^- = \{B \in \mathcal{B} | x'_B < y'_B\}$.

Using these sets we create the complete directed bipartite graph $D_{\mathcal{B}}(x, y)$ on vertices $(\mathcal{B}^+, \mathcal{B}^-)$, where bases $B \in \mathcal{B}^+$ have a supply $x'_B - y'_B$ and bases $B \in \mathcal{B}^-$ have a demand $y'_B - x'_B$. As the total supply equals the total demand, there exists a transshipment t from strategies $B \in \mathcal{B}^+$ to strategies

$B' \in \mathcal{B}^-$, such that, when carried out, we obtain y' from x'. We denote by $t_{(B,B')}$ the amount of load transhipped from $B \in \mathcal{B}^+$ to $B' \in \mathcal{B}^-$.

In the remainder of the proof, we use transshipment t to construct a flow f in graph $D(x,y)$. As the polymatroid is defined by the rank function of a base orderable matroid, for every pair of bases (B, B') there exists a bijective function $g_{B,B'} : B \to B'$ such that both $B - e + g_{B,B'}(e) \in \mathcal{B}$ and $B' + e - g_{B,B'}(e) \in \mathcal{B}$ for all $e \in B$. Note that when $e \in B \cap B'$, $g_{B,B'}(e) = e$. Using the function $g_{B,B'}$, we can decompose the value transhipped from B to B' into a transshipment between resources. For all combinations of resources $(e, e') \in E \times E$ we define:

$$\mathcal{B}^2_{e,e'} := \left\{(B, B') \in \mathcal{B}^+ \times \mathcal{B}^- | e \in B, e' \in B' \text{ and } g_{B,B}(e) = e'\right\}.$$

We define flow f as: $f_{(e,e')} = \sum_{(B,B')\in\mathcal{B}^2_{e,e'}} t_{B,B'}$ for all $(e, e') \in E \times E$. Then f has the following two properties:

1. It satisfies all demands and supplies in $D(x,y)$ as f is created from base decompositions x', y' for strategy profiles x and y.

2. It satisfies capacities $\hat{c}_{x,y}(e, e')$ of $D(x,y)$, as $x + \sum_{(B,B')\in\mathcal{B}^2_{e,e'}} t_{B,B'} \cdot (\chi_{e'} - \chi_e)$ is a convex combination of bases, and thus an element of $P_{d \cdot rk}$. Therefore,
$$f_{(e,e')} = \sum_{(B,B')\in\mathcal{B}^2_{e,e'}} t_{B,B'} < \hat{c}_{x,y}(e, e').$$

Hence, f is a feasible flow in $D(x,y)$, satisfying all supplies and demands. As $x, y \in P_{d \cdot rk}$ were chosen arbitrarily, $P_{d \cdot rk}$ is a bidirectional flow polymatroid. ■

An application of these results can be found in the *spanning tree games*.

Example 2.5.3 (Spanning Tree Games)**.** Consider a finite set of players N and an undirected graph $G = (V, E)$ with non-negative, increasing, differentiable, convex and player specific edge costs functions $c_{i,e}$ for all $e \in E$ and $i \in N$. In a spanning tree game, every player i is associated with a weight d_i and a subgraph G_i of G. A strategy for player i is to divide it's weight along the spanning trees of G_i, to minimize his total costs. If G is a generalized series parallel graph, then $P_{d_i \cdot rk_i}$ is a bidirectional flow polymatroid, where rk_i be the rank function for the graphic matroid on subgraph G_i, (cf. Figure 2.1). Theorem 2.5.2 implies that equilibria will be unique.

For graphic matroids, the generalized series-parallel graph is the maximal graph structure that allows for a bidirectional flow between every pair of strategies.

Theorem 2.5.4 (Korneyenko [42], Nishizeki [57]). *A graph is generalized series-parallel if and only if it does not contain the* K_4 *as a minor.*

Let rk be the rank function for the graphic matroid on the K_4, we show that there exist two conflicting strategies $x, y \in P_{\mathrm{rk}}$, thus there does not exist a flow f in $D(x, y)$.

Example 2.5.5. Polymatroid (E, rk) based on the rank function of the graphic matroid on the K_4 is not a bidirectional flow polymatroid. Let the resources be numbered as in Figure 2.5 and look at the strategies $x = (1, 1, 0, 0, 0, 1)$ and $y = (0, 0, 1, 1, 1, 0)$. Graph $D(x, y)$ is depicted in Figure 2.5. Then there is no flow f in $D(x, y)$ that satisfies all supplies and demands. Resource 1 and 6 have both a supply of 1 and can only exchange load with resource 4 , which only has demand 1. Thus such a flow f does not exist, and (E, rk) is not a bidirectional flow polymatroid.

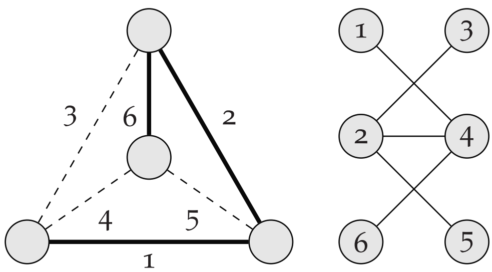

Figure 2.5: Left: the K_4 with two strategies x (thick), y (dashed). Right: $D(x, y)$.

2.6 NON-DIFFERENTIABLE CONVEX FUNCTIONS

So far we assumed cost functions to be differentiable. When this is not the case, the proof we gave in the previous section does not hold. When a function is not differentiable, one speaks about the left derivative $c^-(x)$ and the right derivative $c^+(x)$. In the same way we define $\mu^-_{i,e}(x) = c(x_e) + x_{i,e}c^-(x_e)$ and $\mu^+_{i,e}(x) = c(x_e) + x_{i,e}c^+(x_e)$. For non-differentiable functions, equilibrium condition (2.4.1) generalizes as follows:

Lemma 2.6.1 (Theorem 8.1, [26]). *Let x be a Nash equilibrium in a polymatroid congestion game. If $x_{i,e} > 0$, then for all $e' \in E$ for which there is an $\epsilon > 0$ such that $x_i + \epsilon(\chi_{e'} - \chi_e) \in P_{\rho_i}$, we have $\mu^-_{i,e}(x) \leqslant \mu^+_{i,e'}(x)$.*

The uniqueness proof in the previous section will not hold as equations (2.4) and (2.5) might fail to hold using this new equilibrium condition. An example with multiple equilibria is as follows.

Example 2.6.2. We look at a two player asymmetric game on three resources. Both players have demand $d_1 = d_2 = 1$, and the first player can only use resource 1 or 2, the second player can only use resource 2 or 3. Note that this game is a 1-uniform matroid congestion, and therefore a bidirectional polymatroid. We use the following non-player specific cost functions:

$$c_1(x_1) = 4x_1 \qquad c_2(x_2) = \begin{cases} x_2 & \text{if } x_2 < 1 \\ 10x_2 - 9 & \text{otherwise} \end{cases} \qquad c_3(x_3) = 4x_3.$$

Let $(x_{1,1}, x_{1,2}, x_{2,2}, x_{2,3})$ denote a strategy profile, then both $(\frac{1}{3}, \frac{2}{3}, \frac{1}{3}, \frac{2}{3})$ and $(\frac{2}{3}, \frac{1}{3}, \frac{2}{3}, \frac{1}{3})$ correspond to Nash equilibria. Note that these are two Nash equilibria where the total load on the resources differ.

The same example can be modified to one with symmetric strategy spaces, but player specific costs, by incurring a high cost on the unavailable resources. The question remains unresolved for symmetric player specific cost functions.

2.7 NON-MATROID SET SYSTEMS

We now derive necessary conditions on a given set system $(\mathcal{S}_i)_{i \in N}$ so that any atomic splittable congestion game based on $(\mathcal{S}_i)_{i \in N}$ admits unique equilibria. We show that the *matroid property* is a necessary condition on the players' strategy spaces that guarantees uniqueness of equilibria *without* taking into account how the strategy spaces of different players interweave.[1] To state this property mathematically precisely, we introduce the notion of *embeddings* of $\mathcal{S}_i$ in E. An embedding is a map $\tau := (\tau_i)_{i \in N}$, where every $\tau_i : E_i \to E$ is an injective map from $E_i := \cup_{S \in \mathcal{S}_i} S$ to E. For $X \subseteq E_i$, we denote $\tau_i(X) := \{\tau_i(e), e \in X\}$. Mapping τ_i induces an

1 The term "interweaving" has been introduced by Ackermann et al. [2, 3].

isomorphism $\phi_{\tau_i} : \mathcal{S}_i \to \mathcal{S}'_i$ with $S \mapsto \tau_i(S)$ and $\mathcal{S}'_i := \{\tau_i(S) | S \in \mathcal{S}_i\}$. Thus, isomorphism $\phi_\tau = (\phi_{\tau_i})_{i \in N}$ induces the isomorphic strategy space $\phi_\tau(\mathcal{S}) = (\phi_{\tau_i}(\mathcal{S}_i))_{i \in N}$.

Definition 2.7.1. *A family of set systems $\mathcal{S}_i \subseteq 2^{E_i}$, for $i \in N$ is said to have the* strong uniqueness property *if for all embeddings τ, the induced game with isomorphic strategy space $\phi_\tau(\mathcal{S})$ has unique Nash equilibria.*

Since for bases of matroids any embedding τ_i with isomorphism ϕ_{τ_i} has the property that $\phi_{\tau_i}(\mathcal{S}_i)$ is again a collection of bases of a matroid, we obtain the following immediate consequence of Theorem 2.4.2.

Corollary 2.7.2. *If $(\mathcal{S}_i)_{i \in N}$ consists of bases of a base-orderable matroid M_i, for $i \in N$, then $(\mathcal{S}_i)_{i \in N}$ possess the strong uniqueness property.*

For obtaining necessary conditions we need a certain property of non-matroids stated in the following lemma.

Lemma 2.7.3. *[3, Lemma 16] If $\mathcal{S}_i \subseteq 2^{E_i}$ with $\mathcal{S}_i \neq \emptyset$ is a non-matroid, then there exist $X, Y \in \mathcal{S}_i$ and $\{a, b, c\} \subseteq X \Delta Y := (X \backslash Y) \cup (Y \backslash X)$ such that for each set $Z \in \mathcal{S}_i$ with $Z \subseteq X \cup Y$, either $a \in Z$ or $\{b, c\} \subseteq Z$.*

Theorem 2.7.4. *Let $|N| \geqslant 3$ and assume that for all $i \in N$, $\mathcal{S}_i$ is a non-matroid set system. Then, $(\mathcal{S}_i)_{i \in N}$ does not have the strong uniqueness property.*

Proof. We show that there are embeddings $\tau_i : E_i \to E$, $i \in N$, such that the isomorphic game $\phi_\tau(\mathcal{S}) = (\phi_{\tau_1}(\mathcal{S}_1), \ldots, \phi_{\tau_n}(\mathcal{S}_n))$ admits multiple equilibria.

We can assume w.l.o.g. that each set system $\mathcal{S}_i$ forms an anti-chain (in the sense that $X \in \mathcal{S}_i, X \subset Y$ implies $Y \notin \mathcal{S}_i$) since cost functions are non-negative and strictly increasing. Let us call a non-empty set system $\mathcal{S}_i \subseteq 2^{E_i}$ a *non-matroid* if $\mathcal{S}_i$ is an anti-chain and $(E_i, \{X \subseteq S : S \in \mathcal{S}_i\})$ is not a matroid.

Let $\tilde{E} = \bigcup_{i \in N} \tau_i(E_i)$ denote the set of all resources under the embeddings $\tau_i, i \in N$. The costs on all resources in $\tilde{E} \backslash (\tau_1(E_1) \cup \tau_2(E_2) \cup \tau_3(E_3))$ are set to zero. Also, the demands of all players d_i with $i \in N \backslash \{1, 2, 3\}$ are set to zero. This way, the game is basically determined by the players $1, 2, 3$. We set the demands $d_1 = d_2 = d_3 = 1$.

Let us choose two sets X, Y in $\mathcal{S}_1$ and $\{a, b, c\} \subseteq X \cup Y$ as described in Lemma 2.7.3. Let $e := \tau_1(a), f := \tau_1(b)$ and $g := \tau_1(c)$. We set the costs of all resources in $\tau_1(E_1) \backslash (\tau_1(X) \cup \tau_1(Y))$ to some very large cost M (large enough so that player 1 would never use any of these resources). The cost on all resources in $(\tau_1(X) \cup \tau_1(Y)) \backslash \{e, f, g\}$ is set to zero. This way, player 1

Table 2.1: Cost functions used for constructing a game with multiple equilibria.

	e	f	g
Player 1	$c_{1,e}(x) = x^3$	$c_{1,f}(x) = x + 1$	$c_{1,g}(x) = x + 1$
Player 2	$c_{2,e}(x) = x + 1$	$c_{2,f}(x) = x^3$	$c_{2,g}(x) = x + 1$
Player 3	$c_{3,e}(x) = x + 1$	$c_{3,f}(x) = x + 1$	$c_{3,g}(x) = x^3$

always chooses a strategy $\tau_1(Z) \subseteq \tau_1(X) \cup \tau_1(Y)$ which, by Lemma 2.7.3, either contains e, or it contains both f and g. We apply the same construction for player 2 and 3, only changing the role of e to act as f and g, respectively.

Note that the so-constructed game is essentially isomorphic to the routing game illustrated in Figure 2.6 if we interpret resource e as arc (s_1, t_1), resource f as arc (s_2, t_2), and resource g as arc (s_3, t_3). On every edge there is a player specific cost function, given in Table 2.1.

Every player has two possible paths: the direct path that uses only one edge, or the indirect path that uses two edges. We show that the game where everyone puts all their weight on the direct path is a Nash equilibrium, as is the game where everybody puts their weight on the indirect path.

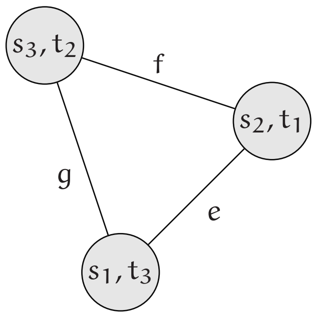

Figure 2.6: Counterexample

If all players put their weight on the direct route, then player 1 cannot deviate to decrease it's costs, as:

$$c_{1,e}(1) + c'_{1,e}(1) \cdot 1 = 1 + 3 \leqslant 2 + 2 = c_{1,f}(1) + c_{1,g}(1).$$

On the other hand, when all players put their weight on the indirect direct route, player 1 can also not deviate, as:

$$c_{1,f}(2) + c'_{1,f}(2) \cdot 1 + c_{1,g}(2) + c'_{1,g}(2) \cdot 1 = 3 + 1 + 3 + 1 \leqslant 8 = c_{1,e}(2).$$

The same inequalities hold for player 2 and 3. And therefore everyone playing the direct route, or everyone playing the indirect route both results in a Nash equilibrium.

■

2.8 A CHARACTERISATION FOR UNDIRECTED GRAPHS

In Section 2.7 we proved that non-matroid set systems in general do not have the strong uniqueness property when there are at least three players, by constructing embeddings τ_i that lead to the counterexample in Figure 2.6. This example also gives new insights in uniqueness of equilibria in network congestion games. In the following, we give a characterization of graphs that guarantee uniqueness of Nash equilibria even when player-specific cost functions are allowed.

Definition 2.8.1. *An undirected graph* G *is said to have the* uniqueness property *if for any atomic splittable network congestion game on* $G = (V, E)$*, equilibria are unique.*

Note that in the above definition, we do not specify how source- and sink vertices are distributed in V. We obtain the following result which is related to Theorem 3 of Meunier and Pradeau [47], where a similar result is given for non-atomic congestion games with player-specific cost functions.

Theorem 2.8.2. *An undirected graph has the uniqueness property if and only if* G *has no cycle of length 3 or more.*

Proof. Let $G = (V, E)$ be the network in an atomic splittable network congestion game. Assume there exists a cycle C in G of length k, with $k \geqslant 3$. Already for three players, we can create a game with multiple equilibria by generalising the previous example visualized in Figure 2.6. Pick three clockwise adjacent vertices v_1, v_2, v_3 in cycle C and create three players which have equal weight 1. Player 1 has source v_1 and sink v_2, player 2 has source v_2 and sink v_3 and player 3 has source v_3 and sink v_1. Let $c_{i,e}(x)$ be the cost function for player i at resource e. Define $c_{i,e}(x)$ as in Table 2.2.

For the same reason as in Figure 2.6 this game has two Nash equilibria: one where all players send their flow clockwise, another where all players send all flow counterclockwise.

On the other hand, assume no cycle of length 3 or more in G exists, then G is a tree with parallel edges. Thus, for every source s and sink t, there

Table 2.2: Cost functions for a game with multiple equilibria, M is sufficiently large.

$c_{i,e}(x)$	(v_1, v_2)	(v_2, v_3)	$C \setminus \{(v_1, v_2), (v_2, v_3)\}$	$e \notin C$
Player 1	x^3	$x+1$	$\frac{1}{k-2}(x+1)$	$x+M$
Player 2	$x+1$	x^3	$\frac{1}{k-2}(x+1)$	$x+M$
Player 3	$x+1$	$x+1$	$\frac{1}{k-2}\,x^3$	$x+M$

is a unique path from s to t in G modulo parallel edges. Therefore, players only have to decide on how to divide their weight over every set of parallel edges they encounter. As the total cost for a player is just the sum of the costs for all resources separately, players compete only in sets of parallel edges. Atomic splittable congestion games on parallel edges with player-specific cost functions are proven to have a unique Nash equilibrium by Orda et al. [59]. Thus when G does not contain cycles of length 3 or more, Nash equilibria are unique. ■

2.9 CONCLUDING REMARKS

In this chapter we studied sufficient and necessary conditions for a pure Nash equilibrium to be unique in atomic splittable congestion games.

As a main result we introduced the class of bidirectional flow polymatroids and proved that whenever the strategy space of each player is equal to the base polytope of an bidirectional flow polymatroid, equilibria are guaranteed to be unique. Then, we demonstrated that bidirectional polymatroids are quite general, as they contain the classes of base-orderable matroids, gammoids, transversal and laminar matroids. We complemented this uniqueness result by showing that when there are at least three players for which set systems $\mathcal{S}_i$ are not bases of matroids, there exists an isomorphic game that admits multiple equilibria.

Secondly, we considered uniqueness of equilibria if the set systems $\mathcal{S}_i$ correspond to paths in undirected graphs. Here, we proved that is we assume at least three players, and we do not specify beforehand which vertices of the graph serve as sources or sins, an undirected graph induces unique equilibria if and only if the graph has no cycle of length at least 3.

APPENDIX 2.A: SUBCLASSES OF BASE ORDERABLE MATROIDS

We give proofs of the inclusions given in Figure 2.1:

- $\subset_1$: A uniform matroid is a partition matroid where the partition contains only one set.
- $\subset_2$: A partition matroid is a laminar matroid if all sets in the laminar family are disjoint.
- $\subset_3$: A partition matroid is a transversal matroid where the sets that need to be traversed are either equal or disjoint.
- $\subset_4$: For the laminar matroid, let $\mathcal{F}$ be the underlying laminar family on ground set S with $S \in \mathcal{F}$. Copy each set X in $\mathcal{F}$ exactly k_X times to create multi set $\mathcal{F}'$, where k_X is the number of elements we are allowed to take from set X. Now create a directed graph $G = (V, A)$, where $V = \mathcal{F}' \cup S$, and $A = \{(X, Y) \in \mathcal{F}' \times \mathcal{F}' | X \subseteq Y\} \cup \{(s, X) \in S \times \mathcal{F}' | s \in X\}$. Let U be the maximal multi set containing only S. Then clearly G with starting points S and endpoints U form a gammoid that corresponds to the laminar matroid.
- $\subset_5$: A transversal matroid is a gammoid according to Corollary 39.5a in [67].
- $\subset_6$: Every binary matroid is a gammoid if and only if it is a graphic matroid on a generalized series-parallel graph [76]. As every graphic matroid is binary [30], the graphic matroid on a generalized series-parallel graph is a binary gammoid, and thus a gammoid.
- $\subset_7$: A gammoid is strongly base orderable according to Theorem 42.12 in [67].
- $\subset_8$: A matroid $\mathcal{M} = (R, \mathcal{I})$ is called *strongly base orderable* (SBO) if for every pair of bases (B, B') there exists a bijective function $g : B \to B'$ such that $B - X + g(X) \in \mathcal{B}$. Take $X = e$ and $X = B \backslash \{e\}$ to obtain the conditions for base orderable matroids.

3 EQUILIBRIUM COMPUTATION FOR GAMES WITH AFFINE COSTS

3.1 INTRODUCTION

As mentioned in Section 1.1, the complexity of computing equilibria is one of the core activities of algorithmic game theory. While the computational tractability of equilibria for unsplittable congestion games has been intensively studied over the last decade (cf. [2, 14, 15, 17, 25, 68]), the equilibrium computation problem for *atomic splittable congestion games* is much less explored.

Regarding polynomial time algorithms for equilibrium computation, we are only aware of two prior works: (1) For affine player-independent cost functions, there exists a convex potential whose global minima are pure Nash equilibria, see Cominetti et al. [18]. Thus, for any $\epsilon > 0$ one can compute an ϵ-approximate equilibrium in polynomial time by convex programming methods. (2) Huang [39] also considered affine player-independent cost functions, and he devised a combinatorial algorithm computing an exact equilibrium for routing games on symmetric s-t graphs that are so-called *well-designed*. This condition is met for instance by series-parallel graphs. His proof technique also uses the convex potential.

3.1.1 *Our Results and Techniques*

ATOMIC SPLITTABLE CONGESTION GAMES. We study atomic splittable *singleton* congestion games with player-specific affine cost functions and develop the first polynomial time algorithm computing a pure Nash equilibrium. Our algorithm is purely combinatorial and computes an *exact* equilibrium. The main ideas and constructions are as follows. By analyzing the first order necessary optimality conditions of an equilibrium, it can be shown that any equilibrium is *rational* as it is a solution to a system of linear equations with rational coefficients (assuming rational input). Using that equilibria are unique for singleton games (see Richmann and Shimkin [62] and Bhaskar et al. [9]), we further derive that the constraint matrix of the equation system is non-singular, allowing for an explicit representation of the equilibrium by Cramer's rule (using determinants of the constraint- and their sub-matrices). This way, we obtain an explicit lower bound on the minimum demand value for any used resource in the equilibrium. We further show that the unique equilibrium is also the unique equilibrium for an associated *integrally-splittable* game in which the players may only distribute the demands in *integer multiples* of a common *packet size* of some value $k^* \in \mathbb{Q}_{>0}$ over the resources. While we are not able to compute k^* exactly, we can efficiently compute some sufficiently small $k_0 \leqslant k^*$ with the property that an equilibrium for the k_0-integrally-splittable game allows us to determine the set of resources on which a player will put a positive amount of load in the atomic splittable equilibrium. Once these *support sets* are known, an atomic splittable equilibrium can be computed in polynomial time by solving a system of linear equations. This way, we can reduce the problem of computing the exact equilibrium for an atomic splittable game to computing an equilibrium for an associated k_0-integrally-splittable game.

The class of integrally-splittable congestion games has been studied before by Tran-Thanh et al. [72] for the case of player-independent convex cost functions and later by Harks et al. [32, 36] (for the more general case of polymatroid strategy spaces and player-specific convex cost functions). In particular, Harks et al. devised an algorithm with running time $n^2 m(\delta/k_0)^3$, where n is the number of players, m the number of resources, and δ is an upper bound on the maximum demand of the players (cf. Corollary 5.2 [32]). As δ is encoded in binary, however, the algorithm is only pseudo-polynomial even for player-specific affine cost functions.

We devise a polynomial time algorithm for integrally-splittable singleton congestion games with player-specific affine cost functions. Our algorithm works as follows. For a game with initial packet size k_0, we start by finding an equilibrium for packet size $k = k_0 \cdot 2^q$ for some q of order $O(\log(\delta/k_0))$, satisfying only a part of the player-specific demands. Then we repeat the following two actions:

1. We half the packet size from k to $k/2$ and construct a $k/2$-equilibrium using the k-equilibrium. Here, a k-equilibrium denotes an equilibrium for an integrally-splittable game with common packet size k. We show that this can be done in polynomial time by repeatedly performing the following operations given a k-equilibrium:
    a) Among players who can improve, we find the player that benefits most by moving one packet of size $k/2$;

b) If necessary, we perform a sequence of backward-shuffles of packets to correct the *load decrease* caused by the first packet movement (this is called a *backward path*);

c) If necessary, we perform a sequence of forward-shuffles of packets to correct the load increase caused by the first packet movement (this is called a *forward path*);

(a)-(c) is iterated until a $k/2$-equilibrium for the currently scheduled demand is reached. For strategy profile x we define $\Delta(x)$ to be a vector that contains the cost for moving one packet to the currently cheapest resource, for each combination of a player and resource. We show that after each iteration $\Delta(x)$ lexicographically increases, which implies that we converge to a $k/2$-equilibrium.

2. For each player i we repeat the following step: if the current packet size k is smaller than the currently unscheduled demand of player i, we add one more packet for this particular player to the game and recompute the equilibrium. This part of the algorithm has also been used in the algorithm by Tran-Thanh et al. [72] and Harks et al. [36].

After q iterations, we have scheduled all demands and obtain an equilibrium for the desired packet size k_0. Polynomial running time of the algorithm is shown by several structural results on the sensitivity of equilibria with respect to packet sizes $2k$ and k. Specifically, we derive bounds on the difference of the resulting global load vectors as well as the individual load vectors of players. We use these insights to show that $\Delta(x)$ reaches a lexicographical maximum in a polynomial number of steps. Overall, compared to the existing algorithms of Tran-Thanh et al. [72] and Harks et al. [36], our algorithm has two main innovations: packet sizes are decreased exponentially (yielding polynomial running time in δ) and k-equilibrium computation for an intermediate packet size k is achieved via a careful construction of a sequence of single packet movements (backward- and forward paths) from a given $2k$-equilibrium (ensuring its polynomial length).

multimarket cournot oligopolies. We then study the equilibrium computation problem for multimarket Cournot oligopolies as introduced in Section 1.2.3. For multimarket oligopolies, we develop a poly-time computable isomorphism mapping a multimarket Cournot competition game to an associated atomic splittable singleton congestion game. The isomorphism is payoff invariant (up to constants) and thus preserves equilibria in either games. As a consequence, we can apply the isomorphism and the polynomial time algorithm for atomic splittable congestion games to efficiently compute Cournot equilibria for models with firm-specific affine price functions and quadratic production costs. In addition, our analysis for integrally-splittable games also implies new bounds on the difference between real and integral Cournot equilibria complementing and extending recent results of Todd [69] obtained for a single market.

3.1.2 *Related Work*

discrete congestion games. As the first seminal work regarding the computational complexity of equilibrium computation in congestions games, Fabrikant et al. [25] showed that the problem of computing a pure Nash equilibrium is PLS-complete for network congestion games. Ackermann et al. [2] strengthened this result to hold even for network congestion games with linear cost functions. On the other hand, there are polynomial algorithms for symmetric network congestion games (cf. Fabrikant et al. [25]), for matroid congestion games with player-specific cost functions (Ackermann et al. [2, 3]), for polymatroid congestion games with player-specific cost functions and polynomially bounded demands (Harks et al. [32, 36]) and for so-called total unimodular congestion games (see Del Pia et al. [22]). Further results regarding the computation of approximate equilibria in congestion games can be found in Caragiannis et al. [14, 15], Chien and Sinclair [17] and Skopalik and Vöcking [68].

atomic splittable congestion games. Atomic splittable network congestion games with player-independent cost functions have been studied (seemingly independent) by Orda et al. [59], Haurie and Marcotte [38] and Marcotte [46]. Both lines of research mentioned that Rosens' existence result for concave games on compact strategy spaces implies the existence of pure Nash equilibria via Kakutani's fixed-point theorem. Cominetti et al. [18] presented the first upper bounds on the price of anarchy in atomic splittable congestion games. These were later improved by Harks [31] and finally shown to be tight by Schoppmann and Roughgarden [65].

For the computation of equilibria, Marcotte [46] proposed four numerical algorithms and showed local convergence results. Meunier and Pradeau [48] developed a pivoting-algorithm (similar to Lemkes algorithm) for nonatomic network congestion games with affine player-specific cost functions. Polynomial running time was, however, not shown and seems unlikely to hold. Gairing et al. [28] considered nonatomic routing games on parallel links with affine player-specific cost functions. They developed a convex potential function that can be minimized within arbitrary precision in polynomial time. Deligkas et al. [23] considered general concave games with compact action spaces and investigated algorithms computing an approximate equilibrium. Roughly speaking, they discretized the compact strategy space and use the Lipschitz constants of utility functions to show that only a finite number of representative strategy profiles need to be considered for obtaining an approximate equilibrium (see also Lipton et al. [45] for a similar approach). The running time of the algorithm, however, depends on an upper bound of the norm of strategy vectors, thus, implying only a pseudo-polynomial algorithm for our setting.

linear complementarity problems. The equilibrium computation problem in atomic splittable congestion games with singleton strategies and affine cost functions can be written as a *linear complementary problem (LCP)*, see Appendix 3.A for the formal derivation. Given a matrix M and a vector q, a linear complementarity problem $\mathrm{LCP}(M, q)$ seeks a vector z such that the following three properties hold: (1) $z \geqslant 0$ (2) $Mz + q \geqslant 0$ and (3) $z^T(Mz + q) = 0$. If matrix M is positive definite, the solution to the $\mathrm{LCP}(M, q)$ can be found within polynomial time for any q. When we write an atomic splittable congestion game as an $\mathrm{LCP}(M, q)$ (see Appendix 3.A), then M is not positive definite.

Matrix M is a P-matrix, i.e., a matrix where every principal minor is positive, if and only if the solution to the $\mathrm{LCP}(M, q)$ is unique for every q [19, 53]. In some specific cases, this solution can be found within polynomial time [77]. Though an atomic splittable congestion game on singleton strategies has a unique solution, M is not a P-matrix in the corresponding $\mathrm{LCP}(M, q)$. When the demands, that only occur in q, are set to zero, multiple solutions to the $\mathrm{LCP}(M, q)$ exist. Ye and Pardalos [78] studied classes of LCP's for which a polynomial time algorithms are known; e.g. when M is co-positive and $q \geqslant 0$, or when M^{-1} is co-positive and $M^{-1}q \leqslant 0$, but the corresponding LCP of our problem does not seem to belong to any of these classes.

multimarket cournot oligopolies. The existence of equilibria in single market Cournot models (beyond quasi-polynomial utility functions) has been studied extensively in the past decades (see Vives [74] for a good survey). For example, Novshek [58] proved that equilibria exists whenever the marginal revenue of each firm is decreasing in the aggregate quantities of the other firms. Then, several works (cf. Topkis [70], Amir [4], Kukushkin [44], Milgrom and Roberts [50], Milgrom and Shannon [51], Topkis [71] and Vives [73]) proved existence of equilibria when the underlying game is supermodular, i.e., when the strategy space forms a lattice and the marginal utility of each firm is increasing in any other firm's output. Using supermodularity, one can obtain existence results without assuming that the utility functions are quasi-convex. Very recently, Todd [69] considered Cournot competition on a single market, where the price functions are linear and cost functions are quadratic. For such games, he proved that equilibria exist and can be computed in time $O(n\log(n))$, where n denotes the number of firms. Additionally, he analyzed the maximum differences of production quantities of real an integral equilibria, respectively.

Abolhassani et al. [1] devised several polynomial time algorithms (partly using algorithms for solving nonlinear complementarity problems) for multimarket Cournot oligopolies. In contrast to our work, they assume that price functions are firm-independent. Bimpikis et al. [10] provided a characterization of the production quantities at the unique equilibrium, when price functions are player-independent and concave, and cost functions are convex. They study the impact of changes in the competition structure on the firm's profit. This framework can be used to either identify opportunities for collaboration and expanding in new markets. Harks and Klimm [35] studied the existence of Cournot equilibria, under the condition that each firm can only sell its items to a limited number of markets simultaneously. They proved that equilibria exist when production cost functions are convex, marginal return functions strictly decrease for strictly increased own quantities and non-decreased aggregated quantities and when for every firm, the firm specific market reaction functions across markets are identical up to market-specific shifts.

3.2 PRELIMINARIES

3.2.1 *Atomic Splittable Singleton Games*

Similar as in Section 1.2.2, an atomic splittable *singleton* congestion game is represented by the tuple:

$$\mathcal{G} := (N, E, (d_i)_{i\in N}, (E_i)_{i\in N}, (c_{i,e})_{i\in N, e\in E_i}),$$

where $E = \{e_1, \dots, e_m\}$ is a finite set of resources and $N = \{1, \dots, n\}$ is a finite set of players. Each player $i \in N$ is associated with a demand $d_i \in \mathbb{Q}_{\geqslant 0}$ and a set of allowable resources $E_i \subseteq E$. A strategy for player $i \in N$ is a (possibly fractional) distribution of the demand d_i over the singletons in E_i. Thus, one can represent the strategy space of every player $i \in N$ by the polytope:

$$\mathcal{S}_i(d_i) := \Big\{ x_i \in \mathbb{R}_{\geqslant 0}^{|E_i|} \mid \sum_{e \in E_i} x_{i,e} = d_i \Big\}.$$

The combined strategy space is denoted by $\mathcal{S} := \times_{i \in N} \mathcal{S}_i(d_i)$ and we denote by $x = (x_i)_{i \in N}$ the overall strategy profile. We define $x_{i,e} := (x_i)_e$ as the load of player i on $e \in E_i$ and $x_{i,e} = 0$ when $e \in E \backslash E_i$. The total load on e is given as $x_e := \sum_{i \in N} x_{i,e}$. Resources have player-specific affine cost functions $c_{i,e}(x_e) = a_{i,e} x_e + b_{i,e}$ with $a_{i,e} \in \mathbb{Q}_{>0}$ and $b_{i,e} \in \mathbb{Q}_{\geqslant 0}$ for all $i \in N$ and $e \in E$. The total cost of player i in strategy distribution x is defined as:

$$\pi_i(x) = \sum_{e \in E_i} c_{i,e}(x_e)\, x_{i,e}.$$

For $i \in N$, we write $\mathcal{S}_{-i}(d_{-i}) = \times_{j \neq i} \mathcal{S}_j(d_j)$ and $x = (x_i, x_{-i})$ meaning that $x_i \in \mathcal{S}_i(d_i)$ and $x_{-i} \in \mathcal{S}_{-i}(d_{-i})$. A strategy profile x is an *equilibrium* if $\pi_i(x) \leqslant \pi_i(y_i, x_{-i})$ for all $i \in N$ and $y_i \in \mathcal{S}_i(d_i)$. A pair $\big(x, (y_i, x_{-i})\big) \in \mathcal{S} \times \mathcal{S}$ is called an *improving move* of player i, if $\pi_i(x_i, x_{-i}) > \pi_i(y_i, x_{-i})$. The *marginal cost* for player i on resource e is defined as:

$$\mu_{i,e}(x) = c_{i,e}(x_e) + x_{i,e} c'_{i,e}(x_e) = a_{i,e}(x_e + x_{i,e}) + b_{i,e}.$$

We obtain the following sufficient and necessary equilibrium condition.

Lemma 3.2.1 (cf. Harks [31]). *Strategy profile x is an equilibrium if and only if the following holds for all $i \in N$: when $x_{i,e} > 0$, then $\mu_{i,e}(x) \leqslant \mu_{i,f}(x)$ for all $f \in E_i$.*

Using that the strategy space is compact and cost functions are convex, Kakutanis' fixed point theorem implies the existence of an equilibrium. Uniqueness is proven by Richmann and Shimkin [62] and Bhaskar et al. [9].

Game $\mathcal{G}$ is called symmetric whenever $E_i = E$ for all $i \in N$. We can project any asymmetric game $\mathcal{G}$ on a symmetric game $\mathcal{G}^*$ by setting:

$$c^*_{i,e}(x_e) = \begin{cases} c_{i,e}(x_e) & \text{if } e \in E_i, \\ x_e + (n+2)(a_{\max})^2 & \text{otherwise,} \end{cases}$$

for all $i \in N$ and $e \in E \backslash E_i$, where:

$$a_{\max} := \max\{a_{i,e}, b_{i,e} \mid i \in N, e \in E_i\}, \{d_i \mid i \in N\}, 1\}.$$

In this case $\mu_{i,e}(0) \geqslant \mu_{i,f}(x_e)$ for any $e \in E \backslash E_i$, $f \in E_i$, $i \in N$ and $x \in \mathcal{S}$. Thus, in an equilibrium y for game $\mathcal{G}^*$ no player i puts load on any resource $e \in E \backslash E_i$. Hence, y is also an equilibrium for game $\mathcal{G}$. In the rest of this paper we project every asymmetric game on a symmetric game using the construction above.

3.2.2 *Integral Singleton Games*

A k-integral game is given by the tuple $\mathcal{G}_k := (N, E, (d_i)_{i \in N}, (c_{i,e})_{i \in N, e \in E})$ with $k \in \mathbb{Q}_{>0}$. Here, players cannot split their load fractionally, but only in multiples of k. Assume d_i is a multiple of k, then the strategy space for player i is the following set:

$$\mathcal{S}_i(d_i, k) := \left\{x_i \in \mathbb{Q}_{\geqslant 0}^{|E|} \mid x_{i,e} = kq, q \in \mathbb{N}_{\geqslant 0}, \sum_{e \in E} x_{i,e} = d_i\right\}.$$

In this game, k is also called the *packet size*. When E, N and $(c_{i,e})_{i \in N, e \in E}$ are clear from the context, we will also refer to the game as $\mathcal{G}_k((d_i)_{i \in N})$. For player-specific affine cost functions the (discrete) marginal increase and decrease are defined as follows:

$$\mu_{i,e}^{+k}(x) = (x_{i,e} + k)c_{i,e}(x_e + k) - x_{i,e}c_{i,e}(x_e), \tag{3.1}$$

$$\mu_{i,e}^{-k}(x) = \begin{cases} x_{i,e}c_{i,e}(x_e) - (x_{i,e} - k)c_{i,e}(x_e - k), & \text{if } x_{i,e} > 0 \\ -\infty, & \text{if } x_{i,e} \leqslant 0. \end{cases} \tag{3.2}$$

Here, $\mu_{i,e}^{+k}(x)$ is the cost for player i to add one packet of size k to resource e and $\mu_{i,e}^{-k}(x)$ is the gain for player i of removing a packet from resource e. Assuming that cost functions are affine, we obtain

$$\mu_{i,e}^{+k}(x) = ka_{i,e}(x_e + x_{i,e} + k) \text{ and } \mu_{i,e}^{-k}(x) = ka_{i,e}(x_e + x_{i,e} - k).$$

Lemma 3.2.2 (cf. Groenevelt [29]). *Strategy profile x is an equilibrium in a k-integral congestion game if and only if for all $i \in N$: when $x_{i,e} > 0$, then $\mu_{i,e}^{-k}(x) \leqslant \mu_{i,f}^{+k}(x)$ for all $f \in E$.*

Define $\mu_{i,\min}^{+k}(x) := \min_{e \in E}\{\mu_{i,e}^{+k}(x)\}$ and $\mu_{i,\max}^{-k}(x) := \max_{e \in E}\{\mu_{i,e}^{-k}(x)\}$. Then strategy profile x is an equilibrium in a k-integral congestion game if and only if $\mu_{i,\max}^{-k}(x) \leqslant \mu_{i,\min}^{+k}(x)$ for all $i \in N$.

3.3 REDUCTION TO INTEGRALLY-SPLITTABLE GAMES

We show that the problem of finding an equilibrium for an atomic splittable game reduces to the problem of finding an equilibrium for a k_0-integral game for some $k_0 \in \mathbb{Q}_{>0}$.

Theorem 3.3.1. *Let x be the unique equilibrium of an atomic splittable singleton game $\mathcal{G}$. Then, there exists a $k^* \in \mathbb{Q}_{>0}$ such that x is also the unique equilibrium for the k^*-integral splittable game $\mathcal{G}_{k^*}$.*

Proof. We define the support set for each player as $I_i := \{e \in E \mid x_{i,e} > 0\}$. Lemma 3.2.1 implies that if x is an equilibrium, and $x_{i,e}, x_{i,f} > 0$, then $\mu_{i,e}(x) = \mu_{i,f}(x)$. Define $p := \sum_{i \in N} |I_i| \leqslant nm$. Then, if the correct support set I_i of each player is known, the equilibrium can be computed by solving the following set of p linear equations on p variables.

1. For every player we have an equation that makes sure the demand of that player is satisfied. Thus, $\sum_{e \in I_i} x_{i,e} = d_i$ for each player $i \in N$.
2. For every $i \in N$, there are $|I_i| - 1$ equations of type $\mu_{i,e}(x) = \mu_{i,f}(x)$ for $e, f \in I_i$, which we write as $a_{i,e}(x_e + x_{i,e}) - a_{i,f}(x_f + x_{i,f}) = b_{i,e} - b_{i,f}$. Note that x_e is not an extra variable, but an abbreviation for $\sum_{i \in N} x_{i,e}$.

From now on we refer to this set of equalities as the system $Ax = b$, where A is a $p \times p$ matrix. Note that as the equilibrium exists and is unique, the corresponding matrix is non-singular. Using Cramer's Rule, the unique solution is given by:

$$x_{i,e} = \det(A_{i,e}) / \det(A) = |\det(A_{i,e})| / |\det(A)|,$$

where $A_{i,e}$ is the matrix formed by replacing the column that corresponds to value $x_{i,e}$ in A by b. We first define the set of input values as

$$Q := \{\{a_{i,e}, b_{i,e} \mid i \in N, e \in E_i\}, \{d_i \mid i \in N\}, 1\}.$$

We define the *greatest common divisor* of Q as:

$$a_{gcd} := \max\{a \in \mathbb{Q}_{>0} \mid \forall q \in Q,\ \exists \ell \in \mathbb{N} \text{ such that } q = a \cdot \ell\}.$$

Then, as all values in A and b depend on adding and subtracting values in Q, $|\det(A_{i,e})|$ is an integer multiple of $(a_{gcd})^p$ and, hence, an integer multiple of $(a_{gcd})^{nm}$. Thus, all player-specific loads are integer multiples of $(a_{gcd})^{nm} / |\det(A)|$ and, hence, if we define $k^* = (a_{gcd})^{nm} / |\det(A)|$, x is an equilibrium for the k^*-integral splittable game. Note that we can compute a_{gcd} in running time $O(nm \log a_{\max})$.

It is left to prove that x is the unique equilibrium for the k^*-integral splittable game. Assume on the contrary that there are two different equilibria x, y, where x is the equilibrium for the atomic splittable game. We define:

$$E^+ := \{e \in E | x_e > y_e\}, \qquad N^+ := \{i \in N | \sum_{e \in E^+} (x_{i,e} - y_{i,e}) > 0\},$$
$$E^- := \{e \in E | x_e \leqslant y_e\}, \qquad N^- := \{i \in N | \sum_{e \in E^-} (x_{i,e} - y_{i,e}) < 0\}.$$

Clearly $N^+ \neq \emptyset$, and as each player distributes the same amount of load in x and y we have $N^+ = N^-$. Choose a player $i \in N^+ = N^-$, then there exist resources e and f such that $x_e > y_e$, $x_{i,e} > y_{i,e}$, $x_f \leqslant y_f$ and $x_{i,f} < y_{i,f}$. Then, we have:

$$\begin{aligned}
&\mu_{i,e}^{+k^*}(y) \\
&< \mu_{i,e}^{+k^*}(x) - (k^*)^2 a_{i,e} && (\text{as } x_e \geqslant y_e + k^* \text{ and } x_{i,e} \geqslant y_{i,e} + k^*) \\
&= k^* \cdot \mu_{i,e}(x) && (\text{by rewriting}) \\
&\leqslant k^* \cdot \mu_{i,f}(x) && (\text{as } x \text{ is the atomic splittable equilibrium}) \\
&= \mu_{i,f}^{-k^*}(x) + (k^*)^2 a_{i,f} && (\text{by rewriting}) \\
&\leqslant \mu_{i,f}^{-k^*}(y). && (\text{as } x_f \leqslant y_f \text{ and } x_{i,f} \leqslant y_{i,f} + k^*)
\end{aligned}$$

This contradicts the fact that y is an equilibrium. Thus, x is the unique k^*-integral splittable equilibrium. ■

Note that we do not know matrix A beforehand, but we do know that $2a_{\max}$ is an upper bound on the values occurring in A. Using Hadamard's inequality we find that $|\det(A)| \leqslant (2a_{\max})^{nm}(nm)^{nm/2}$. Hence, we can find a lower bound of k^*:

$$k^* \geqslant \frac{a_{gcd}^{nm}}{(2a_{\max})^{nm}(nm)^{nm/2}}.$$

For atomic splittable equilibrium x and any k-integral-splittable equilibrium x_k, there exist bounds on $|x_e - (x_k)_e|$ and $|x_{i,e} - (x_k)_{i,e}|$ in terms of k and m. Then, given the equilibrium for some sufficiently small k_0, we are able to compute the correct support set of each player and compute the exact equilibrium by solving system $Ax = b$ as described earlier.

Let x be a Nash equilibrium for an atomic splittable game and x_k an equilibrium for a k-integral splittable game. We show that for each $e \in E$ $|(x_k)_e - x_e| < mk$ (Lemma 3.3.2) and $|(x_k)_{i,e} - x_{i,e}| < m^2 k$ for each $i \in N$ and $e \in E$ (Lemma 3.3.3).

Lemma 3.3.2. *Let x be an equilibrium for an atomic splittable game, and x_k be an equilibrium for a k-integral splittable game. Then $|(x_k)_e - x_e| < mk$ for all $e \in E$.*

Proof. We prove this lemma by contradiction and assume $|(x_k)_e - x_e| \geqslant mk$ for some $e \in E$. There are two cases: $(x_k)_e - x_e \geqslant mk$ or $(x_k)_e - x_e \leqslant -mk$. We discuss why $(x_k)_e - x_e \geqslant mk$ leads to a contradiction. The second case is similar, but with reversed inequalities.

Thus, we assume $(x_k)_e - x_e \geqslant mk$. We then introduce two player sets N_f^+, N_f^- for every resource $f \in E$, where $N_f^+ = \{i \in N | (x_k)_{i,f} > x_{i,f}\}$ and $N_f^- = \{i \in N | (x_k)_{i,f} < x_{i,f}\}$. First note that for every $i \in N_e^+$ we have $(x_k)_e + (x_k)_{i,e} > x_e + x_{i,e} + mk$. Hence,

$$(x_k)_e + (x_k)_{i,e} > x_e + x_{i,e} + k \quad \text{and} \quad (x_k)_{i,e} > x_{i,e} \geqslant 0. \tag{3.3}$$

Then, using the player sets, we obtain:

$$\sum_{i \in N_e^+} ((x_k)_{i,e} - x_{i,e}) + \sum_{i \in N_e^-} ((x_k)_{i,e} - x_{i,e}) = (x_k)_e - x_e.$$

As $\sum_{i \in N_e^-} ((x_k)_{i,e} - x_{i,e}) \leqslant 0$ and $(x_k)_e - x_e \geqslant mk$, we obtain:

$$\sum_{i \in N_e^+} ((x_k)_{i,e} - x_{i,e}) \geqslant mk.$$

The total load distributed by player i does not change, therefore:

$$\sum_{f \neq e} \sum_{i \in N_e^+} (x_k)_{i,f} - x_{i,f} \leqslant -mk.$$

For every resource $f \neq e$ we then obtain:

$$\sum_{f \neq e} \sum_{i \in N_e^+ \cap N_f^-} ((x_k)_{i,f} - x_{i,f}) + \sum_{i \in N_e^+ \cap N_f^+} ((x_k)_{i,f} - x_{i,f}) \leqslant -mk,$$

as $N = N_f^- \cup N_f^= \cup N_+^-$, and for all $i \in N_f^=$ it holds that $(x_k)_{i,f} - x_{i,f} = 0$. By definition of N_f^+, it holds that $\sum_{f \neq e} \sum_{i \in N_e^+ \cap N_f^+} ((x_k)_{i,f} - x_{i,f}) \geqslant 0$. And thus:

$$\sum_{f \neq e} \sum_{i \in N_e^+ \cap N_f^-} ((x_k)_{i,f} - x_{i,f}) \leqslant -mk. \tag{3.4}$$

As $(x_k)_e - x_e \geqslant mk$ and the total load in the system is the same in x and x_k, we have $\sum_{f \neq e}((x_k)_f - x_f) \leqslant -mk$. Therefore

$$\sum_{f \neq e} \sum_{i \in N_e^+ \cap N_f^-} ((x_k)_f - x_f) \leqslant -|N_e^+ \cap N_f^-| mk.$$

We add this to equation (3.4) to obtain

$$\sum_{f \neq e} \sum_{i \in N_e^+ \cap N_f^-} ((x_k)_f - x_f) + ((x_k)_{i,f} - x_{i,f}) \leqslant -(|N_e^+ \cap N_f^-| + 1) mk.$$

By using the pigeonhole principle on the number of resources $f \neq e, f \in E$, there must exists an $f \in E$, $f \neq e$ such that

$$\sum_{i \in N_e^+ \cap N_f^-} ((x_k)_f - x_f) + ((x_k)_{i,f} - x_{i,f}) < -(|N_e^+ \cap N_f^-| + 1) k.$$

Using the pigeonhole principle again on the number of players in $N_e^+ \cap N_f^-$, there must exist an $i \in N_e^+ \cap N_f^-$ such that:

$$(x_k)_f + (x_k)_{i,f} \leqslant x_f + x_{i,f} - k. \tag{3.5}$$

Then:

$$\mu_{i,f}(x) \geqslant_1 \frac{1}{k} \mu_{i,f}^{+k}(x_k) \geqslant_2 \frac{1}{k} \mu_{i,e}^{-k}(x_k) >_3 \mu_{i,e}(x) \tag{3.6}$$

Here $\geqslant_1$ is due to inequality (3.5), $\geqslant_2$ is due to the fact that x_k is a Nash equilibrium and $>_3$ is due to inequality (3.3). Inequality (3.6) now contradicts the fact that x is a Nash equilibrium. ∎

Thus, we established a bound on the difference in total load for an atomic splittable equilibrium and a k-integral splittable equilibrium. We use this bound on the total load to establish bounds on the difference in player-specific load on a resource.

Lemma 3.3.3. *Let x be an equilibrium for an atomic splittable congestion game $\mathcal{G}$, and x_k be an equilibrium for the corresponding k-integral splittable game. Then $|(x_k)_{i,e} - x_{i,e}| < m^2 k$ for all $i \in N$ and $e \in E$.*

Proof. We prove this lemma by contradiction and assume that there exists an $i \in N$ and $e \in E$ such that $|(x_k)_{i,e} - x_{i,e}| \geqslant m^2 k$. Again there are two cases: $(x_k)_{i,e} \geqslant x_{i,e} + m^2 k$ or $(x_k)_{i,e} \leqslant x_{i,e} - m^2 k$. We first discuss why

$(x_k)_{i,e} \geqslant x_{i,e} + m^2k$ leads to a contradiction. The same reasoning holds for $(x_k)_{i,e} \leqslant x_{i,e} - m^2k$, by reversing all inequalities.

Thus, we assume that $(x_k)_{i,e} \geqslant x_{i,e} + m^2k$. From Lemma 3.3.2 we know that $(x_k)_e \geqslant x_e - mk$. Adding both inequalities we obtain

$$(x_k)_e + (x_k)_{i,e} \geqslant x_e + x_{i,e} + m(m-1)k. \tag{3.7}$$

As the total load distributed by player i does not change, and neither does the total load in the system change, we obtain:

$$\sum_{f \neq e} ((x_k)_f + (x_k)_{i,f}) \geqslant \sum_{f \neq e} (x_f + x_{i,f}) - m(m-1)k.$$

As there are $m-1$ remaining resources besides resource e, there must exist at least one resource $f \in E$ such that:

$$(x_k)_f + (x_k)_{i,f} \geqslant x_f + x_{i,f} - mk. \tag{3.8}$$

Note that $x_{i,f} > 0$, as $x_{i,f} = 0$ implies $(x_k)_f < x_f - mk$ which contradicts Lemma 3.3.2. We obtain:

$$\mu_{i,f}(x) \geqslant_1 \frac{1}{k}\mu^{+k}_{i,f}(x_k) \geqslant_2 \frac{1}{k}\mu^{-k}_{i,e}(x_k) >_3 \mu_{i,e}(x). \tag{3.9}$$

Assuming $m \geqslant 2$, $\geqslant_1$ follows from inequality (3.8), $\geqslant_2$ is due to the fact that x_k is a Nash equilibrium and $>_3$ follows from inequality (3.7). Inequality (3.9), combined with $x_{i,f} > 0$, contradicts the fact that x is a Nash equilibrium. Altogether we get $|(x_k)_{i,e} - x_{i,e}| < m^2k$ for all $i \in N$ and $e \in E$. ∎

Thus, if we compute an equilibrium for a sufficiently small k_0, this k_0-integral-splittable equilibrium should be fairly similar to the unique k^*-integral splittable equilibrium. Hence, it enables us to find the correct support sets and compute the exact atomic splittable equilibrium.

Theorem 3.3.4. *Given an atomic splittable congestion game $\mathcal{G}$ and an equilibrium x_{k_0} for k_0-splittable game $\mathcal{G}_{k_0}$, where:*

$$k_0 := \frac{1}{2m^2} \cdot \frac{a_{gcd}^{nm}}{\lceil (2a_{\max})^{nm}(nm)^{nm/2} \rceil},$$

we can compute in $O((nm)^3)$ the exact atomic splittable equilibrium x for game $\mathcal{G}$.

Proof. First note that all demands d_i are integer multiples of k_0, as d_i is an integer multiple of a_{gcd}, and both $2m^2$ and $\lceil (2a_{max})^{nm}(nm)^{nm/2} \rceil$ are integer.

Theorem 3.3.1 implies that there exists a k^* such that the atomic splittable equilibrium is also an equilibrium for the k^*-integral splittable game. In the following we show that there is a load-threshold $m^2 k_0$ that enables us to decide whether or not a resource receives any demand from player i in the equilibrium of the atomic splittable game.

1. If $(x_{k_0})_{i,e} < m^2 k_0$, then $x_{i,e} = 0$. Assume by contradiction that $x_{i,e} > 0$. Remember that the atomic splittable equilibrium is also an k^*-equilibrium and thus, if $x_{i,e} > 0$, then the inequality $x_{i,e} \geqslant k^*$ must hold. We obtain:

$$\begin{aligned} x_{i,e} - (x_{k_0})_{i,e} &> k^* - m^2 k_0 \\ &\geqslant \frac{1}{2} \cdot \frac{a_{gcd}^{nm}}{\lceil (2a_{max})^{nm}(nm)^{nm/2} \rceil} \\ &= m^2 k_0, \end{aligned}$$

   which contradicts Lemma 3.3.3. Thus, $x_{i,e} = 0$.

2. If $(x_{k_0})_{i,e} \geqslant m^2 k_0$, then we prove that $x_{i,e} > 0$. On the contrary, we assume that $x_{i,e} = 0$. In this case we have $(x_{k_0})_{i,e} - x_{i,e} \geqslant m^2 k_0$, which contradicts Lemma 3.3.3. Thus, $x_{i,e} > 0$.

Hence, given an equilibrium (x_{k_0}) for k_0-splittable game $\mathcal{G}_{k_0}$, we can compute the correct support sets I_i for all $i \in N$, where

$$I_i := \{e \in E \mid (x_{k_0})_{i,e} \geqslant m^2 k_0\}.$$

Given the correct support sets, we can easily compute the correct, exact equilibrium by solving the system $Ax = b$ of at most nm linear equations in running time $O((nm)^3)$ using Gaussian elimination [55]. ∎

It is left to compute an equilibrium x_{k_0} for integral game $\mathcal{G}_{k_0}$. Such integral games have been studied in the literature before, see Harks et al. [32, 36]. In particular, [32, Algorithm 1] has running time $O(nm(\delta/k_0)^3)$. Here δ is an upper bound on the player-specific demands. In general, δ is not bounded in k_0, thus, the running time is not polynomially bounded in the size of the input.

3.4 A POLYNOMIAL ALGORITHM FOR INTEGRAL GAMES

The goal of this section is to develop a *polynomial time* algorithm that computes an equilibrium for any k-integral splittable singleton game with player-specific affine cost functions. We use elements of [72, Algorithm 1] and [36, Algorithm 1] to construct a new algorithm that has a running time in $O(n^2m^{14}\log(\delta/k))$. We first introduce some new notation. For two vectors $x_i, y_i \in \mathbb{R}^{|E|}$, we denote their Hamming distance by:

$$H(x_i, y_i) := \sum_{e\in E} |x_{i,e} - y_{i,e}|.$$

For two strategy profiles x and y, we denote their Hamming distance by:

$$H(x, y) := \sum_{e\in E} |x_e - y_e|.$$

For two resources $e^-, e^+ \in E$, we denote $y_i = (x_i)_{e^-\to e^+}$ if $y_{i,e^-} = x_{i,e^-} - k$, $y_{i,e^+} = x_{i,e^+} + k$ and $y_{i,e} = x_{i,e}$ for all $e \in E\backslash\{e^-, e^+\}$. If x is a strategy profile for some game $\mathcal{G}_k$ and $y_i = (x_i)_{e^-\to e^+}$, we denote $(y_i, x_{-i}) = x_{i:e^-\to e^+}$. We define a *restricted best response*:

Definition 3.4.1. *Let x be a strategy profile for game $\mathcal{G}_k((d_i)_{i\in N})$. Assume there exists $e^-, e^+ \in E$ such that $e^- \in \arg\max\{\mu^{-k}_{i,e}(x)\}$, $e^+ \in \arg\min\{\mu^{+k}_{i,e}(x)\}$ and $\mu^{-k}_{i,e^-}(x) > \mu^{+k}_{i,e^+}(x)$. Then, we term strategy $y_i = (x_i)_{e^-\to e^+}$ a* restricted best response *to x for player i.*

Note that when y_i is a restricted best response to x_i, $H(x_i, y_i) = 2k$. We first describe two subroutines, termed Add and Restore.

3.4.1 Add

The first subroutine, Add, is described in Algorithm 2 and consists of lines 4-10 of [36, Algorithm 1]. Given an equilibrium x_k for game $\mathcal{G}_k((d_i)_{i\in N})$, it computes an equilibrium for the game, where the demand for player j is increased by k. First it decides on the best resource f for player j to put her new packet. In effect, the load on resource f increases and only those players with $x_{i,f} > 0$ can potentially decrease their cost by a deviation. In this case, Harks et al. proved in [36, Theorem 3.2] that a best response can be obtained by a restricted best response moving a packet away from f. Thus, only one packet is moved throughout, preserving the invariant that only players using a resource to which the packet is moved may have an incentive to profitably deviate.

Algorithm 2: Subroutine ADD$(x, i, \mathcal{G}_k((d'_i)_{i \in N}))$

Input: equilibrium x_k for $\mathcal{G}_k((d'_i)_{i \in N})$, player i

Output: equilibrium x'_k for $\mathcal{G}_k((d'_i)_{i \in N})$, where $d'_j \leftarrow d'_j + k$;

$x \leftarrow x_k$; $d'_j \leftarrow d'_j + k$; $\mathcal{S}'_j \leftarrow \mathcal{S}_j(d'_j, k)$;

Choose $f \in \arg\min\{\mu^{+k}_{j,e}(x)\}$;

$x_{j,f} \leftarrow x_{j,f} + k$;

while $\exists i \in N$ *who can improve in* $\mathcal{G}_k$ **do**

  Compute a restricted best response $y_i \in \mathcal{S}'_i$;

  $x_i \leftarrow y_i$;

end

$x'_k \leftarrow x$;

return x'_k

3.4.2 RESTORE

The second subroutine, RESTORE, takes as input an equilibrium x_{2k} for packet size $2k$ and game $\mathcal{G}_k((d_i)_{i \in N})$, and constructs an equilibrium for packet size k. This algorithm makes use of two sub-algorithms: Algorithm 3 and Algorithm 6. In Algorithm 3 we create a *backward path of restricted best responses*. In a backward path we are given a resource e^-_1 and a strategy profile x^b_1. In iteration q, we decide if there exists a player i that has a restricted best response from some e^-_{q+1} to e^-_q, and if so, we define $x^b_{q+1} \leftarrow (x^b_q)_{i:e^-_{q+1} \to e^-_q}$. If no player has a restricted best response to resource e^-_q, we check if $(x^b_q)_{e^-_q} > (x_{2k})_{e^-_q} - 2mk$. If so, we end our backward path. Else, we look for a player that has an improving move in which she shifts one packet from some e^-_{q+1} to e^-_q, and then continue the backward path. Note that in each step we preserve the invariant that $H(x^b_1, x^b_q) \in \{0, 2k\}$.

In Algorithm 6 we create a *forward path of restricted best responses*. A forward path is very similar to a backward path, but we change the perspective. Thus, given a resource e^+_q and a strategy profile x^f_q, we check in iteration q if there exists a player that has a restricted best response from e^+_q to some e^+_{q+1}. As this algorithm is very similar to Algorithm 3, it is moved to Appendix 3.B. Both algorithms (back- and forward path) can be seen as a special instantiation of a general restricted best response dynamic (cf. [36, 72]). We are now ready to define subroutine RESTORE.

Algorithm 3: $BP(x_{2k}, x_1^b, e_1^-, \mathcal{G}_k)$: A backward path of restricted best responses.

Input: equilibrium x_{2k} for game $\mathcal{G}_{2k}$, strategy profile x_1^b for game $\mathcal{G}_k$, a resource e_1^- and game $\mathcal{G}_k$.

Output: Strategy profile $x_{q_b}^b$ for game $\mathcal{G}_k$ and resource e_q^-.

Initialize $q \leftarrow 1$;
repeat
  if $(x_q^b)_{e_q^-} \leqslant (x_{2k})_{e_q^-} - 2mk$ **then**
    Find player i and resource e_{q+1}^- satisfying properties:
    B1. $(x_q^b)_{i,e_q^-} < (x_{2k})_{i,e_q^-}$;
    B2. $e_{q+1}^- \in \arg\max\{\mu_{i,e}^{-k}(x_q^b)\}$;
    B3. $\mu_{i,e_q^-}^{+k}(x_q^b) < \mu_{i,e_{q+1}^-}^{-k}(x_q^b)$;
    Define $x_{q+1}^b \leftarrow (x_q^b)_{i:e_{q+1}^- \to e_q^-}$;
    $q \leftarrow q+1$;
  end
  while $\exists i \in N$ *with* $e_q^- = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_q^b)\}$ *and* $\mu_{i,\min}^{+k}(x_q^b) < \mu_{i,\max}^{-k}(x_q^b)$ **do**
    Choose $e_{q+1}^- \in \arg\max\{\mu_{i,e}^{-k}(x_q^b)\}$;
    $x_{q+1}^b \leftarrow (x_q^b)_{i:e_{q+1}^- \to e_q^-}$;
    $q \leftarrow q+1$;
  end
until $(x_q^b)_{e_q^-} > (x_{2k})_{e_q^-} - 2mk$;
return (x_q^b, e_q^-)

We initialize x by equilibrium x_{2k}. While x is not an equilibrium for $\mathcal{G}_k$, we iterate the following. Among players who can improve, we find the player j that benefits most from a restricted best response. We carry out a restricted best response for player j and move a packet from some resource e_1^- to some e_1^+. Then we compute a backward path, starting in resource e_1^-. If the resulting strategy profile has Hamming distance zero with x, we stop this iteration and overwrite x by the resulting strategy profile. Else, we compute a forward path, starting in e_1^+ and overwrite x by the resulting strategy profile. The pseudo-code of subroutine RESTORE can be found in Algorithm 4.

Algorithm 4: Subroutine RESTORE$(x, k, (d'_i)_{i\in N}, \mathcal{G})$

Input: equilibrium x_{2k} for $\mathcal{G}_{2k}((d'_i)_{i\in N})$
Output: equilibrium x_k for $\mathcal{G}_k((d'_i)_{i\in N})$

$x \leftarrow x_{2k}$;
while x *not an equilibrium for* $\mathcal{G}_k((d'_i)_{i\in N})$ **do**
  $j \leftarrow \arg\max_{i\in N}\{\mu^{+k}_{i,\min}(x) - \mu^{-k}_{i,\max}(x)\}$;
  Choose $e^-_1 \in \arg\max\{\mu^{-k}_{j,e}(x)\}$ and $e^+_1 \in \arg\min\{\mu^{+k}_{j,e}(x)\}$;
  $x^b_1 \leftarrow x_{j:e^-_1 \to e^+_1}$;
  $(x^b_{q_b}, e^-_{q_b}) \leftarrow \mathrm{BP}(x_k, x^b_1, e^-_1, \mathcal{G}_k)$;
  if $e^+_1 \neq e^-_{q_b}$ **then**
    $x^f_1 \leftarrow x^b_{q_b}$;
    $(x^f_{q_f}, e^+_{q_f}) \leftarrow \mathrm{FP}(x_k, x^f_1, e^+_1, \mathcal{G}_k)$;
    $x' \leftarrow x^f_{q_f}$;
  else
    $x' \leftarrow x^b_{q_b}$;
  end
  $x \leftarrow x'$;
end
$x_k \leftarrow x$;
return x_k;

3.4.3 PACKETHALVER

Using the subroutines ADD and RESTORE we develop PACKETHALVER, which computes an equilibrium x_{k_0} for the k_0-splittable game $\mathcal{G}_{k_0}((d_i)_{i\in N})$.

In this algorithm we start with an equilibrium x_k for $\mathcal{G}_k((d'_i)_{i\in N})$, where $d'_i = 0$ for all $i \in N$, $k = 2^{q_1}k_0$ and $q_1 = \arg\min_{q\in\mathbb{N}}\{2^q k_0 > \max_{i\in N} d_i\}$. Note that this game has a trivial equilibrium, where $(x_k)_{i,e} = 0$ for all $i \in N$ and $e \in E$. We repeat the following two steps:

- Given an equilibrium x_k for $\mathcal{G}_k((d'_i)_{i\in N})$, we construct an equilibrium for $\mathcal{G}_{k/2}((d'_i)_{i\in N})$ using subroutine RESTORE and set k to $k/2$.
- For each player $i \in N$ we check if $d_i - d'_i \geqslant k$. If so, we increase d'_i by k and recompute equilibrium x_k using subroutine ADD.

After q_1 iterations PACKETHALVER returns an equilibrium x_{k_0} for game $\mathcal{G}_{k_0}((d_i)_{i\in N})$. The pseudo-code of PACKETHALVER can be found in Algorithm 5.

Algorithm 5: Algorithm PacketHalver($\mathcal{G}_{k_0}((d_i)_{i\in N})$)

Input: Integral splittable congestion game
$\mathcal{G}_{k_0} = (N, E, (d_i)_{i\in N}, (c_{i,e})_{i\in N, e\in E})$.
Output: An equilibrium x_{k_0} for $\mathcal{G}_{k_0}$.

Initialize $q_1 = \arg\min_{q\in\mathbb{N}}\{2^q k_0 > \max_{i\in N} d_i\}$; $k \leftarrow 2^{q_1}k_0$; $d'_i \leftarrow 0$; $x_k \leftarrow (0)_{e\in E, i\in N}$;
for $1, \ldots, q_1 - 1$ **do**
    $k \leftarrow k/2$;
    $x_k \leftarrow$ Restore($x_{2k}, \mathcal{G}_k((d'_i)_{i\in N})$);
    for $i \in N$ **do**
        if $d_i - d'_i > k$ **then**
            $x_k \leftarrow$ Add($x_k, i, \mathcal{G}_k((d'_i)_{i\in N})$);
            $d'_i \leftarrow d'_i + k$;
        end
    end
end
return x_k;

3.5 CORRECTNESS

In this section, we prove that PacketHalver indeed returns an equilibrium for game $\mathcal{G}_{k_0}((d_i)_{i\in N})$. In order to do so, we first need to verify that the two subroutines Add and Restore are correct. Subroutine Add is proven to be correct by Harks, Peis, and Klimm [36], thus, it is left to verify correctness of Restore and PacketHalver.

3.5.1 *Correctness* Restore

To verify the correctness of subroutine Restore($x_{2k}, \mathcal{G}_k((d_i)_{i\in N})$), we need to prove that Algorithm 3 and Algorithm 6 are well-defined, and that Restore terminates.

To prove that Algorithm 3 is well-defined, we need to verify that if $(x_q^b)_{e_q^-} \leqslant (x_{2k})_{e_q^-} - 2mk$, there exists a player i and a resource e_{q+1}^- satisfying Property **B1**, **B2** and **B3**.

Lemma 3.5.1. *In Algorithm 3, if* $(x_q^b)_{e_q^-} \leqslant (x_{2k})_{e_q^-} - 2mk$*, then we can find a player* i *and a resource* e_{q+1}^- *satisfying Property* ***B1****,* ***B2*** *and* ***B3****.*

Proof. The idea for this proof is very similar to the proof of Lemma 3.3.2. When $(x^b_q)_{e^-_q} \leqslant (x_{2k})_{e^-_q} - 2mk$, we can find a player that can decrease its cost by moving a packet from some e^-_{q+1} to e^-_q. The maximum cost decrease is attained by choosing $e^-_{q+1} \in \max_{e \in E}\{\mu^{-k}_{i,e}\}$. The full proof can be found in Appendix 3.C. ■

Algorithm 6 is well-defined if we can find a player i and a resource e^+_{q+1} satisfying Property **F1**, **F2** and **F3** whenever $(x^f_q)_{e^+_q} \geqslant (x_{2k})_{e^+_q} + 2mk$.

Lemma 3.5.2. *In Algorithm 6, if* $(x^f_q)_{e^+_q} \geqslant (x_{2k})_{e^+_q} + 2mk$*, then we can find a player* i *and a resource* e^+_{q+1} *satisfying Property* ***F1****,* ***F2*** *and* ***F3****.*

Proof. The proof of this lemma is omitted as it is similar to the proof of Lemma 3.5.1, where all inequalities are reversed. ■

It is left to prove that RESTORE terminates. We define:

$$\begin{aligned}\Delta(x) &:= (\mu^{+k}_{i,\min}(x) - \mu^{-k}_{i,e}(x))_{i \in N; e \in E},\\ \Delta_i(x) &:= (\mu^{+k}_{i,\min}(x) - \mu^{-k}_{i,e}(x))_{e \in E}.\end{aligned}$$

Let $\Delta_{\min}(x)$ be the minimum value in $\Delta(x)$. Note that when all elements in $\Delta(x)$ are non-negative, or, equivalently, when $\Delta_{\min}(x)$ is non-negative, x is an equilibrium. Our goal is to show that after each iteration in the while-loop (lines 2 - 15 of RESTORE) $\Delta(x)$ increases according to a certain lexicographical order defined as follows. Given two vectors $u, v \in \mathbb{R}^n$, we say that v is *sorted lexicographically larger* than u, if there is an $k \in \{1, \dots, n\}$ such that $u_{\phi(i)} = v_{\psi(i)}$ for all $i < k$ and $u_{\phi(k)} < v_{\psi(k)}$, where ϕ and ψ are permutations that sort u and v non-decreasingly. We write $u <_{\text{lex}} v$. If $u_{\phi(i)} = v_{\psi(i)}$ for all $i \in \{1, \dots, n\}$, we write $u =_{\text{lex}} v$.

Proving that $\Delta(x)$ sorted lexicographically increases implies that RESTORE does not cycle, and thus, as the strategy space is finite, terminates. In general, under the hypothesis that $\Delta(x)$ lexicographically increases, we obtain the following (finite) sequence of strategy profiles within a while-loop:

$$x \to x^b_1 \to x^b_2 \to \cdots \to x^b_{q_b} = x^f_1 \to x^f_2 \to \cdots \to x^f_{q_f} = x'. \qquad (3.10)$$

We introduce two types of vectors that help us prove that $\Delta(x) <_{\text{lex}} \Delta(x')$. Let q be the iteration in the backward path where a player moves a packet from e^-_{q+1} to e^-_q. We define values $B^{q,-k}_{i,e}(x)$ and $B^{q,+k}_{i,\min}(x)$ as follows:

$$B^{q,-k}_{i,e}(x) = \begin{cases} \mu^{-k}_{i,e}(x) + k^2 a_{i,e}, & \text{if } e^-_q \neq e^+_1 \text{ and } e = e^-_q, \\ \mu^{-k}_{i,e}(x) - k^2 a_{i,e}, & \text{if } e^-_q \neq e^+_1 \text{ and } e = e^+_1, \\ \mu^{-k}_{i,e}(x), & \text{otherwise.} \end{cases} \tag{3.11}$$

$$B^{q,+k}_{i,\min}(x) = \begin{cases} \mu^{+k}_{i,\min}(x) + k^2 a_{i,e}, & \text{if } e^-_q \neq e^+_1 \text{ and } e^-_q = \arg\min \mu^{+k}_{i,e}(x), \\ \mu^{+k}_{i,\min}(x) - k^2 a_{i,e}, & \text{if } e^-_q \neq e^+_1 \text{ and } e^+_1 = \arg\min \mu^{+k}_{i,e}(x), \\ \mu^{+k}_{i,\min}(x), & \text{otherwise.} \end{cases} \tag{3.12}$$

As k is fixed within Restore, we write $B^{q,-}_{i,e}(x)$ and $B^{q,+}_{i,\min}(x)$ instead. We define:

$$\begin{aligned} B^q_i(x) &= (B^{q,+}_{i,\min}(x) - B^{q,-}_{i,e}(x))_{e \in E}, \\ B^q(x) &= (B^{q,+}_{i,\min}(x) - B^{q,-}_{i,e}(x))_{i \in N; e \in E}. \end{aligned}$$

Let q be the iteration in the forward path, where some player moves a packet from e^+_q to e^+_{q+1}. We define values $F^{q,-k}_{i,e}(x)$ and $F^{q,+k}_{i,\min}(x)$ as follows:

$$F^{q,-k}_{i,e}(x) = \begin{cases} \mu^{-k}_{i,e}(x) - k^2 a_{i,e}, & \text{if } e = e^+_q, \\ \mu^{-k}_{i,e}(x), & \text{otherwise.} \end{cases} \tag{3.13}$$

$$F^{q,+k}_{i,\min}(x) = \begin{cases} \mu^{+k}_{i,\min}(x) - k^2 a_{i,e^f_q}, & \text{if } e^+_q = \arg\min \mu^{+k}_{i,e}(x), \\ \mu^{+k}_{i,\min}(x), & \text{otherwise.} \end{cases} \tag{3.14}$$

Again, as k is fixed, we write $F^{q,-}_{i,e}(x)$ and $F^{q,+}_{i,\min}(x)$ instead. We define:

$$\begin{aligned} F^q_i(x) &= (F^{q,+}_{i,\min}(x) - F^{q,-}_{i,e}(x))_{e \in E}, \\ F^q(x) &= (F^{q,+}_{i,\min}(x) - F^{q,-}_{i,e}(x))_{i \in N; e \in E}. \end{aligned}$$

In order to show $\Delta(x) <_{\text{lex}} \Delta(x')$, we first prove that (3.10) is actually well-defined.

Lemma 3.5.3. *Let x, x^b_q and x^f_q be strategy profiles as described in* Restore, *then $\Delta(x) <_{\text{lex}} B^1(x^b_1) <_{\text{lex}} \cdots <_{\text{lex}} B^{q_b}(x^b_{q_b})$ and $F^1(x^f_1) <_{\text{lex}} \cdots <_{\text{lex}} F^{q_f}(x^f_{q_f})$.*

Proof. We first prove that $\Delta(x) <_{\text{lex}} B^1(x_1^b)$. Remember that there exists an $i \in N$ such that $x_1^b = x_{i:e_1^+ \to e_1^-}$. Using Equation (3.11) and (3.12), for all players $j \neq i$, it holds that:

$$B_{j,e}^{1,-}(x_1^b) = \mu_{j,e}^{-k}(x) \quad \text{and} \quad B_{j,\min}^{1,+}(x_1^b) = \mu_{j,\min}^{+k}(x).$$

So, for all players $j \neq i$, we have $\Delta_j(x) =_{\text{lex}} B_j^1(x_1^b)$. Hence, it is left to prove that $\Delta_i(x) <_{\text{lex}} B_i^1(x_1^b)$. As $\arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\} \neq \emptyset$, at least one of the following is true: (I) $f \in \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$ for some $f \in E \backslash \{e_1^+, e_1^-\}$, (II) $e_1^+ \in \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$ or (III) $e_1^- \in \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$. As all three cases are proven similarly, we only will prove the first case here. Proofs for the other two cases can be found in Appendix 3.D. For the first case, assume that $f \in \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$ for some $f \in E \backslash \{e_1^+, e_1^-\}$. For resource e_2^+ we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e_1^+}^{1,-}(x_1^b) \\
&= \mu_{i,\min}^{+k}(x_1^b) - (\mu_{i,e_1^+}^{-k}(x_1^b) - a_{i,e_1^+}k^2) && \text{(by Definition 3.11 and 3.12)} \\
&> \mu_{i,f}^{+k}(x_1^b) - \mu_{i,e_1^+}^{-k}(x_1^b) && (\text{as } f = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}) \\
&= \mu_{i,f}^{+k}(x) - \mu_{i,e_1^+}^{+k}(x) && (\text{as } \mu_{i,f}^{+k}(x_1^b) = \mu_{i,f}^{+k}(x) \text{ and} \\
& && \mu_{i,e_1^+}^{-k}(x_1^b) = \mu_{i,e_1^+}^{+k}(x)) \\
&> \mu_{i,f}^{+k}(x) - \mu_{i,e_1^-}^{-k}(x) && (\text{as } \mu_{i,e_1^+}^{+k}(x) < \mu_{i,e_1^-}^{-k}(x)) \\
&\geqslant \mu_{i,\min}^{+k}(x) - \mu_{i,e_1^-}^{-k}(x). && (\text{by definition of } \mu_{i,\min}^{+k}(x))
\end{aligned}$$

For resource e_1^- we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e_1^-}^{1,-}(x_1^b) \\
&= \mu_{i,f}^{+k}(x_1^b) - (\mu_{i,e_1^-}^{-k}(x_1^b) + k^2 a_{i,e_1^-}) && \text{(by Definition 3.11 and 3.12)} \\
&= \mu_{i,f}^{+k}(x) - (\mu_{i,e_1^-}^{-k}(x) - k^2 a_{i,e_1^-}) && (\text{as } \mu_{i,e'}^{+k}(x_1^b) = \mu_{i,e'}^{+k}(x) \text{ and} \\
& && \mu_{i,e_1^-}^{-k}(x_1^b) = \mu_{i,e_1^-}^{-k}(x) - 2k^2 a_{i,e_1^-}) \\
&> \mu_{i,\min}^{+k}(x) - \mu_{i,e_1^-}^{-k}(x). && (\text{by definition of } \mu_{i,\min}^{+k}(x))
\end{aligned}$$

For resource $e \in E \backslash e_1^+, e_1^-$ we have:

$$\begin{aligned}
&B^{1,+}_{i,\min}(x_1^b) - B^{1,-}_{i,e}(x_1^b) \\
&= \mu^{+k}_{i,e'}(x_1^b) - \mu^{-k}_{i,e}(x_1^b) && \text{(by Definition 3.11 and 3.12)} \\
&= \mu^{+k}_{i,e'}(x) - \mu^{-k}_{i,e}(x). && \text{(as } \mu^{+k}_{i,e'}(x_1^b) = \mu^{+k}_{i,e'}(x) \\
& && \text{ and } \mu^{-k}_{i,e}(x_1^b) = \mu^{-k}_{i,e}(x))
\end{aligned}$$

Thus when $e' = \arg\min_{e \in E}\{\mu^{+k}_{i,e}(x_1^b)\}$, we have $\Delta_i(x) <_{\text{lex}} B^1_i(x_1^b)$. Hence, $\Delta(x) <_{\text{lex}} B^1(x_1^b)$.

The argumentation above depends on three crucial factors:

- The fact that $x_1^b = x_{i:e_1^- \to e_1^+}$.
- Moving a packet from resource e_1^- to e_1^+ is a restricted improving move for player i, i.e., $\mu^{+k}_{i,e_1^+}(x) < \mu^{-k}_{i,e_1^-}(x)$).
- The relation between $\mu^{-k}_{i,e}(x)$ and $B^{1,-}_{i,e}(x)$ and the relation between $\mu^{+k}_{i,\min}(x)$ and $B^{1,-}_{i,\min}(x)$.

Similar relations hold for $x_q^b, x_{q+1}^b, B^q(x)$ and $B^{q+1}(x)$. Thus, using similar arguments, $B^q(x_q^b) <_{\text{lex}} B^{q+1}(x_{q+1}^b)$ for all $q \in \{1, \dots q_b - 1\}$. The same holds for $x_q^f, x_{q+1}^f, F^q(x)$ and $F^{q+1}(x)$. Hence, $F^q(x_q^f) <_{\text{lex}} F^{q+1}(x_{q+1}^f)$ for all $q \in \{1, \dots q_f - 1\}$. Thus, $\Delta(x) <_{\text{lex}} B^1(x_1^b) <_{\text{lex}} \cdots <_{\text{lex}} B^{q_b}(x_{q_b}^b)$ and $F^1(x_1^f) <_{\text{lex}} \cdots <_{\text{lex}} F^{q_f}(x_{q_f}^f)$. ■

Hence, the backward path and the forward path end after a finite number of steps. We need two more lemmas to connect vectors $\Delta(x), B^q(x)$ and $F^q(x)$.

Definition 3.5.4. *Let $\alpha \in \mathbb{Q}$ and $y \in \mathbb{Q}^{|I|}$ for a finite set I. We define function $\#(\alpha, y)$ to output the number of times that α occurs in y.*

Lemma 3.5.5. *Let x and $x_{q_b}^b$ be as in* RESTORE, *and assume that $e_1^+ \neq e_{q_b}^-$. If $\Delta(x) <_{\text{lex}} B^{q_b}(x_{q_b}^b)$, then $\Delta(x) <_{\text{lex}} F^1(x_{q_b}^b)$ and*

$$\#(\Delta_{\min}(x), F^1(x_{q_b}^b)) \leqslant \#(\Delta_{\min}(x), B^{q_b}(x_{q_b}^b)).$$

Proof. If for all $i \in N$ and $e \in E$ we have:

$$B^{q_b,+}_{i,\min}(x_{q_b}^b) - B^{q_b,-}_{i,e}(x_{q_b}^b) \leqslant F^{1,+}_{i,\min}(x_{q_b}^b) - F^{1,-}_{i,e}(x_{q_b}^b),$$

then $\Delta(x) <_{\text{lex}} B^{q_b}(x^b_{q_b}) \leqslant_{\text{lex}} F^1(x^b_{q_b})$ and the lemma follows. Therefore, assume that for some $i \in N$ and $e \in E$ we have:

$$B^{q_b,+}_{i,\min}(x^b_{q_b}) - B^{q_b,-}_{i,e}(x^b_{q_b}) > F^{1,+}_{i,\min}(x^b_{q_b}) - F^{1,-}_{i,e}(x^b_{q_b}).$$

Using the definitions of $F^{1,-}_{i,e}(x), F^{1,+}_{i,\min}(x), B^{q_b,-}_{i,e}(x)$ and $B^{q_b,+}_{i,\min}(x)$ we obtain:

$$F^{1,-}_{i,e}(x) = \begin{cases} B^{q_b,-}_{i,e}(x) - a_{i,e}k, & \text{if } e = e^-_{q_b}, \\ B^{q_b,-}_{i,e}(x), & \text{otherwise.} \end{cases}$$

$$F^{1,+}_{i,\min}(x) = \begin{cases} B^{q_b,+}_{i,\min}(x) - a_{i,e}k, & \text{if } e^-_{q_b} = \arg\min\{\mu^{+k}_{i,e}(x)\}, \\ B^{q_b,+}_{i,\min}(x), & \text{otherwise.} \end{cases}$$

Therefore, $e^-_{q_b} = \arg\min \mu^{+k}_{i,e}(x^b_{q_b})$, which implies:

$$\begin{aligned} & B^{q_b,+}_{i,\min}(x^b_{q_b}) - B^{q_b,-}_{i,e}(x^b_{q_b}) \\ & \quad > F^{1,+}_{i,e^-_{q_b}}(x^b_{q_b}) - F^{1,-}_{i,e}(x^b_{q_b}) \quad (\text{as } e^-_{q_b} = \arg\min \mu^{+k}_{i,e}(x^b_{q_b})) \\ & \quad \geqslant \mu^{+k}_{i,e^-_{q_b}}(x^b_{q_b}) - \mu^{-k}_{i,e}(x^b_{q_b}) \quad (\text{as } e^-_{q_b} \neq e^+_1) \\ & \quad \geqslant 0. \qquad (\text{as } e^-_{q_b} \text{ is the end of the backward path}) \end{aligned}$$

Thus, if $B^{q_b}(x^b_{q_b}) >_{\text{lex}} F^1(x^b_{q_b})$, it is caused by some positive values in $B^{q_b}(x^b_{q_b})$ decreasing to some smaller positive values in $F^1(x^b_{q_b})$. As x is not an equilibrium, $\Delta(x)$ contains a negative value that is increased by the initial restricted best response. Thus, if $\Delta(x) <_{\text{lex}} B^{q_b}(x^b_{q_b})$, then it holds that $\Delta(x) <_{\text{lex}} F^1(x^b_{q_b})$. Moreover, as $\Delta_{\min}(x) < 0$:

$$\#(\Delta_{\min}(x), F^1(x^b_{q_b})) \leqslant \#(\Delta_{\min}(x), B^{q_b}(x^b_{q_b})).$$

■

Lemma 3.5.6. *Let $x^f_{q_f}$ be as described in* RESTORE. *If $\Delta(x) <_{\text{lex}} F^{q_f}(x^f_{q_f})$, then $\Delta(x) <_{\text{lex}} \Delta(x^f_{q_f})$. Moreover, $\#(\Delta_{\min}(x), \Delta(x^f_{q_f})) \leqslant \#(\Delta_{\min}(x), F^{q_f}(x^f_{q_f}))$.*

Proof. If for all $i \in N$ and $e \in E$ we have:

$$F^{q_f,+}_{i,\min}(x^f_{q_f}) - F^{q_f,-}_{i,e}(x^f_{q_f}) \leqslant \mu^{+k}_{i,\min}(x^f_{q_f}) - \mu^{-k}_{i,e}(x^f_{q_f}),$$

then $\Delta(x) <_{\text{lex}} F^{q_f}(x^f_{q_f}) \leqslant_{\text{lex}} \Delta(x^f_{q_f})$ and the lemma follows. Therefore, assume that for some $i \in N$ and $e \in E$ we have:

$$F^{q_f,+}_{i,\min}(x^f_{q_f}) - F^{q_f,-}_{i,e}(x^f_{q_f}) > \mu^{+k}_{i,\min}(x^f_{q_f}) - \mu^{-k}_{i,e}(x^f_{q_f}).$$

Definition 3.13 and 3.14 imply that in this case $e = e^+_{q_f}$. As $e^+_{q_f}$ is the end of the backward path, for all $i \in N$ we have either (I) $e^+_{q_f} \notin \arg\max\{\mu^{-k}_{i,e}(x^f_{q_f})\}$ or (II) $\mu^{+k}_{i,\min}(x^f_{q_f}) - \mu^{-k}_{i,\max}(x^f_{q_f}) \geqslant 0$. We first assume $e' \in \arg\max\{\mu^{-k}_{i,e}(x^f_{q_f})\}$:

$$\begin{aligned}
&\mu^{+k}_{i,\min}(x^f_{q_f}) - \mu^{-k}_{i,e^+_{q_f}}(x^f_{q_f}) \\
&\quad > \mu^{+k}_{i,\min}(x^f_{q_f}) - \mu^{-k}_{i,e'}(x^f_{q_f}) && (\text{as } e'_{q_f} \in \arg\max\{\mu^{-k}_{i,e}(x^f_{q_f})\}) \\
&\quad \geqslant F^{q_f,+}_{i,\min}(x^f_{q_f}) - F^{q_f,-}_{i,e'}(x^f_{q_f}) && (\text{as } e^+_{q_f} \neq e') \\
&\quad \geqslant \Delta_{\min}(x). && (\text{as } \Delta(x) <_{\text{lex}} F^{q_f}(x^f_{q_f}))
\end{aligned}$$

In the second case, as $0 > \Delta_{\min}(x)$:

$$\mu^{+k}_{i,\min}(x^f_{q_f}) - \mu^{-k}_{i,e^+_{q_f}}(x^f_{q_f}) \geqslant 0 > \Delta_{\min}(x).$$

Thus, if $F^{q_f}(x^f_{q_f}) >_{\text{lex}} \Delta(x^f_{q_f})$, it is caused by some values in $F^{q_f}(x^f_{q_f})$ decreasing to some values in $\Delta(x^f_{q_f})$ that are larger than $\Delta_{\min}(x)$. Thus, if $\Delta(x) <_{\text{lex}} F^{q_f}(x^f_{q_f})$, then $\Delta(x) <_{\text{lex}} \Delta(x^f_{q_f})$. Moreover,

$$\#(\Delta_{\min}(x), \Delta(x^f_{q_f})) \leqslant \#(\Delta_{\min}(x), F^{q_f}(x^f_{q_f})).$$

■

Using Lemma 3.5.3, 3.5.5 and 3.5.6 we prove the following statement.

Lemma 3.5.7. *Let x and x' be defined as in the while-loop (lines 2-15) of* RESTORE. *Then $\Delta(x) <_{\text{lex}} \Delta(x')$, and moreover, $\#(\Delta_{\min}(x), \Delta(x)) > \#(\Delta_{\min}(x), \Delta(x'))$.*

Proof of Lemma 3.5.7. We first prove that $\Delta(x) <_{\text{lex}} \Delta(x')$. Lemma 3.5.3 implies:

$$\Delta(x) <_{\text{lex}} B^1(x^b_1) <_{\text{lex}} \cdots <_{\text{lex}} B^{q_b}(x^b_{q_b}).$$

If $e^+_1 = e^-_{q_b}$, Definition (3.11) and (3.12) imply that $B^{q_b}(x^b_{q_b}) = \Delta(x^b_{q_b})$. Hence, $\Delta(x) <_{\text{lex}} \Delta(x^b_{q_b}) = \Delta(x')$. On the other hand, if $e^+_1 \neq e^-_{q_b}$, we combine the fact that $\Delta(x) <_{\text{lex}} B^{q_b}(x^b_{q_b})$ with Lemma 3.5.5 to obtain: $\Delta(x) <_{\text{lex}} F^1(x^b_{q_b}) = F^1(x^f_1)$. Lemma 3.5.3 implies:

$$F^1(x^f_1) \leqslant_{\text{lex}} \cdots \leqslant_{\text{lex}} F^{q_f}(x^f_{q_f}).$$

Hence, $\Delta(x) <_{\text{lex}} F^{q_f}(x^f_{q_f})$. We use Lemma 3.5.6 to obtain the desired result:

$$\Delta(x) <_{\text{lex}} \Delta(x^f_{q_f}) =_{\text{lex}} \Delta(x').$$

Therefore, if x and x' are as defined in the while-loop, then $\Delta(x) <_{\text{lex}} \Delta(x')$. For the second part of the lemma, we have:

$$\begin{aligned}
&\#(\Delta_{\min}(x), \Delta(x)) \\
&\quad > \#(\Delta_{\min}(x), B^1(x_1^b)) && (\text{as } \min_{i \in N}\{\mu_{j,\min}^{+k}(x) - \mu_{j,\max}^{-k}\} = \Delta_{\min}(x)) \\
&\quad \geqslant \#(\Delta_{\min}(x), B^{q_b}(x_{q_b}^b)) && (\text{as } B^1(x_1^b) <_{\text{lex}} B^{q_b}(x_{q_b}^b)) \\
&\quad = \#(\Delta_{\min}(x), F^1(x_1^f)) && (\text{by Lemma 3.5.5 and as } x_{q_b}^b = x_1^f) \\
&\quad \geqslant \#(\Delta_{\min}(x), F^{q_f}(x_{q_f}^f)) && (\text{as } F^1(x_1^f) <_{\text{lex}} F^{q_f}(x_{q_f}^f)) \\
&\quad \geqslant \#(\Delta_{\min}(x), \Delta(x')). && (\text{by Lemma 3.5.6 and as } x_{q_f}^f = x')
\end{aligned}$$

If $x' = x_{q_b}^b$, then $B^{q_b}(x_{q_b}^b) = \Delta(x')$. Thus, both when $x' = x_{q_b}^b$ and $x' = x_{q_f}^f$:

$$\#(\Delta_{\min}(x), \Delta(x)) > \#(\Delta_{\min}(x), \Delta(x')).$$

■

As $\Delta(x)$ lexicographically increases after each loop and the strategy space $\mathcal{S}$ is finite, RESTORE terminates.

3.5.2 *Correctness* PACKETHALVER

It is left to prove that PACKETHALVER returns an equilibrium for game $\mathcal{G}_{k_0}((d_i)_{i \in N}$.

Theorem 3.5.8. *Given a k_0-integral splittable singleton game with affine player-specific cost functions $\mathcal{G}_{k_0} := (N, E, (d_i)_{i \in N}, ((c_{i,e})_{e \in E})_{i \in N})$,* PACKETHALVER *returns an equilibrium for $\mathcal{G}_{k_0}$.*

Proof. Strategy profile x' is initialized as the all-zero strategy profile, which obviously is an equilibrium for the game $\mathcal{G}_{2^{q_1} k_0}(\vec{0})$. Assume that in iteration q we enter the for-loop in PACKETHALVER with an equilibrium x for game $\mathcal{G}_{2^{q_1 - q + 1} k_0}$ with demands $d_i' = d_i - (d_i \mod 2^{q_1 - q + 1} k_0)$. First algorithm RESTORE computes an equilibrium for packet size $2^{q_1 - q} k_0$ and demands $d_i' = d_i - (d_i \mod 2^{q_1 - q} k_0)$. In lines 5-10 of PACKETHALVER we check for each player $i \in N$ if her unscheduled load satisfies $d_i - d_i' \geqslant 2^{q_1 - q} k_0$. If so, we schedule one extra packet for player i using subroutine ADD. Thus, after the q'th iteration in the for-loop, we obtain an equilibrium for packet size $2^{q_1 - q} k_0$ and demands $d_i' = d_i - (d_i \mod 2^{q_1 - q} k_0)$. Hence, after the q_1'th iteration we obtain an equilibrium for packet size $2^0 k_0 = k_0$ and demands $d_i' = d_i - (d_i \mod k_0) = d_i$, which is an equilibrium for game $\mathcal{G}_{k_0}((d_i)_{i \in N})$. ■

3.6 RUNNING TIME

We prove that the running time of PACKETHALVER is polynomially bounded in n, m, $\log k$ and $\log \delta$, where δ is the upper bound on player-specific demands d_i. For this, we first need to analyze the running time of the two subroutines ADD and RESTORE.

3.6.1 *Running Time* ADD

In [32, Corollary 5.2] Harks et al. proved that it takes time $nm(\delta/k)^2$ to execute ADD. If their algorithm is applied to games with singleton strategy spaces and player-specific affine cost functions, the running time reduces to $O(nm^4)$. The main reason for this is that equilibria are not very sensitive under small changes in demands.

Lemma 3.6.1. *Let x_k be an equilibrium for game $\mathcal{G}_k((d_i)_{i\in N})$ and let x_q be the strategy profile after the q'th iteration of the while-loop described in lines 4-7 of* ADD. *Then $|(x_k)_{i,e} - (x_q)_{i,e}| < 2mk$ for all $i \in N$ and $e \in E$.*

Proof. On the contrary, assume that q is the first iteration in which we obtain $|(x_q)_{i,e} - (x_k)_{i,e}| = 2mk$ for some $i \in N$ and $e \in E$. There are two cases: either (I)$(x_q)_{i,e} - (x_k)_{i,e} = 2mk$ or (II)$(x_k)_{i,e} - (x_q)_{i,e} = 2mk$. We prove that the first case leads to a contradiction. For the second case a contradiction can be obtained in a similar manner.

Harks, Peis and Klimm [36] proved that only the players using a resource whose load increased in the previous iteration may have an improving move, and if so, a best response consists in moving one packet from this resource to another one. This implies that $(x_k)_e \leqslant (x_q)_e \leqslant (x_k)_e + k$ for all $e \in E$. Thus, when assuming $(x_q)_{i,e} = (x_k)_{i,e} + 2mk$, we obtain:

$$(x_q)_e + (x_q)_{i,e} \geqslant (x_k)_e + (x_k)_{i,e} + 2mk. \tag{3.15}$$

Remember that the total load distributed in x_q by player i exceeds the total load distributed in x_k by at most k, hence, $\sum_{f\in E}(x_q)_{i,e} \leqslant k + \sum_{f\in E}(x_k)_{i,e}$. We obtain:

$$\sum_{f\neq e}(x_q)_{i,e} \leqslant \sum_{f\neq e}(x_k)_{i,e} + (1-2m)k < \sum_{f\neq e}(x_k)_{i,e} - 2(m-1)k.$$

The pigeonhole principle implies that there exists a resource $f \in E$ such that $(x_q)_{i,f} < (x_k)_{i,f} - 2k$, and hence, $(x_q)_{i,f} \leqslant (x_k)_{i,f} - 3k$. Combined with inequality $(x_k)_f \leqslant (x_q)_f \leqslant (x_k)_f + k$, this implies:

$$(x_q)_{i,f} + (x_q)_f \leqslant (x_k)_{i,f} + (x_k)_f - 2k. \tag{3.16}$$

As q is the first iteration in which $(x_q)_{i,e} - (x_k)_{i,e} = 2mk$, we have that $x_q = (x_{q-1})_{i:e' \to e}$ for some $e' \in E$. Using inequalities (3.15), (3.16), $m > 1$ and the fact that x_k is an equilibrium for packet size k, we obtain:

$$\mu^{-k}_{i,e}(x_q) > \mu^{+k}_{i,e}(x_k) \geqslant \mu^{-k}_{i,f}(x_k) \geqslant \mu^{+k}_{i,f}(x_q).$$

This, combined with the fact that $(x_q)_{i,e} > (x_k)_{i,e} \geqslant 0$, implies player i can decrease her cost by moving a packet from e to f. This contradicts the fact that in strategy profile x_{q-1} moving a packet to e is a restricted best response for player i. ∎

Lemma 3.6.2. *Algorithm* ADD *has running time* $O(nm^4)$.

Proof. We use the same proof strategy as used in [36, Proof of Theorem 5.1], and give each unit of demand of each player $i \in N$ an identity denoted by i_j, for $j \in \{1, \dots d_i\}$. For a strategy profile x, we define $e(i_j, x) \in E$ to be the resource to which unit i_j is assigned in strategy profile x. Let x_q be the strategy profile after line 5 of the algorithm has been executed for the q'th time, where we use the convention that x_0 denotes the preliminary strategy profile when entering the while-loop. Note that there is a unique resource e_0 such that $(x_0)_{e_0} = x_{e_0} + k$ and $(x_0)_e = x_e$ for all $e \in E \backslash \{e_0\}$. Furthermore, because we choose in Line 5 a restricted best response, a simple inductive argument shows that after each iteration q of the while-loop, there is a unique resource $e_q \in E$ such that $(x_0)_{e_q} = x_{e_q} + k$ and $(x_0)_e = x_e$ for all $e \in E \backslash \{e_q\}$.

For any x_q during the course of the algorithm, we define the marginal cost $\Delta_{i_j}(x_q)$ of packet i_j under profile x_q as:

$$c_{i,e}((x_q)_e)(x_q)_{i,e} - c_{i,e}((x_q)_e - k)((x_q)_{i,e} - k) = \mu^{-k}_{i,e}(x_q),$$

if $e = e(i_j, x_q) = e_q$, and as:

$$c_{i,e}((x_q)_e + k)(x_q)_{i,e} - c_{i,e}((x_q)_e)((x_q)_{i,e} - k),$$

if $e = e(i_j, x_q) \neq e_q$.

In [36, Proof of Theorem 5.1], Harks, Klimm and Peis prove that the sorted vector of marginal costs $\Delta_{i_j}(x)$ lexicographically decreases during the while loop, and hence Algorithm ADD ends within finite time. It is left to prove that its running time is polynomial.

In order to do so, we assume that players move packets according to a *Last In First Out (LIFO)* principle. Thus, whenever player i moves a packet away from e_q, she moves the packet that was placed on this resource last. We keep track of costs Δ_{i_j}, at the moment that packet i_j is moved. Assume that packet i_j is moved in iterations $q_1, \dots q_{p_{i_j}}$, then:

$$
\begin{aligned}
&\Delta_{i_j}(x_{q_1}) \\
&= \mu^{-k}_{i,e_{q_1}}(x_{q_1}) && \text{(by definition of } \Delta_{i_j} \text{ and } e(i_j, x_{q_1}) = e_{q_1}) \\
&> \mu^{+k}_{i,e_{q_1+1}}(x_{q_1}) && \text{(as moving } i_j \text{ is an improving move for } i\text{)} \\
&= \mu^{-k}_{i,e_{q_1+1}}(x_{q_1+1}) && \text{(by construction of } x_{q_1+1}\text{)} \\
&= \mu^{-k}_{i,e_{q_2}}(x_{q_2}) && \text{(as } e_{q_2} = e_{q_1+1} \text{ and } (x_{q_2})_{i,e_{q_2}} = (x_{q_1+1})_{i,e_{q_2}}) \\
&= \Delta_{i_j}(x_{q_2}). && \text{(by definition of } \Delta_{i_j} \text{ and } e(i_j, x_{q_2}) = e_{q_2})
\end{aligned}
$$

Using similar argumentation, we obtain:

$$\Delta_{i_j}(x_{q_1}) > \Delta_{i_j}(x_{q_2}) > \cdots > \Delta_{i_j}(x_{q_{p_{i_j}}}).$$

Note that in iterations $q_1, \dots q_{p_{i_j}}$, $\Delta_{i_j}(x_{q_\ell})$ does not depend on the aggregated load $(x_{q_\ell})_{e_{q_\ell}}$, as $(x_{q_\ell})_{e_{q_\ell}} = (x_0)_{e_{q_\ell}} + k$ for each $\ell \in \{1, \dots, p_{i_j}\}$. Instead it only depends on the player-specific load $(x_{q_\ell})_{i,e_{q_\ell}}$. Lemma 3.6.1 implies that each player $i \in N$ will move at most $2m$ packets from each resource and hence there will occur at most $4m$ different values of $(x_{q_\ell})_{i,e_{q_\ell}}$. Thus, $\Delta_{i_j}(x_{q_\ell})$ can take $4m$ different values and each packet visits each resource at most $4m$ times.

Hence, each player i moves at most $2m^2$ packets, and each packet visits each resource (m resources) at most $4m$ times. Therefore the running time of ADD is bounded by $O(nm^4)$. ■

3.6.2 *Running Time* RESTORE

We analyze the running time of RESTORE. The crucial idea is that for each strategy profile y (for a game with packet size k) obtained during the execution of RESTORE, we have both $|(y_e - (x_{2k})_e| \leqslant 2mk$ (Lemma 3.6.3 and Lemma 3.6.4) and $|y_{i,e} - (x_{2k})_{i,e}| < 2m^2k$ (Lemma 3.6.5) for all $i \in N$ and $e \in E$.

The first lemma follows trivially from the fact that $H(x, x^b_{q_1}) \in \{0, 2k\}$ and $H(x, x^f_{q_2}) \in \{0, 2k\}$.

Lemma 3.6.3. *Let x_{2k} be an equilibrium for game $\mathcal{G}_{2k}$ and let x, $x^b_{q_1}$ and $x^f_{q_2}$ be as described in* RESTORE *for $q_1 \in \{1, \ldots, q_b\}$ and $q_2 \in \{1, \ldots, q_f\}$. If for all $e \in E$ $|x_e - (x_{2k})_e| < 2mk$, then for all $e \in E$ $|(x^b_{q_1})_e - (x_{2k})_e| \leqslant 2mk$ and $|(x^f_{q_2})_e - (x_{2k})_e| \leqslant 2mk$.*

Lemma 3.6.3 is only useful when $|x'_e - (x_{2k})_e| < 2mk$ for all x' obtained through the while loop. Thus, for x' we prove a slightly stronger result.

Lemma 3.6.4. *Let x_{2k} be an equilibrium for game $\mathcal{G}_{2k}$ and let x and x' be as described in* RESTORE*. If $|x_e - (x_{2k})_e| < 2mk$ for all $e \in E$, then for all $e \in E$ $|x'_e - (x_{2k})_e| < 2mk$.*

Proof. If $x'_e = (x_{2k})_e$ for all $e \in E$ the lemma follows trivially. Thus, assume that there exists an e such $x'_e \neq (x_{2k})_e$. By construction of x', we have $x'_{e^-_{q_b}} = x_{e^-_{q_b}} - k$, $x'_{e^+_{q_f}} = x_{e^+_{q_f}} + k$ and $x'_e = (x_{2k})_e$ for all $e \in E \backslash \{e^-_{q_b}, e^+_{q_f}\}$. Thus, we only need to check that:

1. $x'_{e^-_{q_b}} > (x_{2k})_{e^-_{q_b}} - 2mk$.
2. $x'_{e^+_{q_f}} < (x_{2k})_{e^+_{q_f}} + 2mk$.

For the first case we note that $x'_{e^-_{q_b}} = (x^b_{q_b})_{e^-_{q_b}}$. Then, Algorithm 3 implies $(x^b_{q_b})_{e^-_{q_b}} > (x_{2k})_{e^-_{q_b}} - 2mk$. Hence $x'_{e^-_{q_b}} > (x_{2k})_{e^-_{q_b}} - 2mk$. For the second case $x'_{e^-_{q_b}} = (x^f_{q_f})_{e^+_{q_f}}$. Algorithm 6 implies $(x^f_{q_f})_{e^+_{q_f}} < (x_{2k})_{e^+_{q_f}} + 2mk$. Hence, $x'_{e^+_{q_f}} < (x_{2k})_{e^+_{q_f}} + 2mk$. Thus, if $|x'_e - (x_{2k})_e| < 2mk$ for all $e \in E$, then $|x'_e - (x_{2k})_e| < 2mk$. ■

Moreover, the difference in player-specific load on a resource between any strategy obtained during RESTORE and equilibrium x_{2k} is bounded by $2m^2k$.

Lemma 3.6.5. *Let x_{2k} be an equilibrium for game $\mathcal{G}_{2k}$ and let x, x^b_q and x^f_q be as described in* RESTORE*. If $|x_e - (x_{2k})_e| < 2mk$ and $|x_{i,e} - (x_{2k})_{i,e}| < 2m^2k$ for all $i \in N$ and $e \in E$, then $|(x^b_q)_{i,e} - (x_{2k})_{i,e}| < 2m^2k$ for all $i \in N$, $e \in E$ and $q \in \{1, \ldots, q_b\}$, and $|(x^f_q)_{i,e} - (x_{2k})_{i,e}| < 2m^2k$ for all $i \in N$, $e \in E$ and $q \in \{1, \ldots, q_f\}$.*

Proof. As the proofs for both statements are almost identical, we only prove the first statement here, which we do by using induction on q. We define $e_0^- := e_1^+$ and $x_0^b := x$, then this statement is trivially true for $q = 0$.

Thus, assume that $|(x_q^b)_{i,e} - (x_{2k})_{i,e}| < 2m^2k$ for all $e \in E$, then we prove that $|(x_{q+1}^b)_{i,e} - (x_{2k})_{i,e}| < 2m^2k$ for each $i \in N$ and $e \in E$. On the contrary, assume that there exists an $i \in N$ and an $e \in E$ such that $|(x_{q+1}^b)_{i,e} - (x_{2k})_{i,e}| = 2m^2k$. There are two cases that we need to check: (a) $(x_{2k})_{i,e_{q+1}^-} - (x_{q+1}^b)_{i,e_{q+1}^-} = 2m^2k$ and (b) $(x_{q+1}^b)_{i,e_q^-} - (x_{2k})_{i,e_q^-} = 2m^2k$.

We first show that case (a) leads to a contradiction. Property **B2** and line 7 in Algorithm 3 imply that $e_{q+1}^- \in \arg\max\{\mu_{i,e}^{-k}(x_q^b)\}$ for all $q \in \{0, \dots q_b - 1\}$. We use Lemma 3.6.3 to obtain:

$$(x_{q+1}^b)_{e_{q+1}^-} + (x_{q+1}^b)_{i,e_{q+1}^-} \leqslant (x_{2k})_{e_{q+1}^-} + (x_{2k})_{i,e_{q+1}^-} - 2m(m-1)k. \quad (3.17)$$

As the total load distributed by player i is the same in x_{q+1}^b and x_{2k}, and so is the total load in the system, we obtain:

$$\sum_{e \neq e_{q+1}^-} ((x_{q+1}^b)_e + (x_{q+1}^b)_{i,e}) \geqslant \sum_{e \neq e_{q+1}^-} ((x_{2k})_e + (x_{2k})_{i,e}) + 2m(m-1)k.$$

By the pigeonhole principle, there must exist at least one resource $f \in E$ such that:

$$(x_{q+1}^b)_f + (x_{q+1}^b)_{i,f} \geqslant (x_{2k})_f + (x_{2k})_{i,f} + 2mk. \quad (3.18)$$

Note that $(x_{q+1}^b)_{i,f} > 0$, as $(x_{q+1}^b)_{i,f} = 0$ implies $(x_{q+1}^b)_f > (x_{2k})_f + 2mk$, which contradicts the fact that $|(x_{2k})_f - (x_{q+1}^b)_f| \leqslant 2mk$. We obtain:

$$\begin{aligned}
&\mu_{i,e_{q+1}^-}^{-k}(x_q^b) \\
&= \mu_{i,e_{q+1}^-}^{+k}(x_{q+1}^b) && \text{(as } (x_{q+1}^b) = (x_q^b)_{e_{q+1}^- \to e_q^-}\text{)} \\
&< \frac{1}{2}\mu_{i,e_{q+1}^-}^{-2k}(x_{2k}) && \text{(as } m \geqslant 2, \text{ this follows from equation (3.17))} \\
&\leqslant \frac{1}{2}\mu_{i,f}^{+2k}(x_{2k}) && \text{(as } x_{2k} \text{ is an equilibrium for } \mathcal{G}_{2k}\text{)} \\
&< \mu_{i,f}^{-k}(x_{q+1}^b) && \text{(as } m \geqslant 2, \text{ this follows from equation (3.18))} \\
&\leqslant \mu_{i,f}^{-k}(x_q^b). && \text{(as } f \neq e_{q+1}^-\text{)}
\end{aligned}$$

This contradicts the fact that $e_{q+1}^- = \arg\max\{\mu_{i,e}^{-k}(x_q^b)\}$.

For the second case we prove that $(x^b_{q+1})_{i,e} - (x_{2k})_{i,e} = 2m^2k$ leads to a contradiction. Note that property **B1** implies that x^b_{q+1} is not obtained from x^b_q through lines 3-5 (of Restore), but in lines 6-10 instead. Hence, $e^-_q = \arg\min\{\mu^{+k}_{i,e}(x^b_q)\}$. Using similar argumentation as in the first case, we are able to show that there exists an $f \in E \backslash \{e^-_q\}$ with $\mu^{+k}_{i,f}(x^b_q) < \mu^{+k}_{i,e^-_q}(x^b_q)$, contradicting the fact that $e^-_q = \arg\min\{\mu^{+k}_{i,e}(x^b_q)\}$. ∎

The bounds on the total and player-specific load enable us to prove that Algorithm 3 and Algorithm 6 run in polynomial time.

Lemma 3.6.6. *Let x_{2k} be an equilibrium for packet size $2k$. And let x be a strategy profile for packet size k such that $|(x_{2k})_e - x_e| < mk$ for all $e \in E$. Then Algorithm 3 and Algorithm 6 both have running time $O(nm^6)$.*

Proof. Similar as algorithm Add, we again work with a sequence of restricted best responses. As mentioned before, in [36, Proof of Theorem 5.1], Harks, Klimm and Peis prove that the sorted vector of marginal costs as defined in Lemma 3.6.2 lexicographically decreases during the while loop of improving moves.

Lemma 3.6.3, 3.6.4 and 3.6.5 imply that for each player $i \in N$ and each resource $e \in E$, at most $4m^2$ different marginal cost values Δ_{i_j} can occur whenever a packet i_j is moved within a path of restricted best responses. Using the same argumentation as in Lemma 3.6.2, this implies that each unit of demand for player i ($m \cdot 2m^2$ units) visits each resource (m resources) at most $4m^2$ times. Therefore the running time of both Algorithm 3 and Algorithm 3.5.6 is bounded by $O(nm^6)$. ∎

We combine all previous results to prove Lemma 3.6.7.

Lemma 3.6.7. Restore *has running time $O(n^2m^{14})$.*

Proof. The running time of Restore is dominated by the number of times we enter the while-loop, and the running time of Algorithm 3 and Algorithm 6. Using Lemma 3.6.6 we know that in each iteration, the running time of Algorithm 3 and Algorithm 6 is $O(nm^6)$. Hence, the running time of a complete iteration is $O(nm^6)$.

Note that Lemma 3.6.3 and Lemma 3.6.5 imply that on each resource $e \in E$ at most $O(m^3)$ different values $\mu^{-k}_{i,e}(\cdot)$ can occur and $O(m^4)$ different values $\mu^{+k}_{i,\min}(\cdot)$. Thus, for each player at most $O(m^7)$ different values $\mu^{+k}_{i,\min}(\cdot) - \mu^{-k}_{i,e}(\cdot)$ can appear on a resource, thus $O(nm^8)$ different values in total. In Lemma 3.5.7 we prove that $\#(\Delta_{\min}(x), \Delta(x))$ decreases after each iteration in the while-loop, Hence we enter the while-loop at most $O(nm^8)$ times.

As we enter the while-loop at most $O(nm^8)$ times, and each iteration runs in $O(nm^6)$, PacketHalver runs in $O(n^2m^{14})$. ■

3.6.3 *Running Time* PacketHalver

Finally, we prove the following theorem.

Theorem 3.6.8. PacketHalver *runs in time* $O(n^2m^{14}\log(\delta/k_0))$.

Proof. Note that we picked $q_1 \in \mathbb{N}$ to be the smallest number such that $2^{q_1}k_0 > d_i$ for all player-specific demands d_i. This implies that q_1 is bounded in $O(\log(\delta/k_0))$, where δ is an upper bound on the player-specific demands. Thus, we execute lines 3-10 $O(\log(\delta/k_0))$ times. In line 4 we call Restore, which runs in $O(n^2m^{14})$. In line $5-9$ we execute Add (which runs in $O(nm^6)$) at most n times. Thus, the computation time of lines $5-10$ is $O(n^2m^6)$. This implies that it takes time $O(n^2m^{14})$ to go through a complete iteration in the for loop. Thus, PacketHalver runs in time $O(n^2m^{14}\log(\delta/k_0))$. ■

It is left to show that in an atomic splittable game $\mathcal{G}$, $\log(1/k_0)$ is polynomially bounded in the input.

$$
\begin{aligned}
&O\left(\log\left(\frac{1}{k_0}\right)\right) \\
&= O\left(\log\left(\frac{2m^2(2a_{\max})^{nm}(nm)^{nm/2}}{a_{\mathrm{gcd}}^{nm}}\right)\right) \\
&= O\left(\log m + \log\left(\det(A)\right) + \log\left(\frac{1}{a_{\mathrm{gcd}}^{nm}}\right)\right) \\
&= O\left(\log m + \log((2a_{\max}\sqrt{nm})^{nm}) + \log\left(\prod_{i\in N, e\in E_i} \overline{d_i}\cdot\overline{a_{i,e}}\cdot\overline{b_{i,e}}\right)\right) \\
&= O\left(nm\log(nma_{\max}) + \sum_{i\in N, e\in E_i}\left(\log(\overline{d_i}) + \log(\overline{a_{i,e}}) + \log(\overline{b_{i,e}})\right)\right).
\end{aligned}
$$

Which is indeed polynomial in the size of the input. Remember that if we are computing an atomic splittable equilibrium, we first compute the k_0

splittable equilibrium using the algorithm above. Second, we compute the exact equilibrium in time $O((nm)^3)$.

Corollary 3.6.9. *Given game $\mathcal{G}$, we can compute an atomic splittable equilibrium for $\mathcal{G}$ in running time:*

$$O\left((nm)^3 + n^2m^{14}\log\left(\frac{\delta}{k_0}\right)\right).$$

3.7 MULTIMARKET COURNOT OLIGOPOLY

In this section, we derive a strong connection between atomic splittable singleton congestion games and multimarket Cournot oligopolies (defined in Section 1.2.3). Recall that such a game is compactly represented by the tuple

$$\mathcal{M} = (N, E, (E_i)_{i\in N}, (p_{i,e})_{i\in N, e\in E_i}, (C_i)_{i\in N}),$$

where N is a set of n firms and E a set of m markets. Each firm i only has access to a subset $E_i \subseteq E$ of the markets. Each market e is endowed with firm-specific, decreasing, affine price functions, $p_{i,e}(t) = s_{i,e} - r_{i,e}t$, where $i \in N$. In a strategy profile, a firm chooses a non-negative production quantity $x_{i,e} \in \mathbb{R}_{\geqslant 0}$ for each market $e \in E_i$. We denote a strategy profile for a firm by $x_i = (x_{i,e})_{e\in E_i}$, and a joint strategy profile by $x = (x_i)_{i\in N}$. The production costs of a firm are of the form $C_i(t) = c_i t^2$ for some $c_i \geqslant 0$. The goal of each firm $i \in N$ is to maximize its utility, which is given by:

$$u_i(x) = \sum_{e\in E_i} p_{i,e}(x_e)x_{i,e} - C_i\Big(\sum_{e\in E_i} x_{i,e}\Big),$$

where $x_e := \sum_{i\in N} x_{i,e}$. In the rest of this section we prove that several results that hold for atomic splittable equilibria and k-splittable equilibria carry over to multimarket oligopolies.

A strategic game $\mathcal{G} = (N, (X)_{i\in N}, (u_i)_i \in N)$ is defined by a set of players N, a set of feasible strategies X_i for each player $i \in N$ and a pay-off function $u_i(x)$ for each $i \in N$, where $x \in \bigtimes_{i\in N} X_i$.

Definition 3.7.1. *Let $\mathcal{G} = (N, (X_i)_{i\in N}, (u_i)_{i\in N})$ and $\mathcal{H} = (N, (Y_i)_{i\in N}, (v_i)_{i\in N})$ be two strategic games with identical player set N. Then, $\mathcal{G}$ and $\mathcal{H}$ are called* isomorphic, *if for all $i \in N$ there exists a bijective function $\phi_i : X_i \to Y_i$ and $A_i \in \mathbb{R}$ such that:*

$$u_i(x_1, \dots x_n) = v_i(\phi_1(x_1), \dots, \phi_n(x_n)) + A_i.$$

Let $\mathcal{G} = (N, (X_i)_{i \in N}, (u_i)_i \in N)$ and $\mathcal{H} = (N, (Y_i)_{i \in N}, (v_i)_i \in N)$ be isomorphic games. Then, $(x_i)_{i \in N}$ is an equilibrium of game $\mathcal{G}$ if and only if $(\phi_i(x_i))_{i \in N}$ is an equilibrium of game $\mathcal{H}$. This implies that $(x_i)_{i \in N}$ is the unique equilibrium of game $\mathcal{G}$ if and only if $(\phi_i(x_i))_{i \in N}$ is the unique equilibrium of game $\mathcal{H}$.

We prove that for each multimarket oligopoly, there exists an isomorphic atomic splittable game. Moreover, we can construct the isomorphism in polynomial time.

Theorem 3.7.2. *Given a multimarket oligopoly $\mathcal{M}$, there exists an atomic splittable game $\mathcal{G}$ that is isomorphic to $\mathcal{M}$.*

Proof. Given multimarket oligopoly $\mathcal{M}$, we construct an atomic splittable singleton game $\mathcal{G}$ as follows. For every firm $i \in N$ we create a player i and we define the demand d_i for this player as an upper bound on the maximal quantity that firm i will produce, that is,

$$d_i := \sum_{e \in E_i} \max\{t \mid p_{i,e}(t) = 0\}.$$

Note that if we limit the strategy space for each player $i \in N$ in game $\mathcal{M}$ to strategies x satisfying $\sum_{e \in E_i} x_{i,e} \leqslant d_i$, all equilibria are preserved. Then, for every player i we introduce a special resource e_i, and we define the set of allowable resources for this player as:

$$\tilde{E}_i = E_i \cup \{e_i\} \quad \text{with } e_i \neq e_j \text{ for } i \neq j.$$

The cost on these special resources e_i are defined as:

$$c_{i,e_i}(t) := c_i(t - 2d_i) \text{ for all } i \in N.$$

The cost on resources $e \in E_i$ are defined as:

$$c_{i,e}(t) := -p_{i,e}(t) = r_{i,e}t - s_{i,e} \text{ for all } i \in N.$$

In order to guarantee affine cost functions with only non-negative constant terms, one can add a large positive constant $c_{\max}$ to every cost function. We define:

$$c_{\max} = \max\left\{\{s_{i,e} \mid \text{ for all } i \in N, e \in E_i\} \cup \{2c_i d_i \mid \text{ for all } i \in N\}\right\}.$$

Note that adding $c_{\max}$ to every cost function does not change the equilibrium, it only adds $d_i c_{\max}$ to the total cost of each player. The total cost of a strategy x for player i in game $\mathcal{G}$ is:

$$\pi_i(x') = \sum_{e \in \tilde{E}_i} c_{i,e}(x'_e)x'_{i,e},$$

which is equal to:

$$\pi_i(x') = \sum_{e \in E_i} -p_{i,e}(x'_e)x'_{i,e} + x'_{i,e_i}c_i(x'_{i,e_i} - 2d_i). \tag{3.19}$$

Note that the payoff function of player i in x' is $\nu_i(x') = -\pi_i(x')$. It is left to prove that game $\mathcal{G}$ is isomorphic to game $\mathcal{M}$. Let x be a feasible strategy in game $\mathcal{M}$. For each $i \in N$, we define the bijective function $\phi_i : E_i \to \tilde{E}$ as:

$$\phi_i(x_{i,1}, \dots, x_{i,m}) = (x_{i,1}, \dots, x_{i,m}, d_i - \sum_{e \in E_i} x_{i,e}) =: (x'_{i,1}, \dots, x'_{i,m}, x'_{i,m+1}).$$

As we limited the strategy space for each player $i \in N$ in game $\mathcal{M}$ to strategies x where $\sum_{e \in E_i} x_{i,e} \leqslant d_i$, $x' := \phi(x)$ is a feasible strategy in $\mathcal{G}$. For each feasible strategy x for game $\mathcal{M}$, and for each $i \in N$, we have:

$$\begin{aligned}
u_i(x) &= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - C_i\Big(\sum_{e \in E_i} x_{i,e}\Big) \\
&= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - c_i\Big(\sum_{e \in E_i} x_{i,e}\Big)^2 \\
&= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - c_i\Big(d_i - \sum_{e \in E_i} x_{i,e}\Big)\Big(-d_i - \sum_{e \in E_i} x_{i,e}\Big) - c_i d_i^2 \\
&= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - c_i\Big(d_i - \sum_{e \in E_i} x_{i,e}\Big)\Big(d_i - \sum_{e \in E_i} x_{i,e} - 2d_i\Big) - c_i d_i^2 \\
&= -\pi_i(\phi_1(x_1), \dots, \phi_1(x_n)) - c_i d_i^2 \\
&= \nu_i(\phi_1(x_1), \dots, \phi_1(x_n)) - c_i d_i^2.
\end{aligned}$$

Thus, game $\mathcal{M}$ and $\mathcal{G}$ are isomorphic. ■

Remark 3.7.3. *Given a multimarket oligopoly $\mathcal{M}$, one can construct an atomic splittable singleton game isomorphic to $\mathcal{M}$ within running time $O(nm)$.*

One of the main results of this paper is our polynomial time algorithm that finds the unique equilibrium for atomic splittable singleton congestion games with linear cost functions within polynomial time. As for each multimarket oligopoly there exists an atomic splittable game isomorphic to it, we are also able to construct this unique equilibrium within polynomial time.

Theorem 3.7.4. *Given a multimarket oligopoly $\mathcal{M}$, an equilibrium can be computed within running time:*

$$O\left(n^{16}m^{14}\log\left(\frac{\delta}{k_0}\right)\right).$$

Proof. This theorem follows directly from the fact that we can construct an atomic splittable singleton game $\mathcal{G}$ isomorphic to $\mathcal{M}$ (Theorem 4.4.1) and the fact that $x = (x_i)_{i \in N}$ is an equilibrium in $\mathcal{G}$ if and only if $x = (\phi_i(x_i))_{i \in N}$ is an equilibrium in $\mathcal{M}$. Note that if in $\mathcal{M}$, firms compete over m markets, the isomorphic atomic splittable singleton game $\mathcal{G}$ has $m + n$ resources. For such a game, Corollary 3.6.9 implies that an equilibrium can be found in

$$O\left(n^3(m+n)^3 + n^2(m+n)^{14} \log\left(\frac{\delta}{k_0}\right)\right) = O\left(n^{16} m^{14} \log\left(\frac{\delta}{k_0}\right)\right).$$

■

In an *integral multimarket oligopoly* players sell indivisible goods. Thus, players can only produce and sell integer quantities, i.e., $x_{i,e} \in \mathbb{N}_{\geqslant 0}$ for each $i \in N$ and $e \in E_i$. For these games, we can construct an isomorphic 1-splittable congestion game.

Theorem 3.7.5. *Given an integral multimarket oligopoly $\mathcal{M}$, we can construct an 1-splittable congestion game $\mathcal{G}$ isomorphic to $\mathcal{M}$ within running time* $O(nm)$.

Proof. We define

$$d_i := \sum_{e \in E_i} \lfloor \max\{t \mid p_{i,e}(t) = 0\} \rfloor.$$

Then, the theorem follows using the same construction as in Theorem 4.4.1. ■

Corollary 3.7.6. *Given an integral multimarket oligopoly $\mathcal{M}$, an integral equilibrium can be computed within*

$$O\left(n^{16} m^{14} \log\left(\frac{\delta}{k_0}\right)\right).$$

Proof. Theorem 3.7.5 implies that we can construct an atomic splittable singleton game $\mathcal{G}$ isomorphic to $\mathcal{M}$. Note that if in $\mathcal{M}$, n firms compete over m markets, the isomorphic atomic splittable singleton game has $m + n$ resources. For such a game, Theorem 3.6.8 implies the desired running time. ■

Lastly, we extend a very recent result by Todd [69], where the total and individual production in one market in an integer equilibrium and a real equilibrium are compared.

Theorem 3.7.7. *Given an multimarket oligopoly $\mathcal{M}$, with real equilibrium $(x_i)_{i\in N}$. Then, for any integer equilibrium $(y_i)_{i\in N}$:*

- $|x_e - y_e| \leqslant m + n$.
- $|x_{i,e} - y_{i,e}| \leqslant (m+n)^2$.

Proof. Assume that in game $\mathcal{M}$, n firms compete over m markets. According to Theorem 4.4.1, there exists an atomic splittable congestion game $\mathcal{G}$ on $m + k$ resources that is isomorphic to $\mathcal{M}$ using bijection ϕ.

Let $x = (x_i)_{i\in N}$ be an atomic splittable equilibrium of $\mathcal{M}$ and $y = (y_i)_{i\in N}$ a 1-splittable equilibrium of $\mathcal{M}$. Then $x' := (\phi_i(x_i))_{i\in N}$ is an atomic splittable equilibrium of $\mathcal{G}$ and $y' := (\phi_i(y_i))_{i\in N}$ is a 1-splittable equilibrium of $\mathcal{G}$. According to Lemma 3.3.2 and 3.3.3 we know that for any real equilibrium x' and 1-splittable equilibrium y' it holds that $|x'_e - y'_e| < (m+n)$ and $|x'_{i,e} - y'_{i,e}| < (m+n)^2|$ for all $i \in N$ and $e \in E_i$. Then, using the bijection ϕ described in (4.9), we obtain that $|x_e - y_e| < (m+n)$ and $|x_{i,e} - y_{i,e}| < (m+n)^2$. ∎

Todd [69] showed that the total production in a 1-splittable equilibrium is in the worst-case at most $n/2$ away from that in the real equilibrium, and the individual firm's choice can be more that $(n-1)/4$ away from her choice in the real equilibrium. Our bounds a larger than Todd's, yet, they hold for a more general model – multiple markets and firm-specific price functions. We pose as an open question, whether or not our bounds are tight or can be further improved.

3.8 CONCLUDING REMARKS

In this chapter we studied atomic splittable singleton congestion games with player-specific affine cost functions and developed the first polynomial time algorithm computing the unique pure Nash equilibrium. In order to do so, we first developed an algorithm that computes a Nash equilibrium for an associated k-integral splittable congestion game within polynomial time. Then, we showed that we can bound the differences between the atomic splittable equilibrium and an k-integral splittable equilibrium in terms of the number of resources and packet size k. We found a k such that the support set of the k-splittable equilibrium would allow us to compute the support set for the atomic splittable equilibrium. Once this support set is known, the atomic splittable equilibrium could be found easily by solving a set of linear equations.

Then, we then developed a polynomial time computable transformation mapping a multimarket Cournot competition game with firm-specific affine price functions and quadratic costs to an associated atomic splittable congestion game with affine cost functions. As this transformation preserves equilibria, our previous algorithm also computes Cournot equilibria in this setting within polynomial time. Lastly, our analysis for integrally splittable games lead to new bounds on the differences between real and integer Cournot equilibria and, hence, generalize the bounds for single markets obtained by Todd [69].

APPENDIX 3.A ATOMIC SPLITTABLE CONGESTION GAMES AS LCP

Assume we are given an atomic splittable congestion game $\mathcal{G}$ as defined in Section 3.2.1. Given M, q as below, any solution for the linear complementarity problem $LCP(M, q)$ leads to an equilibrium in atomic splittable congestion game $\mathcal{G}$. Here, M is an $n(m+1) \times n(m+1)$ matrix given by:

$$M := \begin{bmatrix} 2A_1 & A_1 & \dots & A_1 & -V_1 \\ A_2 & 2A_2 & \dots & A_2 & -V_2 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ A_N & A_N & \dots & 2A_N & -V_N \\ H_1 & H_2 & \dots & H_M & 0 \end{bmatrix},$$

where for all $i \in N$ we define A_i, H_i and V_i to be $m \times n$ matrices. Here A_i is defined as:

$$A_i := \begin{bmatrix} a_{i,1} & 0 & \dots & 0 \\ 0 & a_{i,2} & \dots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \dots & a_{i,n} \end{bmatrix},$$

Matrix V_i contains ones in column i and zero's in all other colums:

$$(V_i)_{p,q} := \begin{cases} 1 & \text{when } p = i \\ 0 & \text{otherwise.} \end{cases}$$

Matrix H_i contains ones in row i and zero's in all other colums:

$$(H_i)_{p,q} := \begin{cases} 1 & \text{when } q = i \\ 0 & \text{otherwise.} \end{cases}$$

We define q as:

$$q^T := \left[\begin{array}{ccc:c:ccc:ccc} b_{1,1} & \dots & b_{1,m} & \dots & b_{n,1} & \dots & b_{n,m} & -d_1 & \dots & -d_n \end{array} \right].$$

We denote any solution z that is a solution to $LCP(M, q)$ as:

$$z^T := \left[\begin{array}{ccc:c:ccc:ccc} x_{1,1} & \dots & x_{1,m} & \dots & x_{n,1} & \dots & x_{n,m} & \lambda_1 & \dots & \lambda_n \end{array} \right].$$

The first mn equations in $Mz + q$ model exactly the optimality conditions, and the last n inequalities model the demand constraints. Then,

$(x_{i,e})_{i\in N, e\in E}$ forms an atomic splittable equilibrium for game $\mathcal{G}$, and $(\lambda_i)_{i\in N}$ are the KKT multipliers. In any solution to this LCP, λ_i will be the marginal cost for player i on all resources for which $x_{i,e} > 0$.

Note that when the cost functions are player independent, matrix M is positive definite, any algorithm that can solve quadratic programs, e.g. Lemke's algorithm, can be used to find a solution to the LCP within polynomial time

APPENDIX 3.B FORWARD PATH OF RESTRICTED BEST RESPONSES

Algorithm 6: FP$(x_{2k}, x_1^f, e_1^+, \mathcal{G}_k)$: A forward path of restricted best responses.

Input: equilibrium x_{2k} for game $\mathcal{G}_{2k}$, strategy profile x_1^f for game $\mathcal{G}_k$, a resource e_1^+ and game $\mathcal{G}_k$.
Output: Strategy profile $x_{q_f}^f$ for game $\mathcal{G}_k$ and resource e_q^+.

Initialize $q \leftarrow 1$;
repeat
  if $(x_q^f)_{e_q^+} \geqslant (x_{2k})_{e_q^+} + 2mk$ **then**
    Find player i and resource e_{q+1}^+ satisfying properties:
    F1. $(x_q^b)_{i,e_q^+} > (x_{2k})_{i,e_q^+}$;
    F2. $e_{q+1}^+ \in \arg\min\{\mu_{i,e}^{+k}(x_q^f)\}$;
    F3. $\mu_{i,e_q^+}^{-k}(x_q^f) > \mu_{i,e_{q+1}^+}^{+k}(x_q^f)$;
    $x_{q+1}^f \leftarrow (x_q^f)_{i:e_q^+ \to e_{q+1}^+}$;
    $q \leftarrow q+1$;
  end
  while $\exists i \in N$ *with* $e_q^+ = \arg\max_{e\in E}\{\mu_{i,e}^{-k}(x_q^f)\}$ *and* $\mu_{i,\min}^{+k}(x_q^f) < \mu_{i,\max}^{-k}(x_q^f)$ **do**
    Choose $e_{q+1}^+ \in \arg\min\{\mu_{i,e}^{+k}(x_q^f)\}$;
    $x_{q+1}^f \leftarrow (x_q^f)_{e_q^+ \to e_{q+1}^+}$;
    $q \leftarrow q+1$;
  end
until $(x_q^f)_{e_q^+} < (x_{2k})_{e_q^+} + 2mk$;
return (x_q^f, e_q^+)

APPENDIX 3.C FULL PROOF OF LEMMA 3.5.1

Full proof of Lemma 3.5.1.

Full proof of Lemma 3.5.1. We introduce two player sets N_e^+, N_e^- for every resource $e \in E$, where:

$$N_e^+ = \{i \in N | (x_{2k})_{i,e} > (x_q^b)_{i,e}\} \text{ and } N_e^- = \{i \in N | (x_{2k})_{i,e} < (x_q^b)_{i,e}\}.$$

Note that for every $i \in N_e^+$ we have:

$$(x_q^b)_{e_q^-} + (x_q^b)_{i,e_q^-} < (x_{2k})_{e_q^-} + (x_{2k})_{i,e_q^-} - 2mk. \tag{3.20}$$

Using the player sets, we obtain:

$$\begin{aligned}\sum_{i \in N^+_{e_q^-}} ((x_{2k})_{i,e_q^-} - (x_q^b)_{i,e_q^-}) &+ \sum_{i \in N^-_{e_q^-}} ((x_{2k})_{i,e_q^-} - (x_q^b)_{i,e_q^-}) \\ &= (x_{2k})_{e_q^-} - (x_q^b)_{e_q^-} \\ &\geqslant 2mk.\end{aligned}$$

As $\sum_{i \in N^-_{e_q^-}} ((x_{2k})_{i,e_q^-} - (x_q^b)_{i,e_q^-}) \leqslant 0$, we have:

$$\sum_{i \in N^+_{e_q^-}} ((x_{2k})_{i,e_q^-} - (x_q^b)_{i,e_q^-}) \geqslant 2mk.$$

The total load distributed by a player does not change, therefore:

$$\sum_{f \neq e_q^-} \sum_{i \in N^+_{e_q^-}} ((x_{2k})_{i,f} - (x_q^b)_{i,f}) \leqslant -2mk.$$

For every resource $f \in E \backslash \{e_q^-\}$ we split $N^+_{e_q^-}$ in two parts: $N^+_{e_q^-} \cap N_f^-$ and $N^+_{e_q^-} \cap N_f^+$:

$$\sum_{f \neq e_q^-} \sum_{i \in N^+_{e_q^-} \cap N_f^-} ((x_{2k})_{i,f} - (x_q^b)_{i,f}) + \sum_{i \in N^+_{e_q^-} \cap N_f^+} ((x_{2k})_{i,f} - (x_q^b)_{i,f}) \leqslant -2mk.$$

Using the definition of N_f^+, we obtain:

$$\sum_{f \neq e_q^-} \sum_{i \in N^+_{e_q^-} \cap N_f^-} ((x_{2k})_{i,f} - (x_q^b)_{i,f}) \leqslant -2mk. \tag{3.21}$$

As $(x_{2k})_{e_q^-} - (x_q^b)_{e_q^-} \geqslant 2mk$, we have $\sum_{f \neq e_q^-}((x_{2k})_f - (x_q^b)_f) \leqslant -2mk$. Therefore:

$$\sum_{f \neq e_q^-} \sum_{i \in N^+_{e_q^-} \cap N_f^-} ((x_{2k})_f - (x_q^b)_f) \leqslant -|N^+_{e_q^-} \cap N_f^-| 2mk.$$

We add this to equation (3.21) to obtain the following:

$$\begin{aligned}\sum_{f \neq e_q^-} \sum_{i \in N^+_{e_q^-} \cap N_f^-} & ((x_{2k})_f - (x_q^b)_f) + ((x_{2k})_{i,f} - (x_q^b)_{i,f}) \\ & \leqslant -(|N^+_{e_q^-} \cap N_f^-| + 1) 2mk.\end{aligned}$$

By using the pigeonhole principle on the number of resources $f \in E \backslash \{e_q^-\}$, there exists an $f \in E \backslash \{e_q^-\}$ such that:

$$\sum_{i \in N^+_{e_q^-} \cap N_f^-} ((x_{2k})_f - (x_q^b)_f) + ((x_{2k})_{i,f-(x_q^b)_{i,f}}) < -(|N^+_{e_q^-} \cap N_f^-| + 1) 2k.$$

Using the pigeonhole principle again on the number of players in $N^+_{e_q^-} \cap N_f^-$, there exists an $i \in N^+_{e_q^-} \cap N_f^-$ such that

$$((x_{2k})_f - (x_q^b)_f) + ((x_{2k})_{i,f} - (x_q^b)_{i,f}) \leqslant -3k. \tag{3.22}$$

We combine Equation (3.20), Equation (3.22) and the fact that x is an equilibrium for packet size k to obtain:

$$\mu^{+k}_{i,e_q^-}(x_q^b) < \frac{1}{2}\mu^{-2k}_{i,e_q^-}(x_{2k}) \leqslant \frac{1}{2}\mu^{+2k}_{i,f}(x_{2k}) \leqslant \mu^{-k}_{i,f}(x_q^b). \tag{3.23}$$

Define $e^-_{q+1} = \arg\max_{e \in E}\{\mu^{-k}_{i,e}(x_q^b)\}$, then we have found a player i and a resource e^-_{q+1} for which **B1** holds as $i \in N^+_{e_q^-}$, **B2** holds by definition of e^-_{q+1} and **B3** holds as equation (3.23) implies

$$\mu^{+k}_{i,e_q^-}(x_q^b) < \mu^{-k}_{i,f}(x_q^b) \leqslant \mu^{-k}_{i,e^-_{q+1}}(x_q^b).$$

■

APPENDIX 3.D REMAINING PROOF OF LEMMA 3.5.3

Remaining Proof of Lemma 3.5.3. In the second case we assume that $e_1^+ \in \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$. For resource e_1^+ we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e_1^+}^{1,-}(x_1^b) \\
&=_1 (\mu_{i,\min}^{+k}(x_1^b) - a_{i,e_1^+}k^2) - (\mu_{i,e_1^+}^{-k}(x_1^b) - a_{i,e_1^+}k^2) \\
&=_2 \mu_{i,e_1^+}^{+k}(x_1^b) - \mu_{i,e_1^+}^{-k}(x_1^b) \\
&>_3 0.
\end{aligned}$$

Here, $=_1$ is true by Definition (4.1), $=_2$ as $e_1^+ \in \arg\min_{e\in E}\{\mu_{i,e}^{+k}(x_1^b)\}$ and $>_3$ by Definition (3.11) and (3.12). For resource e_1^- we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e_1^-}^{1,-}(x_1^b) \\
&\quad =_1 (\mu_{i,e_1^+}^{+k}(x_1^b) - k^2 a_{i,e_1^+}) - (\mu_{i,e_1^-}^{-k}(x_1^b) + k^2 a_{i,e_1^-}) \\
&\quad =_2 (\mu_{i,e_1^+}^{+k}(x) + k^2 a_{i,e_1^+}) - (\mu_{i,e_1^-}^{-k}(x) - k^2 a_{i,e_1^-}) \\
&\quad >_3 \mu_{i,\min}^{+k}(x) - \mu_{i,e_1^-}^{-k}(x).
\end{aligned}$$

Here, $=_1$ is true by Definition (4.1), $=_2$ as $\mu_{i,e_1^+}^{+k}(x_1^b) = \mu_{i,e_1^+}^{+k}(x) + 2k^2 a_{i,e_1^+}$ and $\mu_{i,e_1^-}^{-k}(x_1^b) = \mu_{i,e_1^-}^{-k}(x) - 2k^2 a_{i,e_1^-}$, and $>_3$ by definition of $\mu_{i,\min}^{+k}(x)$. For resource $e \in E\backslash\{e_1^+, e_1^-\}$ we have:

$$\begin{aligned}
&B_{i,\min}^{i,+}(x_1^b) - B_{i,e}^{i,-}(x_1^b) \\
&\quad =_1 (\mu_{i,\min}^{+k}(x_1^b) - a_{i,e_1^-}) - \mu_{i,e}^{-k}(x_1^b) \\
&\quad =_2 (\mu_{i,e_1^+}^{+k}(x) + k^2 a_{i,e_1^+}) - \mu_{i,e}^{-k}(x) \\
&\quad >_3 \mu_{i,\min}^{+k}(x) - \mu_{i,e}^{-k}(x).
\end{aligned}$$

Here, $=_1$ is true by Definition (4.1), $=_2$ as $\mu_{i,e_1^+}^{+k}(x_1^b) = \mu_{i,e_1^+}^{+k}(x) + 2k^2 a_{i,e_1^+}$ and $\mu_{i,e}^{-k}(x_1^b) = \mu_{i,e}^{-k}(x)$, and $>_3$ by definition of $\mu_{i,\min}^{+k}(x)$. Thus, when $e_1^+ = \arg\min_{e\in E}\{\mu_{i,e}^{+k}(x_1^b)\}$, we have:

$$B_{i,\min}^{1,+}(x_1^b) - B_{i,e}^{1,-}(x_1^b) > \min\{\mu_{i,\min}^{+k}(x) - \mu_{i,e}^{-k}(x), 0\}.$$

Thus we have $\Delta_i(x) <_{\text{lex}} B_i^1(x_1^b)$ and hence, $\Delta(x) <_{\text{lex}} B^1(x_1^b)$. For the third case we assume that $e_1^- = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(y')\}$. For resource e_1^+ we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e_1^+}^{1,-}(x_1^b)\\
&\quad =_1 (\mu_{i,\min}^{+k}(x_1^b) + k^2 a_{i,e_1^-}) - (\mu_{i,e_1^+}^{-k}(x_1^b) - k^2 a_{i,e_1^+})\\
&\quad >_2 \mu_{i,e_1^-}^{+k}(x_1^b) - \mu_{i,e_1^+}^{-k}(x_1^b)\\
&\quad =_3 \mu_{i,e_1^-}^{-k}(x) - \mu_{i,e_1^+}^{+k}(x)\\
&\quad >_4 0.
\end{aligned}$$

Here, $=_1$ is true by Definition 3.11 and 3.12, $=_2$ as $e_1^- = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$, $>_3$ as player i moved a packet from e_1^- to e_1^+ and $>_4$ as $\mu_{i,e_1^-}^{-k}(x) > \mu_{i,e_1^+}^{+k}(x)$. For resource e_1^- we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e_1^-}^{1,-}(x_1^b)\\
&\quad =_1 (\mu_{i,\min}^{+k}(x_1^b) + k^2 a_{i,e_1^-}) - (\mu_{i,e_1^-}^{-k}(x_1^b) + k^2 a_{i,e_1^-})\\
&\quad =_2 \mu_{i,e_1^-}^{+k}(x_1^b) - \mu_{i,e_1^-}^{-k}(x_1^b)\\
&\quad >_3 0.
\end{aligned}$$

Here, $=_1$ is true by Definition 3.11 and 3.12, $=_2$ as $e_1^- = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$, and $>_3$ by Definition 4.1. For resource $e \in E \backslash \{e_1^+, e_1^-\}$ we have:

$$\begin{aligned}
&B_{i,\min}^{1,+}(x_1^b) - B_{i,e}^{1,-}(x_1^b)\\
&\quad =_1 (\mu_{i,\min}^{+k}(x_1^b) + k^2 a_{i,e_1^-}) - \mu_{i,e}^{-k}(x_1^b)\\
&\quad >_2 \mu_{i,e_1^-}^{+k}(x_1^b) - \mu_{i,e}^{-k}(x)\\
&\quad =_3 \mu_{i,e_1^-}^{-k}(x) - \mu_{i,e}^{-k}(x)\\
&\quad >_4 \mu_{i,e_1^+}^{+k}(x) - \mu_{i,e}^{-k}(x)\\
&\quad \geqslant_5 \mu_{i,\min}^{+k}(x) - \mu_{i,e}^{-k}(x).
\end{aligned}$$

Here, $=_1$ is true by Definition 3.11 and 3.12, $=_2$ as $e_1^- = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$, $=_3$ as player i moves a packet from e_1^- to e_1^+, $>_4$ as $\mu_{i,e_1^-}^{-k}(x) > \mu_{i,e_1^+}^{+k}(x)$ and $\geqslant_5$ by definition of $\mu_{i,\min}^{+k}(x)$.

Thus, when $e_1^- = \arg\min_{e \in E}\{\mu_{i,e}^{+k}(x_1^b)\}$:

$$B_{i,\min}^{1,+}(x_1^b) - B_{i,e}^{1,-}(x_1^b) \geqslant \min\{\mu_{i,\min}^{+k}(x) - \mu_{i,e}^{-k}(x), 0\}.$$

As $\mu^{+k}_{i,\min}(x) - \mu^{-k}_{i,e_1^-}(x) < 0 < B^{1,+}_{i,\min}(x^b_1) - B^{1,-}_{i,e_1^-}(x^b_1)$, we have $\Delta_i(x) <_{\text{lex}} B^1_i(x^b_1)$ and hence, $\Delta(x) <_{\text{lex}} B^1(x^b_1)$. ■

4 EQUILIBRIUM COMPUTATION FOR GAMES WITH CONVEX COSTS

4.1 INTRODUCTION

In Chapter 3, we studied the computation of pure Nash equilibria in atomic splittable congestion games with singleton strategy spaces and affine cost functions, and developed an algorithm that constructs a pure Nash equilibrium within polynomial time. In this chapter, we generalize this model and consider cost functions which are convex and increasing.

4.1.1 *Our Results and Techniques*

ATOMIC SPLITTABLE CONGESTION GAMES We study atomic splittable polymatroid congestion games with player-specific, convex cost functions and show how to compute pure Nash equilibria within pseudo-polynomial time. As equilibria in these games are not guaranteed to be rational, we look for ϵ-approximate equilibria. Here, we say a strategy profile is an ϵ-approximate equilibrium when no player can deviate from her current strategy and decrease her cost by at least ϵ. Similar to Chapter 3, the idea is to compute a pure Nash equilibrium for an associated *integrally splittable* congestion game.

Recall that the class of integrally splittable *singleton* congestion games has been studied before by Tran-Thanh et al. [72] for the case of player-independent convex cost functions. Later, this problem was studied by Harks et al. [32, 36] for the more general case of polymatroid strategy spaces and player-specific convex cost functions. In particular, Harks et al. prove that the algorithm by Tran-Thanh et al. has a running time which is pseudo-polynomial in the aggregated load of the players (cf. Corollary 5.2 [32]). In Section 4.2 we discuss the algorithm by Harks et al. [32] in some more detail.

In Section 4.3.2 we prove that for each $\epsilon > 0$, we can compute a packet size k_ϵ such that the k_ϵ-integral equilibrium is guaranteed to be an ϵ-approximate equilibrium. Then, by using the algorithm by Harks et al. [32, 36], one is able to compute an ϵ-approximate equilibrium within pseudo-polynomial time.

MULTIMARKET OLIGOPOLIES In Section 4.4 we study the problem of computing equilibria in multimarket oligopolies, which were already introduced in Section 1.2.3 and discussed in Section 3.7. For multimarket oligopolies with firm-specific, concave price functions and quadratic production costs, we prove that there exists a polynomial time transformation to atomic splittable congestions games with player-specific, convex costs. Hence, using the same techniques we used to find ϵ-approximate equilibria for atomic splittable congestion games, we are able to compute an ϵ-approximate equilibrium for multimarket oligopolies with concave pricing functions and quadratic production costs within pseudo-polynomial time.

4.1.2 *Related Work*

Most of the related work on this topic is already mentioned in Section 2.1.2 and Section 3.1.2. Additionally, one should be aware of the following two results.

The problem of computing approximate equilibria in atomic splittable congestion games with increasing, convex cost functions has been studied very recently by Bhaskar and Lolakapuri [8]. They devise two algorithms computing an ϵ-approximate equilibrium that runs in exponential time in either the number of resources or the number of players, respectively. Both algorithms are based on guessing the *marginal costs* of the players at an equilibrium. These marginal costs appear to have several monotonicity properties, which they exploit using a high-dimensional binary search algorithm.

Deligkas et al. [24] study the computation of ϵ-approximate equilibria in general concave games with compact strategy spaces and Lipschitz continuous cost functions. In their paper, they decide on a number k, discretize the strategy space and only consider k-uniform points, i.e., vectors where all elements are integer multiples of k. Then, as for each player only finitely many of these vectors exist, they enumerate all feasible k-uniform strategy profiles, and pick the best candidate (see also Lipton et al. [45] for a similar approach). This method results in an algorithm that finds ϵ-approximate equilibria in exponential time.

4.2 PRELIMINARIES

POLYMATROIDS. The definition of a polymatroid and a polymatroid congestion game were discussed in Section 2.2. The definitions found in this chapter differ slightly from the definitions given in Section 2.2, as we introduce d as the *rank* of the polymatroid.

Let $E = \{e_1, \ldots, e_m\}$ be a finite set of resources and $\rho : 2^E \to \mathbb{R}$ be (1) submodular, i.e, $\rho(U) + \rho(V) \geqslant \rho(U \cup V) + \rho(U \cap V)$ for all $U, V \subseteq E$, (2) monotone, i.e., $\rho(U) \leqslant \rho(V)$ for all $U \subseteq V$ and (3) normalized, i.e, $\rho(\emptyset) = 0$. Then, the pair (E, ρ) is called a *polymatroid* and the associated polyhedron is defined as:

$$\mathcal{P}_\rho := \left\{x \in \mathbb{R}^E_{\geqslant 0} \mid x(U) \leqslant \rho(U)\ \forall U \subseteq E \right\},$$

where $x(U) := \sum_{e \in U} x_e$ for all $U \subseteq E$. Given a polyhedron $\mathcal{P}_\rho$ and a rational $d \in \mathbb{Q}_{>0}$ with $d \leqslant \rho(E)$, a polymatroid base polytope of rank d is defined as:

$$\mathcal{P}_\rho(d) := \left\{x \in \mathbb{R}^E_{\geqslant 0} \mid x(U) \leqslant \rho(U)\ \forall U \subseteq E,\ x(E) = d\right\}.$$

ATOMIC SPLITTABLE POLYMATROID CONGESTION GAMES. An atomic splittable polymatroid congestion game is represented by the tuple:

$$\mathcal{G} := (N, E, (d_i)_{i \in N}, (\rho_i)_{i \in N}, (c_{i,e})_{i \in N, e \in E}).$$

Here, $N = \{1, \ldots, n\}$ is a finite player set and with every $i \in N$, we associate a player-specific polymatroid (E, ρ_i). The strategy space of player $i \in N$ is defined as the (player-specific) polymatroid base polytope

$$\mathcal{P}_{\rho_i}(d_i) := \left\{x_i \in \mathbb{R}^E_{\geqslant 0} \mid x_i(U) \leqslant \rho_i(U)\ \forall U \subseteq E,\ x_i(E) = d_i\right\}.$$

The combined strategy space is denoted by $\mathcal{P} := \prod_{i \in N} \mathcal{P}_{\rho_i}(d_i)$ and we denote by $x = (x_i)_{i \in N}$ the overall strategy profile. The entry $x_{i,e}$ of the vector x_i is the load of player i on $e \in E$ and $x_e := \sum_{i \in N} x_{i,e}$ is defined as the total load on e. Resources have player-specific cost functions $c_{i,e}(x_e)$, where $c_{i,e}(x_e)$ is non-decreasing, non-negative, differentiable and convex. We further assume that all cost functions $c_{i,e}(x)$ and their derivatives $c'_{i,e}(x)$ are Lipschitz continuous:

Definition 4.2.1 (Lipschitz continuity). *A function $c : \mathbb{R} \to \mathbb{R}$ is called Lipschitz continuous if there exists a constant $L > 0$ such that for all $x, y \in R$: $|c(y) - c(x)| \leqslant L|y - x|$. Here L is called the Lipschitz constant.*

The total cost of player i in strategy distribution x is defined as

$$\pi_i(x) = \sum_{e \in E} c_{i,e}(x_e)\, x_{i,e}.$$

The goal of each player is to choose a strategy x_i such that her personal cost $\pi_i(x)$ is minimized. As pure Nash equilibria in these games are not guaranteed to be rational, the notion of ϵ-approximate equilibria are suitable.

Definition 4.2.2 (ϵ-approximate equilibrium)**.** *Strategy x is an ϵ-approximate equilibrium for game $\mathcal{G}$, if for each $i \in N$, and every strategy $y_i \in \mathcal{P}_i$:*

$$\pi_i(y_i, x_{-i}) \geqslant \pi_i(x_i, x_{-i}) - \epsilon.$$

A pair $\big(x, (y_i, x_{-i})\big) \in \mathcal{P} \times \mathcal{P}$ is called an *improving move* of player i, if $\pi_i(x_i, x_{-i}) > \pi_i(y_i, x_{-i})$. Given $x_{-i} \in \mathcal{P}_{-i}(d_{-i})$, a strategy $x_i \in \mathcal{P}_i(d_i)$ is called a *best response* of player i to x_{-i} if $\pi_i(x_i, x_{-i}) \leqslant \pi_i(y_i, x_{-i})$ for all $y_i \in \mathcal{P}_i(d_i)$.

When $E, N, (\rho_i)_{i \in N}$ and $(c_{i,e})_{i \in N, e \in E}$ are clear from the context, we refer to the game as $\mathcal{G}((d_i)_{i \in N})$, and write $\mathcal{P}_i(d_i)$ instead of $\mathcal{P}_{\rho_i}(d_i)$. For each $i \in N$, we write $\mathcal{P}_{-i}(d_{-i}) = \prod_{j \neq i} \mathcal{P}_j(d_j)$ and $x = (x_i, x_{-i})$ meaning that $x_i \in \mathcal{P}_i(d_i)$ and $x_{-i} \in \mathcal{P}_{-i}(d_{-i})$.

INTEGRAL POLYMATROID CONGESTION GAMES. A k-integral polymatroid congestion game is given by the tuple

$$\mathcal{G}_k := (N, E, (d_i)_{i \in N}, (\rho_i)_{i \in N}, (c_{i,e})_{i \in N, e \in E}, k)\,.$$

Such games are very similar to atomic splittable polymatroid congestion games, except for the fact that we are given a *packet size* $k \in \mathbb{Q}_{>0}$, where $d_i/k \in \mathbb{N}$ for all $i \in N$. Here, players cannot choose any strategy in the polymatroid base polytope, but only k-integral points. Thus, the strategy space of player i is defined by the (player-specific) k-integral polymatroid base polytope $P^k_{\rho_i}(d_i)$:

$$\left\{x_i \in \mathbb{R}^E_{\geqslant 0} \mid x_i(U) \leqslant \rho_i(U)\ \forall U \subseteq E,\ x_i(E) = d_i \text{ and } x_{i,e} = kq, q \in \mathbb{N}_{\geqslant 0}\right\}.$$

When $E, N, (\rho_i)_{i \in N}$ and $(c_{i,e})_{i \in N, e \in E}$ are clear from the context, we refer to the game as $\mathcal{G}_k((d_i)_{i \in N})$ and to the strategy spaces as $P^k_i(d_i)$.

Similar to atomic splittable congestion games, the complete strategy space of the game is defined by $\mathcal{P}^k := \prod_{i \in N} \mathcal{P}^k_i(d_i)$. For all $i \in N$, we write $\mathcal{P}^k_{-i}(d_{-i}) = \prod_{j \neq i} \mathcal{P}^k_j(d_j)$ and $x = (x_i, x_{-i})$ meaning that $x_i \in \mathcal{P}^k_i(d_i)$ and

$x_{-i} \in \mathcal{P}^k_{-i}(d_{-i})$. A strategy profile x is an *equilibrium* if $\pi_i(x) \leqslant \pi_i(y_i, x_{-i})$ for all $i \in N$ and $y_i \in \mathcal{P}^k_i(d_i)$. A pair $\big(x, (y_i, x_{-i})\big) \in \mathcal{P}^k \times \mathcal{P}^k$ is called an *improving move* of player i, if $\pi_i(x_i, x_{-i}) > \pi_i(y_i, x_{-i})$. Given $x_{-i} \in \mathcal{P}^k_{-i}(d_{-i})$, a strategy $x_i \in \mathcal{P}^k_i(d_i)$ is called a *best response* of player i to x_{-i} if

$$\pi_i(x_i, x_{-i}) \leqslant \pi_i(y_i, x_{-i})$$

for all $y_i \in \mathcal{P}^k_i(d_i)$.

Harks et al. [32] developed an algorithm that computes an exact Nash equilibrium for k-integral splittable polymatroid congestion games with non-negative, increasing and convex cost functions (see Algorithm 7). The running time of this algorithm is pseudo-polynomial in the aggregated demand of the players. Their algorithm [32, Algorithm 1] starts with an equilibrium for the game where the demand for each player is set to zero: $d'_i = 0$. This game has a unique equilibrium, where $x_{i,e} = 0$ for each $i \in N$ and $e \in E$. Then, they repeatedly look for a player for whom $d'_i < d_i$. For this player, they increase d'_i by k and a preliminary equilibrium with respect to the current demands d'_i is recomputed by following a sequence of best responses of the players. The running time of Algorithm 7 is $O(nm(\delta/k)^3)$, where δ is an upper bound on the demands d_i.

Algorithm 7: Computing a k-integral PNE.

Input: $\mathcal{G}_k := (N, E, (\rho_i)_{i \in N}, (c_{i,e})_{i \in N, e \in E}, k)$
Output: A pure Nash equilibrium x

$d'_i \leftarrow 0, \mathcal{P}'_i \leftarrow \mathcal{P}^k_i(0)$ and $x_{i,e} \leftarrow 0$ for all $i \in N$, $e \in E$;
for $q = 1, \dots, \left(\sum_{i \in N} d_i\right)/k$ **do**
  Choose $i \in N$ with $d'_i < d_i$;
  $d'_i \leftarrow d'_i + k; \mathcal{P}'_i \leftarrow \mathcal{P}^k_i(d'_i)$;
  Choose a best response $y_i \in \mathcal{P}'_i$ with $||y_i - x_i||_1 = k$;
  $x_i \leftarrow y_i$;
  while $\exists i \in N$ *who can improve in* $\mathcal{G}'_k := \big(N, E, (\mathcal{P}'_i)_{i \in N}, (c_{i,e})_{i \in N, e \in E}\big)$ **do**
    Compute a best response $y_i \in \mathcal{P}_i$ with $||y_i - x_i||_1 = 2k$;
    $x_i \leftarrow y_i$;
  end
end
Return x;

4.3 LIPSCHITZ CONTINUITY AND APPROXIMATE EQUILIBRIA

As a main result in this chapter, we prove that when all cost functions and their derivatives are *Lipschitz continuous* with Lipschitz constant L, we can decide on a packet size k_ϵ such that the exact k_ϵ-integral equilibrium is an ϵ-approximate atomic splittable equilibrium.

In Section 4.3.1 we first discuss a sufficient and necessary conditions for a strategy x to be either an atomic splittable equilibrium or a k-integral equilibrium. Then, in Section 4.3.2, we use these conditions to prove that for any $\epsilon > 0$, we can find a packet size k_ϵ such that the exact k_ϵ-integral equilibrium is an ϵ-approximate atomic splittable equilibrium. As we can find a k_ϵ-integral equilibrium within pseudo-polynomial time using 7 by Harks, Peis and Klimm, this implies that we can find ϵ-approximate equilibria within pseudo-polynomial time.

4.3.1 *Equilibrium Conditions in* k*-Integral Games*

In atomic splittable congestion games, the *marginal cost* for player i on resource e is defined as:

$$\mu_{i,e}(x) = c_{i,e}(x_e) + x_{i,e}c'_{i,e}(x_e).$$

Intuitively, the marginal cost represents the cost increase of player i when she would increase her load on resource e. For k-integral games, the marginal costs are defined as follows:

$$\mu^{+k}_{i,e}(x) = (x_{i,e} + k)c_{i,e}(x_e + k) - x_{i,e}c_{i,e}(x_e), \tag{4.1}$$

$$\mu^{-k}_{i,e}(x) = \begin{cases} x_{i,e}c_{i,e}(x_e) - (x_{i,e} - k)c_{i,e}(x_e - k), & \text{if } x_{i,e} > 0, \\ -\infty, & \text{if } x_{i,e} \leqslant 0. \end{cases} \tag{4.2}$$

Here, $\mu^{+k}_{i,e}(x)$ is the cost for player i to add one packet of size k to resource e and $\mu^{-k}_{i,e}(x)$ is the gain for player i for removing a packet of size k from resource e.

Fujishige studied the properties of any optimal strategy that minimizes a seperable convex cost function $\sum_{e \in E} c_e(x_e)$ over a submodular system. In this context, he showed in Theorem 8.1 [26] that for any strategy that can be obtained from the optimal strategy by a single exchange from e to e', the right derivative of $c(x_{e'})$ is at least the left derivative of $c(x_e)$. When we apply Theorem 8.1 [26] to k-integral games, we obtain the following lemma:

Lemma 4.3.1. *Given an atomic splittable congestion game $\mathcal{G}$, and a packet size k such that $\rho_i(U)/k \in \mathbb{N}$ for all $U \subseteq E, i \in N$. Given a strategy $x_{-i} \in \mathcal{P}^k_{-i}(d_{-i})$ for k-integral splittable game $\mathcal{G}_k$, then strategy $x_i \in P^k_i(d_i)$ is a best response for player i if and only if for every pair $(e, e') \in E^2$ the following holds. If there exists an $\alpha > 0$ such that: $x_i + \alpha(\chi_e - \chi_{e'}) \in \mathcal{P}_i(d_i)$, then we have: $\mu^{+k}_{i,e}(x) \geqslant \mu^{-k}_{i,e'}(x)$.*

Proof. Assume there exists an $\alpha > 0$ such that: $x_i + \alpha(\chi_e - \chi_{e'}) \in \mathcal{P}_i(d_i)$. Then, as $\rho_i(U)/k \in \mathbb{N}$ for all $U \subseteq E$ and all $i \in N$, we can define $\alpha' = k\lceil \frac{1}{k}\alpha \rceil$, for which it holds that $x_i + \alpha'(\chi_e - \chi_{e'}) \in \mathcal{P}^k_i(d_i)$. Then, Theorem 8.1 [26] implies that $\mu^{+k}_{i,e}(x) \geqslant \mu^{-k}_{i,e'}(x)$. ∎

When the cost functions $c_{i,e}(x)$ and their derivatives $c'_{i,e}(x)$ are Lipschitz continuous with Lipschitz constant L, we obtain the following relations between marginal costs in atomic splittable games and k-integral games:

Lemma 4.3.2. *For any feasible strategy x for game $\mathcal{G}_k$, for $\delta = \max_{i \in N}\{d_i\}$, we have:*

1. $\frac{1}{k}\mu^{+k}_{i,e}(x) \leqslant \mu_{i,e}(x) + kL(\delta + 1)$,
2. $\frac{1}{k}\mu^{-k}_{i,e}(x) \geqslant \mu_{i,e}(x) - kL(\delta + 1)$ *whenever* $x_{i,e} > 0$.

Proof. We start by proving $\mu^{+k}_{i,e}(x)/k \leqslant \mu_{i,e}(x) + kL(d_i + 1)$. We obtain:

$$\begin{aligned}
\mu^{+k}_{i,e}(x)/k &= \frac{(x_{i,e} + k)c_{i,e}(x_e + k) - x_{i,e}c_{i,e}(x_e)}{k} \\
&= x_{i,e}\frac{c_{i,e}(x_e + k) - c_{i,e}(x_e)}{k} + c_{i,e}(x_e + k) \\
&\leqslant_1 x_{i,e}c'(x_e + k) + c_{i,e}(x_e + k) \\
&\leqslant_2 x_{i,e}(c'(x_e) + kL) + c_{i,e}(x_e) + kL \\
&= x_{i,e}c'(x_e) + c_{i,e}(x_e) + (x_{i,e} + 1)kL \\
&\leqslant \mu_{i,e}(x) + kL(\delta + 1).
\end{aligned}$$

Inequality $\leqslant_1$ holds as $c_{i,e}$ is convex and increasing, which implies $c'(x_e + k) \geqslant (c_{i,e}(x_e + k) - c_{i,e}(x_e))/k$. We obtain inequality $\leqslant_2$ using Lipschitz constant L, as it implies $c(x_e + k) \leqslant c(x_e) + kL$ and $c'(x_e + k) \leqslant c'(x_e) + kL$. The second inequality can be obtained in a similar way. ∎

Lemma 4.3.3. *Given an atomic splittable congestion game $\mathcal{G}$, and a packet size k such that $\rho_i(U)/k \in \mathbb{N}$ for all $U \subseteq E$ and all $i \in N$. Let $x_i \in \mathcal{P}_i^k(d_i)$ be a strategy for player i in an k-integral game. Then, if there exists an $\alpha > 0$ such that:*

$$x_i + \alpha(\chi_{e'} - \chi_e) \in \mathcal{P}_i(d_i),$$

it holds that:

$$\mu_{i,f}(x) - \mu_{i,e}(x) \geqslant -2kL(\delta + 1).$$

Proof. First note that Lemma 4.3.1 states that:

$$\mu_{i,e}^{-k}(x) \leqslant \mu_{i,f}^{+k}(x).$$

As $x_i + \alpha(\chi_{e'} - \chi_e) \in \mathcal{P}_i(d_i)$, we have $x_{i,e} > 0$. We combine this with Lemma 4.3.2 to obtain that:

$$\mu_{i,e}(x) - kL(\delta + 1) \leqslant \tfrac{1}{k}\mu_{i,e}^{-k}(x) \leqslant \tfrac{1}{k}\mu_{i,f}^{+k}(x) \leqslant \mu_{i,f}(x) + kL(\delta + 1).$$

We rewrite the previous inequality, and obtain the desired statement:

$$\mu_{i,f}(x) - \mu_{i,e}(x) \geqslant -2kL(\delta + 1).$$

■

4.3.2 *Atomic Splittable and k-Integral Equilibria*

In this section we focus on finding ϵ-approximate equilibria for atomic splittable congestion games with increasing, non-negative, differentiable and convex cost functions, where both the original function and its derivative are bounded by Lipschitz constant L. In order to do so, for each $\epsilon > 0$ we define a k_ϵ, and prove that the k_ϵ-integral equilibrium will be an ϵ-approximate atomic splittable equilibrium.

Theorem 4.3.4. *Given an atomic splittable game $\mathcal{G}$, where all cost functions and their derivatives are bounded by a Lipschitz constant L. Then, for any $\epsilon > 0$, there exists a $k_\epsilon > 0$ such that an exact equilibrium x for the k_ϵ-integral splittable game $\mathcal{G}_{k_\epsilon}$ is an ϵ-equilibrium for $\mathcal{G}$.*

Proof. As π_i is convex, for any alternative strategy $y_i \in \mathcal{P}_i(d_i)$, we have that:

$$\pi_i(y_i, x_{-i}) \geqslant \pi_i(x_i, x_{-i}) + \nabla_i \pi_i(x_i, x_{-i}) \cdot (y_i - x_i). \tag{4.3}$$

Thus, our goal is to determine a k_ϵ that bounds $\nabla_i \pi_i(x_i, x_{-i}) \cdot (y_i - x_i)$ from below by ϵ. We define

$$E^{i,+} := \{e \in E | y_{i,e} > x_{i,e}\} \text{ and } E^{i,-} := \{e \in E | x_{i,e} > y_{i,e}\}.$$

Note that:

$$\nabla_i \pi_i(x_i, x_{-i}) \cdot (y_i - x_i) = \sum_{e \in E^{i,+}} \mu_{i,e}(x)(y_{i,e} - x_{i,e}) + \sum_{e \in E^{i,-}} \mu_{i,e}(x)(y_{i,e} - x_{i,e}). \tag{4.4}$$

Consider the complete, directed, bipartite graph $G(x_i, y_i)$ on node sets $E^{i,-}$ and $E^{i,+}$, where each node $e \in E^{i,-}$ has a supply of $x_{i,e} - y_{i,e}$ and each node $e \in E^{i,+}$ has a demand of $y_{i,e} - x_{i,e}$. The edges of $G(x_i, y_i)$ are directed from $E^{i,-}$ to $E^{i,+}$ and the capacity $c_{e,e'}(x_i, y_i)$ of edge $e, e' \in E^{i,-} \times E^{i,+}$ is defined as:

$$c_{e,e'}(x_i, y_i) = \max\{\alpha | x_i + \alpha(\chi_{e'} - \chi_e) \in \mathcal{P}_i(d_i)\}.$$

Then, as y_i and x_i are points in the polymatroid base polytope $\mathcal{P}_i(d_i)$, there exists a transshipment t in $G(x_i, y_i)$ from resources in $E^{i,-}$ to resources in $E^{i,+}$ that exactly satisfies all supplies, demands and capacities (Lemma 2.3.2). We denote by $t_{e,f}$ the amount of load transshipped from resource e to resource f in t, thus: $\sum_{f \in E^{i,+}} t_{e,f} = x_{i,e} - y_{i,e}$ if $e \in E^{i,-}$ and $\sum_{f \in E^{i,-}} t_{f,e} = y_{i,e} - x_{i,e}$ if $e \in E^{i,+}$ Using this transshipment, we rewrite (4.4) in terms of t. Hence,

$$\begin{aligned}
&\nabla_i \pi_i(x_i, x_{-i}) \cdot (y_i - x_i) \\
&= \sum_{e \in E^{i,+}} \mu_{i,e}(x)(y_{i,e} - x_{i,e}) + \sum_{e \in E^{i,-}} \mu_{i,e}(x)(y_{i,e} - x_{i,e}) \\
&= \sum_{e \in E^{i,+}} \mu_{i,e}(x) \left(\sum_{f \in E^{i,-}} t_{f,e} \right) - \sum_{e \in E^{i,-}} \mu_{i,e}(x) \left(\sum_{f \in E^{i,+}} t_{e,f} \right) \\
&= \sum_{(f,e) \in E^{i,-} \times E^{i,+}} \mu_{i,e}(x) t_{f,e} - \sum_{(e,f) \in E^{i,-} \times E^{i,+}} \mu_{i,e}(x) t_{e,f} \\
&= \sum_{(e,f) \in E^{i,-} \times E^{i,+}} \mu_{i,f}(x) t_{e,f} - \sum_{(e,f) \in E^{i,-} \times E^{i,+}} \mu_{i,e}(x) t_{e,f} \\
&= \sum_{(e,f) \in E^{i,-} \times E^{i,+}} (\mu_{i,f}(x) - \mu_{i,e}(x)) t_{e,f}. \qquad (4.5)
\end{aligned}$$

Note that, in order to use Lemma 4.3.3, we need a packet size k such that $\rho_i(U)/k \in \mathbb{N}$ for all $U \subseteq E$ and all $i \in N$. Note that $\rho_i(U) \in \mathbb{Q}_{\geqslant 0}$ for all $U \subseteq E$ and all $i \in N$, hence, we define

$$\rho_{\mathrm{gcd}} := \max\{a \in \mathbb{Q}_{>0} \mid a \leqslant 1 \text{ and } \forall i \in N, U \subseteq E,\ \exists \ell \in \mathbb{N} \text{ s.t. } \rho_i(U) = a \cdot \ell\}.$$

Given any $\epsilon > 0$, we define

$$k_\epsilon = \frac{\rho_{\mathrm{gcd}}}{\left\lceil \frac{2m^2 L\delta(\delta+1)}{\epsilon} \right\rceil}. \tag{4.6}$$

Note that k_ϵ has the following two properties: (1) as $\rho_{\mathrm{gcd}} \leqslant 1$, we know that $k_\epsilon \leqslant \frac{\epsilon}{2m^2 L\delta(\delta+1)}$; (2) as $\rho_{\mathrm{gcd}}/k_\epsilon \in \mathbb{N}$, we know that $\rho_i(U)/k_\epsilon \in \mathbb{N}$ for all $U \subseteq E$ and all $i \in N$.

We prove that the k_ϵ-integral equilibrium is also an ϵ-approximate equilibrium for the corresponding atomic splittable game. Using Lemma 4.3.3, we know that if there exists an $\alpha > 0$ such that: $x_i + \alpha(\chi_{e'} - \chi_e) \in \mathcal{P}_i(d_i)$, we have that:

$$\mu_{i,f}(x) - \mu_{i,e}(x) \geqslant -2k_\epsilon L(\delta + 1) \geqslant -\frac{\epsilon}{m^2\delta}. \tag{4.7}$$

By the choice of transshipment t, we have that

$$x_i + t_{e,f}(\chi_f - \chi_e) \in \mathcal{P}_i(d_i).$$

We combine Equation (4.7) with Equation (4.5) and obtain:

$$\nabla_i \pi_i(x_i, x_{-i}) \cdot (y_i - x_i) \geqslant -\left(\sum_{e,f \in E^{i,-} \times E^{i,+}} \frac{\epsilon}{m^2\delta} t_{e,f}\right). \tag{4.8}$$

Note that $t_{e,f} \leqslant \delta$ and $|E^{i,-} \times E^{i,+}| < m^2$. Hence:

$$\nabla_i \pi_i(x_i, x_{-i}) \cdot (y_i - x_i) > -\epsilon.$$

Using Equation (4.3) we obtain:

$$\pi_i(y_i, x_{-i}) > \pi_i(x_i, x_{-i}) - \epsilon.$$

Thus, player i cannot gain more than ϵ by playing an alternative strategy y_i. As player i was chosen arbitrarily, x is an ϵ-approximate equilibrium. ∎

Corollary 4.3.5. *Given an atomic splittable polymatroid congestion game, where the cost functions are non-negative, increasing, differentiable, convex, and where both the original function as its derivative are bounded by a Lipschitz constant, we can compute an* ϵ*-approximate equilibrium within a running time*

$$O\left(nm\left(\frac{\delta\left\lceil\frac{2m^2L\delta(\delta+1)}{\epsilon}\right\rceil}{\rho_{gcd}}\right)^3\right)=O\left(nm^7\delta^7\left(\frac{L}{\epsilon\cdot\rho_{gcd}}\right)^3\right).$$

Proof. Assume we are given $\epsilon > 0$. Using Theorem 4.3.4, we can then find a packet size k_ϵ such that for any $k \leqslant k_\epsilon$, any k-splittable equilibrium is an ϵ-approximate equilibrium. Using the Algorithm [36, Algorithm 1] by Harks, Peis and Klimm, we can compute a k-splittable equilibrium within running time $O(nm(\delta/k)^3)$. Thus, using the definition of k_ϵ in (4.6), we can find ϵ-approximate equilibria within the required running time. ∎

Corollary 4.3.6. *Given an atomic splittable singleton game, where the cost functions are non-negative, increasing, differentiable, convex, and where both the original function as its derivative are bounded by a Lipschitz constant, we can compute an* ϵ*-approximate equilibrium within a running time*

$$O\left(nm^7\delta^7\left(\frac{L}{\epsilon\cdot\rho_{gcd}}\right)^3\right).$$

4.4 MULTIMARKET COURNOT OLIGOPOLY

In this section, we derive a strong connection between atomic splittable singleton congestion games with convex cost functions and multimarket Cournot oligopolies with concave, decreasing and differentiable price functions and quadratic costs. Such a game is compactly represented by the tuple

$$\mathcal{M} = (N, E, (E_i)_{i\in N}, (p_{i,e})_{i\in N, e\in E_i}, (C_i)_{i\in N}),$$

where N is a set of n firms and E a set of m markets. Each firm i only has access to a subset $E_i \subseteq E$ of the markets and each market e is endowed with firm-specific, non-increasing, differentiable and concave price functions $p_{i,e}(t) : \mathbb{R} \to \mathbb{R}$, for all $i \in N$. In a strategy profile, a firm $i \in N$ chooses a non-negative production quantity $x_{i,e} \in \mathbb{R}_{\geqslant 0}$ for each market $e \in E_i$. We denote a strategy profile for a firm by $x_i = (x_{i,e})_{e\in E_i}$, and a joint strategy profile by $x = (x_i)_{i \in N}$. The production costs of a firm are of the form $C_i(t) = c_i t^2$ for some $c_i \geqslant 0$. The goal of each firm $i \in N$ is to maximize its utility, which is given by:

$$u_i(x) = \sum_{e \in E_i} p_{i,e}(x_e) x_{i,e} - C_i\Big(\sum_{e \in E_i} x_{i,e}\Big),$$

where $x_e := \sum_{i \in N} x_{i,e}$. Note that a connection between Cournot games with affine price functions and atomic splittable games with affine cost functions has already been made in Section 3.7. In the rest of this section we generalize the connection stated in Section 3.7 and prove that several results that hold for atomic splittable equilibria and k-splittable equilibria in games with convex cost functions carry over to multimarket oligopolies with concave price functions.

More precisely, we prove that for each multimarket oligopoly with concave, decreasing and differentiable price functions and quadratic costs, there exists an isomorphic atomic splittable game with convex, increasing and differentiable costs. Moreover, we can construct the isomorphism in polynomial time.

Theorem 4.4.1. *Given a multimarket oligopoly $\mathcal{M}$ with concave, decreasing and differentiable price functions and quadratic costs, there exists an atomic splittable singleton game $\mathcal{G}$ with convex, increasing and differentiable costs that is isomorphic to $\mathcal{M}$. Moreover, if the price functions in $\mathcal{M}$ are Lipschitz continuous with Lipschitz constant L, then the cost functions in $\mathcal{G}$ are Lipschitz continuous with Lipschitz constant* $\max\{L, c_i\}$.

Proof. This proof generalizes a similar transformation stated in [34, Theorem 7.2]. Given multimarket oligopoly $\mathcal{M}$, we construct an atomic splittable singleton game $\mathcal{G}$ as follows. For every firm $i \in N$ we create a player i and we define the demand d_i for this player as an upper bound on the maximal quantity that firm i will produce, that is,

$$d_i := \sum_{e \in E_i} \max\{t \mid p_{i,e}(t) = 0\}.$$

Note that a rational player i would never produce more than d_i, as this implies that she charges a negative price in at least one of the markets in E_i. Thus, we can limit the strategy space for each player $i \in N$ in game $\mathcal{M}$ to strategies x satisfying $\sum_{e \in E_i} x_{i,e} \leqslant d_i$, and preserve all equilibria. Then, for every player i we introduce a special resource e_i, and we define the set of allowable resources for this player as:

$$\tilde{E}_i = E_i \cup \{e_i\} \text{ with } e_i \neq e_j \text{ for } i \neq j.$$

The cost on these special resources e_i is defined as:

$$c_{i,e_i}(t) := c_i(t - 2d_i) \text{ for all } i \in N,$$

which is affine and increasing, and hence differentiable, convex and Lipschitz continuous with Lipschitz constant c_i. The cost on resources $e \in E_i$ is defined as:

$$c_{i,e}(t) := -p_{i,e}(t) \text{ for all } i \in N.$$

Note that as $p_{i,e}(t)$ is concave, differentiable, decreasing and Lipschitz continuous with constant L, $-p_{i,e}(t)$ is convex, differentiable, increasing and Lipschitz continuous with constant L. Note that all const functions are lipschitz continuous with constant

$$L' := \max\{\{L\} \cup \{c_i\}_{i \in N}\}.$$

In order to guarantee that all cost functions are non-negative, one can add a large positive constant $c_{\max}$ to every cost function. We define:

$$c_{\max} = \max\left\{\{p_{i,e}(0) \mid \text{ for all } i \in N, e \in E_i\} \cup \{2c_i d_i \mid \text{ for all } i \in N\}\right\}.$$

Note that adding $c_{\max}$ to every cost function does not change the equilibrium, it only adds $d_i c_{\max}$ to the total cost of each player. The total cost of a strategy x for player i in game $\mathcal{G}$ is:

$$\pi_i(x') = \sum_{e \in \tilde{E}_i} c_{i,e}(x'_e) x'_{i,e},$$

which is equal to:

$$\pi_i(x') = \sum_{e \in E_i} -p_{i,e}(x'_e) x'_{i,e} + x'_{i,e_i} c_i(x'_{i,e_i} - 2d_i). \tag{4.9}$$

Note that the utility function of player i in x' is $\nu_i(x') = -\pi_i(x')$. It is left to prove that game $\mathcal{G}$ is isomorphic to game $\mathcal{M}$. Let x be a feasible strategy in game $\mathcal{M}$. For each $i \in N$, we define the bijective function $\phi_i : E_i \to \tilde{E}$ as:

$$\phi_i(x_{i,1}, \dots, x_{i,m}) = (x_{i,1}, \dots, x_{i,m}, d_i - \sum_{e \in E_i} x_{i,e}) =: (x'_{i,1}, \dots, x'_{i,m}, x'_{i,m+1}).$$

As we limited the strategy space for each player $i \in N$ in game $\mathcal{M}$ to strategies x where $\sum_{e \in E_i} x_{i,e} \leqslant d_i$, $x' := \phi(x)$ is a feasible strategy in $\mathcal{G}$. For each feasible strategy x for game $\mathcal{M}$, and for each $i \in N$, we have:

$$\begin{aligned}
u_i(x) &= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - C_i\Big(\sum_{e \in E_i} x_{i,e}\Big) \\
&= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - c_i\Big(d_i - \sum_{e \in E_i} x_{i,e}\Big)\Big(-d_i - \sum_{e \in E_i} x_{i,e}\Big) - d_i^2 \\
&= \sum_{e \in E_i} p_{i,e}(x_e)x_{i,e} - c_i\Big(d_i - \sum_{e \in E_i} x_{i,e}\Big)\Big(d_i - \sum_{e \in E_i} x_{i,e} - 2d_i\Big) - d_i^2 \\
&= -\pi_i(\phi_1(x_1), \dots, \phi_1(x_n)) - d_i^2 \\
&= \nu_i(\phi_1(x_1), \dots, \phi_1(x_n)) - d_i^2.
\end{aligned}$$

Thus, game $\mathcal{M}$ and $\mathcal{G}$ are isomorphic. ■

Remark 4.4.2. *Given a multimarket oligopoly $\mathcal{M}$, one can construct an atomic splittable singleton game isomorphic to $\mathcal{M}$ within running time $O(nm)$.*

Combining Corollory 4.3.5 and Theorem 4.4.1 we obtain the following result.

Theorem 4.4.3. *Given a multimarket oligopoly $\mathcal{M}$, where all cost functions and their derivatives are Lipschitz continuous with Lipschitz constant L, one can compute an ϵ-approximate equilibrium within a running time that is pseudo-polynomial in $\max\{d_i \mid i \in N\}$, L and $\frac{1}{\epsilon}$.*

4.5 CONCLUDING REMARKS

In this chapter we studied the construction of ϵ-approximate Nash equilibria in atomic splittable polymatroid congestion games with player-specific, convex cost functions. As our main result, we proved that for any $\epsilon > 0$, we can find a packet size k_ϵ such that any k_ϵ-splittable equilibrium is an ϵ-approximate atomic splittable equilibrium. Note that it was known that this k_ϵ-splittable equilibrium can be found in pseudo-polynomial time using Algorithm 1 by Harks, Peis and Klimm [32].

Then, we considered multimarket oligopolies with decreasing, concave price functions and quadratic production costs and showed that there exists a polynomial time transformation to atomic splittable congestion games. Using our first result, this implies that we can compute ϵ-approximate

Cournot-Nash equilibria for multimarket oligopolies with player specific, decreasing, concave pricing functions and quadratic cost functions within pseudo-polynomial time.

BIBLIOGRAPHY

[1] Melika Abolhassani, Mohammad Hossein Bateni, MohammadTaghi Hajiaghayi, Hamid Mahini, and Anshul Sawant. „Network Cournot Competition". In: *Web and Internet Economics*. Ed. by Tie-Yan Liu, Qi Qi, and Yinyu Ye. Cham: Springer International Publishing, 2014, pp. 15–29. ISBN: 978-3-319-13129-0.

[2] H. Ackermann, H. Röglin, and B. Vöcking. „On The Impact of Combinatorial Structure on Congestion Games". In: *J. ACM* 55.6 (2008), pp. 1–22.

[3] H. Ackermann, H. Röglin, and B. Vöcking. „Pure Nash equilibria in player-specific and weighted congestion games". In: *Theoret. Comput. Sci.* 410.17 (2009), pp. 1552–1563.

[4] R. Amir. „Cournot Oligopoly and the Theory of Supermodular Games". In: *Games Econom. Behav.* 15.2 (1996), pp. 132–148.

[5] R. Aumann and S. Hart, eds. *Handbook of Game Theory with Economic Applications*. Vol. 3. 3. Amsterdam, Netherlands: Elsevier, 2002.

[6] Robert J. Aumann. „What is game theory trying to accomplish?" In: *Frontiers of economics*. Oxford: Basil Blackwell, 1985, pp. 28–76.

[7] M. Beckmann, C. McGuire, and C. Winsten. *Studies in the Economics and Transportation*. New Haven, CT, USA: Yale University Press, 1956.

[8] U. Bhaskar and P. R. Lolakapuri. „Equilibrium Computation in Atomic Splittable Routing Games with Convex Cost Functions". In: *arXiv* abs/1804.10044 (2018).

[9] U. Bhaskar, L. Fleischer, D. Hoy, and C.-C. Huang. „Equilibria of Atomic Flow Games are not Unique". In: *Math. Oper. Res.* 40.3 (2015), pp. 634–654.

[10] Kostas Bimpikis, Shayan Ehsani, and Rahmi Ilkiliç. „Cournot Competition in Networked Markets". In: *Proceedings of the Fifteenth ACM Conference on Economics and Computation*. EC '14. Palo Alto, California, USA: ACM, 2014, pp. 733–733. ISBN: 978-1-4503-2565-3.

[11] R.A. Brualdi. „Exchange systems, matchings, and transversals". In: *J. Comb. Theory* 5.3 (1968), pp. 244–257.

[12] R.A. Brualdi and E.B. Scrimger. „Comments on bases in dependence structures". In: *Bull, Austral. Math. Soc.* 1 (1969), pp. 161–167.

[13] J. Bulow, J. Geanakoplos, and P. Klemperer. „Multimarket Oligopoly: Strategic Substitutes and Complements". In: *J. Polit. Econ.* 93.3 (1985), pp. 488–511.

[14] Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. „Efficient Computation of Approximate Pure Nash Equilibria in Congestion Games". In: *FOCS 2011, Palm Springs, CA, USA, 2011*. 2011, pp. 532–541.

[15] Ioannis Caragiannis, Angelo Fanelli, Nick Gravin, and Alexander Skopalik. „Approximate Pure Nash Equilibria in Weighted Congestion Games: Existence, Efficient Computation, and Structure". In: *ACM Trans. Economics and Comput.* 3.1 (2015), p. 2.

[16] X. Chen, X. Deng, and S.-H. Teng. „Settling the complexity of computing two-player Nash equilibria". In: *J. ACM* 56.3 (2009), 14:1–14:55.

[17] S. Chien and A. Sinclair. „Convergence to approximate Nash equilibria in congestion games". In: *Games Econom. Behav.* 71.2 (2011), pp. 315–327.

[18] R. Cominetti, J. R. Correa, and N. E. Stier-Moses. „The Impact of Oligopolistic Competition in Networks". In: *Operations Research* 57.6 (2009), pp. 1421–1437.

[19] R. Cottle, J. Pang, and R. Stone. *The Linear Complementarity Problem*. Society for Industrial and Applied Mathematics, 2009.

[20] A. A. Cournot. *Recherches sur les principes mathématiques de la théorie des richesses/par Augustin Cournot*. L. Hachette, 1838.

[21] Constantinos Daskalakis, Paul W. Goldberg, and Christos H. Papadimitriou. „The Complexity of Computing a Nash Equilibrium". In: *SIAM J. Comput.* 39.1 (2009), pp. 195–259.

[22] Alberto Del Pia, Michael Ferris, and Carla Michini. „Totally Unimodular Congestion Games". In: *Proc. 28th Annual ACM-SIAM Sympos. on Discrete Algorithms*. 2017.

[23] Argyrios Deligkas, John Fearnley, and Paul G. Spirakis. „Lipschitz Continuity and Approximate Equilibria". In: *SAGT, Liverpool, UK, 2016*. 2016, pp. 15–26.

[24] Argyrios Deligkas, John Fearnley, and Paul Spirakis. „Lipschitz Continuity and Approximate Equilibria". In: *Algorithmic Game Theory: 9th International Symposium, SAGT 2016, Liverpool, UK, September 19–21, 2016, Proceedings*. Ed. by Martin Gairing and Rahul Savani. Berlin, Heidelberg: Springer Berlin Heidelberg, 2016, pp. 15–26. ISBN: 978-3-662-53354-3.

[25] A. Fabrikant, C. Papadimitriou, and K. Talwar. „The complexity of pure Nash equilibria". In: *STOC 2004*. Ed. by László Babai. 2004, pp. 604–612.

[26] S. Fujishige. *Submodular functions and Optimization*. Elsevier, 2005.

[27] S. Fujishige, M.X. Goemans, T. Harks, and B. Peis. „Matroids are Immune to Braess Paradox". http://arxiv.org/abs/1504.07545F. 2015.

[28] M. Gairing, B. Monien, and K. Tiemann. „Routing (Un-)Splittable Flow in Games with Player-Specific Linear Latency Functions". In: *ACM Trans. Algorithms* 7.3 (2011), pp. 1–31.

[29] H. Groenevelt. „Two algorithms for maximizing a separable concave function over a polymatroid feasible region". In: *Eur. J. Oper. Res.* 54.2 (1991), pp. 227 –236.

[30] Frank Harary and Dominic Welsh. „Matroids versus graphs". English. In: *The Many Facets of Graph Theory*. Ed. by G. Chartrand and S.F. Kapoor. Vol. 110. Lecture Notes in Mathematics. Springer Berlin Heidelberg, 1969, pp. 155–170. ISBN: 978-3-540-04629-5.

[31] T. Harks. „Stackelberg Strategies and Collusion in Network Games with Splittable Flow". In: *Theory Comput. Syst.* 48 (4 2011), pp. 781–802.

[32] T. Harks, M. Klimm, and B. Peis. „Sensitivity Analysis for Convex Separable Optimization over Integral Polymatroids". In: *Siam J. of Optimization* forthcoming (2018).

[33] T. Harks, T. Oosterwijk, and T. Vredeveld. „A Logarithmic Approximation for Polymatroid Matroid Congestion Games". Unpublished manuscript. 2014.

[34] T. Harks and V. Timmermans. „Equilibrium Computation in Resource Allocation Games". In: *https://arxiv.org/pdf/1612.00190.pdf* (2016).

[35] Tobias Harks and Max Klimm. „Multimarket Oligopolies with Restricted Market Access". In: *Algorithmic Game Theory*. Ed. by Ron Lavi. Berlin, Heidelberg: Springer Berlin Heidelberg, 2014, pp. 182–193. ISBN: 978-3-662-44803-8.

[36] Tobias Harks, Max Klimm, and Britta Peis. „Resource Competition on Integral Polymatroids". In: *WINE 2014*. Ed. by Tie-Yan Liu, Qi Qi, and Yinyu Ye. 2014, pp. 189–202.

[37] Tobias Harks and Veerle Timmermans. „Equilibrium Computation in Atomic Splittable Singleton Congestion Games". In: *Integer Programming and Combinatorial Optimization*. Ed. by Friedrich Eisenbrand and Jochen Koenemann. Cham: Springer International Publishing, 2017, pp. 442–454. ISBN: 978-3-319-59250-3.

[38] A. Haurie and P. Marcotte. „On the relationship between Nash-Cournot and Wardrop equilibria". In: *Networks* 15 (1985), pp. 295–308.

[39] C.-C. Huang. „Collusion in Atomic Splittable Routing Games". In: *Theory Comput. Syst.* 52.4 (2013), pp. 763–801.

[40] S. Kakutani. „A generalization of Brouwer's fixed point theorem". In: *Duke Mathematics Journal* 8.3 (1941), pp. 457–458.

[41] Y. Korilis, A. Lazar, and A. Orda. „Capacity Allocation Under Noncooperative Routing". In: *IEEE Trans. on Aut. Contr.* 42.3 (1997), pp. 309–325.

[42] N.M. Korneyenko. „Combinatorial algorithms on a class of graphs". In: *Discrete Appl. Math.* 54.2-3 (1994), pp. 215–217.

[43] Elias Koutsoupias and Christos Papadimitriou. „Worst-case equilibria". In: *Proc. 16th Internat. Sympos. Theoretical Aspects of Comput. Sci.* Ed. by C. Meinel and S. Tison. Vol. 1563. LNCS. 1999, pp. 404–413.

[44] N. Kukushkin. „A fixed-point theorem for decreasing mappings". In: *Econ. Lett.* 46 (1994), pp. 23–26.

[45] Richard J. Lipton, Evangelos Markakis, and Aranyak Mehta. „Playing large games using simple strategies". In: *Proceedings 4th ACM Conference on Electronic Commerce (EC-2003), San Diego, California, USA, June 9-12, 2003*. 2003, pp. 36–41.

[46] Patrice Marcotte. „Algorithms for the Network Oligopoly Problem". In: *Journal of the Operational Research Society* 38.11 (1987), pp. 1051–1065.

[47] Frédéric Meunier and Thomas Pradeau. „The uniqueness property for networks with several origin-destination pairs". In: *Eur. J. Oper. Res.* 237.1 (2012), pp. 245–256.

[48] Frédéric Meunier and Thomas Pradeau. „A Lemke-Like Algorithm for the Multiclass Network Equilibrium Problem". In: *WINE 2013*. 2013, pp. 363–376.

[49] I. Milchtaich. „Topological Conditions for Uniqueness of Equilibrium in Networks". In: *Math. Oper. Res.* 30.1 (2005), pp. 225–244.

[50] P. Milgrom and J. Roberts. „Rationalizability, learning, and equilibrium in games with strategic complementarities". In: *Econometrica* 58 (1990), pp. 1255–1277.

[51] P. Milgrom and C. Shannon. „Monotone Comparative Statics". In: *Econometrica* 62.1 (1994), pp. 157–80.

[52] K. Murota. *Discrete Convex Analysis*. SIAM, 2003.

[53] Katta G. Murty. „On the number of solutions to the complementarity problem and spanning properties of complementary cones". In: *Linear Algebra and its Applications* 5.1 (1972), pp. 65 –108. ISSN: 0024-3795.

[54] J. Nash. „Non-cooperative games". PhD thesis. Princeton, 1950.

[55] George L. Nemhauser and Laurence A. Wolsey. *Integer and Combinatorial Optimization*. New York, NY, USA: Wiley-Interscience, 1988. ISBN: 0-471-82819-X.

[56] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani. *Algorithmic Game Theory*. Cambridge, UK: Cambridge University Press, 2007.

[57] T. Nishizeki and N. Chiba. *Planar Graphs: Theorems and Algorithms*. North-Holland, 1988.

[58] W. Novshek. „On the existence of Cournot equilibrium". In: *Rev. Econ. Stud.* 52.1 (1985), pp. 85–98.

[59] A. Orda, R. Rom, and N. Shimkin. „Competitive Routing in Multi-User Communication Networks". In: *IEEE/ACM Trans. Networking* 1 (1993), pp. 510–521.

[60] A. Pigou. *The Economics of Welfare*. London, UK: Macmillan, 1920.

[61] .S. Pym and H. Perfect. „Submodular function and independence structures". In: *J. of Mathematical Analysis and Applications* 30.1 (1970), pp. 1–31.

[62] O. Richman and N. Shimkin. „Topological Uniqueness of the Nash Equilibrium for Selfish Routing with Atomic Users". In: *Math. Oper. Res.* 32.1 (2007), pp. 215–232.

[63] J. Rosen. „Existence and Uniqueness of equilibrium points in concave n-player games". In: *Econometrica* 33.3 (1965), pp. 520–534.

[64] R. Rosenthal. „A class of games possessing pure-strategy Nash equilibria". In: *Internat. J. Game Theory* 2.1 (1973), pp. 65–67.

[65] Tim Roughgarden and Florian Schoppmann. „Local smoothness and the price of anarchy in splittable congestion games". In: *J. Econom. Theory* 156 (2015). Computer Science and Economic Theory, pp. 317 –342. ISSN: 0022-0531.

[66] David Schmeidler. „Equilibrium points of nonatomic games". In: *Journal of Statistical Physics* 7.4 (1973), pp. 295–300.

[67] Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*. Vol. 24. Springer, 2003.

[68] A. Skopalik and B. Vöcking. „Inapproximability of pure Nash equilibria". In: *Proc. 40th Annual ACM Sypos. Theory Comput.* 2008, pp. 355–364.

[69] Michael J. Todd. „Computation, Multiplicity, and Comparative Statics of Cournot Equilibria in Integers". In: *Mathematics of Operations Research* 41.3 (2016), pp. 1125–1134.

[70] D. Topkis. „Equilibrium points in nonzero n-person submodular games". In: *SIAM J. Control Optim.* 17 (1979), pp. 773–787.

[71] D. Topkis. *Supermodularity and Complementarity*. Princeton, NJ, USA: Princeton University Press, 1998.

[72] L. Tran-Thanh, M. Polukarov, A. Chapman, A. Rogers, and N. Jennings. „On the Existence of Pure Strategy Nash Equilibria in Integer-Splittable Weighted Congestion Games". In: *SAGT 2011*. Ed. by G. Persiano. 2011, pp. 236–253.

[73] X. Vives. „Nash equilibrium with strategic complementarities". In: *J. Math. Econom.* 19.3 (1990), pp. 305–321.

[74] Xavier Vives. „Games with strategic complementarities: New applications to industrial organization". In: *Int. J. Ind. Organ.* 23.7-8 (2005), pp. 625–637.

[75] J. Wardrop. „Some theoretical aspects of road traffic research". In: *Proc. Inst. Civil Engineers* 1.Part II (1952), pp. 325–378.

[76] D.J.A. Welsh. *Matroid Theory*. Dover books on mathematics. Dover Publications, 2010. ISBN: 9780486474397.

[77] Y. Ye. „A Further Result on the Potential Reduction Algorithm for the P-Matrix Linear Complementarity Problem". PhD thesis. Dept. of Management Sciences, Univ. of Iowa, Iowa City, 1988.

[78] Yinyu Ye and Panos M. Pardalos. „A class of linear complementarity problems solvable in polynomial time". In: *Linear Algebra and its Applications* 152 (1991), pp. 3 –17. ISSN: 0024-3795.

NEDERLANDSE SAMENVATTING

Speltheorie is een gebied in de wiskunde dat de strategische interactie tussen verschillende beslissers (of *spelers*) analyseert en probeert te voorspellen. In dit proefschrift bestuderen we strategische spellen, waarin een beperkte hoeveelheid goederen moet worden verdeeld onder de verschillende deelnemers van het spel. Zulke spellen zijn er in vele varianten te vinden, en in deze thesis bekijken we twee soorten spellen in het bijzonder: atomisch deelbare congestiespellen en multimarkt Cournot competities. In deze spellen zijn er een aantal essentiële elementen: een verzameling van spelers, een verzameling van mogelijke strategieën voor iedere speler, en een functie voor iedere speler die een waarde toekent aan de gespeelde combinatie van strategieën. We leggen we de atomisch deelbare congestiespellen en Cournot competities iets preciezer uit.

In een atomisch deelbaar congestiespel is er een verzameling spelers en een verzameling hulpbronnen. Daarnaast heeft iedere speler een 'vraag' naar de hulpbronnen, en kan ze deze vraag verdelen over verschillende toegestane combinaties van hulpbronnen. De prijs die moet worden betaald voor het gebruiken van een hulpbron hangt af van de totale vraag naar deze specifieke hulpbron. Het doel van iedere speler is om egoïstisch haar eigen kosten te minimaliseren. Een voorbeeld van een atomisch deelbaar congestiespel vinden we in het olie transport. Stel dat de overheid een netwerk van oliepijpen heeft liggen, en verschillende bedrijven willen olie transporteren door dit netwerk. Dan zijn er mogelijk meerdere routes waardoor de olie getransporteerd kan worden. Een bedrijf kan er voor kiezen om al haar olie door dezelfde route te transporteren, of dit te verdelen over verschillende routes. Om een pijp te mogen gebruiken moet er een vergoeding worden betaald aan de overheid, en die vergoeding hangt af van de totale hoeveelheid olie die alle bedrijven samen door dit netwerk willen transporteren. Des te groter de totale vraag, des te hoger de prijs. Andere toepassingen van atomisch deelbare congestiespellen zijn te vinden in het verkeer, het overbrengen van data over het internet en het opslaan van data in datacentra.

In een Cournot competitie spreken we over een verzameling bedrijven die allemaal een vergelijkbaar product willen verkopen op verschillende markten. Daarnaast kan het zijn dat niet ieder bedrijf toegang heeft tot iedere markt. De prijs die bedrijven voor hun product in een bepaalde markt kunnen vragen hangt af van het totale aanbod op die markt. Daarnaast moet ieder bedrijf rekening houden met de kosten die ze maken voor het produceren van hun product. Het doel van ieder bedrijf is om haar eigen winst te maximaliseren. In dit geval kan je denken aan de 77 verschillende cola merken. Niet alle 77 merken zijn beschikbaar in ieder land. Zo wordt Amrat Cola wel verkocht in Pakistan, maar niet in Nederland. In een strategie bepaalt een bedrijf hoeveel cola ze in totaal moet produceren en hoeveel cola er in welke landen verkocht moet worden.

Een belangrijk concept binnen de speltheorie is het *Nash-evenwicht*. Een Nash-evenwicht is een collectie van strategieën waarin geen enkele speler kan profiteren van een wijziging in haar eigen strategie, wanneer ze aanneemt dat alle andere spelers hun gekozen strategie behouden. Zo'n evenwicht wordt ook wel een stabiele oplossing genoemd. Merk op dat een Nash-evenwicht meestal hogere kosten met zich meebrengt dan een strategie profiel dat is bepaald door een centrale autoriteit. In multimarkt Cournot competities spreken we van een Cournot-Nash evenwicht. Dit betekent dat een evenwicht wordt gegeven in de optimale verkoopcijfers voor ieder land.

In dit proefschrift vind u drie artikelen die Nash-evenwichten analyseren in atomisch deelbare congestiespellen en multimarkt Cournot competities. We focussen ons op twee aspecten: het vinden van een Nash-evenwicht en bepalen aan welke criteria een spel moet voldoen zodat het Nash-evenwicht uniek is.

HOOFDSTUK 2 We bestuderen voor welke verzamelingen van strategieën een Nash-evenwicht gegarandeerd uniek is. Het uniek zijn van evenwichten is een fundamentele eigenschap van een strategisch spel. Deze eigenschap maakt het namelijk mogelijk om van te voren de uitkomst van het spel te voorspellen. Wanneer er meerdere evenwichten mogelijk zijn, is het niet duidelijk welk evenwicht zal worden geselecteerd door de spelers.

Het belangrijkste resultaat dat we in dit hoofdstuk bewijzen is gebaseerd op de theorie over polymatroïden. We introduceren een speciale klasse van polymatroïden en noemen die *tweezijdige stroom-polymatroïden*. Het Nash-evenwicht van een spel is uniek wanneer voor iedere speler geldt dat haar mogelijke strategieën samen een basis vormen van een tweezijdige stroom-polymatroïde. Daarnaast bewijzen we dat de klasse van tweezijdige stroom-polymatroïden de bekende klasse van basisorderbare matroïden bevat, welke regelmatig voorkomen in de praktijk.

We vullen onze resultaten aan met de theorie dat matroïden noodzakelijk zijn om te garanderen dat het Nash-evenwicht uniek is. In andere woorden, voor elk atomisch deelbaar congestiespel met ten minste drie spelers geldt dat als er een speler bestaat waarvoor de strategieruimte niet gelijk is aan de basis van een matroïde, er een isomorfisme bestaat naar een een spel dat meerdere evenwichten heeft. Er is nog niets bekend over spellen waar de staretegie ruimtes van de spelers wel matroïden zijn, maar geen basisorderbare matroïden.

Daarnaast bestuderen we ook voor welke kostenfuncties van de hulpbronnen we kunnen garanderen dat er een uniek Nash-evenwicht is. Uit eerder onderzoek was bekend dat wanneer alle kostenfuncties polynomen zijn met een graad van maximaal drie, de evenwichten uniek zijn. We generaliseren dit bewijs en introduceren een grotere klasse van kostenfuncties waarvoor hetzelfde resultaat geldt.

HOOFDSTUK 3 In dit hoofdstuk ontwikkelen we het eerste algoritme dat Nash-evenwichten construeert binnen polynomiale tijd, voor atomisch deelbare congestiespellen met affiene kostenfuncties en wanneer de strategieën bestaan uit slechts één element. Ons algoritme is puur combinatorisch en construeert een exact evenwicht wanneer de invoer rationeel is. Het idee is om een puur Nash-evenwicht te vinden voor een geassocieerde geheeltallig-deelbare polymatroïde, waar spelers hun vraag alleen kunnen verdelen in geheeltallige veelvouden van een gemeenschappelijke pakketgrootte. Hoewel geheeltallig-deelbare congestiespellen al eerder zijn bestudeerd, is er nog geen algoritme bekend dat Nash-evenwichten in polynomiale tijd construeert. Ook voor deze klasse ontwikkelen we het eerste algoritme dat werkt in polynomiale tijd en gebruiken deze als bouwsteen voor ons belangrijkste resultaat van dit hoofdstuk.

Daarnaast ontwikkelen we ook een transformatie, die werkt binnen polynomiale tijd, van een multimarkt Cournot competitie met bedrijfs-specifieke affiene prijsfuncties en kwadratische productiekosten naar een atomisch deelbaar congestiespel. Deze transformatie behoudt de evenwichten, en daarom kunnen we via ons algoritme voor atomisch deelbare congestiespellen ook voor deze spellen de Nash-evenwichten vinden binnen polynomiale tijd. Tenslotte volgen uit onze analyse ook nieuwe grenzen op het verschil tussen atomisch deelbare en geheeltallig-deelbare Cournot-evenwichten. Deze nieuwe grenzen kunnen worden gezien als een generalisatie van de recente grenzen voor Cournot-competities op een enkele markt door Todd [69].

HOOFDSTUK 4 We construeren we ϵ-benaderbare Nash-evenwichten voor atomisch deelbare congestiespellen met convexe kostenfuncties, en strategieruimtes die gelijk zijn aan de basis van een polymatroïde. Net als in Hoofdstuk 3, pakken we dit aan door een puur Nash-evenwicht te vinden voor een geassocieerde geheeltallig-deelbare polymatroïde. We weten dat een Nash-evenwicht voor een geheeltallig-deelbaar congestiespel geconstrueerd kan worden binnen pseudo-polynomiale tijd. In dit hoofdstuk beslissen we voor iedere $\epsilon > 0$ over een pakketgrootte k, zodat het geassocieerde k-deelbare evenwicht een ϵ-benaderbaar evenwicht is voor het originele atomisch deelbare congestiespel.

VALORIZATION

All knowledge is connected to all other knowledge. The fun is in making the connections.

— Arthur Aufderheide

For research to be valuable, one should be able to explain its utilisation and its impact on society. Thus, in this chapter, we discuss the economic and social relevance of this thesis and identify groups for which, in addition to the academic community, these results are of interest. We start with the use of game theory in general, and its shortcomings. Then, we discuss how the field of Algorithmic Game Theory (AGT) arose from these shortcomings, and we discuss some real-life examples in which the AGT-community made a difference. This thesis focusses on the uniqueness and computation of Nash equilibria in resource allocation games, and hence, we also discuss the relevance of these specific type of games and how the results obtained in this thesis can be used in practice. Lastly, we discuss the relevance of this research for the field of Transportation Science, and the steps that are made to collaborate with other communities in order to face the challenges that are arising in network design and autonomous vehicles.

Game theory, also referred to as decision theory, models strategic interaction of multiple rational decision makers. In contrast to the field of classical optimization, game theorists assume that there is no central authority that is making these decisions, but instead, players behave selfishly and aim to maximize their private utility. Applications of this field include a large number of economic and political phenomena and approaches, such as auctions, voting systems, fair division, duopolies, oligopolies and social network formation. Even in biology, game theory has been used as a model to understand many different phenomena, like the stable approximate 1:1 sex ration in most species, animal communication, fighting behavior and territoriality.

Research in game theory usually focusses on equilibria or 'solution concepts', and the most famous of these is the Nash equilibrium. A set of strategies is called a Nash equilibrium when no player can unilaterally deviate from her current strategy and increase her utility. Though game theorists are particularly interested in these Nash equilibria, they do not consider the complexity issues that might arise when one actually wants to compute them, which would be one of the main concerns of any computer scientist. Algorithmic game theory arose due to this discrepancy between game theory and computer science.

Nisan et al. write in their book *Algorithmic Game Theory* [56] that "if an equilibrium concept is not efficiently computable, much of its credibility as a prediction of the behavior of rational agents is lost". We should be able to simulate each decision maker by a machine and as Kamal Jain said: "If your laptop cannot find the equilibrium, neither can the market". Therefore, one of the core activities of the AGT-community is to find efficient algorithms that compute Nash equilibria. Besides computing them, this community also studies the existence of equilibria, the uniqueness of equilibria and the complexity of computing them. Lastly, another popular topic is the convergence of best responses. Given an arbitrary state of the game, people might be able to increase their utility by unilaterally changing their strategy. The best improvement that can be made is called a *best response*. A change of strategy might cause a reaction by other players, and we obtain a sequence of best responses. Does such a sequence converge to a steady state (a Nash equilibrium)? Or will it cycle?

Results on the existence and uniqueness of equilibria are currently used in network design. For example, it is seen in practice that adding a road to an existing network might cause an increased overall journey time, a phenomenon called the Braess paradox. In Seoul, South Korea, congestion was reduced when a motorway was removed as part of a restoration project and in 1990, closing 42nd street in New York improved traffic in that area. When we are able to efficiently compute equilibria, one can identify such roads and close them to reduce the overall journey times. Furthermore, if one knows that best responses will converge to a Nash equilibrium, any set of arbitrary paths chosen by the cars will after some time converge to a steady traffic flow. Thus, applications of being able to efficiently compute equilibria can be found, among others, in transportation science and network design.

In this thesis, I study the uniqueness and complexity of computing equilibria in resource allocation games. These games play a key role in a wide range of applications including traffic networks and telecommunication networks. In particular, we looked at atomic splittable congestion games and multimarket oligopolies. The practical application for multimarket oligopolies seems clear. Given a set of companies and markets, a strategy of a company is to decide how many goods it wants to sell in each market. Atomic splittable congestion games also occur often, but they might be a bit harder to spot. Besides the obvious applications in traffic, we find one in queueing theory: assume we are given a set of M/M/1 queues served in a first-come-first-serve fashion, and a finite set of companies, each sending packets to the queues with some company-specific arrival rate. Each queue has a single server with an exponentially distributed service time. Note that when more packets are assigned to a queue, the average sojourn time will increase. Each company needs to find a fractional distribution of her packets over the queues that minimizes the average sojourn time of her packets.

In Chapter 2 we study the uniqueness of Nash equilibria in strategic games. This property is key to actually predict the outcome of distributed resource allocation: if there are multiple equilibria, it is not clear upfront which equilibrium will be selected by the players. This issue has been raised explicitly by Aumann [5]: "...it is by no means clear how the players would arrive at an equilibrium, why they should play equilibrium strategies, and how a specific equilibrium would be chosen from among the set of all equilibria". In this chapter, we investigate strategy spaces that have the property that Nash equilibria are unique, no matter how the strategy spaces of the different players are interweaved. As a result, we introduced the class of bidirectional flow polymatroids. Such a result can be useful in the design of new games. For example: assume that given a network, each player needs to connect all her nodes in a connected subgraph of this network, e.g., players need to divide their demand over spanning trees of a connected network. Then, whenever this network is known to be generalized series-parallel, equilibria are known to be unique.

In Chapter 3 and Chapter 4 we study the computation of Nash equilibria on networks with only parallel edges or, more generally, strategy spaces where each strategy consists of a single resource. Note that it is already known that equilibria are unique in this setting. The main idea is to construct an equilibrium in a corresponding integral splittable congestion game. This can be done within polynomial time in case cost functions are affine, and within pseudo-polynomial time whenever the cost functions are increasing, non-negative and convex. Furthermore, we showed that multimarket oligopolies with affine price functions and quadratic costs can be modelled as atomic splittable congestion games, and hence, also compute Nash equilibria for this setting. The impact of these results is that whenever a game can be modelled as an atomic splittable congestion game, players are able to compute the unique Nash equilibrium of this game and play accordingly.

In order to make a real change, collaboration with companies as TomTom, Google Maps and the Transportation Science community is necessary. Though the Algorithmic Game Theory community can find exact equilibria in many situations, the models that are used are merely a simplification of reality. Through collaboration and discussions, we should settle on new basic assumptions that make the current models more realistic. Currently, one of the most interesting collaborations in this direction is the Dagstuhl seminar on Dynamic Traffic Flow models in Transportation Science. So far, this seminar has been organized twice, first in Oktober 2016 and a second time in March 2018. This seminar brings together researchers from three different communities: Simulations (SIM), Dynamic Traffic Assignment (DTA) and Algorithmic Game Theory (AGT). Among other points, the seminar initiated a systematic study of the complexity of equilibrium computation for DTA models – which is the core task when resolving dynamic traffic assignment problems. As equilibrium computation and its complexity are one of the core topics of this thesis, I gave a talk at both seminars and in this way contributed to the valuable discussions.

CURRICULUM VITAE

Veerle Timmermans was born in Nijmegen, the Netherlands on August 17, 1990. In 2008, she received her Gymnasium diploma cum laude from the MaasWaal College in Wijchen. In September of the same year she started studying Mathematics at the Radboud University in Nijmegen. During this time, she spend a month in Kampala, Uganda, showing local teachers how to use computers. Under the supervision of prof. dr. Wieb Bosma and dr. Tjark Vredeveld she received her Master's degree cum laude in July 2014.

From July 2014 until June 2018 she was a PhD student at Maastricht University, under the supervision of prof. dr. Tobias Harks and prof. dr. Stan van Hoesel. During this time, she mainly studied the uniqueness and computation of pure Nash equilibria in resource allocation games. Parts of the results of her research are presented in this thesis. Veerle presented her work at various international conferences and published multiple results in this thesis in international refereed academic journals.

In July 2018, she will start a post-doctoral fellowship at the RWTH Aachen.

This thesis considers resource allocation games and studies whether or not Nash equilibria are unique and can be computed efficiently.

Chapter 1 presents an introduction to algorithmic game theory and explains why equilibrium computation is one of its core activities. Then, it discusses resource allocation games, which contain the classes of atomic splittable congestion games and multimarket oligopolies.

Chapter 2 studies sufficient and necessary properties for the uniqueness of Nash equilibria in atomic splittable congestion games.

Chapter 3 is dedicated to computing exact equilibria for atomic splittable congestion games with singleton strategies and affine cost functions. The results of this chapter are then translated to multimarket oligopolies.

Chapter 4 complements the results of Chapter 3 as it considers the computation of approximate equilibria in atomic splittable congestion games with convex cost functions and multimarket oligopolies with concave price functions.

The thesis ends with a Dutch summary, a valorisation and academic curriculum vitae.