Tolls For Atomic Congestion Games

by

Alexander Stoll

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Master of Mathematics
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2021

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

In games where selfish players compete for resources, they often arrive at equilibria that are less desirable than the social optimum. To combat this inefficiency, it is common for some central authority to place tolls on the resources in order to guide these players to a more advantageous result. In this thesis, we consider the question of how to add tolls to atomic unsplittable congestion games in order to enforce a specific flow as the unique equilibrium. We consider this question in the context of both routing games and matroid congestion games. In the former case, we show that for the class of series-parallel graphs the nonatomic tolls suffice, and investigate examples for which nonatomic tolls fail. In the latter case, we show that the nonatomic tolls can also be used to impose flows in atomic laminar matroid games.

Acknowledgements

An enormous thank you to my advisor, Chaitanya Swamy. Without your invitation to join the C&O program, my graduate school experience would have been entirely different. Thank you for helping me find a problem to work on that we were both passionate about. I have grown tremendously as a mathematician thanks to you; without your constant feedback I would not have been able to explain my ideas precisely in a way that others can easily understand. Most importantly, thank you for always believing in me, even when I did not.

To Joseph Cheriyan and Ricardo Fukasawa, thank you for reading my thesis and providing valuable feedback.

To Melissa, thank you for answering every administrative question I had, and for always being available, especially when timelines were tight.

To my friends Andrew, Sabrina, Matt, and Tina, thank you for making me feel welcome in your offices, and for making my graduate experience both on and off campus a memorable one. I am blessed to have met you.

And of course, a special thank you to my mom and dad, for supporting me unconditionally every step of the way.

To perserverance.

Table of Contents

List of Figures

Chapter 1

Introduction

Congestion games [18] are a class of games in which players selfishly select a subset of resources and incur cost based on the number of players selecting those resources. Congestion games are a popular topic of study in game theory because they model important real-world problems, such as how to route traffic in a large transportation or communication network that has no central authority.

Congestion games are split into several different models, based on the influence each individual player has on the game. In an *atomic* game, there are finitely many players who each choose a single strategy to route their own traffic. In a *nonatomic* game, there are infinitely many players, each controlling a negligible amount of traffic. In between these two definitions, there also exist *atomic splittable* games, in which the finitely many players are allowed to divide their traffic fractionally among their available strategies.

Congestion games are further classified by the environment in which they are played. A *network routing game* is played on a directed graph $G = (V, E)$, and each player selects an $s \rightsquigarrow t$ path to route his portion of the flow from a sink s to a source t. By doing so, he incurs costs from each edge he selects that depend on the congestion on those edges. In a *matroid congestion game*, instead of choosing a path through a network, each player instead chooses a basis of the matroid $M = (E, \mathcal{I})$. Similarly, he incurs cost based on how many other players have selected the same elements of his basis. We give a fuller description of the different types of congestion games in Chapter 2.

The key problem that arises in almost every game is that left to their own devices, even rational players will often arrive at an outcome that is inferior to the social optimum. For a simple example of this, consider the following nonatomic routing game, first considered by Pigou [17]:

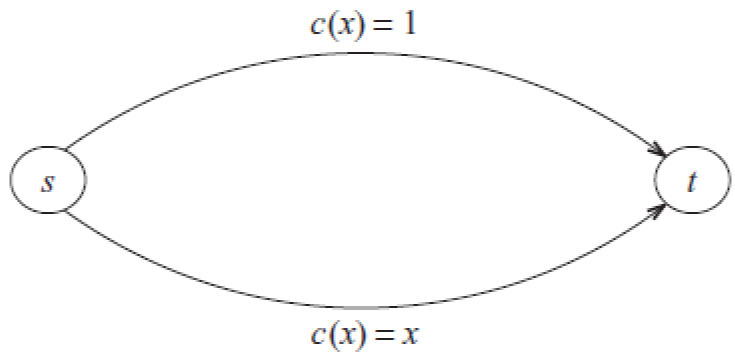

Figure 1.1: Pigou's example.

In this graph, the cost function $c(x)$ for an edge depends on the amount of flow x routed through that edge. Suppose infinitely many players each control a very small portion of one total unit of flow. Then the bottom route is the dominant strategy, since it will never cost more than the top route, even if every player chooses it. Notice that if any player chooses the top route, then the bottom route will be less expensive. Therefore, the situation in which every player chooses the bottom route (and so incurs a cost of 1) is an equilibrium.

However, this solution is not socially optimal. The social optimum of the game is to send half of the players through the top route, and the other half through the bottom. In this case, half of the players incur a cost of 1, and the others incur a cost of 1/2. On average, these players' costs are 3/4 as expensive compared to before.

The ratio of the total cost between the equilibrium induced by selfish players and the cost of the social optimum is known as the *price of anarchy* (PoA), first introduced by Papadimitriou [16]. Given that the inefficiency of these equilibria is not desirable, it makes sense to introduce a central authority or network manager whose job is to impose extra constraints or incentive to push players towards a socially optimal equilibrium.

Perhaps the most natural way to influence players is through the use of tolls, in which the network manager places a constant toll τ_e on each resource e. In the previous example, adding a toll of 1/2 to the bottom edge induces the social optimum as the equilibrium. If more than half of the users selected the bottom edge, then its cost would be greater than 1, giving incentive to use the top edge instead.

In general, the goal is to find tolls to place on each resource so that the resulting equilibrium's cost (including the tolls) is relatively close to the cost of the social optimum. Tolls are extremely well studied: they were first proposed by Pigou [17] in 1920, and since then there have been a multitude of results showing how to use tolls in congestion games. The nonatomic case is by far the most understood; under mild conditions, it is possible to compute tolls that induce the optimal flow as the equilibrium [11].

However, much less is understood about tolls in the atomic setting. One hurdle that bars progress in this setting is that while nonatomic equilibria are unique, atomic equilibria are not [15]. Therefore, one would need to use tolls to not only induce some flow as an atomic equilibrium, but the unique one in order to show a positive result mirroring what is known in the nonatomic setting. In addition, it is important to be careful in how we define equilibria; the same flow can often arise from different sets of player strategies. Of these, it is possible that some are equilibria while others are not. To combat this complication, we define the notion of a "profile-decomposition independent" equilibrium, or a flow that maintains the properties of an equilibrium regardless of how it is decomposed into player strategies. We elaborate on these complications and the resulting definition in Section 2.5.

1.1 Our Results and Organization of the Thesis

We show that we can compute tolls to enforce a flow f as the unique profile-decomposition independent (PDI) atomic equilibrium in a variety of settings.

In chapter 2, we discuss some preliminaries, and precisely define what it means for an equilibrium to be profile-decomposition independent. This gives the terminology and notation to state the central problem of this thesis precisely.

In chapter 3, we show that in series-parallel graphs the tolls that induce f as a nonatomic equilibrium also suffice to impose f as the unique PDI atomic equilibrium (Theorem 3.1). We then give an example that shows that the nonatomic tolls do not suffice in general, but also show how to modify these tolls a small amount to create valid atomic tolls (Proposition 3.3). However, we then show that while a small perturbation of the nonatomic tolls does not always suffice (Proposition 3.5), we can still compute atomic tolls that make any f enforceable in our example, given stricter constraints on the cost functions.

In chapter 4, we develop the relationship between polymatroids and matroid congestion games. We show in Theorem 4.7 that if the polymatroid associated with a matroid congestion game exhibits the properties in Conjecture 4.6, then the nonatomic tolls will also suffice in the atomic setting. We then show that Conjecture 4.6 holds for both partition (Theorem 4.9) and laminar (Theorem 4.14) matroid congestion games.

In chapter 5, we conclude the thesis by surveying the remaining classes of graphs and matroids for which our techniques are currently insufficient. In both cases, we discuss what parts of the structure of our graph and matroid classes were most essential in allowing us to develop our tools, and consider what aspects of the structure are simplest to remove first for the purposes of generalizing our results.

1.2 Related Work

The idea of a nonatomic congestion game was first created by Wardrop [24] as a model for road traffic. The first price of anarchy results for nonatomic routing games were found by Roughgarden and Tardos [21], who gave tight bounds for games with linear latency functions. Roughgarden [19, 20] then expanded these results for many other classes of latency functions.

Tolls were first considered in the literature by Pigou [17]. One issue with adding tolls to a network is that different users may have a different trade-off between time and money (in particular, let each player a have a function $\alpha(a)$ that converts latency into money), and so the same toll may be experienced differently by different players. Beckman, McGuire, and Winsten [2] showed that adding marginal cost tolls to a Nash flow yields the social optimum when players' α functions are identical.. Dafermos [9] and Smith [22] provided early results on how to compute tolls when each player's α value is known, and Cole, Dodis, and Roughgarden [6] found tolls for single commodity networks. As mentioned previously, Fleisher, Jain and Mahdian [11] generalize this work and give a complete characterization of flows that are enforceable by tolls in the nonatomic setting.

In atomic splittable routing, various bounds have been developed on the price of anarchy [3, 7, 12]. Work has also been done to relate the cost of Nash equilibria in atomic splittable games to their nonatomic counterparts [3, 7, 13]. In addition, Yang and Zhang [25] have shown that optimal tolls exist in the atomic splittable setting, provided the users are homogeneous with the same α value. These results were generalized by Swamy [23], who gave a complete characterization of flows that can be induced as a Nash equilibrium via tolls.

For atomic unsplittable routing games, some PoA bounds are given by Awerbuch, Azar and Epstein [1], and Christodoulou and Koutsoupias [5], who show that these bounds can be worse than they are in the other two settings. Caragiannis et al. [4] have shown in their study of parallel-link graphs with linear cost functions that optimal tolls do not always exist in atomic unsplittable games, and gave various ranges on the PoA that is achievable using tolls.

Chapter 2

Preliminaries

2.1 Nonatomic Congestion Games

A *nonatomic congestion game* consists of a tuple $\mathcal{N} = (K, E, (\mathcal{S}_i)_{i \in K}, (c_e)_{e \in E}, (d_i)_{i \in K})$, where $K = \{1, \ldots, k\}$ is a nonempty, finite set of populations, and $E = \{e_1, \ldots, e_m\}$ is a nonempty, finite set of resources. Players are infinitesimally small, but each population i together is responsible for d_i total *volume* or *flow*. Each player in population $i \in K$ chooses a strategy P from the strategy set $\mathcal{S}_i$, which is a nonempty, finite collection of subsets of E.

We keep track of the aggregate choices of the players by defining $f_P \geq 0$ to be the total volume of players using strategy P. Thus, the f_P that encode the choices of population i must satisfy $\sum_{P \in \mathcal{S}_i} f_P = d_i$, and the choices of the entire population must satisfy this constraint for all $i \in K$. We will employ a common abuse of notation here, and for convenience use f_e to denote the total volume of players choosing resource e as part of their strategy; that is, $f_e = \sum_{i \in K} \sum_{P \in \mathcal{S}_i : e \in P} f_P$.

We refer to f_e as the amount of congestion on e, and call $f = (f_e)_{e \in E}$ as the *congestion vector*. As is clear from the definition, the congestion vector depends on the aggregate strategy profile $(f_P)_{P \in \bigcup_{i \in K} \mathcal{S}_i}$ of the players.

Each element $e \in E$ has a corresponding cost function $c_e(f_e)$, where the cost of a resource depends on its congestion. Nonatomic congestion games involving various conditions on the c_e's have been studied in the literature. For our purposes, we will assume that each c_e is strictly increasing.

A strategy profile $f = (f_P)_{P \in \bigcup \mathcal{S}_i}$ of a nonatomic congestion game is said to be a *Nash equilibrium* if the following holds:

$$\forall i \in K, \ \ \forall P \in \mathcal{S}_i \text{ s.t. } f_P > 0, \ \ \forall Q \in \mathcal{S}_i, \quad \sum_{e \in P} c_e(f_e) \leq \sum_{e \in Q} c_e(f_e).$$

In words, every strategy chosen is a minimum cost strategy under the congestion vector $f = (f_e)_{e \in E}$. Another well known definition of Nash equilibria is contained the the following lemma, used early on by Dafermos and Sparrow [10]:

Lemma 2.1. *Let $\mathcal{N} = (K, E, (\mathcal{S}_i)_{i \in K}, (c_e)_{e \in E}, (d_i)_{i \in K})$ be a nonatomic congestion game with nondecreasing and continuous edge costs. Then f is a Nash equilibrium iff f is a min-cost flow under the $\{c_e(f_e)\}$ edge costs, i.e. $\sum_{e \in E} c_e(f_e) f_e \leq \sum_{e \in E} c_e(f_e) g_e$, for every feasible congestion vector g.*

Proof. By definition, f is a nonatomic Nash equilibrium iff

$$\forall i \in K, \ \ \forall P \in \mathcal{S}_i \text{ s.t. } f_P > 0, \ \ \forall Q \in \mathcal{S}_i, \quad \sum_{e \in P} c_e(f_e) \leq \sum_{e \in Q} c_e(f_e).$$

Since every strategy in a nonatomic equilibrium is a minimum cost strategy, it must be that

$$\sum_{P \in \bigcup \mathcal{S}_i} c_P(f) f_P \leq \sum_{P \in \bigcup \mathcal{S}_i} c_P(f) g_P.$$

for every feasible flow g. Then, since $c_P(f) = \sum_{e \in P} c_e(f_e)$ (the cost of a strategy is the sum of the costs of the elements chosen by that strategy), we can rewrite this inequality as

$$\sum_{P \in \bigcup \mathcal{S}_i} [f_P \sum_{e \in P} c_e(f_e)] \leq \sum_{P \in \bigcup \mathcal{S}_i} [g_P \sum_{e \in Q} c_e(f_e)].$$

Reversing the order of summation, we obtain

$$\sum_{e \in E} c_e(f_e) f_e \leq \sum_{e \in E} c_e(f_e) g_e$$

as desired. □

This inequality is sometimes called the variational characterization of Nash equilibria. Under mild assumptions, NE exist and can be efficiently computed (within any desired accuracy). Furthermore, if the latency functions are strictly increasing, then there is a unique Nash equilibrium. To see this, suppose f and f' are two distinct NE. Then $\sum_e (f_e - f'_e)c_e(f_e) \le 0$, and also $\sum_e (f'_e - f_e)c_e(f'_e) \le 0$. Adding, we get that $\sum_e (f_e - f'_e))(c_e(f_e) - c_e(f'_e)) \le 0$. But since the edge costs are strictly increasing, each term in the summation is nonnegative, so each term must be 0, and this can only happen if $f_e = f'_e$ for all e.

2.2 Atomic Congestion Games

Atomic congestion games differ from nonatomic congestion games in that our players are no longer infinitesimally small. An *atomic congestion model* is given by $\mathcal{A} = (K, E, (\mathcal{S}_i)_{i\in K}, (c_e)_{e\in E})$. Instead of k populations, $K = \{1, \ldots, k\}$ now denotes k players, each responsible for 1 unit of flow. Each player i selects a single strategy P_i from their corresponding strategy set $\mathcal{S}_i$, and sends his flow along P_i, thereby loading each resource $e \in P_i$. The congestion on resource e is simply $f_e = |\{i : e \in P_i\}|$, the number of players selecting that resource. The cost of each resource e is $c_e(f_e)$ as before, and so the total cost for player i is $\sum_{e\in P_i} c_e(f_e)$, the sum of the costs for each resource used by the player. The chief difference between an atomic and nonatomic game is that each player in an atomic game has a noticeable effect on the congestion of the resources he uses. This correspondingly changes the definition of a Nash equilibrium.

In an atomic congestion game, a strategy profile $\{P_i\}_{i\in K}$ is a pure Nash Equilibrium if no single player is able to improve his cost by switching to some other strategy. More formally,

$$\forall i \in K, \;\; \forall Q_i \in \mathcal{S}_i, \quad \sum_{e\in P_i} c_e(f_e) \le \sum_{e\in P_i\cap Q_i} c_e(f_e) + \sum_{e\in Q_i\setminus P_i} c_e(f_e+1).$$

Here, the "+1" term comes from the fact that when player i deviates to strategy Q_i, he increases the congestion of the resources in $Q_i \setminus P_i$ by 1. [1] The (pure) Nash equilibria of atomic congestion games are much less understood, and enjoy fewer nice properties compared to the equilibria of nonatomic congestion games. It is known that they always

[1]We have described what are called unweighted atomic congestion games, wherein all players control the same amount of flow. More generally, player i may control d_i units of flow, and increases the congestion of all resources in his chosen strategy by d_i. In this more general setting, pure Nash equilibria need not even exist [15], so we do not discuss this model.

exist, via a potential function argument (see Section 2.4), and can be computed efficiently in certain settings (e.g. in an atomic routing game, where each player's strategy set is the collection of $s \rightsquigarrow t$ paths). However, unlike the nonatomic setting, equilibria need not be unique. Moreover, an additional complication is that a congestion vector $(f_e)_{e \in E}$ does not convey all the information about the players' strategies, and there could be two different strategy profiles mapping to the same congestion vector, where one them is an NE and the other is not (see Section 2.5).

2.3 Tolls

One common issue that arises in congestion games with self-interested players is that equilibria can be inefficient, meaning that an equilibrium can be much more costly than an optimal congestion vector. One measure that has been studied in the literature to mitigate this inefficiency is the use of tolls to modify the cost functions in a controlled way so as to impose the optimal congestion vector (or some other desired congestion vector) as the equilibrium.

Let $\tau \in \mathbb{R}_+^E$ be the toll vector for a congestion game. These tolls add an additional constant cost to each resource $e \in E$, and so the cost of each resource e is now $c_e(f_e) + \tau_e$.

The use of tolls to enforce a desired congestion vector in the nonatomic setting is very well understood. In particular, Fleisher, Jain and Mahdian [11] show that a congestion vector f is enforceable as a nonatomic equilibrium iff f is minimal, meaning that if the same game admits another congestion vector f' such that $f'_e \leq f_e$, for all $e \in E$, then $f' = f$.

Theorem 2.2. *A feasible congestion vector f is enforceable as a nonatomic equilibrium via tolls iff f is minimal.*

Here, feasible simply means that the congestion vector f arises from a strategy profile of the game. This result has immediate consequences in the atomic setting, since a nonatomic equilibrium is also an atomic equilibrium. However, as noted earlier, atomic equilibria are not unique. This motivates the central question of this thesis: for a given congestion vector f, do there exist tolls that enforce f as the unique atomic equilibrium? We will state this question more precisely in Section 2.5.

2.4 Potential Functions

A useful tool for analyzing congestion games is a *potential function*, which is a single global function that expresses the incentive of the game's players to change their strategies. A potential function ϕ of a game is a real-valued function defined on the strategy profiles. It has the property that it "tracks" the change in cost of a deviating player for any single player's deviation, and therefore pure Nash equilibria correspond to local minimizers of ϕ.

The potential function for an atomic congestion game is the following [15]:

$$\phi(f) = \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i).$$

The potential function ϕ is usually viewed as a function mapping congestion vectors to reals. It has the property that if single player's deviation from some starting congestion vector f results in the congestion vector f', then $\phi(f') - \phi(f)$ is precisely the change in the deviating player's total cost.

Lemma 2.3. *Let $\mathcal{A} = (K, E, (\mathcal{S}_i)_{i \in K}, (c_e)_{e \in E})$ be an atomic congestion game. Let f be the congestion vector created by some strategy profile $\{P_i\}_{i \in K}$. Suppose player i for some $i \in K$ changes his strategy from P to P', resulting in a new congestion vector f'. Then $\phi(f') - \phi(f) = \sum_{e \in P' \setminus P} c_e(f'_e) - \sum_{e \in P \setminus P'} c_e(f_e)$.*

Proof.

$$\begin{aligned}
\phi(f') - \phi(f) &= \sum_{e \in E} \sum_{i=1}^{f'_e} c_e(i) - \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i) \\
&= \sum_{e \in P \setminus P'} \sum_{i=1}^{f'_e} c_e(i) + \sum_{e \in P' \setminus P} \sum_{i=1}^{f'_e} c_e(i) - \sum_{e \in P \setminus P'} \sum_{i=1}^{f_e} c_e(i) - \sum_{e \in P' \setminus P} \sum_{i=1}^{f_e} c_e(i) \\
&= \sum_{e \in P' \setminus P} c_e(f'_e) - \sum_{e \in P \setminus P'} c_e(f_e)
\end{aligned}$$

The second equality follows from the fact that $f_e = f'_e$ for all $e \in P \cap P'$ and $e \notin P \cup P'$, and the final equality follows from the fact that $f'_e = f_e + 1$ when $e \in P' \setminus P$, and $f_e = f'_e + 1$ when $e \in P \setminus P'$. $\square$

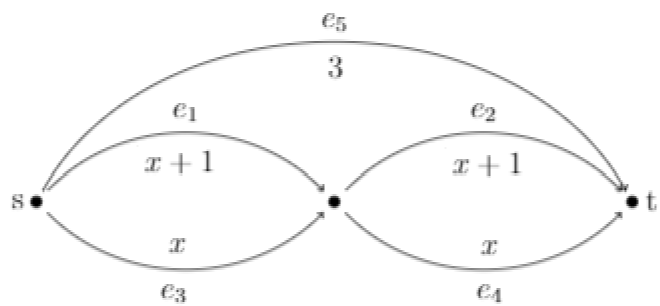

Figure 2.1: A directed graph where flow can be decomposed in multiple ways.

Lemma 2.3 immediately shows that an atomic congestion game always has a pure Nash equilibrium. If f^* is the minimizer of the potential function, and $(P_1, \ldots, P_k)$ is a strategy profile yielding f^* as the congestion vector, then no deviation by any player can lower the potential function, and then by the lemma no deviation can decrease a player's cost.

2.5 Profile-Decomposition Independent Nash Equilibria of Atomic Congestion Games

One of the complications that arises in dealing with equilibria of atomic congestion games is that the congestion vector $f = (f_e)_{e \in E}$ need not map to a unique strategy profile. In particular, there could be two strategy profiles that map to the same congestion vector, where one is a Nash equilibrium, and the other is not. Consider the following example of an atomic routing game – a congestion game arising from a directed graph – where we have 2 players whose strategy is to pick an $s \rightsquigarrow t$ path. In this graph, each edge has an associated cost function $c_e(x)$ that depends on the amount of flow x through that edge.

One Nash Equilibrium for the game occurs when both players opt to avoid the expensive e_5 edge and instead pick crossing routes of total cost 3, where one player selects e_1 and e_4 while the other selects e_2 and e_3. Now, consider instead what would happen if we routed the first player along e_1 and e_2 and the other player through e_3 and e_4. Clearly this is not a Nash equilibrium, as the first player incurs a cost of 4, and has incentive to reroute via e_5. However, both of these strategy profiles correspond to the exact same congestion vector $f = (f_{e_1}, f_{e_2}, f_{e_3}, f_{e_4}, f_{e_5}) = (1, 1, 1, 1, 0)$, even though only one is a Nash Equilibrium. Therefore, even if a strategy profile is a NE, it's possible that the corresponding f is not a local minimizer of ϕ.

In terms of the potential function ϕ (which only depends on the congestion vector) this means that we cannot quite use ϕ to reason about Nash equilibria (other than noting that the global minimizer of ϕ does correspond to a NE). In particular, define the *neighborhood* of a congestion vector f, denoted $NBD(f)$, to be all congestion vectors that arise via a single player's deviation in strategies, where one considers all strategy profiles resulting in the congestion vector f. This definition allows us to define more precisely what it means to be a local minimizer of ϕ. We say that f is a local minimizer of ϕ if it is the lowest potential of all of its neighbors, meaning

$$\phi(f) \leq \phi(f'), \forall f' \in NBD(f).$$

While Lemma 2.3 would seem to indicate that a NE is a local minimum of ϕ, the above example shows that this is not true. In that example, $f = (1,1,1,1,0)$ was an NE, but $f' = (0,0,1,1,1)$ lies in $NBD(f)$ and has smaller ϕ value. To circumvent these issues, and keeping in mind that our goal is to impose a target congestion vector as the unique equilibria, we define the following special type of equilibrium.

Definition 2.4. *f is a Profile-decomposition independent Nash equilibrium (PDI NE) if every strategy profile that admits f as a congestion vector is a Nash equilibrium.*

PDI equilibria are much more mathematically convenient to work with since they allow us to work with the congestion vector and not worry about how it is decomposed as a strategy profile. In particular, we can now use the potential function to identify and reason about such equilibria since PDI equilibria correspond precisely to local minimizers of ϕ.

Lemma 2.5. *f is a PDI Nash Equilibrium for an atomic congestion game iff f is a local minimum of ϕ.*

Proof. By definition f is a PDI Nash Equilibrium for an atomic congestion game with potential function ϕ iff every strategy profile that gives rise to f is a NE. Equivalently, no player in any of these strategy profiles can deviate and decrease his cost (and therefore decrease ϕ), meaning that f is a local minimum of ϕ. □

Notice that our argument from earlier proving that the global minimizer f^* of ϕ corresponds to a NE shows in fact that f^* is a PDI equilibrium. Thus every atomic congestion game admits at least one PDI equilibrium.

Since our goal is to find tolls to impose a given target congestion vector f as the unique atomic equilibrium (modulo player renaming), this implies that f must be a PDI equilibrium, and so the PDI property arises as a natural necessary property in this context. We can now state the question that we study in this thesis precisely; given a target congestion vector f, find tolls so that after imposing these: (i) f is a PDI equilibrium; and (ii) there is no other PDI equilibrium.

While both (i) and (ii) mention f being a PDI NE, it is important to note that if f is the unique atomic equilibrium, then it must also be PDI, since the global minimum of the potential function is always a PDI equilibrium. However, this does not mean that the use of PDI in (ii) is redundant; it is entirely possible to have both f as the only PDI atomic equilibrium and some other non-PDI atomic equilibria that are distinct from f.

In fact, in the earlier example $f' = (0, 0, 1, 1, 1)$ is the only PDI atomic equilibrium, yet $f = (1, 1, 1, 1, 0)$ is an atomic equilibrium. These are the only two atomic equilibria. To see this, observe that if both players choose e_5, then there is incentive to deviate to the bottom route, and if both players select the same e_i for $i = 1, 2, 3, 4$, then at least one player will have a cost of 4 and prefer to deviate to the top route. Therefore, no edge is used twice. If neither player chooses e_5, and the two players choose disjoint routes, then the resulting congestion vector is f. The remaining congestion vectors are those in which one player opts to choose e_5, and the other player chooses a subset of the remaining edges. In all of these cases, if the latter player does not choose exactly e_3 and e_4 (therefore creating the congestion vector f'), he can reduce his cost by doing so. f' only admits one decomposition, so it is necessarily PDI.

2.6 Network Routing Games

In this thesis, we study two structured classes of congestion games, one of which is network routing games. A *network routing game* $(K, G, (c_e)_{e \in E}, (d_i)_{i \in K})$ is a congestion game that arise from a directed graph $G = (V, E)$. In this game, the i^{th} player (or population of players, in the nonatomic setting) is associated with a source-sink vertex pair (s_i, t_i), and its strategy set is the set of all $s_i \rightsquigarrow t_i$ paths. The edges are the resources, and as before, each edge $e \in E$ has an associated strictly increasing cost function $c_e(f_e)$ that depends on the congestion of that edge, and the total cost incurred by a player is simply the sum of the costs of the edges he chooses. In the atomic case, we only consider the case where all players have the same source-sink pair, and control one unit of flow (i.e., $d_i = 1$).

Observe that in the nonatomic setting, a strategy profile $f = f_P$ is a nonatomic equilibrium if all paths chosen by players are shortest paths under the $\{c_e(f_e)\}$ edge costs.

In the atomic setting, a strategy profile is an atomic equilibrium if no player can reduce his own cost by choosing another path. Thus, since all players have the same strategy set (consisting of $s \rightsquigarrow t$ paths) in the atomic setting we consider, this means that in an equilibrium, all players incur the same cost.

2.7 Matroid Congestion Gamess

Let E be a finite set and let $\mathcal{I}$ be a family of subsets of E, called *independent sets*. We say that $M = (E, \mathcal{I})$ is a *matroid* if the following three axioms hold [8]:

1. The empty set is independent: $\emptyset \in \mathcal{I}$.
2. Subsets of independent sets are independent: if $S \subseteq T \in \mathcal{I}$, then $S \in \mathcal{I}$.
3. The *exchange property holds*: if $S, T \in \mathcal{I}$ and $|S| < |T|$, then there exists an $e \in T \setminus S$ such that $S \cup \{e\} \in \mathcal{I}$.

If B is a maximal subset of A, then we call it a *basis* of A. The bases of E are also called the bases of M, and the set of all these bases is $\mathcal{B}$. The *rank* of a subset $S \subset E$ is the maximum size of an independent subset of A. Using this definition, the rank function $r(S) : 2^E \mapsto \mathbb{R}$ of a matroid maps sets of elements to their ranks.

We can also play games using the structure of a matroid. An atomic matroid congestion game is specified by a tuple $\mathcal{M} = (k, M, (c_e)_{e \in E})$. There are k players, $M = (E, \mathcal{I})$ is a matroid defined over a ground set E of "resources", and each player's strategy is to choose a basis B of M. Each element e of the ground set has an associated strictly increasing cost function $c_e(f_e)$ that depends on the congestion of resource e. A strategy profile $(B_1, \ldots, B_k)$, where B_i is the basis of M chosen by player i, yields the congestion vector $f = (f_e)_{e \in E}$, where $f_e = |\{i : e \in B_i\}|$. The total cost of player i is then $\sum_{e \in B_i} c_e(f_e)$. We sometimes overload notation and use f_B for a basis B to denote the number of players who choose basis B.

2.7.1 Polymatroids

Our study of matroid congestion games will crucially utilize a polytope associated with matroids, and more generally *submodular functions* called a polymatroid. A function

$\rho : 2^E \to \mathbb{R}$ is called *submodular* if it satisfies $\rho(U) + \rho(V) \geq \rho(U \cup V) + \rho(U \cap V)$ for all $U, V \subset E$. The function ρ is *monotone* if $\rho(U) \leq \rho(V)$ for all $U \subseteq V$, and *normalized* if $\rho(\emptyset) = 0$.

Given a function $r : 2^E \mapsto \mathbb{R}_+$ that is submodular, monotone, and normalized, the pair (E, r) is called a *polymatroid*, and the corresponding *polymatroid base polytope* is

$$P_r = \{x \in \mathbb{R}_+^E \,|\, x(U) \leq r(U), \forall U \subset E,\, x(E) = r(E)\}$$

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

Chapter 3

Network Routing Games

Let $\mathcal{N} = (K, G, (\mathcal{S}_i)_{i\in K}, (c_e)_{e\in E}, (d_i)_{i\in K})$ be an atomic routing game played on the directed graph $G = (V, E)$. In this chapter, we only consider games that are symmetric (there is only one source s and one sink t, and every player chooses a strategy $P \in \mathcal{S}$, the set of all $s \rightsquigarrow t$ paths) and each player controls one unit of flow ($d_i = 1$, for all $i \in K$). Therefore, we shorten our notation to $\mathcal{N} = (K, G, \mathcal{S}, (c_e)_{e\in E})$.

Each edge $e \in E$ has an associated strictly increasing cost function $c_e(f_e) + \tau_e$ that depends on the congestion of edge e and the constant toll on the edge. The total cost for player i choosing some $P \in \mathcal{S}$ is $\sum_{e\in P}[c_e(f_e) + \tau_e]$, the sum of the costs of all of the edges in the path chosen by the player. We call $f = (f_e)_{e\in E}$ the congestion vector induced by the players' strategy profile. We will overload our notation for f and and also allow it to be indexed by the players' available strategies when convenient (see Section 3.2).

When brevity is important, we will use $c_f(e)$ instead of $c_e(f_e)$ to mean the cost of edge e under flow f, and use $c_f(P) = \sum_{e\in P} c_e(f_e)$ to mean the total cost incurred by a player choosing path P under f.

By Theorem 2.2, we know that given a desired minimal congestion vector f we can enforce it as the nonatomic Nash equilibrium via tolls. In this chapter, we will show that for some classes of network routing games we can also find tolls that induce f as the unique PDI atomic equilibrium, and compare these tolls to the nonatomic setting.

3.1 Series-Parallel Graphs

Directed *Series-parallel* (SP) with source s and sink t are a class of graphs that are generated recursively via composition operations. The smallest SP graph consists of just a source s and sink t that are connected by a single edge. Any two SP graphs can be combined in either series or parallel to generate another SP graph.

Let G_1 and G_2 be two SP graphs, with source-sink pairs (s_1, t_1) and (s_2, t_2) respectively. To combine G_1 and G_2 in series, take the disjoint union of G_1 and G_2 and then merge t_1 and s_2. The resulting SP graph G has source s_1 and sink t_2. To combine G_1 and G_2 in parallel, again take the disjoint union of G_1 and G_2. Now, merge the sources of the two graphs, and rename it s, and merge the sinks of the two graphs, and rename it t. G has source-sink pair (s, t).

Suppose we have a nonatomic congestion game $(K, G, (c_e)_{e \in E}, (d_i)_{i \in K})$ on a series-parallel graph G where the tolls τ induce the nonatomic equilibrium flow f. Again, assume that the edge costs are strictly increasing.

Theorem 3.1. *Let $\mathcal{N} = (K, G, \mathcal{S}, (c_e)_{e \in E})$ be an atomic network routing game on the directed series-parallel graph $G = (V, E)$ with source s and sink t. If τ is a set of tolls that induces f as the nonatomic equilibrium, then τ also induces f as the unique PDI atomic equilibrium.*

Before proving the theorem, we will first show the following claim about series-parallel graphs.

Lemma 3.2. *Let $G = (V, E)$ be a directed SP graph with source s and sink t. Let f, f' be two distinct $s - t$ flows routing D, D' units of flow respectively, with $D \geq D', D > 0$. Then, there is some $s - t$ path P such that for every $e \in P$, $f_e > 0$ and $f_e > f'_e$.*

Proof. We'll proceed by induction. In the base case for SP graphs, the graph is simply two vertices connected by an edge, so the claim is clearly true. Now, consider the series-parallel graph G which is the combination of two smaller graphs G_1 and G_2 combined in either series or parallel.

By induction, we know that the claim holds for both G_1 and G_2. Therefore, there exists paths P and Q where P satisfies the inductive hypothesis for G_1, and Q satisfies the inductive hypothesis for G_2. In the case that G was created by combining G_1 and G_2 in series, concatenating P and Q immediately gives a path with the desired properties. Now, assume that G was created by combining G_1 and G_2 in parallel. We know that since f

routes at least at many units of flow as f', it must route at least as many units of flow through either G_1 or G_2, otherwise we'd have a contradiction. Assume without loss of generality that f routes at least as many units of flow through G_1 compared to f'. Then P satisfies the properties we need by the inductive hypothesis. Both cases have provided a suitable path, so the claim holds for all series-parallel graphs by induction. $\square$

It's important to note that this lemma is nearly identical to Claim 3.3 from Swamy [23]. The only difference is that the Swamy's results only require a path where $f_e \geq f'_e$. However, we can extract strict inequality from the proof technique he describes with no extra effort. We can now prove the theorem. In the following proof, for a vector $x \in \mathbb{R}_+^E$, define $supp(x) := \{e : x_e > 0\}$. Therefore, for a congestion vector f, $supp(f)$ is the set of edges that are used in at least one player's strategy.

Proof. Since the given set of tolls τ induce a nonatomic equilibrium, we have that for any given path P used by the flow f,

$$\sum_{e \in P} [c_e(f_e) + \tau_e] \leq \sum_{e' \in Q} [c_{e'}(f_{e'}) + \tau_{e'}]$$

where Q is an arbitrary path.

Now, to show that τ induces f as an atomic equilibrium, we need that no player has incentive to deviate. Assume for contradiction that player i can reduce his cost by switching from P to Q. Thus,

$$\sum_{e' \in Q} [c_{e'}(f_{e'} + 1) + \tau_{e'}] < \sum_{e \in P} [c_e(f_e) + \tau_e]$$

However, observe that

$$\sum_{e \in P} [c_e(f_e) + \tau_e] \leq \sum_{e' \in Q} [c_{e'}(f_{e'}) + \tau_{e'}] \leq \sum_{e' \in Q} [c_{e'}(f_{e'} + 1) + \tau_{e'}]$$

where the first inequality follows from f being a nonatomic equilibrium, and the second from the fact that the edge costs are nondecreasing. Combining these inequality chains, we see that $\sum_{e \in P} [c_e(f_e) + \tau_e] < \sum_{e \in P} [c_e(f_e) + \tau_e]$, a contradiction. Therefore, f is indeed an atomic equilibrium.

It remains to be shown that f is the unique PDI atomic equilibrium. Assume for contradiction that some other PDI atomic equilibrium flow f' exists. Note that we can always impose some very large tolls on the edges not in $supp(f)$ and eliminate equilibria using edges not in $supp(f)$. So we may assume that $supp(f') \subseteq supp(f)$. Since f and f' are distinct flows, by Lemma 3.2 there exists a path P such that $f_e > 0$ and $f_e > f'_e$ for every edge $e \in P$. Since f is PDI, there exists a decomposition of f into an atomic equilibrium where at least one player uses P. Since both f, f' are atomic equilibria they must be integer flows, and so we have more specifically that $f_e \geq f'_e + 1$ for each $e \in P$. Similarly, we can use Lemma 3.2 again to find a path Q such that $f''_e > 0$ for all $e' \in Q$ (which means that $f'_e > 0$ since f' sends flow through a subset of the edges used by f) where $f_e + 1 \leq f'_e$ for all $e \in Q$. Again, since f' is also PDI, there exists a decomposition of f' that is an atomic equilibrium using Q. Now, consider these two decompositions using P and Q. Since f is a nonatomic equilibrium and we have strictly increasing cost functions, we have that

$$\sum_{e \in P} c_e(f_e) = \sum_{e' \in Q} c_{e'}(f_{e'}) < \sum_{e' \in Q} c_{e'}(f_{e'} + 1) \leq \sum_{e' \in Q} c_{e'}(f'_{e'}).$$

Similarly, since f' is an atomic equilibrium, we have that

$$\sum_{e' \in Q} c_{e'}(f'_{e'}) \leq \sum_{e \in P} c_e(f'_e + 1) \leq \sum_{e \in P} c_e(f_e).$$

And thus we see that we have both $\sum_{e \in P} c_e(f_e) < \sum_{e' \in Q} c_{e'}(f'_{e'})$ and $\sum_{e' \in Q} c_{e'}(f'_{e'}) \leq \sum_{e \in P} c_e(f_e)$, a contradiction. Thus, f is the unique PDI atomic equilibrium. □

3.2 The Four-Link Graph

While the tolls in the nonatomic case can be extended to the atomic case for SP graphs, this is not true in general. One of the simplest counterexamples can be shown in the following network routing game. Consider an atomic routing game played on the following "four-link" directed graph $G = (V, E)$ on six vertices with edge set $E = \{e_1, \ldots, e_8\}$.

This graph is very similar to the one found in Braess' Paradox; in fact, simply delete edge e_5 and contract e_3 and e_6 to recover the famous example.

Now, we will consider a simple atomic network routing game played on this graph. Let there be 4 players, and let the edge costs be $c_e(f_e) = f_e$ for every $e \in E$ (no tolls have been

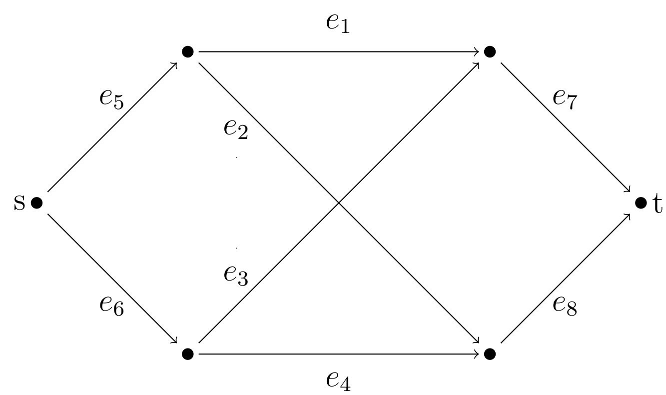

Figure 3.1: The four-link graph.

imposed yet). Notice that there are four $s \rightsquigarrow t$ paths, and each of these paths uses exactly one of e_1, e_2, e_3, or e_4. We will call these edges *identifying* edges, and label the paths that use them as P_1, P_2, P_3, and P_4 respectively.

Consider the path-decomposition $f_P = (f_{P_1}, f_{P_2}, f_{P_3}, f_{P_4}) = (f_{e_1}, f_{e_2}, f_{e_3}, f_{e_4}) = (1, 1, 1, 1)$. In this four-link example, the amount of flow through these four identifying edges is exactly the same as the flow on their respective paths. Therefore, the flow through the remaining four edges is entirely determined by the flow through the identifying edges. For example, $f_{e_5} = f_{e_1} + f_{e_3}$, since e_5 in an edge in paths P_1 and P_3. Under $f = (1, 1, 1, 1)$, each of the four paths cost 5, since there are 2 units of traffic on the outgoing edges of s and incoming edges to t, and 1 unit of traffic on the other edges. Therefore, zero tolls establish $(1, 1, 1, 1)$ as a nonatomic equilibrium.

However, zero tolls do not establish a unique equilibrium in the atomic setting. In the atomic case, in addition to $(1, 1, 1, 1)$, the congestion vectors $(2, 0, 0, 2)$ and $(0, 2, 2, 0)$ are also atomic equilibria. However, we can impose a small toll to enforce $(1, 1, 1, 1)$ as the unique equilibrium.

Proposition 3.3. *Let $\mathcal{N} = (K, G, (\mathcal{S}), (c_e)_{e \in E})$ be a four player atomic routing game on the four-link graph. Let the cost functions be $c_e(x) = ax + b_e$ for all $e \in E$. Let $\tau_7 = \epsilon > 0$ be the only nonzero toll, and let ϵ be arbitrarily small. Then $f = (1, 1, 1, 1)$ is the unique PDI atomic equilibrium.*

Proof. We will first argue that $(1, 1, 1, 1)$ is still an atomic equilibrium. Not including the arbitrarily small toll on e_7, every cost function is integral, so if there is incentive to deviate, it must be by an integral amount. Therefore, an ϵ toll will not provide incentive to deviate, meaning that $(1,1,1,1)$ is an atomic equilibrium.

Next, we show that neither $(2,0,0,2)$ nor $(0,2,2,0)$ are atomic equilibria. In the former case, a player using P_1 can decrease his cost by ϵ by switching to P_2. In the latter case, a player using P_3 can decrease his cost by ϵ by switching to P_4.

Finally, recall that every other strategy profile had an integral incentive to deviate before the toll was added. Therefore, since the toll was arbitrarily small, there is still incentive to deviate in each of these strategy profiles, and so $(1,1,1,1)$ is the unique atomic equilibrium, and is therefore PDI. □

While this set of tolls used in Proposition 3.3 are different from the nonatomic setting, the two sets of tolls are nearly identical. This begs the question of whether the tolls that induce f as the unique atomic equilibrium can be created by only modifying some set of nonatomic tolls by a small amount. More formally, we have:

Question 3.4. *Let f be a nonatomic equilibrium under a set of tolls τ. Does there exist τ', which is an arbitrarily small perturbation of τ, such that f is the unique PDI atomic equilibrium under τ'?*

We can provide a negative answer to this question, even while remaining in the realm of the four-link example.

Proposition 3.5. *Let $\mathcal{N} = (K, G, (\mathcal{S}), (c_e)_{e\in E})$ be a four player atomic routing game on the four-link graph. Let $c_e(x) = 2x$ for e_5, e_6, e_7, e_8 and $c_e(x) = x$ otherwise. Let τ be a set of tolls that induces $f = (1,1,1,1)$ as the nonatomic equilibrium. Then no arbitrarily small perturbation of τ exists in which $(1,1,1,1)$ is the unique PDI atomic equilibrium.*

Proof. We will first show that the proposition holds for a specific set of tolls, and then prove the result in general. Let $\tau = \vec{0}$. Then every path has the same cost (9), so zero tolls enforce $(1,1,1,1)$ as the nonatomic equilibrium, and also as an atomic equilibrium. Now, consider the flow $(2,0,0,2)$. Here, every player incurs a cost of 10. If they switch paths, in the best case scenario they choose either P_2 or P_3 and incur cost 11. If they switch from P_1 to P_4 or vice versa, they will experience a much higher cost (15). Therefore, we cannot create arbitrarily small tolls on any edge to encourage one of these players to switch, as the difference in cost between a player's current path and the cost they would incur after deviating is integral. Therefore, $(2,0,0,2)$ remains an atomic equilibrium regardless of how we construct arbitrarily small tolls.

Now, consider any other set of tolls that makes $(1,1,1,1)$ the nonatomic equilibrium. If $(1,1,1,1)$ is the nonatomic equilibrium, then this necessitates that the cost of all four paths must be the same. Therefore, when we examine the $(2,0,0,2)$ setting, the difference in costs between the four paths is identical to the zero tolls case. Thus, it is impossible to remove $(2,0,0,2)$ as an atomic equilibrium by perturbing our tolls by an arbitrarily small amount in the general case.

□

Before proving the main result, we first mention a few important properties of congestion games played on the four-link graph.

- Every congestion vector corresponds to a unique path-decomposition. In particular, the number of players choosing path P_i is f_{e_i}. This means that if f is an atomic equilibrium, it is PDI, since f can only be decomposed one way, modulo player renaming.
- Each path P_i has three edges. In addition, paths P_i and P_j either share 0 edges or 1 edge, when $i \neq j$.
- Each path P_i has edge e_i which is unique to it. Therefore, we can use tolls to adjust the cost of any path without affecting the cost of the other three paths. This means that we can induce f as a nonatomic equilibrium using only τ_1, τ_2, τ_3, and τ_4, with zero tolls elsewhere.
- Given a set of tolls τ, we can find another set of tolls $\tau' = (\tau'_1, \tau'_2, \tau'_3, \tau'_4, 0, 0, 0, 0)$ in which the cost of each path under τ' is the same as it was under τ. To construct τ', let $\tau'_i = \sum_{e_j \in P_i} \tau_j$, for $i = 1, 2, 3, 4$. In this sense, we can always condense our tolls onto the identifying edges for convenience.

We will now show a lemma that uses the structure of the four-link graph to quickly describe the situations where a player has incentive to deviate. In this lemma, we will use the shorthand $c_f(P)$ to mean the cost of path P under the flow f. Additionally, $c_{f+1}(P)$ is used to mean the cost of P when each of its edges have one extra unit of congestion compared to f.

Lemma 3.6. *Let $\mathcal{N} = (K, G, (\mathcal{S}), (c_e)_{e \in E})$ be a atomic routing game on the four-link graph. For a given flow f, and edge costs $c_e(f_e) = af_e + b_e$, for all $e \in E$, a player has incentive to deviate from P_i to P_j iff $c(P_i) > c(P_j) + a \cdot |P_j \setminus P_i|$.*

Proof. There is an incentive to deviate from P_i to P_j iff

$$\begin{aligned} c_f(P_i) &> c_{f+1}(P_j \setminus P_i) + c_f(P_j \cap P_i) \\ &= \sum_{e \in P_j \setminus P_i} [a(f_e+1)+b_e] + \sum_{e \in P_j \cap P_i} [af_e + b_e] \\ &= \sum_{e \in P_j \setminus P_i} a + \sum_{e \in P_j \setminus P_i} c_f(e) + \sum_{e \in P_j \cap P_i} c_f(e) \\ &= c(P_j) + a \cdot |P_j \setminus P_i| \end{aligned}$$

□

This lemma gives a threshold where players will prefer to deviate from P_i to P_j. As mentioned before the lemma, in the four-link graph two paths share either one edge or no edges. Therefore, incentive to deviate exists when the difference in path costs differ by more than $3a$ or $2a$, respectively.

Recall that in Proposition 3.3 we showed than an ϵ toll on e_7 was enough to make $f = (1,1,1,1)$ the unique PDI atomic equilibrium, provided our cost functions are of the form $c_e(x) = ax + b_e$. We will now expand on this result, and show that $(1,1,1,1)$ is still the unique atomic equilibrium even if we allow negative flow. In particular, we will show that there is incentive to deviate away from any congestion vector $g = (w,x,y,z)$ when $w + x + y + z = 4$.

Lemma 3.7. *Let $\mathcal{N} = (K, G, (\mathcal{S}), (c_e)_{e \in E})$ be a four player atomic routing game on the four-link graph. Let the cost functions be $c_e(x) = ax + b_e$ for all $e \in E$. Let $\tau_7 = \epsilon > 0$ be the only nonzero toll, and let ϵ be arbitrarily small. Then $f = (1,1,1,1)$ is a PDI atomic equilibrium, and no $g = (w,x,y,z)$ is an atomic equilibrium if $w+x+y+z = 4$.*

Proof. We have already shown in Proposition 3.3 that $f = (1,1,1,1)$ is an atomic equilibrium. We have also shown that $g = (w,x,y,z)$ is not an atomic equilibrium if $w+x+y+z = 4$ and $w,x,y,z \in \mathbb{Z}_{\geq 0}$.

Let some entry of g be negative. Since g is integral, that entry is at most -1. Assume for now that $w \leq -1$. We calculate $c_f(P_1)$ and $c_g(P_1)$, which are

$$c_f(P_1) = 3a + a + a + b_1 + b_3 + b_7 + \tau_7$$

and

$$c_g(P_1) = 3aw + ax + ay + b_1 + b_3 + b_7 + \tau_7$$

since each unit of flow through P_1 contributes $3a$ to the cost of P_1, and each unit of flow through P_2 and P_3 contributes a to the cost of P_1 (since P_2 and P_3 share exactly one edge with P_1). Now, the difference in the cost of P_1 under g compared to f is

$$c_g(P_1) - c_f(P_1) = 3a(w-1) + a(x-1) + a(y-1)$$

The minimum decrease in cost of P_1 can be found by maximizing this function subject to $w \leq -1$ and $w+x+y+z = 4$. This minimum decrease is $4a$, and happens when $w = -1$ and $(x-1)+(y-1) = 2$. To verify this, notice that decreasing w' further will only serve to create a larger cost discrepancy, because its $3a$ coefficient outweighs the combination of the positive terms. These positive terms are maximized when $(x-1)+(y-1) = -(w-1) = 2$, as the combined flow increase through P_2 and P_3 can be at most the flow decrease through P_1 (since $w + x + y + z = 4$).

We now argue that after this $4a$ decrease in cost of P_1, some player will have incentive to deviate to P_1 from one of the other three paths. Recall that before the epsilon toll was added to e_7, f was a nonatomic equilibrium for $\mathcal{N}$, so with the toll, $c_f(P_i) - c_f(P_1)$ is 0 if $i = 3$, since P_3 also uses e_7 and also incurs the epsilon toll, and $-\epsilon$ otherwise.

Now, given that $c_g(P_1) - c_f(P_1) \leq -4a < 0$, it is impossible that all of the following inequalities also hold:

- $c_g(P_2) - c_f(P_2) < 0$
- $c_g(P_3) - c_f(P_3) < 0$
- $c_g(P_4) - c_f(P_4) < 0$

If this were the case, then remove the ϵ toll on e_7, and see that every path is lower cost in g compared to f, contradicting the fact that f is the nonatomic equilibrium under the original tolls. Thus it must be that at least one of P_2, P_3, P_4 is more expensive under g than f.

- If $c_g(P_2) - c_f(P_2) \geq 0$ then
$$c_g(P_2) \geq c_f(P_2) = c_f(P_1) - \epsilon \geq c_g(P_1) + 4a - \epsilon$$
.

- If $c_g(P_3) - c_f(P_3) \geq 0$ then

$$c_g(P_3) \geq c_f(P_3) = c_f(P_1) \geq c_g(P_1) + 4a$$

.

- If If $c_g(P_4) - c_f(P_4) \geq 0$ then

$$c_g(P_4) \geq c_f(P_4) = c_f(P_1) - \epsilon \geq c_g(P_1) + 4a - \epsilon$$

.

In every case, there exists some $i \in \{2, 3, 4\}$ such that

$$\begin{aligned} c_g(P_i) &\geq c_g(P_1) + 4a - \epsilon \\ &> c_g(P_1) + a \cdot |e(P_1 \setminus P_i)| \end{aligned}$$

because $|e(P_1 \setminus P_i)| = 3$ if $i = 4$, and $|e(P_1 \setminus P_i)| = 2$ otherwise. Therefore, by Lemma 3.6, there is incentive for a player to deviate, and so g is not an atomic equilibrium. Recall that all of this analysis was done under the assumption that $w \leq -1$, meaning that P_1 was a path with significantly less flow in g compared to f. However, under symmetry P_1, P_2, P_3, and P_4 all function identically to each other in the four-link graph, except that only P_1 and P_3 are affected by the epsilon toll on e_7. The exact same argument can be immediately used for P_3 to show that there is always incentive to deviate to P_3 if $y \leq -1$.

Nearly identical analysis can be used for P_2; the only difference is that P_2 does not incur the extra epsilon cost. In particular, under the ϵ toll, $c_f(P_i) - c_f(P_2)$ is 0 if $i = 4$, since P_4 also does not incur the epsilon toll, and ϵ otherwise. This has a cascading effect – in the final step, we instead finish with

$$\begin{aligned} c_g(P_i) &\geq c_g(P_2) + 4a \\ &> c_g(P_2) + a \cdot |e(P_2 \setminus P_i)| \end{aligned}$$

for some $i \in \{1, 3, 4\}$. The analysis for P_4 is the same, as it is symmetric to P_2. Therefore in all cases we can show that g is not an atomic equilibrium. □

The reason we extend the results of Proposition 3.3 is that our strategy for the more general case will be to use tolls to mimic the relative costs of the paths in the 4 player game. We will expand upon this idea in the next lemma.

Lemma 3.8. *Let $\mathcal{N} = (K, G, (\mathcal{S}), (c_e)_{e\in E})$ be a atomic routing game on the four-link graph. Let the edge costs be $c_e(x) = ax + b_e$. Consider two distinct flows f, f', and suppose that the relative costs of each of the paths in $\mathcal{S}$ are identical, meaning $c_f(P_i) - c_f(P_j) = c_{f'}(P_i) - c_{f'}(P_j)$, for all $P_i, P_j \in \mathcal{S}$. Then f is an atomic equilibrium iff f' is an atomic equilibrium.*

Proof. f is an atomic equilibrium iff for every $P_i \in supp(f)$ and any $P_j \in \mathcal{S}$,

$$\begin{aligned} c_f(P_i) &\leq c_{f+1}(P_j \setminus P_i) + c_f(P_j \cap P_i) \\ &= c_f(P_j) + a \cdot |P_j \setminus P_i| \end{aligned}$$

by Lemma 3.6. Rearranging terms, we have

$$c_f(P_i) - c_f(P_j) \leq a \cdot |P_j \setminus P_i|$$

and then since the relative costs of the paths are identical,

$$c_{f'}(P_i) - c_{f'}(P_j) \leq a \cdot |P_j \setminus P_i|.$$

Therefore,

$$\begin{aligned} c_{f'}(P_i) &\leq c_{f'}(P_j) + a \cdot |P_j \setminus P_i| \\ &= c_{f'+1}(P_j \setminus P_i) + c_{f'}(P_j \cap P_i) \end{aligned}$$

and so f' is an atomic equilibrium, as desired. ☐

Lemmas 3.7 and 3.8 yield a very powerful result. Consider some f for an atomic routing game $\mathcal{N}$ on the four-link graph with edge costs $c_e(f_e) = af_e + b_e$. Now, construct $\mathcal{N}'$, a game that is identical to $\mathcal{N}$ except it is only played with 4 players. If we can find a congestion vector f' that imposes the same relative costs on the paths in $\mathcal{N}'$ as f in $\mathcal{N}$, then f' is an atomic equilibria iff f is. In our main result, we show that for every feasible f admitted by $\mathcal{N}$, we can find an f' for $\mathcal{N}'$ that will tell us whether or not f is an atomic equilibrium.

Theorem 3.9. *Let $\mathcal{N} = (K, G, (\mathcal{S}), (c_e)_{e\in E})$ be a atomic routing game on the four-link graph. Let $w, x, y, z \in \mathbb{Z}_{\geq 0}$. Then for any flow $f = (w, x, y, z)$, there exist tolls τ that induce f as the unique PDI atomic equilibrium.*

Proof. Using Theorem 2.2, we can find a set of tolls τ that induce f as a nonatomic equilibrium in the nonatomic setting. This only guarantees that the costs of each path with nonzero flow are the same, so if any P_i has zero flow, adjust the toll τ_i so that this path also experiences the same cost as the others (if this toll would need to be negative, simply increase the tolls on the identifying edges of the other routes instead). Since f is a nonatomic equilibrium, it is also an atomic equilibrium.

Let $\mathcal{N}' = (K' = \{1, 2, 3, 4\}, G, (\mathcal{S}), (c_e)_{e\in E})$ be a nearly identical atomic routing game to $\mathcal{N}$, with the only change being that this game is played by only four players. Again, we can find tolls τ' that enforce $f' = (1, 1, 1, 1)$ as a nonatomic, and therefore atomic, equilibrium. Note that since both f and f' are nonatomic equilibria of their respective games, $c_f(P_i) - c_f(P_j) = c_{f'}(P_i) - c_{f'}(P_j) = 0$, for all $P_i, P_j \in \mathcal{S}$.

For convenience, update each b_e in the cost functions of both games to $b_e + \tau_e$ and $b_e + \tau'_e$ respectively, and discard the tolls. Now, add an $\epsilon > 0$ toll to e_7 in both games. In both cases, $f = (w, x, y, z)$ and $f' = (1, 1, 1, 1)$ both remain atomic equilibria for their respective games. To see this, assume that this is not the case and some player had incentive to deviate. Since the cost functions are strictly increasing, this player's new strategy must cost more than his old one did before the toll was added. Therefore, we can always make ϵ small enough to make sure this incentive to deviate doesn't exist. Note that since we have only added the same single additive toll to both games, we have maintained that $c_f(P_i) - c_f(P_j) = c_{f'}(P_i) - c_{f'}(P_j)$, for all $P_i, P_j \in \mathcal{S}$.

Now, we will show that any congestion vector $g = (w', x', y', z') \neq f$ is not an atomic equilibrium of $\mathcal{N}$. Construct $g' = f' + (g - f)$. Observe that since our cost functions are linear and $g' = f' + (g - f)$, then $c_{g'}(P_i) = c_{f'}(P_i) + c_g(P_i) - c_f(P_i)$, for all $P_i \in \mathcal{S}$. Similarly, we can write $g = f + (g - f)$ and so $c_g(P_i) = c_f(P_i) + c_g(P_i) - c_f(P_i)$, for all $P_i \in \mathcal{S}$. Then

$$\begin{aligned} c_{g'}(P_i) - c_{g'}(P_j) &= c_{f'}(P_i) + c_g(P_i) - c_f(P_i) - c_{f'}(P_j) - c_g(P_j) + c_f(P_j) \\ &= c_f(P_i) + c_g(P_i) - c_f(P_i) - c_f(P_j) - c_g(P_j) + c_f(P_j) \\ &= c_g(P_i) - c_g(P_j) \end{aligned}$$

for all $P_i, P_j \in \mathcal{S}$. The second step follows from the fact that $c_f(P_i) - c_f(P_j) = c_{f'}(P_i) - c_{f'}(P_j)$, for all $P_i, P_j \in \mathcal{S}$. Therefore, since $c_{g'}(P_i) - c_{g'}(P_j) = c_g(P_i) - c_g(P_j)$, by Lemma 3.8 if g' is not an atomic equilibrium of $\mathcal{N}'$, then g is not an atomic equilibrium of $\mathcal{N}$.

Now, since $g \neq f$, $g' \neq f' = (1,1,1,1)$. For $g' = (g'_{e_1}, g'_{e_2}, g'_{e_3}, g'_{e_4})$, we can confirm that $g'_{e_1} + g'_{e_2} + g'_{e_3} + g'_{e_4} = 4$, since $f'_{e_1} + f'_{e_2} + f'_{e_3} + f'_{e_4} = 1+1+1+1 = 4$, $g' = f' + (g - f)$, and g and f have the same total amount of flow since they are from the same game. Then by Lemma 3.7 g' is not an atomic equilibrium of $\mathcal{N}'$, and so g is not an atomic equilibrium of $\mathcal{N}$. Therefore, f is the only atomic equilibrium, and is therefore PDI.

□

Chapter 4

Matroid Congestion Games

Recall the definition of an atomic matroid congestion game $\mathcal{M} = (k, M, (c_e)_{e \in E})$. We are given a matroid $M = (E, \mathcal{I})$ over a ground set E of "resources", there are k players, and each player's strategy involves choosing a basis of the matroid M. Each element e of the ground set has an associated strictly increasing cost function $c_e(f_e)$ that depends on the congestion of resource e. The total cost for player i is $\sum_{e \in B} c_e(f_e)$, the sum of the costs of all of the elements in the basis chosen by the player. We call $f = (f_e)_{e \in E}$ the congestion vector induced by the players' strategy profile.

From Theorem 2.2, we know that given a desired minimal congestion vector f we can enforce it as a nonatomic Nash equilibrium via tolls. In this chapter, we will show that for certain classes of matroid congestion games, f can also be enforced as the unique PDI atomic equilibrium.

4.1 Relating Polymatroids to Matroid Congestion Games

Let $r(A) : 2^E \mapsto \mathbb{R}$ be the rank function of the matroid M. It is well known that r is monotone, submodular, and normalized. Now, observe that we can scale the rank function by a constant k and maintain the three properties. Therefore, we can scale the polymatroid (E, r) by k, creating the following polymatroid base polytope.

$$P_{k,r} = \{x \in \mathbb{R}^E_+ \,|\, x(U) \leq kr(U), \forall U \subseteq E,\, x(E) = kr(E)\}$$

4.1.1 Strategy Profiles and Polymatroids

We will show that congestion vectors of a matroid congestion game correspond precisely to the integral points of its associated polymatroid.

Theorem 4.1. *Let $\mathcal{M} = (k, M, (c_e)_{e\in E})$ be a matroid congestion game, and $P_{k,r}$ be the associated polymatroid. Then every strategy profile f of $\mathcal{M}$ corresponds to an integer point in $P_{k,r}$.*

Proof. We first show the easy direction. Let f be a strategy profile of $\mathcal{M}$. Clearly f is integral. Let f correspond to the strategy profile $(B_1, \dots, B_k)$, where each B_i is a basis of $\mathcal{M}$. Then, $f(E) = \sum_{i=1}^{k} |B_i| = kr(E)$, and for any set U, we have $f(U) = \sum_{i=1}^{k} |B_i \cap U| \leq kr(U)$ where the inequality is because $B_i \cap U$ is an independent set.

For the other direction, we will rely on the Matroid Intersection Theorem.

Lemma 4.2. [8] For matroids M_1, M_2 defined on a ground set S, with rank functions r_1, r_2

$$max\{|J| : J \in \mathcal{I}_1 \cap \mathcal{I}_2\} = min\{r_1(A) + r_2(\overline{A}) : A \subseteq S\}.$$

We will use matroid intersection to show that integer points $f \in P_{k,r}$ are strategy profiles. First, create two new matroids. Let $E_1, \dots, E_k$ be k disjoint copies of the ground set E, and $E' = \bigcup E_i$. Let M_i be a copy of M on the ground set E_i, for $i = 1, \dots, k$. Let N_1 be the direct sum $M_1 \bigoplus \cdots \bigoplus M_k$, which has ground set E'. Let N_2 be the matroid on E', where a set $J \subseteq E'$ is independent iff it contains at most f_e copies of each $e \in E$

Thus, a basis of N_1 corresponds to each of the k players choosing a basis of M, and a basis of N_2 contains exactly f_e copies of each $e \in E$. Therefore, if we can find a common independent set of size $kr(E)$ of N_1 and N_2, which therefore is a common basis of N_1 and N_2 $f \in P_{k,r}$ can be decomposed into k bases and therefore corresponds to a strategy profile.

Let r_1 and r_2 be the rank functions of N_1 and N_2 respectively. For $e \in E$, let $C_e \subseteq E'$ denote the set of all copies of e. Observe that $r_1(A) = \sum_{i=1}^{k} r(A \cap E_i)$, and $r_2(A) = \sum_{e\in E} min(f_e, |C_e \cap A|)$.

Therefore, using Lemma 4.2 we want to show that

$$max\{|J| : J \in \mathcal{I}_1 \cap \mathcal{I}_2\} = min\{r_1(A) + r_2(\overline{A}) : A \subseteq E'\} \geq kr(E).$$

Note that showing inequality is enough, as the maximum size of an independent set in N_1 is clearly $kr(E)$. First, we will argue that for any $A \subset E'$, we can find $A' \subset E'$ such that $r_1(A) + r_2(\overline{A}) \geq r_1(A') + r_2(\overline{A'})$, where A' has the special property that for every $e \in E$, A' chooses either all k copies of e or no copy of e.

Essentially, we construct $\overline{A'}$ as follows. For any element e, if $\overline{A}$ contains at most f_e copies of e, then we remove all copies of e from $\overline{A}$ (so now they belong to A'); otherwise, we expand $\overline{A}$ to include all k copes of e (so $\overline{A'}$ now contains C_e). In the former case, $r_1(A')$ may increase by at most $|C_e \cap A|$ but $r_2(\overline{A'})$ decreases by precisely $|C_e \cap A|$, and in the latter case, $r_1(A')$ cannot increase and $r_2(\overline{A'})$ remains unchanged. Formally, we have $A' = \bigcup_{e \in E : |C_e \cap A| \geq k - f - e} C_e$, and we have established above that $r_1(A) + r_2(\overline{A}) \geq r_1(A') + r_2(\overline{A'})$.

What we have shown above is that in the min-expression, we can restrict to sets A that are of the form $\bigcup_{e \in S} C_e$ for some set $S \subseteq E$. That is, we have

$$\min\{r_1(A) + r_2(\overline{A}) : A \subseteq E'\} = \min\{r_1(A') + r_2(\overline{A'}) : A' = \bigcup_{e \in S} C_e, S \subseteq E\}.$$

Now, consider a set $A^* = \bigcup_{e \in S^*} C_e$ that minimizes the above expression. We have $r_1(A^*) = kr(S^*) \geq f(S^*)$, where the last inequality is because f lies in the polymatroid, and $r_2(\overline{A^*}) = \sum_{e \notin S^*} f_e = f(\overline{S^*})$, and so $r_1(A^*) + r_2(\overline{A^*}) \geq f(E) = kr(E)$.

Therefore, by Lemma 4.2 there exists a common independent set of size $kr(E)$ between N_1 and N_2. Since the size of this set is $kr(E)$, it is a basis of N_1, so f is decomposable into k bases and is therefore a strategy profile. □

4.1.2 Results on Strategy Profiles

Relating strategy profiles of matroid congestion games to their polymatroid base polytope allows us to draw from results in polymatroid theory. One property of matroids that is especially useful to us is Strong Basis Exchange.

Lemma 4.3. [8] For $B_1, B_2 \in \mathcal{B}$, and any $x \in B_1 \setminus B_2$, there exists $y \in B_2 \setminus B_1$ such that $(B_1 \setminus \{x\}) \cup \{y\} \in \mathcal{B}$, and $(B_2 \cup \{x\}) \setminus \{y\} \in \mathcal{B}$.

In the context of polymatroid theory, strong basis exchange manifests itself in the following analogues of Lemma 4.3 and are given by Murota.

Lemma 4.4. [14] Let $x, y \in P_{k,r}$, and suppose $x_u > y_u$ for some $u \in E$. Then there exists $v \in E$ such that $x_v < y_v$ and $\epsilon > 0$ such that:

$$x + \epsilon(\chi_v - \chi_u) \in P_{k,r} \text{ and } y + \epsilon(\chi_u - \chi_v) \in P_{k,r}.$$

Corollary 4.5. [14] Let x, y be integer points in $P_{k,r}$, and suppose $x_u > y_u$ for some $u \in E$. Then there exists $v \in E$ such that $x_v < y_v$ and

$$x - \chi_u + \chi_v \in \mathcal{Q} \text{ and } y + \chi_u - \chi_v \in \mathcal{Q}.$$

Using these results, we can take existing strategy profiles and find additional ones by taking a "step" from one strategy profile in the direction of another. As we show in Section 4.2, it turns out that our central question of whether a target integral congestion vector f (imposed as a nonatomic equilibrium via tolls) can be imposed as a PDI atomic equilibrium, is closely related to the following key question: if we take two strategy profiles g and g' that are only one step apart, can we can find decompositions of them that isolate this small difference into only one basis, while the remainder of their decompositions stay identical? The following conjecture states precisely that this always holds.

Conjecture 4.6. *Let g and g' be integer points in $P_{k,r}$, where $g' = g + \chi_x - \chi_y$, for some $x, y \in E$. Then for the strategy profiles g and g', $g' \in NBD(g)$.*

Recall that $g' \in NBD(g)$ means that there exist decompositions of g and g' that differ in exactly one basis. In the next section, we will show that if this conjecture holds, we can show positive results about the enforeability of atomic equilibria.

4.2 Atomic Equilibria in Matroid Congestion Games

We already know that a given desired congestion vector f can be imposed as a nonatomic equilibrium via tolls. We will now show that if f is integral and Conjecture 4.6 holds, these same tolls are sufficient to also enforce f as the unique PDI atomic Nash equilibrium.

Theorem 4.7. *Let $\mathcal{M}$ be an atomic matroid congestion game with matroid $M = (E, \mathcal{I})$ and associated polymatroid base polytope $P_{k,r}$. If f is an integral nonatomic Nash Equilibrum and Conjecture 4.6 holds, then f is also the unique PDI atomic Nash Equilibrium.*

Proof. Assume for contradiction that some other PDI atomic Nash Equilibrium $g \neq f$ exists. Because g is an atomic equilibrium, it is integral and corresponds to a strategy profile, so $g \in P_{k,r}$. Similarly, f being a nonatomic equilibrium implies $f \in P_{k,r}$. Since $f \neq g$, there exists $x \in E$ such that $f_x > g_x$ (since $f(E) = g(E)$). Therefore, by Corollary 4.5, there exists $y \in E$ such that $f_y < g_y$, and both $f' = f + \chi_y - \chi_x$ and $g' = g + \chi_x - \chi_y$ are integral points in $P_{k,r}$.

We assume that Conjecture 4.6 holds, so $g' \in NBD(g)$. Since g is an atomic equilibrium, it is a local minimum of the potential function, and since $g' \in NBD(g)$ we have

$$\phi_a(g) - \phi_a(g') \leq 0.$$

or equivalently,

$$c_y(g_y) - c_x(g'_x) \leq 0$$

since the difference in potential between two neighboring congestion vectors is exactly the cost incurred by the single deviating player by switching from resource y to resource x.

Similarly, since f is a nonatomic equilibrium, by Lemma 2.1 f and f' satisfy

$$\sum_{e \in E} c_e(f_e)(f_e - f'_e) \leq 0$$

which when simplified evaluates to

$$c_x(f_x) \leq c_y(f_y).$$

Now, we have

$$c_x(f_x) \leq c_y(f_y) < c_y(g_y) \leq c_x(g'_x) = c_x(g_x + 1) \leq c_x(f_x).$$

The second inequality here follows from the fact that our cost functions are strictly increasing, and $f_y < g_y$. The equality comes from the definition of g', and the last inequality holds because f and g are integral, and $f_x > g_x$. Thus we have that $c_x(f_x) < c_x(f_x)$, a contradiction, proving the theorem. □

The rest of the chapter is dedicated to proving Conjecture 4.6 for certain classes of matroids. We first provide a proof for the relatively straightforward case of partition matroids (Section 4.3), and then continue with the more general laminar matroid case (Section 4.4).

4.3 Proof of Conjecture 4.6 for Partition Matroids

A partition matroid $M = (E, \mathcal{I})$ has its independent sets defined by $\mathcal{I} = \{I : |I \cap E_i| \leq r(E_i) \text{ for all } i = 1, \dots, m\}$ where m is an integer and $E_1, \dots, E_m$ form a partition of E. One thing to note about partition matroids is that every basis B of the matroid contains exactly $r(E_i)$ elements from each E_i. If B contained fewer than $r(E_i)$ elements from E_i, we could simply add any element in $E_i \setminus B$ to get a larger independent set.

Observe that proving Conjecture 4.6 entails finding a decomposition of g containing a basis B such that: (i) $y \in B, x \notin B$, and (ii) $B' = B \cup \{x\} \setminus \{y\}$ is also a basis. If (i) and (ii) hold, then replacing B in the decomposition of g with B' yields a decomposition of g' where only one basis has changed, showing that $g' \in NBD(g)$.

We first show that for any matroid, one can always find a decomposition of g containing a basis B for which (i) holds. From there, our main task is to show that one can suitably modify B (and the decomposition of g) while maintaining (i) so that (ii) also holds.

Lemma 4.8. *Let $M = (E, \mathcal{I})$ be a matroid with corresponding polymatroid base polytope P_ρ. Let $g, g' \in P_{k,r}$ be integral strategy profiles satisfying $g' = g + \chi_x - \chi_y$, for some $x, y \in E$. Then there exists a decomposition $\mathcal{D}$ of g containing a basis B_j such that $y \in B_j, x \notin B_j$.*

Proof. Observe that since $g' = g + \chi_x - \chi_y$ and strategy profiles are nonnegative vectors, $g_y > g'_y \geq 0$. Since $g_y > 0$ and g is integral, $g_y \geq 1$, so in every decomposition of g there is some basis B containing y.

Take an arbitrary decomposition $\mathcal{D}$ of g and corresponding basis B. If $x \notin B$, we are done, so assume that $x \in B$. We have $g_x < g'_x \leq k$, so there is some basis $B' \in \mathcal{D}$ such that $x \notin B'$.

Since $x \in B \setminus B'$, there must also exist some resource $z \in B' \setminus B$. Note that $z \neq y$, since $y \in B$. Now, we can use strong basis exchange to create two new bases $\bar{B} = B - x + z$ and $\bar{B}' = B' + x - z$. Notice that $\chi_{\bar{B}} + \chi_{\bar{B}'} = \chi_B + \chi_{B'}$, so we can replace B and B' in $\mathcal{D}$ with $\bar{B}$ and $\bar{B}'$ to obtain another decomposition of g. In the new decomposition, $y \in \bar{B}$, $x \notin \bar{B}$, as desired. □

Now, to complete the proof for partition matroids, we need to show that this basis B from a decomposition of g can be swapped for a suitable B'_j to obtain a neighboring decomposition of g'.

Theorem 4.9. *Let $\mathcal{M}$ be an atomic matroid congestion game with partition matroid $M = (E, \mathcal{I})$. Then Conjecture 4.6 holds, i.e. if g and g' are integral strategy profiles satisfying $g' = g + \chi_x - \chi_y$ for some $x, y \in E$, then $g' \in NBD(g)$.*

Proof. By Lemma 4.8, there exists a decomposition $\mathcal{D}$ of g containing B where $y \in B, x \notin B$. We now argue that x and y must be in the same partition, say E_i. Suppose not, and $x \in E_i$, but $y \notin E_i$. Then, we have $g'(E_i) = g(E_i) + 1 = kr(E_i) + 1$, where the last equality is because every basis must contain $r(E_i)$ elements from E_i. But then $g' \notin P_{k,r}$, which yields a contradiction.

If $x, y \in E_i$ for the same i, then create $B' = B + x - y$, and clearly $B' \in \mathcal{B}$ because it has the same number of elements as B in every partition. Now, create $\mathcal{D}' := (\mathcal{D} \setminus B) \cup B'$, which is a decomposition of g'. $\mathcal{D}$ and $\mathcal{D}'$ only differ with respect to B and B', so $g' \in NBD(g)$ as desired. □

4.4 Proof of Conjecture 4.6 for Laminar Matroids

A family of sets $\mathcal{F}$ is called a *laminar family* if for any two sets $A, B \in \mathcal{F}$, either $A \cap B = \emptyset$, $A \subseteq B$, or $B \subseteq A$. A *laminar matroid* $M = (E, \mathcal{I})$ over a ground set E is defined in terms of a laminar family $\mathcal{F}$ over E, *capacities* $\{cap(A)\}_{A \in \mathcal{F}}$, and has independent sets given by

$$\mathcal{I} = \{S \subseteq E : |S \cap A| \leq cap(A), \forall A \in \mathcal{F}\}.$$

We say that a set S exceeds the capacity of $A \in \mathcal{F}$ if $|S \cap A| > cap(A)$. Conversely, if $|S \cap A| < cap(A)$, then we say S does not utilize the full capacity of A. In the remainder of this section, $M = (E, \mathcal{I})$ is fixed to be the laminar matroid defined by the laminar family $\mathcal{F}$ and capacities $\{cap(A)\}_{A \in \mathcal{F}}$.

For sets $A, B, C \in \mathcal{F}$, if $C \subset A$, then we say C is a descendant of A. If no B exists so that $C \subset B \subset A$, then C is a child of A. Note that by definition S is only independent if it doesn't violate the capacity of A, so it follows that $cap(A) \geq r(A)$ for sets $A \in \mathcal{F}$. However, while rank and capacity are similar for sets $A \in \mathcal{F}$, it's not the case that $cap(A) = r(A)$, since it may be impossible to fully utilize the capacity of A without violating the capacity constraints of its children. An element $a \in A$ is called a *ring element* of A if none of the children of A contain a.

To aid us in the upcoming proofs, we will introduce another general definition that exists outside of the laminar framework. We say a set $A \subseteq E$ is *g-tight* if $g(A) = kr(A)$. One key observation to make is that if $A \subseteq E$ is g-tight, then if $x \in A$, $y \in A$. Otherwise, in g' we'd have $g'(A) = g(A) + 1 > kr(A)$, meaning that g' is not decomposable.

Before showing that Conjecture 4.6 holds for laminar matroids, we first state a few properties of circuits found in Cook et al.

Proposition 4.10. [8] Let J be an independent set, and $e \notin J$. Then $J \cup \{e\}$ contains at most one circuit.

Proposition 4.11. [8] Let J be a dependent set containing a unique circuit C. Then, for any $e \in C$, we have that $J \setminus \{e\}$ is independent.

We now introduce a few tools admitted by the structure of laminar matroids that will aid us in the proof.

Lemma 4.12. *Let $C = B \cup \{e\}$, where B is a basis of the laminar matroid $M = (E, \mathcal{I})$, and $e \notin B$. Let T be a minimal set of $\mathcal{F}$ such that $|C \cap T| > cap(T)$. Then, $C \cap T$ is a circuit of T.*

Proof. We first observe that T is well defined since C is a dependent set. We need to show that: (i) $C \cap T$ is a circuit, and (ii) $|C \cap T| = r(T) + 1$. Note that (ii) follows easily from our assumptions. We have $|C \cap T| > \cap(T) \geq r(T)$, and since C is a basis plus an element, we also have $|C \cap T| \leq r(T) + 1$. Hence, we have $|C \cap T| = r(T) + 1$.

We now argue that $C \cap T$ is a circuit, i.e., the deletion of an arbitrary $e' \in C \cap T$ yields an independent set. Let $B' = C \cap T \setminus \{e'\}$. Consider any set A of the laminar family.

- If $A \supseteq T$, then $|B' \cap A| < |C \cap A| \leq r(A) + 1$, where the last inequality is because C consists of a basis plus an element. Hence, $|B' \cap A| \leq r(A) \leq cap(A)$.
- If $A \subset T$, then $|C \cap A| \leq cap(A)$ by the minimality of T, so $|B' \cap A| \leq cap(A)$.
- If $A \cap T = \emptyset$, then $B' \cap A = \emptyset$ since $B' \subseteq T$.

So $|B' \cap A| \leq cap(A)$ for all $A \in \mathcal{F}$, which means that B' is independent. Hence, $C \cap T$ is a circuit. □

Let $M_{|S}$ be the matroid M restricted to S, meaning that the independent sets of $M_{|S}$ are the independent sets of M that are contained in S. In the upcoming proofs, we will slightly abuse notation and say that a set $A \subseteq E$ is independent in $M_{|T}$ to mean that $A \cap T$ is independent in $M_{|T}$, i.e., that $A \cap T$ is independent in M.

Lemma 4.13. *Let $C \subseteq E$, and let $T \in \mathcal{F}$ be such that $C \cap T$ is a circuit of T. Let $B \in \mathcal{I}$ satisfy $|B \cap T| < r(T)$. Let $e \in C \cap T$. Then there exists $b \in ((C \setminus B) \cap T) \setminus \{e\}$ such that $C \setminus \{b\}$ is a basis of $M_{|T}$, and $B \cup \{b\}$ is independent in $M_{|T}$.*

Proof. First, note that since $C \cap T$ is a circuit of T, we have $|C \cap T| = r(T) + 1$ and therefore $|(C \cap T) \setminus \{e\}| = r(T)$. B is independent in M, so B is also independent in $M_{|T}$. Since $|B \cap T| < r(T) = |(C \cap T) \setminus \{e\}|$, so we can use the properties of independent sets to find $b \in ((C \cap T) \setminus \{e\}) \setminus (B \cap T)$ such that $B \cup \{b\}$ remains independent in $M_{|T}$. Since $C \cap T$ is a circuit, and $b \in C \cap T$, we have that $(C \cap T) \setminus \{b\} = (C \setminus \{b\}) \cap T$ is independent. By definition, this means that $C \setminus \{b\}$ is independent in $M|_T$. $|(C \setminus \{b\}) \cap T| = r(T)$, so $C \setminus \{b\}$ is a basis of $M_{|T}$, as desired.

As an additional remark, if C was initially a basis plus one element, $B' := C \setminus \{b\}$ is not only a basis of $M_{|T}$, but also a basis of M. If C is a basis plus an element, then C contains a unique circuit, and this circuit is therefore precisely $C \cap T$. Since $b \in C \cap T$, it follows that B' is independent. Also, $|B'| = r(E)$ since $|C| = r(E) + 1$. Hence, B' is a basis of M.

□

To prove Conjecture 4.6 for laminar matroids, we will again lean on Lemma 4.8. However, unlike the partition matroid case, an arbitrary decomposition of g does not necessarily provide a suitable basis B containing y but not x such that $(B \cup \{x\}) \setminus \{y\}$ is also a basis. Our alternative strategy will be to take a starting decomposition of g, and create a decomposition of $g^* = g + \chi_x = g' + \chi_y$ by adding x to B_j to create some B^* containing both x and y. We will modify this decomposition of g^* algorithmically until we arrive at a point where it is possible to delete a copy of x or y from B^* to return valid neighboring decompositions of g and g' respectively.

Theorem 4.14. *Let $\mathcal{M}$ be an atomic matroid congestion game with laminar matroid $M = (E, \mathcal{I})$. Then Conjecture 4.6 holds, i.e. if g and g' are integral strategy profiles satisfying $g' = g + \chi_x - \chi_y$ for some $x, y \in E$, then $g' \in NBD(g)$.*

Proof. Let $g^* = g + \chi_x = g' + \chi_y$. By Lemma 4.8, there exists a decomposition of g into k bases, one of which is B_j satisfying $y \in B_j$, $x \notin B_j$. Add x to B_j to create C_0. Now, replacing B_j with C_0 in the decomposition of g yields a decomposition of g^* containing $k-1$ bases and C_0, which is a basis plus an element.

We will provide an algorithmic method that maintains a decomposition of g^*, denoted $\mathcal{D}^*$, satisfying various invariants. The decomposition $\mathcal{D}^*$ will contain sets with two special properties. First, at the beginning of the i^{th} iteration, $\mathcal{D}_i^*$ contains $k-1$ bases and one other set C_i, which is a basis plus one element. Second, there exists $B_i^* \in \mathcal{D}_i^*$ containing both x and y. Note that C_i and B_i^* are not necessarily distinct; initially, we have $C_0 = B_0^*$.

In the following algorithm, we give an iterative process that will eventually terminate, returning neighboring decompositions $\mathcal{D}$ and $\mathcal{D}'$ of g and g' respectively. At each step in the algorithm, we will consider a set $T_i \in \mathcal{F}$, and perform certain actions depending on whether $y \in T_i$. The set T_i is obtained via Lemma 4.12: it is a minimal set of $\mathcal{F}$ such that $|C_i \cap T_i| > cap(T_i)$.

1. **Looping.** As long as $y \notin T_i$, we perform the following steps.
   (a) T_i is not g-tight, so find $B_i \in \mathcal{D}_i^*$ such that $|B_i \cap T_i| < r(T_i)$.
   (b) Use Lemma 4.13 to find appropriate b to send from $C_i \cap T_i$ to B_i. If $x \in C_i$, set $e = x$ in the lemma to ensure that $b \neq x$. Otherwise, e can be an arbitrary element in $C_i \cap T_i$.
   (c) Create $C_{i+1} := B_i \cup \{b\}$, and $B_b = C_i \setminus \{b\}$. Set $B_{i+1}^* := B_i^*$, unless $B_i^* = C_i$, in which case set $B_{i+1}^* := B_i^* \setminus \{b\}$. Create $\mathcal{D}_{i+1}^*$ by replacing B_i and C_i in $\mathcal{D}^*{}_i$ with B_b and C_{i+1}.
   (d) C_{i+1} is a basis plus an element, so use Lemma 4.12 with C_{i+1} to find $T_{i+1} \in \mathcal{F}$ such that $C_{i+1} \cap T_{i+1}$ is a circuit of T_{i+1}.
2. **Termination.** If $y \in T_i$, then we construct neighboring decompositions of g and g'.
   (a) If $B_i^* = C_i$, construct $B_x = B_i^* \setminus \{x\}$ and $B_y = B_i^* \setminus \{y\}$. Return $\mathcal{D} = (\mathcal{D}_i^* \setminus \{B_i^*\}) \cup \{B_x\}$, $\mathcal{D}' = (\mathcal{D}_i^* \setminus \{B_i^*\}) \cup \{B_y\}$.
   (b) If $B_i^* \neq C_i$, find element $r \in C_i \setminus B_i^*$ such that the unique circuit in $B_i^* \cup \{r\}$ is contained in T_i or some ancestor of T_i. Create $B_x := (B_i^* \cup \{r\}) \setminus \{x\}$, $B_y := (B_i^* \cup \{r\}) \setminus \{y\}$, and $B_r := C_i \setminus \{r\}$.
   Return $\mathcal{D} = (\mathcal{D}_i^* \setminus \{C_i, B_i^*\}) \cup \{B_r, B_x\}$, and $\mathcal{D}' = (\mathcal{D}_i^* \setminus \{C_i, B_i^*\}) \cup \{B_r, B_y\}$.

There are several things that remain to be shown. We argue all of the following:

1. The looping process preserves our invariants.

- If $\mathcal{D}_i^*$ is a decomposition of g^*, so is $\mathcal{D}_{i+1}^*$. During the i^{th} iteration, B_i and C_i were the only sets modified, and we only exchanged elements between them to create B_b and C_{i+1}. Therefore, the number of copies of each element in the decomposition remains unchanged, so $\mathcal{D}_{i+1}^*$ is a decomposition of g^*.
- We maintain $k-1$ bases and one basis plus an element in $\mathcal{D}_{i+1}^*$. B_i was originally a basis, and had an element added to it to create C_{i+1}, a basis plus an element. C_i was initially a basis plus an element, and by Lemma 4.13, $B_b = C_i \setminus \{b\}$ is a basis.
- B_{i+1}^* still contains both x and y. Note that B_{i+1}^* is B_i^* or $B_i^* \setminus \{b\}$. In the latter case, we have $b \neq x$, since we invoke Lemma 4.13 taking $e = x$; also $b \neq y$, since $b \in T_i$, and $y \notin T_i$.

2. The looping process terminates.

   - We argue that $T_{i+1} \supset T_i$. Notice that $b \in T_i$, and b must be present in the circuit $T_{i+1} \cap C_{i+1}$, since $C_{i+1} \setminus \{b\} = B_i$, a basis. So $b \in T_{i+1}$, implying that $T_{i+1} \cap T_i \neq \emptyset$. Additionally, $T_{i+1} \not\subseteq T_i$, since C_{i+1} is independent in $M_{|T_i}$ and $C_{i+1} \cap T_{i+1}$ is a circuit. Therefore, since $T_i, T_{i+1} \in \mathcal{F}$, $T_{i+1} \supset T_i$. As long as we continue to loop, we will find larger and larger T_i until we eventually find some T_i containing y.

3. We can find a satisfactory element r in the last termination case 2(b).

   - If $B_i^* \neq C_i$, it is a basis. We claim that $C_i \setminus B_i^*$ contains a ring element of some set R that is T_i or a descendant of T_i. $C_i \cap T_i$ is a circuit and B_i^* is a basis, so $|C_i \cap T_i| > |B_i^* \cap T_i|$. Partition T_i into its children and $ring(T_i)$. By the pigeonhole principle, C_i will have a larger intersection with one of these sets. If that set is $ring(T_i)$, we are done. Otherwise, take the child that C_i has larger intersection with, and consider its ring elements and children. Continuing this process will either directly lead to a set of ring elements that C_i has larger intersection with, or lead to a set with no descendants, in which case all of its elements are ring elements.
   - Select an abitrary ring element $r \in C_i \setminus B_i^*$ from R, and create $B_i^* \cup \{r\} =: C_r$. By Proposition 4.10, C_r contains a unique circuit because B_i^* is a basis. We claim that this circuit is not contained in any descendant of T_i. To see this, observe that for every $S \in \mathcal{F}$ satisfying $R \subseteq S \subset T_i$, $|C_i \cap S| > |B_i^* \cap S|$. By Lemma 4.12, T_i is a minimal set of $\mathcal{F}$ such that $|C_i \cap T_i| > cap(T_i)$. Therefore, $|C_r \cap S| \leq |C_i \cap S| \leq cap(S)$ for every such set S, showing that that $C_r \cap S$

is independent. Since the circuit must contain r and r is a ring element of R, the circuit also cannot be contained in the descendants of R, and therefore the circuit in C_r must be contained in T_i or some ancestor of T_i.

- Both x and y are present in the unique circuit contained in C_r. By Lemma 4.12 4.10, there is some $A \in \mathcal{F}$ such that $C_r \cap A$ is a circuit. We have argued already that A is T_i or some ancestor of T_i. Therefore, since $\{x, y\} \subseteq C_r \cap T$, it follows that x and y belong to the unique circuit in C_r. Now create $(B_i^* \cup \{r\}) \setminus \{x\} = C_r \setminus \{x\} =: B_x$ and $(B_i^* \cup \{r\}) \setminus \{y\} = C_r \setminus \{y\} =: B_y$. By construction, both B_x and B_y are bases since removing either x or y from C_r creates an independent set, and B_x and B_y are the same size as B_i^*. Similarly, B_r is also a basis since r is obviously present in the unique circuit contained in C_r.

4. $\mathcal{D}$ and $\mathcal{D}'$ are decompositions of g and g' into k bases that neighbor each other.

   - $\mathcal{D}$ and $\mathcal{D}'$ contain k bases. If we terminate in step 2(a) with $B_i^* = C_i$, $\mathcal{D}_i^*$ contains $k-1$ bases and B_i^*, which is exchanged for either the basis B_x or B_y, so $\mathcal{D}$ and $\mathcal{D}'$ contain k bases. If we terminate in step 2(b), $\mathcal{D}_i^*$ contains $k-2$ bases, along with B_i^* and C_i. Both of these sets are swapped for bases, so again we have k bases.
   - $\mathcal{D}$ and $\mathcal{D}'$ are decompositions of g and g', respectively. Note that $\mathcal{D}$ contains one less copy of x compared to the decomposition $\mathcal{D}_i^*$, and $\mathcal{D}'$ contains one less copy of y compared to $\mathcal{D}^*$. Since $\mathcal{D}^*$ is a decomposition of $g^* = g + \chi_x = g' + \chi_y$, it follows that $\mathcal{D}$ and $\mathcal{D}'$ are decompositions of g and g' respectively.
   - In both cases, $\mathcal{D} \setminus \{B_x\} = \mathcal{D}' \setminus \{B_y\}$, so the decompositions neighbor each other.

□

Chapter 5

Conclusions

In this thesis, we studied the question of how to add tolls to atomic unsplittable congestion games to induce a desired flow f as the unique atomic Nash equilibrium. To make this question precise, we introduced the idea of profile-decomposition independent equilibria, which are congestion vectors that maintain their equilibrium properties regardless of how they are decomposed into player strategies. Therefore, the precise question that we sought to answer is "given a target flow vector f, can we find tolls that impose f not only as a PDI NE, but also as the unique PDI NE?"

We obtained results for both network routing games and matroid congestion games. For network routing games, we show that in series-parallel graphs, the exact same tolls that suffice in the nonatomic setting also induce f as the unique PDI NE in the atomic setting. However, we supply an example (a generalization of Braess's Paradox) that shows that the nonatomic tolls do not suffice in general, and show that for general linear cost functions, even a perturbation of the nonatomic tolls is not always sufficient. At the same time, we provide positive results on how to compute tolls in this example. Whether or not tolls exist (and can be computed) for general atomic network routing games remains an open question. One key structural aspect of our four-link example is that each path contains an edge not utilized in any other path. Crucially, this guarantees that tolls can adjusted to change the cost of any one particular path while maintaining the cost of all other strategies. The class of graphs exhibiting this property may be a natural target for expanding our results.

For matroid congestion games, we show that our question regarding tolls can be reduced to a question about the structure of the polymatroid associated with the game. Essentially, if neighboring points in the polymatroid have neighboring basis decompositions, then the nonatomic tolls will suffice for the atomic game as well. We show that this polymatroid structure does indeed hold for laminar matroids, leading to positive results on tolls for laminar matroid games. While some of our tools do work in the general matroid setting, the existence and computability of tolls in general matroid congestion games remains an open question.

References

[1] B. Awerbuch, Y. Azar, and A. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 57–66, 2005.

[2] M. Beckman, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[3] U. Bhaskar, L. Fleicher, and C. Huang. The price of collusion in series-parallel networks. In *Proceedings of the 14th IPCO*, pages 313–326, 2010.

[4] I. Caragiannis, C. Kaklamanis, and Kanellopoulos. P. Taxes for linear atomic congestion games. In *Proceedings of the 14th Annual European Symposium on Algorithms*, pages 184–195, 2006.

[5] G. Christodoulou and E. Koutsoupias. The price of anarchy in finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing*, pages 67–73, 2005.

[6] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. *STOC 2003*, pages 521–530, 2003.

[7] R. Cominetti, J. R. Correa, and N. Stier-Moses. The impact of oligopolistic competition in networks. *Operations Research*, 57(6):1421–1437, 2009.

[8] William J. Cook, William H. Cunningham, William R. Pulleyblank, and Alexander Schrijver. *Combinatorial Optimization*. John Wiley & Sons, Inc., New York, NY, USA, 1998.

[9] S. C. Dafermos. Toll patterns for multiclass-user transportation networks. *Transportation Science*, 7:211–223, 1973.

[10] S. C. Dafermos and F. T. Sparrow. The traffic assignment problem for a general network. *Journal of Research of the U.S. National Bureau of Standards 73B*, pages 91–118, 1969.

[11] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proceedings of the 45th Annual IEEE Symposium on Foundations of Computer Science, held 2004 in Rome, Italy*, pages 277–285, 2004.

[12] T. Harks. Stackelberg strategies and collusion in network games with splittable flow. In *Proceedings of the 6th WAOA*, pages 133–146, 2008.

[13] Ara Hayrapetyan, É. Tardos, and T. Wexler. The effect of collusion in congestion games. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 89–98, 2006.

[14] K. Murota. *Discrete Convex Analysis*. Society for Industrial and Applied Mathematics, 2003.

[15] Noam Nisan, Tim Roughgarden, Éva Tardos, and Vijay V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.

[16] C. H. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.

[17] A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.

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

[19] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[20] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[21] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49:236–259, 2002.

[22] M. J. Smith. The marginal cost taxation of a transportation network. *Transportation Research Ser. B*, 13:237–242, 1979.

[23] C Swamy. The effectiveness of stackelberg strategies and tolls for network congestion games. *ACM Transactions on Algorithms*, 8(4):36, 2012.

[24] J. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Part II, 1*, pages 325–378, 1952.

[25] H. Yang and X. Zhang. Existence of anonymous link tolls for system optimum on networks with mixed equilibrium behaviors. *Transportation Research B*, 42:99–112, 2008.