
**Best Response Dynamics in
Simultaneous and Sequential Network
Design Games**

Matthew Milan Radoja

Thesis submitted for the degree of
Doctor of Philosophy



Department of Informatics
University of Leicester

February 2017

Abstract

This thesis is concerned with the analysis of best response dynamics in simultaneous and sequential network design games with fair-cost sharing, for both capacitated and uncapacitated networks. We address questions related to the evolution of stable states through selfish updates. First, we examine in general what effects such updates can have, from various perspectives, on the quality of the solutions to a game. From this, we move on to a more specific analysis of updates which begin from an optimal profile, providing insight to the price of stability measure of network efficiency, from the perspective of the user incurring the highest cost in the game. Finally, we investigate the process of updates beginning from an empty strategy profile, and make some observations about the quality of the resultant profile in such situations.

Acknowledgements

First and foremost I thank my supervisor, Thomas Erlebach, without whose support and guidance this thesis would not be. Secondly I thank my parents, without whom I would not be. I also express my gratitude to the coffee growers and tobacco farmers of the world, for their part in reversing my tendency for inertia. Last of all, but by no means least of all, I thank Maria Psyllou, for being there for me and for making me smile.

Contents

1	Introduction	1
1.1	Game Theory	2
1.2	Network Design	4
1.3	Literature Review	5
1.3.1	Algorithmic Game Theory	5
1.3.2	Network Design Games	6
1.4	Our Contribution	8
1.5	Outline of Thesis	10
2	Background	11
2.1	Model	11
2.2	Definitions	12
2.3	Network Topology	14
3	Best Response Dynamics	17
3.1	Introduction	17
3.2	Preliminary and Known Results	19
3.2.1	Effect on the Sum Cost of a Profile	21
3.2.2	General Games	23
3.2.3	Symmetric Uncapacitated Networks	23
3.3	Convergence to Nash Equilibria	24
3.3.1	Uncapacitated Games	26
3.3.2	General Games	28
3.4	Effect on the Maximum Cost of a Profile	28
3.4.1	Uncapacitated Games	29
3.4.2	General Games	30
3.5	Effect on the Cost to an Individual	36
3.5.1	Uncapacitated Games	36
3.5.2	General Games	38
3.6	Concluding Remarks	40
4	Price of Stability With Respect to Max Cost	42
4.1	Introduction	42
4.2	Preliminary and Known Results	46
4.3	General Games	49

4.3.1	Asymmetric Games	50
4.3.2	Rooted Games	57
4.4	Uncapacitated Games	59
4.5	Concluding Remarks	60
5	Sequential Network Design Games	62
5.1	Introduction	62
5.2	Existence of Reachable Nash Equilibria	65
5.2.1	General Games	65
5.3	Price of Anarchy for Sequential Games	66
5.3.1	General Games	66
5.3.2	Uncapacitated Games	67
5.4	Price of Stability for Sequential Games	70
5.4.1	General Games	70
5.4.2	Uncapacitated Games	70
5.5	Concluding Remarks	84
6	Conclusion and Future Research	86
6.1	Summary of the Contribution	86
6.2	Directions for Future Research	88

List of Figures

1.1	Pay-off matrix for the prisoners' dilemma.	3
3.1	Symmetric game Δ where BRD increases sum-cost by a factor $O(\log n)$	22
3.2	Two-player game Δ_u where BRD convergence is unbounded in n	27
3.3	Rooted general game where effect of BRD on max cost $\Omega(n)$	30
3.4	Asymmetric game Δ where BRD on max cost is $n \log n$	32
3.5	Rooted game Δ_u where an individual player experiences an increase in cost of a factor n	38
3.6	Asymmetric game Δ where $m = 3$ and an individuals cost increases by an arbitrary factor w.r.t n	39
4.1	Examples showing lower bounds for PoA_{sc} in Δ and Δ_u	48
4.2	Asymmetric game Δ showing $\text{PoS}_{mc}(\Delta) \geq \Omega(n \log n)$	52
4.3	Rooted game Δ_u showing $\text{PoS}_{mc}(\Delta_u) \geq \Omega(n)$	59
5.1	Symmetric game Δ where no feasible sequential solution exists.	66
5.2	Symmetric sequential game Δ where PoA w.r.t sum-cost and max- cost unbounded in n	67
5.3	Rooted game Δ_u for $n = 2$ where sequential PoA is $8/5$	68
5.4	Asymmetric sequential game Δ_u where PoS_{sc} is $\Omega(\sqrt{n})$	71
5.5	Illustration of sets of intervals which are agreeable and nested.	74
5.6	Example for the construction of the tree T used in the proof of The- orem 16.	79
5.7	Illustration of the sum of distances between a set of intervals which is agreeable	80
5.8	Illustration of the sum of distances between a set of intervals which is nested	80

List of Tables

1.1	Summary of results related to best response dynamics.	8
1.2	Summary of results for the price of stability with regard to the maximum cost to any player.	9
1.3	Summary of results for the sum-cost sequential price of anarchy.	9
3.1	Worst case increase of sum cost as a result of BRD	21
3.2	Worst case convergence of BRD to Nash Equilibria.	26
3.3	Worst case effect of BRD on the max-cost of a profile.	29
3.4	Worst case effect of BRD on the cost of an individual player.	37
4.1	Results related to the Prices of Anarchy and Stability.	44
5.1	Summary of results related to PoA_{sc} for Sequential Network Design Games	65

Chapter 1

Introduction

In this thesis we are interested in the application of game theoretical methods to the modelling and analysis of the behaviour of self motivated agents operating within the framework of a network.

Arguably one of the most influential ideas of the twentieth century, game theory has become the standard modelling technique for situations involving cooperation and/or conflict. The publication, in 1944, of *Games and Economic Behaviour* [50], laid the foundations for both Macroeconomics and Utility theory, and has become the basis for much of modern economic, political, and social science, with ever expanding influence in disciplines ranging from biology to philosophy. In short, game theory can be applied to any situation where more than one agents interact.

Networks are pervasive in society, and the rise of the communication age has propelled humanity forward at a speed far greater than that of the industrial revolution. The many benefits of communication networks include the dissemination of information, pushing society ever closer to the capitalist goal of perfect information.

In this thesis we are interested in applying game theoretical concepts and tools to the analysis of networks. The abstract notion of a network can be applied to many settings in the real world, the more obvious being communication networks such as the internet, and transport systems such as a road or rail network, on which numerous selfish agents compete with the aim of reducing their travel time. Any one of these can be modelled as a collection of agents selfishly operating on a graph, it's edges representing choices, their costs representing the pay-offs of each choice.

Of course, the more complex the situation the more complex the model, and the more intractable the analysis of these situations become. Much of current research in the intersection of game theory and network design is interested in bounding the inefficiency of anarchic situations. The concept of an equilibrium point [50] , allows us to find situations where no agent has incentive to change their strategy. It is these states we are interested in.

1.1 Game Theory

With its origins in economics, it is not surprising that game theory is based on the idea of quantifiable outcomes to decision problems. The general idea is that one can, in any situation where two or more entities are in some sort of competition, create a framework within which to not only analyse the best strategy for an individual, but to mathematically predict which actions will lead to a stable and mutually beneficial situation.

An often used example of the application of game theory to every day situations is the *prisoners' dilemma*, which we will now briefly describe.

Prisoners' Dilemma

Following a recent spate of bank robberies, the police have arrested their two prime suspects, Bonnie (B) and Clyde (C), and proceed to interrogate each separately. While the police have no concrete evidence of the pair's involvement in the crime, their resistance to the arrest and possession of a firearm means they now also face lesser charges with the possibility of up to five years incarceration. Knowing that the maximum sentence for armed robbery is 25 year, each prisoner is confronted with the following dilemma: *confess* and implicate their partner, with the hope of lenient sentencing for themselves, or *deny* all involvement, accepting the lesser charge. The specifics of the offer are that

- If B betrays C , while C refuses to cooperate, B receives just 2 years which C serves the full 25, and vice versa.
- If both choose the betray their accomplice, both receive a sentence of 10 years.

		<i>Bonnie</i>	
		<i>confess</i>	<i>deny</i>
<i>Clyde</i>	<i>confess</i>	10	25
	<i>deny</i>	2	5

Figure 1.1: Pay-off matrix for the prisoners' dilemma.

- If both refuse to talk, the lesser charges of possessing a deadly weapon and the assault of the arresting officer will be pressed, and both will serve 5 years.

The question now, at least from the perspective of the accused, is “Which of these two choices leads to the best outcome?”.

Observe that the pay-offs for the two players of the game depends on the actions of their accomplice. Figure 1.1 shows the pay-offs for each player for all possible outcomes. First consider the jail time associated with the possible outcomes for Bonnie. She has no control over the actions of Clyde, so must choose whether to confess or deny her crimes with the aim of minimising her sentence regardless of his actions. Notice that if Clyde confesses, the best outcome for Bonnie is to also confess and serve 10 year, rather than deny and serve the maximum sentence of 25 year. Notice also that, in the case where Clyde remains silent, the best outcome for Bonnie is still to confess, as doing so will result in just 2 years jail time while a denial will lead to the lesser charges of assault being pressed, the punishment for which is 5 year. By analysing the game in this way, Bonnie can deduce that a confession will in all cases lead to a lesser sentence than denying involvement. As Clyde is faced with the same two options, the pay-offs for which are identical as they were for Bonnie's options, he too has a best choice in confessing.

This example aptly illustrates the power of game theory when it comes to analysing interactive scenarios. For the individual player of the game, it provides a framework by which to decide their strategy; for the external observer it gives predictive power over the outcomes to interactions between multiple agents with

competing and sometimes conflicting interests; and for the designer of the game, it provides the insight to allow for the promotion of whatever objective they choose.

1.2 Network Design

The daily commute, treading a well worn route perfected over successive journeys, is a prime example of selfishness in action. You know from experience which roads to avoid and at what times to do so, and understand that the shortest route may not be the quickest. Every turn you make has one goal, which is to minimise the time spent in your vehicle. As an autonomous user of the road network you have complete control over your actions, and will react to the current situation in a way which most benefits yourself, as do your fellow travellers. But how does this behaviour effect the performance of the overall system? Being selfish one might not give much thought to the experience of other users of the road, but this it to overlook the fundamental nature of the network on which you operate, which is that as your actions effect others, so the actions of others have the power to both shorten and increase your journey time.

It is natural to ask, when considering the fact that each action we make will ripple through the system, what tsunamis might result.

Of course, as we are now addressing questions related to computer science a more apt example for the application of game theory to network design might be a communication network such as the internet, but for an illustration of a situation involving individual agents a concrete example serves better than an abstract one.

Game theory provides tool which allow a much better understanding of the real world performance of networks, and as such provided invaluable insight into one of the most important and widespread developments of modern society. By applying game theoretic techniques to the analysis of networks we have immense power to promote efficiency. In modelling the actions of individual users, we are able to better understand how the system as a whole will behave, and from this gain insight into how to avoid situations where selfishness has a high social cost.

In the following we aim to further the understanding of a small subsection of network design, namely, the analysis of the behaviour of agents in a game which

utilizes the well studied concept of fair cost sharing first introduced by Shapley and Mondorer.

1.3 Literature Review

1.3.1 Algorithmic Game Theory

For a general outline of the applications of game theory in computer science, we direct the reader to [53]. An invaluable resource for the fundamentals, the aforementioned discusses many of the topics explored in this thesis, and so provides comprehensive background reading. We now outline the contents of some of the more relevant chapters, making mention of the original papers on which they are based:

Basic Solution Concepts and Computational Issues [53, Chapter 1]. A more detailed definition of a game than the one provided in this thesis, followed by some simple examples illustrating the applicability of the game concept to every day scenarios. This chapter covers much of what was first introduced here [51], and expanded on here [9].

Complexity of Finding Equilibria [53, Chapter 2]. A discussion of the computational issues arising when searching for stable solutions to games. Some early works on which this chapter is based include [33, 35], while some more recent examinations of the complexity of computing equilibria in multi-player games can be found here [55, 56]. Also of interest is work on reducibility among equilibrium problems [34].

Introduction to Mechanism Design [53, Chapter 9]. Discusses some of the issues related to finding equilibria through social choice. Of interest in the context of Best Response Dynamics, and Reachable Equilibria. For further reading on this, see [52, 8, 43].

Cost Sharing [53, Chapter 15]. Explores some of the properties (fairness, robustness) by which one can judge the efficacy of cost-sharing methods, and provides a detailed description of the Shapley value, on which the cost sharing method used in our model is based. One may find interesting reading on this in [38, 10].

Introduction to the Inefficiency of Equilibria [53, Chapter 17]. This chapter introduces measures of inefficiency for games, providing background on the motivation for doing so, before illustrating these concepts in various models. Of particular relevance to this thesis, we suggest [39, 54, 58, 60] for some background.

Routing Games [53, Chapter 18]. A more detailed examination of the inefficiency of equilibria for routing games, exploring proof techniques used in bounding these measures.

1.3.2 Network Design Games

Inefficiency of Equilibria. The quantification of the inefficiency of Nash equilibria has received considerable attention in recent years. The concept of the *price of anarchy*, measuring the inefficiency of the worst Nash equilibrium (NE) of a given game compared to a social optimum, was introduced by Koutsoupias and Papadimitriou [39], who called it the *coordination ratio*. The *price of stability*, measuring the inefficiency of the best NE of a given game, was first studied by Schulz et al. [60], under the name *optimistic price of anarchy*. Games for which these measures have been studied include scheduling games [39], routing games [59], network design games [6], and capacitated network design games [29].

Fair Cost Sharing. Network design games with fair cost sharing, where the cost of an edge is distributed to all players using the edge in equal shares, were first studied by Anshelevich et al. [6]. They observe that these games are potential games, meaning it is possible to track the changes in cost to an individual when making a change in strategy [49], and therefore always have a NE in pure strategies. They also observe that the process of iterative updates in strategy, or best response dynamics, must converge to such a NE.

Best Response Dynamics. Apart from the study of the inefficiency of NE, one is also interested in the convergence time of best response dynamics (BRD), i.e., the process that starts with an arbitrary strategy profile and iteratively allows one of the players to update her strategy to one that optimises her cost given the current strategies of all the other players.

It is known that in symmetric network design games where there is no limitation on the number of players which may use a particular edge (uncapacitated), BRD

converge to a NE in at most n steps, as the best response for the first player will also be the best response for all other players. In the general case, convergence of symmetric games can be as bad as $n^{3/2}$ given a particular order of updates [29].

Further to this, it has been shown by Anshelevich et al. [6] that convergence in asymmetric uncapacitated games can be exponential in the number of players, but as in general games this only holds if players make their improving moves in a certain order [6].

Price of Stability. The price of stability was first examined in this setting here [6]. For asymmetric, uncapacitated network design games on directed graphs, they show that the price of stability with respect to the total cost to all players (sum-cost) is at most $H(n) = \Theta(\log n)$, where $H(n)$ is the n^{th} harmonic number. They prove this result by considering a potential function that decreases with every improving move of a player and using it to show that BRD from an optimal strategy profile must lead to a NE whose sum-cost is at most $H(n)$ times the sum-cost of the starting profile. We will use the same potential function several times in this paper. They also show that the upper bound of $H(n)$ on the price of stability for sum-cost holds for several generalisations, including capacitated network design games. The price of stability of uncapacitated network design games with respect to sum-cost for *undirected* networks is still open. The best known lower bounds are constant and the best known upper bound is $(1 - \Theta(1/n^4))H(n)$, showing that the maximum price of stability for undirected networks is smaller than it is for directed networks [21].

As already noted in [29], it is easy to see that the price of stability is 1 for both sum-cost and max-cost for symmetric network design games without capacities, since the strategy profile where all players choose the same minimum-cost path from the common source to the common destination is a NE and also the social optimum.

Feldman and Ron [29] present a comprehensive study of symmetric capacitated network design games in undirected networks. They show that the price of anarchy in general networks is unbounded when considering both the total costs paid by all players, and the maximum cost paid by all players, but is bounded by $O(n)$ for parallel links and series-parallel networks. For the price of stability with respect to sum-cost, they show a bound of $O(\log n)$ that is tight even for parallel links. For the price of stability with respect to max-cost, they give tight bounds of $O(n)$ for

	General Games	Uncapacitated Games
Convergence to NE	unbounded	unbounded
Effect on the maximum cost to all players	$\Theta(n \log n)$	$\Theta(n)$
Effect on the cost to an individual	unbounded	$\Theta(n)$

Table 1.1: Summary of results related to best response dynamics.

parallel links and series-parallel networks, but for arbitrary networks their upper bound of $O(n \log n)$ leaves a gap to the lower bound of $\Omega(n)$.

1.4 Our Contribution

We now outline the contributions made in this thesis.

Best Response Dynamics. We begin by examining a process of strategy updating known as best response dynamics (BRD), and show that in the context of our model convergence time to NE is unbounded in the number of players (except in the case of uncapacitated symmetric networks, for which the tight bound of n has existed for some time). We then move on to the analysis of the negative impacts of BRD, i.e. how does selfish behaviour effect the cost of solutions by a number of metrics. We examine it's effects on the maximum cost of a solution, showing respective tight bounds of $\Theta(n \log n)$ and $\Theta(n)$ for the general and uncapacitated cases. From a more individualistic perspective we show that no players' cost can increase by a factor of more than $\Theta(n)$ in the uncapacitated case, while in the general case it is possible for an individual to experience an arbitrary (with respect to the cost of their initial profile) increase in cost as a result of selfish updates. We summarise these results in Table 1.1.

Price of Stability w.r.t Max-cost. Having completed our analysis of BRD, we turn our attention to a measure of efficiency known as the price of stability (PoS), which for a given class of games is the worst case ratio between the best NE of a particular game and the optimum solution for that game. When calculating the cost of solutions we are interested in the maximum cost of a profile. We show that, in the uncapacitated case the price of stability with respect to max-cost is tightly $\Theta(n)$ in both asymmetric and rooted games. For the general case, we find that the

	General Games	Uncapacitated Games
Routed	$\Theta(n)$	\vdots
Asymmetric	$\Theta(n \log n)$	$\Theta(n)$

Table 1.2: Summary of results for the price of stability with regard to the maximum cost to any player. In this context, a rooted game is one where all players share a common source node, and may have distinct end points, and an asymmetric game is one where all players may have a unique source destination pair.

	General Games	Uncapacitated Games
Symmetric	\vdots	1
Routed	\vdots	8/5 for $n = 2$
Asymmetric	unbounded	$\Omega(\sqrt{n}), O(\sqrt{n} \log^4 n)$

Table 1.3: Summary of results for the sum-cost sequential price of anarchy.

price of stability is tightly $\Theta(n \log n)$ for asymmetric games, and $\Theta(n)$ for rooted and symmetric games. These findings are show in Table 1.2.

Sequential Network Design Games. When examining the sequential version of network design games, where the game is initialised with an empty strategy set, and players chooses their joining strategy in turn, so that each player joins with a best response to the existing solution, we find a feasibility issue for general games (in that games exist where there is no solution by which all players form an $s-t$ connection). In uncapacitated games, there must be some solution which can be arrived at by the described method. In the analysis of the efficiency of stable states in such games, we find that in the general case, both the price of anarchy and price of stability measures cannot be bounded by any function on the number of players. In the uncapacitated case, we show an upper bound of $O(\sqrt{n} \log^4 n)$ for the price of anarchy w.r.t maximum cost, coupled with an asymmetric game where the reachable price of stability is $\Omega(\sqrt{n})$. We also show results for games with two players, with a tight bound of 8/5 for both the price of stability and price of anarchy measures. Table 1.3 summarises these findings.

1.5 Outline of Thesis

The remainder of this thesis is structured as follows. In Chapter 2 we give some background information on what is already known, and make some observations about the setting which we use throughout the thesis. In Chapter 3 we investigate the process of BRD, in particular its convergence to NE, and its effect on the quality of the solution with respect to the sum-cost, the max-cost, and the cost to any individual player of the game. In Chapter 4 we examine the Price of Stability of our games, both general and uncapacitated, with respect to the max-cost to all players of the game.

In Chapter 5 we study sequential network design games, exploring both the existence of solutions in this setting, and the quality of the solutions they produce.

We conclude with a summary of the results, and a discussion of their implications for further research in chapter 6.

Chapter 2

Background

In this chapter we will give detailed definitions of the terminology and concepts discussed throughout this thesis. We start by defining the class of games we examine, and discuss the main points related to them. We then categorise networks based on their topology, and discuss the applicability of results to these different categories. This is followed by a literature review, first outlining some of the more general related work, and then discussing those specifically relating to the topics explored in this thesis.

2.1 Model

Capacitated Network Design Games. We consider capacitated network design games, also known as capacitated cost sharing (CCS) games [29] and referred to as *general games* throughout this thesis. These games are discrete. All players, who we will also refer to as agents, have perfect knowledge of the choices available to them, as well as the associated cost of each of these choices with respect to the current choices of their opponents. The aim of the game is as follows: for some directed or undirected graph $G = (V, E)$, each player i in a set of n must establish a connection between their *source* and *sink* nodes, represented as s_i and t_i respectively. Every edge $e \in E$ has cost $p(e) \in \mathbb{R}^{\geq 0}$ and capacity $c(e) \in \mathbb{N}$. We also write p_e for $p(e)$ and c_e for $c(e)$. Let $[n]$ denote the set $\{1, 2, \dots, n\}$. This represents the set of all players. We will sometimes refer to individual players by a letter, e.g. player 2 as

player b , when doing so is clearer. The game can be represented as the tuple

$$\Delta = \langle n, G = (V, E), \{s_i\}_{i \in [n]}, \{t_i\}_{i \in [n]}, \{p_e\}_{e \in E}, \{c_e\}_{e \in E} \rangle.$$

Uncapacitated Network Design Games. A special case of this model, where the capacity of each $e \in E$ is infinite (or at least as great as the number of players), has been widely studied. As the properties of these games differ from the general model defined above, it is necessary to differentiate between the two; we will refer to these games as *uncapacitated games*, and use the notation Δ_u to refer to the subset of Δ where $c_e \geq n$ for all $e \in E$.

2.2 Definitions

Strategies. Any agent i may have several options when forming their s - t path in G ; we will refer to their chosen path as their *strategy*, denoting this connection S_i ; the set of all available paths for this player we denote by Σ_i and refer to as their *strategy set*. In general we will refer to the strategy of an agent by the vertices it crosses, although occasionally we will talk about the edges this sequence consists of if doing so is clearer. As such, a strategy S_i crossing vertices x and y may be written as (s_i, x, y, t_i) or $((s_i, x), (x, y), (y, t_i))$. The strategies of all n players is the *strategy profile*, and represents a solution to the game. Profiles are denoted as S , where $S = (S_1, S_2, \dots, S_n)$. The strategy set of all players, which is the set of all possible combinations of s - t paths for all players is written as Σ . When talking about solutions it is sometimes necessary to know the actions or choices of a player's opponents. For some profile S , the actions of all players but i is shown as S_{-i} , and the strategy set of this player's opponents is written as Σ_{-i} .

Cost of Strategy. The price of an edge is shared evenly between all players in whose strategy it falls. This fair cost division scheme is derived from the Shapley value, and is one of the most widely studied protocols [49]. The price of an individual's strategy S_i , with respect to the strategy profile S , is defined as

$$p_i(S) = \sum_{e \in S_i} \frac{p_e}{x_e(S)}, \quad (2.1)$$

where $x_e(S) = |\{i : e \in S_i\}|$ denotes the number of agents that use e in their path.

Cost of Strategy Profile. We consider two social cost functions: the *sum-cost*

of a profile S , denoted by $sc_\Delta(S) = \sum_{i \in [n]} p_i(S)$, is the sum of all edges used and the total cost to all agents in S , while the *max-cost* of a profile S , denoted by $mc_\Delta(S) = \max_{i \in [n]} p_i(S)$, is the maximum cost of any agent in S . We omit the subscript Δ if the game is clear from the context.

Feasible Games. As we are considering capacitated networks, a feasibility issue arises. A strategy profile $S = (S_1, \dots, S_n)$ is *feasible* if $x_e(S) \leq c_e$ for all $e \in E$. Throughout this thesis we will only consider feasible games, i.e., games for which there is at least one feasible strategy profile.

Updates in strategy. The game theoretic take on this model assumes that all players will change their strategy if doing so benefits them. We refer to an agent's change in strategy as an *update*. Players being selfish we assume no collaboration, and say that updates in strategy must be sequential. For a given profile S a player may have several paths which represent an improving move. We assume that a rational self-motivated individual would always choose the best, that is, cheapest, option available to them. An update by i , from this profile S , will yield a new strategy profile, and we can represent this change with the notation $S \xrightarrow{i} S'$, where S' is the strategy profile arrived at after i 's update: $S' = S_{-i} \cup S'_i$.

Best Response Dynamics. We assume agents have full knowledge of the paths available to them, as well as their opponents' strategies, so they know the cost of all alternatives with respect to S_{-i} . Being self-motivated, players will update their strategies to the cheapest path available at any given point in what is known as *best response dynamics* (BRD). We do not specify the order in which updates are made, only that they are sequential and that the choice of strategy of the player making the update must be the best response to her opponents' current strategies.

Nash Equilibria. The concept of an equilibrium is key to game theory, and the analysis of these states is central in current research. A profile S is said to be a *Nash equilibrium* (NE) if no agent can improve their cost by a unilateral deviation from the profile, that is, for every player i we have that for all s_i - t_i paths S'_i , it holds that $p_i(S) \leq p_i(S'_i, S_{-i})$.

Existence of Nash Equilibria. The CCS games we consider fall into the class of *congestion games* studied by Monderer and Shapley [49], who show that all such games have pure Nash equilibria. They do this by defining a potential function Φ ,

which in the context of our model is

$$\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{i} \quad [57]. \quad (2.2)$$

Note that $\Phi(S)$ is bounded by $H(n)$ times the sum-cost of S , where $H(n) = \sum_{i=1}^n 1/i$ denotes the n -th harmonic number, and is asymptotically equal to $\log n$. As players only make improving moves, best response dynamics will strictly reduce the potential of the solution with each step, meaning a profile cannot be revisited. As the strategy space of a game is finite, any sequence of updates will terminate at a profile where no player can make a unilateral improvement, which must be a NE.

Quality of Nash Equilibria. For games where feasible solutions exist, when measuring the quality of a NE we will compare its cost, by either the sum-cost or max-cost objective, to that of the optimal solution. The ratio between the objective value of the worst NE and the optimal objective value is called the *price of anarchy*, while the ratio between the objective value of the cheapest NE and the optimal objective value is called the *price of stability*, abbreviated to PoA and PoS, respectively.

We refer to the optimal objective value with respect to max-cost as OPT_{mc} , and that with respect to sum-cost as OPT_{sc} . Furthermore, we write $PoS_{mc}(\Delta)$ for the price of stability with respect to max-cost, and similarly for the other cases. For a particular CCS game Δ whose set of Nash equilibria is denoted by $NE(\Delta)$, the prices of anarchy and stability with respect to max-cost are defined as

$$PoA_{mc}(\Delta) = \frac{\max_{S \in NE(\Delta)} mc_{\Delta}(S)}{OPT_{mc}(\Delta)} \quad PoS_{mc}(\Delta) = \frac{\min_{S \in NE(\Delta)} mc_{\Delta}(S)}{OPT_{mc}(\Delta)}$$

with analogous calculations for sum-cost.

2.3 Network Topology

The behaviour of the agents within a game is dictated by the topology of the network on which it is played. We will classify games by the properties of the underlying network.

Undirected Graphs. Throughout this thesis we examine games in which the underlying graph is undirected, i.e. edges may be traversed in both directions. We

refer to these constructions as *undirected graphs*. A *directed graph* is one where each edge can only be used in one direction.

Note that an undirected graph G can be transformed into an equivalent directed graph G' using the following well known construction: Every undirected edge $\{u, v\}$ of G is replaced by the directed edges (u, x_1) , (v, x_1) , (x_1, x_2) , (x_2, u) , and (x_2, v) , where x_1 and x_2 are two new nodes created for the transformation of $\{u, v\}$. The capacity and cost of (x_1, x_2) are set equal to those of $\{u, v\}$, the remaining edges have infinite capacity and cost 0.

As a consequence of this transformation, any construction of undirected general games establishing a lower bound on the price of stability (or on the convergence time of BRD) automatically yields an equivalent construction of directed general games. Similarly, any upper bound on the price of stability proved for directed general games automatically yields the same upper bound for undirected general games. When it is clear from the context that we are considering undirected graphs, we also write undirected edges in the form (u, v) instead of $\{u, v\}$.

Asymmetric Games. The most general networks we examine are those where every player may have a distinct source and destination. We refer to these games as *asymmetric*, in reference to the asymmetry of the players' strategy sets.

Rooted Games. We refer to games where all players have a shared sink node as *rooted*. In rooted games each player may have a distinct destination. Note that all rooted games are also asymmetric games, i.e. rooted games form a subset of all asymmetric games.

Symmetric Games. We call games where all players share the same source and destination nodes *symmetric*. Note that, all symmetric games are also rooted (the former being a subset of the latter), and that in these games the strategy set for each player is identical.

Symmetric to Rooted to Asymmetric. Any symmetric network can be transformed into a rooted network by adding distinct sources for each player, and connecting these to the common source with a single zero cost edge. Any rooted network can be made asymmetric by adding distinct destinations for each player and connecting each of these to the common destination with zero cost edges.

Observation 1. *A lower bound for a game where the underlying network is symmetric applies to the rooted case.*

Observation 2. *A lower bound for a game where the underlying network is rooted applies to the asymmetric case.*

An uncapacitated network can be seen as a special case of a capacitated network, and so the class of uncapacitated games is a subset of capacitated games. Any uncapacitated network can be made capacitated, by giving each edge a capacity equal to the number of players.

Observation 3. *An upper bound for the class of capacitated games applies to all uncapacitated games, as the latter is a subset of the first.*

Observation 4. *A lower bound in the uncapacitated case will apply to the capacitated case.*

Chapter 3

Best Response Dynamics

In this chapter we will examine a strategy updating rule known as *best response dynamics*, or BRD, and explore some of the open questions related to the number of updates possible before a stable state is reached, as well as quantifying the negative effects of such updates from a global and individual perspective.

3.1 Introduction

A key aspect of the game-theoretic take on network design and analysis is the study of the mechanisms through which the games, or the solutions they give, evolve. One such mechanism, known as best response dynamics, has been the focus of a large amount of research and has produced invaluable insight into the field overall. BRD is the process through which agents make unilateral updates in strategy, in no particular order, with the sole proviso that any update must be the best response to the current profile. We will assume that players have complete knowledge of the paths available to them, that is, all paths in their strategy set, as well as the cost associated with each relative to the current paths of their opponents. This gives us a very natural model in which to examine the effects of selfish behaviour, without restricting the actions of agents within the game.

Limitations. While BRD give some insight into the evolution of stable solutions, it relies on agents responding only to the current state, and so could be considered a myopic approach to NE discovery. That said, while an individual player may have

perfect knowledge of the choices available to their opponent, using this information to develop a long term strategy can be very complicated.

Questions. In congestion games such as ours, the study of BRD is primarily concerned with the discovery of NE. In the following chapter we divide our research into these two questions:

1. What is the worst case convergence of BRD?
2. What are the negative impacts (with regard to price) of BRD?

The motivation for the first question is simple: are BRD a viable tool for discovering NE? The second question is directly related to measures of network efficiency, in that, a bound on the negative impacts of BRD allows us to bound the *Price of Stability* for any given game.

Outline of the remainder of the chapter. In this chapter we examine the questions outlined above.

We start in Section 3.2 by giving some background information and exploring what is already known about BRD. In Section 3.2.1 we outline what is known about the effects of BRD on the sum cost of a profile. This is followed by a brief discussion of the general setting in Section 3.2.2. We conclude with some observations about the effects of BRD for symmetric uncapacitated games in Section 3.2.3

In Section 3.3 we address the question of BRD convergence, and show that in all but the symmetric uncapacitated case convergence cannot be bounded by any function on n , by showing an uncapacitated rooted construction where this is the case.

We then move on to the effects of BRD on the cost of a solution. In Section 3.4 we examine the effects of BRD on the maximum cost of a profile. In the general setting, we show that it is possible for the max-cost of a profile to increase by a factor of $n \log n$ times the cost of the initial profile, showing that the existing upper bound is tight. While this does not directly mean that the Price of Stability for these games is tightly $n \log n$, as it does not show that there is no better equilibria, its existence raises an interesting question which we further explore in Chapter 4. In the case of uncapacitated games, we introduce an upper bound of $O(n)$ (improving on the

previous upper bound of $n \log n$), and an improved lower bound of $\Omega(n)$ (previously linear), thus showing a tight bound of $\Theta(n)$.

This is followed by an examination of the worst possible effects on the cost to an individual player of the game in Section 3.5. For this, we show a tight bound of $\Theta(n)$ for the uncapacitated case. In the general case, we find that it is possible for an individual player to experience an arbitrary increase in cost, relative to their cost in the initial profile.

The chapter is concluded in Section 3.6 with some possible directions for future research.

3.2 Preliminary and Known Results

When analysing networks from a game-theoretic perspective we are interested in stable solutions, and so BRD are a natural place to start when asking questions such as 'how bad is the best stable state?', as they allow us to measure the effect updates can have on a solution.

Bounding the negative effects of BRD has implications to the price of stability measure of network efficiency. If we know that BRD cannot increase the cost of a solution by more than a factor x , we know that there must be a NE which is at most a factor of x times more expensive than the optimum, as BRD from OPT converges to NE.

As an update by one player may increase the cost to their opponents, it is interesting to ask how costs can evolve as a result of BRD from an arbitrary start profile to NE. To answer this, it is necessary to consider both the number of updates required before a NE is reached, and the effect that each update can have.

The maximum number of updates required before BRD converge to NE has been explored in our model, first by Anshelevich et al. [6] for the special case where all edges have a capacity of n , and then in the general case by Feldman & Ron [29].

Anshelevich et al. [6] show that in the asymmetric uncapacitated setting convergence can be exponential in the number of players. They do this by constructing a game where each player has the choice of two paths, and so can be made to represent a bit in a binary counter. They then show that if players make their updates in a

particular order, BRD will iterate through a sequence of profiles which simulates the counter incrementing from an all zero state to an all one state. They conjecture that it may be possible to schedule updates in a way that guarantees convergence in polynomial time. They also show that, when all players have the same source and destination, there cannot be more than n updates before a NE is reached. We discuss symmetric uncapacitated games in Section 3.2.1.

For the general setting, Feldman & Ron [29] show a symmetric construction where convergence can take as many as $n^{3/2}$ steps, although this again relies on a specific ordering of updates.

The impact of BRD on cost has been well studied in the context of the overall cost of a solution, or the sum-cost. In the uncapacitated setting, we again look to Anshelevich et al. [6] for the most significant results. They use the potential function for the game to give an upper bound of $O(\log n)$ to the increase in the total cost of the game. They show a matching lower bound for directed network design games. For undirected networks, there are no known examples where BRD results in an increase to the total cost of the solution by more than a constant factor. In the general setting, Feldman & Ron [29] show that the sum-cost can increase by a factor of $\log n$, which coupled with the upper bound for uncapacitated games give a tight bound on the worst case increase. There has been much further research into BRD on the sum-cost of a solution, and we outline some of the most interesting results in Section 3.2.1.

Another interesting question regarding the effect that BRD can have on cost is that of the maximum cost paid by all players. From the perspective of a central authority, promoting a solution which has a low overall cost would seem a natural goal, however, if we take a more individual viewpoint, a solution where one player pays far more than all others could be considered unfair. For our setting, an upper bound of $n \log n$ and a lower bound of n for both the general and uncapacitated case is shown by Feldman & Ron [29].

The question of max-cost is motivated by a more individual (and therefore more selfish) perspective. Bounds for the previous measures can be applied in some cases, but there has been little research into the effects of BRD on an individual's cost.

3.2.1 Effect on the Sum Cost of a Profile

One of the most widely studied measures of the efficiency of a solution looks at the sum cost of a solution, which is the total cost to all players. This is a very natural measure, as from the point of view of a central authority a solution with a lower overall cost would appear more appealing than a more expensive one. Of course, there are instances where a low sum-cost comes at the cost of a high price for a particular agent, and it is these cases which we explore in the following sections.

In this section we will outline what is already known about the effect of BRD on sum cost. Table 3.1 gives a summary of this.

	General	Uncapacitated
symmetric	$\Theta(\log n)$ [29]	1
asymmetric	$\Theta(\log n)$ [29]	$\Omega(1), O(\log n)$ [7]

Table 3.1: Worst case increase of sum cost as a result of BRD

Existing Upper Bound

Network design games fall into the category of congestion games, and as such have a potential function which can be used to track the change in cost to players as a result of updates. In the case of our model, which uses fair-cost allocation as a means of sharing the cost of edges between those using them, the potential function is defined as

$$\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{i},$$

where $x_e(S)$ represents the number of agents in whose strategy the edge e appears.

As the potential for a solution cannot be more than the sum cost of the solution times the n^{th} harmonic number (where n is the number of players in the solution), i.e $\Phi(S) \leq H(n) \cdot sc(S)$, and that the sum cost of the solution can never exceed the potential of the solution ($sc(S) \leq \Phi(S)$), it follows that the change in sum-cost as a result of best responses cannot be more than factor of $\log n$ times the initial profile [6]. This upper bounds all general and uncapacitated games.

Existing Lower Bounds

General Games. In the general case the upper bound of $O(\log n)$ is tight [29]. Consider the n player game played on a symmetric network with n parallel links e_1, \dots, e_n , where the cost of an edge e_i for $i \in [2, \dots, n]$ is $1/i$, its capacity 1, the cost of the edge e_1 is $1 + \varepsilon$, and the capacity of that edge is n . Figure 3.1 shows this network. Consider the profile where all players share the top edge (e_1). The sum cost of this profile is $1 + \varepsilon$, and each player pays one n^{th} of this. Observe that the bottom edge represents an improving move for all players, but as it has capacity 1 only one player may benefit from using this path, however, their deviation from the initial profile will increase the cost share of those players still using the top edge, which will in turn trigger another update in strategy by one player to the second edge from the bottom. Observe that there are in total $n - 1$ alternative connections to the top edge, with costs $\frac{1}{n}, \frac{1}{n-1}, \dots, \frac{1}{3}, \frac{1}{2}$ and capacity 1, and that when the cheapest i of these are in use the cost to each of the remaining $n - i$ players on the top edge is $(1 + \varepsilon)/(n - i)$. We therefore have a stable solution where each player uses a disjoint path, the sum-cost of which is $\sum_{i=1}^n \frac{1}{i}$. For ε approaching 0 we have a symmetric network where the effect of BRD on sum-cost is $\Omega(\log n)$.

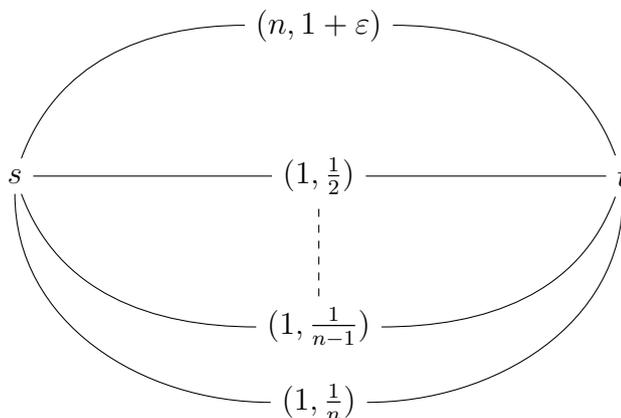


Figure 3.1: Symmetric game Δ where the effect of BRD on sum-cost is $O(\log n)$. Edges are labelled in the form $(capacity, cost)$. Note that the $n - 4$ single capacity edges with costs $\frac{1}{n-2}, \frac{1}{n-3}, \dots, \frac{1}{3}$ have been omitted from the illustration and are instead represented by the dashed line.

Uncapacitated Games. In the case where no edge in the underlying network has a capacity less than the number of players, there is no example where BRD increases

the sum-cost of a solution by more than a constant factor. Results for BRD on sum-cost are generally used in the pursuit of bounding measures related to the quality of solutions, such as the *Price of Stability*. For games with n players, Christodoulou et al. [14] show a lower bound of $42/23 \approx 1.8261$. We omit detailed discussion on existing lower bounds from this chapter, and direct the reader to Section 4.2 for more information, where we outline existing results for the Price of Stability.

3.2.2 General Games

We will now consider the most general class of games, where edges can have capacity less than the number of players. We will now show that, in the context of BRD analysis, it is not necessary to distinguish between symmetric and asymmetric networks. We do this by showing that any asymmetric game can be made symmetric without changing the paths available to each player during BRD.

Take any asymmetric game, and consider a player i who must choose a path between s_i and t_i . If we introduce a node s which will become the common source of all players, we can connect this to s_i with a zero cost single capacity edge. Now do the same for a common destination t . If we consider the profile where each player i uses the edge connecting s to s_i , and t to t_i , we can see that the choices she has to complete her $s-t$ path are those paths she had in the asymmetric version of the game. Further to this, as each connection from the common source and sink has capacity 1, and all of these will be in use, no player has any alternative than to travel to what was in the asymmetric version their original source and sink. So long as the initial profile has each i using s, s_i and t, t_i , she will have the same choices available to her as in the asymmetric version of the game.

Observation 5. *In the context of BRD, it is not necessary to differentiate between symmetric, rooted, and asymmetric games.*

3.2.3 Symmetric Uncapacitated Networks

We will now make some observations regarding the properties of uncapacitated symmetric games. Unlike the most general case, where there may be a limit on the number of players who can use a given path, the strategy set is the same for all

players of uncapacitated symmetric games. This fact allows us to make a number of observations about best responses in these games, giving us tight bounds for this case for the remainder of this chapter.

First let us examine the maximum number of updates needed to reach a stable state.

Consider an uncapacitated symmetric game, and let the first update be player i updating their strategy from a path P to a path Q , which they choose over all alternatives, which we denote as R . By X_i we denote the cost player i would pay to use the path X , which gives us $P_i > Q_i < R_i$.

For the second update, we now consider the actions of some player i' , and can say that $P_{i'} > P_i$ and that if $Q \cap R \neq \emptyset$ and $R_{i'} = R_i - x$ for some $x > 0$, we can also say that $Q_{i'} \leq Q_i - x$ giving $Q_{i'} < R_{i'}$.

The path Q therefore also represents a best response for the second player.

Observation 6. *For any Δ_u where all players share the same source destination pair, no player will ever make more than one update before BRD converges to NE.*

Observation 7. *In symmetric Δ_u , BRD is guaranteed to converge to NE after n steps.*

Now consider the change in costs which each player experiences. A player's cost will only increase if we reach a stable state where they pay for a greater portion of their path than in the initial profile. As all players have the same best response, we cannot reach a stable state where all players do not share the same path P . Any player joining P will decrease her cost, and all subsequent updates by her opponent will further decrease her cost.

Observation 8. *In symmetric Δ_u , BRD from an initial profile S will terminate at a stable profile where all players pay no more than they did in S .*

Observation 9. *In symmetric Δ_u , BRD benefit all players.*

3.3 Convergence to Nash Equilibria

Best response dynamics are of interest both as a method for discovering Nash Equilibria and for the effect they can have on any player's cost. Knowing that BRD

iterates through combinations of individual strategies, and that each step can affect all players, it is natural to ask what limits there are to the number of updates required before a stable solution is reached.

In the uncapacitated case, Ashelevich et al. [6] show for a specific ordering of updates it is possible for convergence to be as high as 2^n , but also conjecture that an ordering of updates may exist where convergence is polynomial in the number of players.

The convergence of BRD to NE has also been studied in scheduling and routing game, and we direct the reader to the works of Feldman and Tamir [30], Fotakis [32], and Even-Dar et al. [26].

Convergence has been studied in our setting [29] and it has been shown that it is possible for there to be as many as $n^{3/2}$ updates, even in the symmetric case.

Interestingly, the fact that uncapacitated games are a subset of our general games, which we show in Observation 4, means that there is a general asymmetric game where convergence can take up to 2^n steps. We can take this lower bound further and apply it to the general symmetric case, as shown in Observation 5.

These facts are however somewhat trivial, as they do not address the underlying question regarding the guaranteed worst case convergence time. If we ask what the best worst case is, there is no example where BRD need more than n steps to converge.

In this section we show that convergence can in fact be unbounded in the number of players, no matter the order in which players make their updates. We show this for all cases except symmetric uncapacitated games, by first showing an asymmetric uncapacitated game where this is the case, and then extending this result to general games. We do this for the general setting by giving an asymmetric construction for two players which converges in $\Omega(|V|)$ steps. This implies matching results for the general case.

A summary of our results is shown in bold in Table 3.2, alongside previous results. We also present the result showing that convergence is unbounded in the general case here [25].

	General		Uncapacitated	
symmetric	\vdots	\vdots	$\Theta(n)[6]$	
rooted	$\Omega(n^{3/2})[29]$	\vdots	$\Omega(n)[6]$	\vdots
asymmetric	$\Omega(2^n)[6]$	unbounded	$\Omega(2^n)[6]$	unbounded

Table 3.2: Worst case convergence of BRD to Nash Equilibria. Results in bold are proved within this thesis.

3.3.1 Uncapacitated Games

Theorem 1. *For rooted Δ_u , the number of updates required to reach a stable state cannot be bounded by any function on the number of players.*

To prove the above we will show that there is a two player rooted game where convergence is unbounded in n .

Theorem 2. *There exists a two-player uncapacitated rooted game Δ_u where BRD converge to NE in $\Omega(|V|)$ steps, regardless of the order in which players make their updates.*

We define Δ_u as a game with two players, whom we refer to as player a and player b , where the underlying graph is defined as

$$\begin{aligned}
 V &= \{s_a, s_b, t\} \cup \{x_i : 0 < i \leq m\} \\
 E &= \{(s_a, x_i) : i \text{ even}\} \cup \{(s_b, x_i) : i \text{ odd}\} \cup \{(x_i, x_{i-1}) : 1 \leq i \leq m\} \\
 p(e) &= \begin{cases} 2/3 & \text{if } e = (x_i, x_{i-1}) \text{ for } 0 < i \leq m \\ m - \lfloor \frac{i}{2} \rfloor (1 + \varepsilon) & \text{if } e = (s_a, x_i) \text{ or } e = (s_b, x_i) : 0 \leq i \leq m \end{cases}
 \end{aligned}$$

for some even $m > 4$, and $\varepsilon \rightarrow 0$. Players a and b have s_a and s_b as their respective source nodes, with x_0 being the destination for both. See Figure 3.2 for an illustration of the above structure, noting that we use the label t for the node x_0 .

Definition 1. *A direct path is one which does not cross the player's opponent's source, and is of the form $s_i, x_j, x_{j-1}, \dots, x_1, x_0, t$.*

Definition 2. *An indirect path is one which crosses the player's opponent's source, and is of the form $s_i, x_j, \dots, x_k, s_l, x_l, \dots, x_0, t$, for $k > l$.*

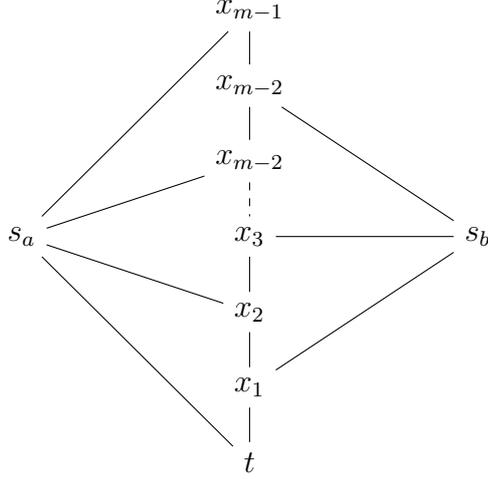


Figure 3.2: Two-player game Δ_u where BRD convergence is unbounded in n .

Lemma 1. *If their opponent uses a direct path, player i 's best response must also be a direct path.*

Proof. If i 's opponent uses a direct path, the cost to player i of an indirect path will be at least the full cost of (s_i, x_j) and $(s_{i'}, x_k)$, plus half of the cost of $(s_{i'}, x_l)$. We will refer to the costs of these edges as J, K, L respectively. As $L \geq \min(J, K)$, we can say that $J + K + L/2 \geq 5/2 \cdot \min(J, K)$, meaning that any indirect path will cost at least $5/2 \cdot (m - (\lfloor m/2 \rfloor (1 + \varepsilon)))$, which is $5m/4 + 5\varepsilon m/4$. As both players have a direct path costing at most $m - (\lfloor 1/2 \rfloor (1 + \varepsilon))$, they will never have incentive to use an indirect path. \square

Lemma 2. *The best response for a player i , while i' uses the direct path $s_{i'}, x_j \rightarrow x_0, t$, is $s_i, x_{j+1} \rightarrow x_0, t$.*

Proof. Lemma 1 shows that i 's best response will be a direct path, and will be of the form $s_i, x_k \rightarrow x_j \rightarrow x_{k'} \rightarrow x_0, t$ or $s_i, x_{k'} \rightarrow x_0, t$, for some $k > j, k' < j$. Note that i will pay the full cost of any connection (x_{k+1}, x_k) , and half of the cost of any connection $(x_{k'}, x_{k'-1})$.

The cost of (s_i, x_k) is the cost of (s_i, x_{k+2}) plus $1 + \varepsilon$, and the cost of the edges $(x_{k+2}, x_{k+1}), (x_{k+1}, x_k)$ is $2 \cdot 2/3 = 4/3$. The cheapest path connecting to some $x_k : k > j$ will be $s_i, x_{j+1} \rightarrow x_j \rightarrow x_0, t$.

The cost of $(s_i, x_{k'})$ is the cost of $(s_i, x_{k'+2})$ plus $1 + \varepsilon$, and the cost of the edges $(x_{k'+2}, x_{k'+1}), (x_{k'+1}, x_{k'})$ is $2 \cdot 1/3 = 2/3$. The cheapest path connecting to some $x_{k'} : k' < j$ will be $s_i, x_{j-1} \rightarrow x_0, t$.

Now compare the costs of these two paths. As they both pass through the node x_{j-1} we can compare the cost of reaching this node. The cost of (s_i, x_{j-1}) is $p(s_i, x_{j+1}) + 1 + \varepsilon$, and the cost of the edges $(x_{j+1}, x_j), (x_j, x_{j-1})$ is $2/3 + 1/3$, so the path $s_i, x_{j+1} \rightarrow x_0, t$ is i 's best response. \square

Proof of Theorem 2. Take the start profile $S = ((s_a, x_0, t), (s_b, x_1, x_0, t))$. Lemma 2 shows that player b already uses the cheapest path available to them, and that player a 's best response is to use the path s_a, x_2, x_1, x_0, t . Each time a player updates, their opponent will have a best response, until we arrive at the profile $S = ((s_a, x_m, t), (s_b, x_{m-1}, x_0, t))$. As each node $x_m \dots x_2$ corresponds to an update, we have a game for two players where BRD converges to NE in $\Theta(|V|)$ steps, and so cannot be bounded by any function on the number of players. \square

3.3.2 General Games

Theorem 3. *In general games, convergence of BRD to NE is arbitrary in the number of players even in the symmetric case.*

Proof. As an uncapacitated game can be viewed as a general game where all edges have a capacity of n , the example given for asymmetric uncapacitated games also applies to general games. This is shown by Observation 4 in Section 2.3. As we have an asymmetric game Δ where convergence is unbounded in the number of players, we also have a symmetric game Δ where this is the case, shown by Observation 5. \square

3.4 Effect on the Maximum Cost of a Profile

In this section we will examine a measure of efficiency which uses the maximum cost to all players when comparing solutions. We are interested in bounding the worst possible increase in cost to the most expensive path in the profile after a series of best responses.

The upper bound of $O(n \log n)$ follows from the upper bound on the increase in sum-cost shown by [6]. As the sum-cost cannot increase by a factor of more than $H(n)$ times the sum-cost of the initial profile, and the initial max-cost could have been no less than $1/n$ of this, maximum cost cannot increase by more than $nH(n)$ times the maximum cost in the optimal profile.

A summary of our contributions can be seen in Table 3.3, alongside previously known results.

	General			Uncapacitated		
symmetric	⋮	⋮	⋮	1		
rooted	⋮	⋮	⋮	⋮	⋮	⋮
asymmetric	$\Omega(n)$ [29]	$\Theta(n \log n)$	$O(n \log n)$ [6]	$\Omega(1)$ [6]	$\Theta(n)$	$O(n \log n)$ [6]

Table 3.3: Worst case effect of BRD on the max-cost of a profile. Results in bold are show within this thesis.

3.4.1 Uncapacitated Games

Theorem 4. *For rooted Δ_u the worst case effect of BRD on the max-cost of a solution is tightly $\Theta(n)$.*

We prove this by giving matching upper and lower bounds.

Lemma 3. *In Δ_u , BRD cannot lead to a NE where the max-cost is more than n times the max-cost in the initial profile.*

Proof. No player will ever pay more than the raw edge cost of the cheapest path in their strategy set, and will never pay less than $1/n$ of this. It therefore follows that no player could ever experience a cost increase of a factor more than n . \square

Lemma 4. *There exists a rooted game Δ_u where best responses increase the max-cost by a factor $\Omega(n)$.*

Proof. Let Δ_u be a game with the underlying graph G and its associated edge costs defined as

$$\begin{aligned}
 V &= \{s_a, s_i, x, t\} \\
 E &= \{(s_a, x), (s_i, x), (x, t), (s_i, t)\}
 \end{aligned}
 \quad
 p(e) = \begin{cases} n & \text{if } e = (s_a, x) \\ 2 & \text{if } e = (s_i, t) \\ n^2 & \text{otherwise} \end{cases},$$

which we depict in Figure 3.3. Player a has the source node s_a , while all other players have the source node s_i . All players have the common destination t .

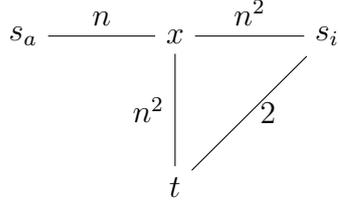


Figure 3.3: Rooted general game where effect of BRD on max cost $\Omega(n)$.

Give the profile S where all players travel through x , player a will pay $2n$, while all other players pay $n + (n^2)/(n - 1)$. At this point player a cannot improve by passing through s_i , as to do so would cost $2n + 2$. It is easy to see that players $2..n$ all have an improving move by travelling directly to t : the first to deviate will pay 2 , with each subsequent deviation further reducing the cost of this edge. Once all players $2..n$ use the edge s_i, t , the cost to player a will be $n^2 + n$, and a will have no incentive to update as to reach the node s_i will also cost $n^2 + n$.

As all players in the initial profile have a cost of roughly $2n$, this represents an increase in max-cost of a factor $(n^2 + n)/(2n)$, which is roughly $n/2$, and so we have that BRD on max-cost is $\Omega(n)$. \square

3.4.2 General Games

In the general case, where edges may have capacity less than the number of players, we show a symmetric game where the max cost of the Nash Equilibrium reached after BRD has $O(n \log n)$ times the max cost of the initial profile.

Theorem 5. *There exists a game Δ for which BRD results in an increase in max cost by a factor arbitrarily close to $nH(n)$.*

The following construction uses parameters $m \in \mathbb{N}$ and $\varepsilon > 0$ where m is sufficiently large, and ε is sufficiently small, e.g. $\varepsilon < 0.1$, yet still satisfies the inequality $\varepsilon > \frac{im(im+1)}{(i-1)2^{m-1}}$ for any $1 < i \leq n$. To show that such a value for ε exists we need only note that for very large m , $\frac{im(im+1)}{(i-1)2^{m-1}}$ approaches 0. It is useful to think of m approaching infinity.

$$\begin{aligned}
V &= \{s_1, t_1, x_{1,m}, z_{1,m}\} \cup \\
&\quad \{s_i, t_i, z_{i,0} \mid 1 < i \leq n\} \cup \\
&\quad \{x_{i,j}, y_{i,j}, z_{i,j} \mid 1 < i \leq n, 1 \leq j \leq m\} \\
E &= \left\{ \begin{array}{l} (x_{i,j}, x_{i,j-1}), (x_{i,j}, y_{i,j}), \\ (y_{i,j}, z_{i,j}), (y_{i,j}, z_{i,j-1}), (z_{i,j}, z_{i,j-1}) \end{array} \mid 1 < i \leq n, 1 < j \leq m \right\} \cup \\
&\quad \{(s_1, x_{1,m}), (t_1, z_{n,m}), (x_{1,m}, z_{1,m})\} \cup \\
&\quad \left\{ \begin{array}{l} (z_{i,0}, z_{i-1,m}), (z_{i,0}, z_{i,1}), (z_{i,0}, y_{i,1}), (x_{i,1}, y_{i,1}), \\ (z_{i,1}, y_{i,1}), (x_{1,m}, x_{i,1}), (z_{i,0}, s_i), (s_i, z_{1,m}), (t_i, x_{i,m}) \end{array} \mid 1 < i \leq n \right\}
\end{aligned}$$

An example of the above is shown in Figure 3.4. We will define the capacities and costs as follows.

$$c(e) = \begin{cases} n & \text{if } e = (x_{1,m}, z_{1,m}) \\ 2 & \text{if } e = (x_{i,j}, y_{i,j}) : 1 < i \leq n, 1 \leq j \leq m \\ 1 & \text{otherwise} \end{cases}$$

$$p(e) = \begin{cases} 1 + \varepsilon & \text{if } e = (x_{1,m}, z_{1,m}) \\ \frac{1}{i2^{j-1}} & \text{if } e = (x_{i,j}, y_{i,j}) : 1 < i \leq n, 1 \leq j \leq m \\ H(i) - \frac{3}{i2^j} + \frac{\varepsilon}{im+j} & \text{if } e = (y_{i,j}, z_{i,j}) : 1 < i \leq n, 1 \leq j \leq m \\ \frac{\varepsilon}{i+j} & \text{if } e = (y_{i,j}, z_{i,j-1}) : 1 < i \leq n, 1 \leq j \leq m \\ 0 & \text{otherwise} \end{cases}$$

Throughout this proof we will refer to player 1 as player a , to s_1 as s_a , and to t_1 as t_a . When talking about paths we will use the notation $z_{i,j} \rightarrow z_{i',j'}$ to mean the sequence of contiguous z vertices between these two vertices, with similar abbreviations for paths along sequences of x nodes, e.g. $x_{i,j} \rightarrow x_{i,j+2}$ would refer to the path $(x_{i,j}, x_{i,j+1}), (x_{i,j+1}, x_{i,j+2})$.

We will use the following notation to refer to choices available to player a .

Definition 3. $A^0 = s_a, x_{1,m}, z_{1,m} \rightarrow z_{n,m}, t_a$.

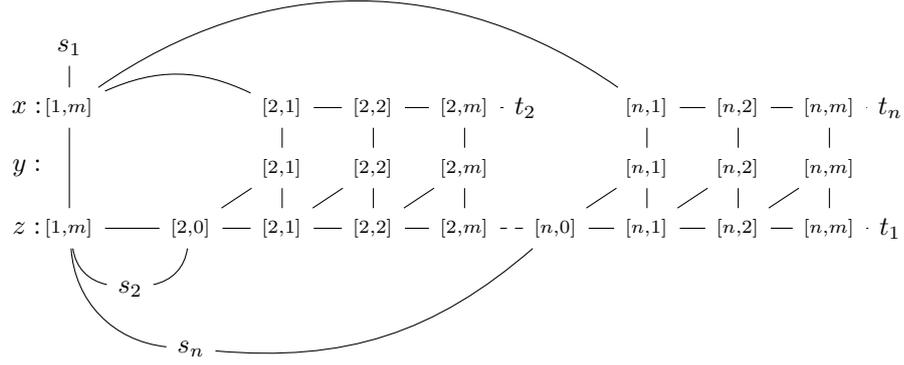


Figure 3.4: Asymmetric game Δ where BRD increases max cost by a factor of $n \log n$. Node $[i, j]$ in row x is the node $x_{i,j}$, and similarly for row y and z .

Definition 4. $A_j^i = s_a, x_{1,m}, x_{i,1} \rightarrow x_{i,j}, y_{i,j}, z_{i,j} \rightarrow z_{n,m}, t_a$.

We now introduce the notation by which we refer to the choices available to a player $i \in \{2..n\}$.

Definition 5. $P_i^0 = s_i, z_{1,m}, x_{1,m}, x_{i,1} \rightarrow x_{i,m}, t_i$.

Definition 6. $P_i^j = s_i, z_{i,0} \rightarrow z_{i,j-1}, y_{i,j}, x_{i,j} \rightarrow x_{i,m}, t_i$.

First we will establish some inequalities which will be used throughout the proof.

Since $p(x_{i,j}, y_{i,j}) = 1/i2^{j-1}$ we have that $p(x_{i,j}, y_{i,j}) = 2 \cdot p(x_{i,j+1}, y_{i,j+1})$. As $p(y_{i,j}, z_{i,j-1}) = \varepsilon/(i+j)$, we have

$$p(x_{i,j}, y_{i,j})/2 + p(y_{i,j}, z_{i,j-1}) > p(x_{i,j+1}, y_{i,j+1}) + p(y_{i,j+1}, z_{i,j}) \quad (3.1)$$

for all $2 \leq i \leq n$ and $1 \leq j < m$.

As $p(y_{i,j}, z_{i,j}) = H(i) - 3/i2^j + \varepsilon/(im+j)$ we have that $p(y_{i,j+1}, z_{i,j+1}) = p(y_{i,j}, z_{i,j}) + 3/(i2^{j+1}) - \varepsilon/((im+j+1)(im+j))$. We can also say that $p(x_{i,j+1}, y_{i,j+1})/2 + p(y_{i,j+1}, z_{i,j+1}) = 1/i2^{j+1} + p(y_{i,j}, z_{i,j}) + 3/i2^{j+1} - \varepsilon/((im+j+1)(im+j)) = p(y_{i,j}, z_{i,j}) + 1/i2^{j-1} - \varepsilon/((im+j+1)(im+j))$. We therefore have that

$$p(x_{i,j}, y_{i,j}) + p(y_{i,j}, z_{i,j}) > p(x_{i,j+1}, y_{i,j+1})/2 + p(y_{i,j+1}, z_{i,j+1}) \quad (3.2)$$

We define the start profile S as all players travelling through the edge $x_{1,m}, y_{1,m}$ (which they reach by the most direct route, and then continuing to their destination using only zero cost edges (those along the top or bottom of any grid)).

$$S = (A^0, P_2^0, P_3^0, \dots, P_n^0)$$

In this profile, all players pay $(1 + \varepsilon)/n$.

Lemma 5. *Player a has no improving move from the profile S .*

Proof. In the profile S player a uses the path $s_a, x_{1,m}, z_{1,m} \rightarrow z_{n,m}, t_a$, paying $1/n$ of the cost of the edge $(x_{1,m}, z_{1,m})$, and 0 for all other edges in this path. As all the edges $(x_{1,m}, x_{i,1})$ are in use and have a capacity of 1, player a must use the edge $(x_{1,m}, z_{1,m})$, and so cannot improve her path. \square

Lemma 6. *If not using the edge $(x_{1,m}, z_{1,m})$, player a must travel through some edge $(y_{i,j}, z_{i,j})$.*

Proof. If not using $(x_{1,m}, z_{1,m})$, player a must use an edge $(x_{1,m}, x_{i,1})$ and then travel between some $x_{i,j}$ and $z_{i,j}$ using the edges $(x_{i,j-1}, x_{i,j})$ and $(z_{i,j}, z_{i,j+1})$, leaving the i^{th} grid via $(z_{i,m}, z_{i+1,0})$. Player i must travel between $z_{i,j-1}$ and $x_{i,j+1}$, and so must travel along either $(z_{i,j-1}, z_{i,j})$ or $(z_{i,j-1}, y_{i,j})$. Player a cannot therefore use both of these edges, and so must reach $z_{i,j}$ via $(y_{i,j}, z_{i,j})$. \square

Lemma 7. *Player a will not have an improving move until she pays the full cost of the edge $(x_{1,m}, z_{1,m})$.*

Proof. The cheapest alternative to the path using the edge $(x_{1,m}, z_{1,m})$ uses the edge $(y_{2,1}, z_{2,1})$ and costs $H(2) - 3/4 + \varepsilon/(2m + 1)$, which is more than half the cost of $(x_{1,m}, z_{1,m})$. \square

Lemma 8. *From the profile S , player n has a best response of P_n^1 .*

Proof. The cost of P_n^0 is $(1 + \varepsilon)/n$, while the cost of P_n^1 is $1/n + \varepsilon/(n + 1)$, so the latter represents an improving move. To show that there is no alternative cheaper than P_n^1 , it is only necessary to note that if not using the path P_n^0 or P_n^1 , player n must use some edge $(y_{n,j}, z_{n,j})$ for a cost of at least $H(n) - \frac{3}{2n}$. \square

Lemma 9. *While player a uses the path A^0 , and each player $j \in \{n, \dots, i + 1\}$ uses the path P_j^1 , player i has the best response of P_i^1 .*

Proof. The cost to i of the path P_i^0 will be at least $(1 + \varepsilon)/i$, as players $i + 1 \dots n$ do not use the edge $(z_{1,m}, x_{1,m})$. The cost of P_i^1 being $1/i + \varepsilon/(i + 1)$, this path represents an improving move. As any alternative to P_i^0, P_i^1 will use some edge $(y_{i,j}, z_{i,j})$ costing at least $3/4$, P_i^1 is the best response. \square

We have shown that from the profile S , BRD will lead to the profile

$$S' = (A^0, P_2^1, P_3^1, \dots, P_n^1)$$

where player a pays $1 + \varepsilon$ and each player $i \in \{2..n\}$ pays $1/i + \varepsilon/(i + 1)$.

Lemma 10. *From the profile S' , player a 's best response is the path $A^{2,1}$.*

Proof. If not using the path A^0 , they will pay the full cost of some edge $(y_{i,j}, z_{i,j})$, the cheapest of which is $(y_{2,1}, z_{2,1})$, which costs $H(2) - 3/4 + \varepsilon/(2m + 1)$. The cheapest path containing this edge is $A^{2,1}$, the price of which is the cost of the edge $(y_{2,1}, z_{2,1})$ plus half of the cost of the edge $(x_{2,1}, y_{2,1})$, which is $1 + \varepsilon/(2m + 1)$. \square

We will now show a sequence of best responses by which player a iterates through paths $A^{2,1}$ to $A^{n,m}$.

Lemma 11. *Given the profile $S^* = (A_i^j, P_2^m, \dots, P_{i-1}^m, P_i^j, P_{i+1}^1, \dots, P_n^1)$, player i has the best response of P_i^{j+1} .*

Proof. The cost to player i of their current path is $1/(i2^j)$. Note that, while player a uses the path A_i^j , a path $P_i^{j'}$ is only feasible for $j' = j$ or $j' = j + 1$. Inequality (3.1), the path P_i^{j+1} represents an improving move for player i . We now show that it is the best response to the profile S^* . Consider alternatives available to i . If not using a path of the form $P_i^{j'}$, they may reach their destination through the edge $(s_i, z_{1,m})$, in which case they will either pay for $(z_{1,m}, x_{1,m})$, which, costing $1 + \varepsilon$, is more expensive than their current path, or pay for some edge $(y_{i',j'}, z_{i',j'})$ for some $i' < i$ and an arbitrary j' , all of which cost more than $1/i$ and therefore more than their current path. Player i cannot improve their cost with any path starting with the edge $(s_i, z_{1,m})$. If they instead chose to start with the edge $(s_i, z_{i,0})$, the only choices available which are not of the form $P_i^{j'}$ use must use some edge $(y_{i,j'}, z_{i,j'})$. As any of these edges costs more than $1/i$ (player i does not have an improving move which is not of the form $P_i^{j'}$). The path P_i^{j+1} must therefore be their best response. \square

Lemma 12. *Given the profile $S^* = (A_i^{j-1}, P_2^m, \dots, P_{i-1}^m, P_i^j, P_{i+1}^1, \dots, P_n^1)$, player a has the best response of A_i^j .*

Proof. The cost to player a of their current path is $p(x_{i,j-1}, y_{i,j-1}) + p(y_{i,j-1}, z_{i,j-1}) = H(i) - 1/i2^{j-2} + \varepsilon/(im + j - 1)$. We show that player a has an improving move in

the path A_i^j by Inequality (3.2). Now consider alternatives to this path. As player i uses the edges $z_{i,0} \rightarrow z_{i,j-1}$, any path for a which passes through a node $x_{i,j'} : j' < j$ will also use the edge $(y_{i,j'}, z_{i,j'})$. It therefore follows that of the paths available to a which pass through the node $z_{i,j-1}$, A_i^{j-1} is the cheapest. Now consider paths for a which do not pass through the node $z_{i,j-1}$, excluding P_i^j . As no path $P_i^{j'} : j < j'$ is feasible we need only examine paths $P_{i'}^{j'} : i' > i$, all of which must cost at least $H(i)$, and so could not be an improving move. The path A_i^j is the only improving move for a and therefore their best response to the profile S^* . \square

Lemma 13. *Given the profile $S^* = (A^{i-1,m}, P_2^m, \dots, P_{i-1}^m, P_i^1, \dots, P_n^1)$, player a has a best response in the path $A^{i,1}$.*

Proof. The cost of the profile S^* to players a and $i-1$ is $H(i-1) - 1/(i-1)2^{m-1} + \varepsilon/im$ and $1/(i-1)2^m + \varepsilon/(i-1)m$ respectively. Player a has the choice of the direct paths $A^{i-1,m-1}, A^{i,1}, \dots, A^{n,1}$. The path $A^{i-1,m-1}$ costs the full price of the edges $(x_{i-1,m-1}, y_{i-1,m-1})$ and $(y_{i-1,m-1}, z_{i-1,m-1})$, which is $H(i-1) - 1/(i-1)2^{m-1} + \varepsilon/(im-1)$ and is $\varepsilon/((im)^2 - im)$ more expensive than a 's current path, so cannot be an improving move. To use the path $A^{i,1}$, a will pay the full price of the edge $(y_{i,1}, z_{i,1})$, and half the price of the edge $(x_{i,1}, y_{i,1})$, which is $H(i-1) + \varepsilon/(im+1)$. For this to be an improving move we need that $H(i-1) - 1/(i-1)2^{m-1} + \varepsilon/im > H(i-1) + \varepsilon/(im+1)$, i.e. that $\varepsilon/im - \varepsilon/(im+1) > 1/(i-1)2^{m-1}$. As this is satisfied by our initial requirement that $\varepsilon > im(im+1)/(i-1)2^{m-1}$, the path A_i^1 is the only improving move for player a , and therefore her best response to the profile S^* . \square

By the above Lemmas, we have a start profile S which will lead to a profile $\bar{S} = (A^{n,m}, P_2^m, \dots, P_i^m, \dots, P_n^m)$. We conclude our proof by showing that this profile is in fact stable.

Lemma 14. *Given the profile $\bar{S} = (A^{n,m}, P_b^m, \dots, P_i^m, \dots, P_n^m)$, no player has an improving move and so $\bar{S} \in NE(\Delta)$.*

Proof. No player $i \in \{2, \dots, n\}$ can improve by using any $P_i^j : j < m$, by (3.1). Neither can they improve by using any path containing an edge $(y_{i',j}, z_{i',j-1})$, as this would require them to also use the edge $(z_{i',j-1}, z_{i',j})$ or $(z_{i',j-1}, z_{i',j-2})$, both of which have capacity 1 and are used by player i' , or the edge $(z_{i',j-1}, y_{i',j-1})$, which

has a cost of at least $3/4$. So, there is no improving move for any player $2 \dots n$. Player a does not have an improving move as, by Lemma 13, their current path is the best-response to their opponents' paths. \square

Proof of Theorem 5. We have shown a game with a start profile S , where $mc(S) = (1 + \varepsilon)/n$, and a sequence of best-responses leading to a stable profile \bar{S} . In this NE where player a pays the full cost of the edge $(y_{n,m}, z_{n,m})$, and half the cost of the edge $(x_{n,m}, y_{n,m})$, which is $H(n) - 3/(n2^m) + \varepsilon/(nm + m) + 1/(n2^{m-1})$ which, as $m \rightarrow \infty$, approaches $H(n)$. As the potential of the start profile was $H(1) + \varepsilon$, it follows that player a has the max cost of the resultant NE. We therefore have a game where the max cost increases by a factor approaching $H(n) / ((1 + \varepsilon)/n)$, which is $nH(n)$. \square

3.5 Effect on the Cost to an Individual

Much of the research into the effects of BRD on cost has been motivated by the search for bounds on the Price of Stability, and so has focused on it's effects on the sum cost of the solution. We now ask the question, given an arbitrary player of a game, what is the worst case increase in cost that they can experience. As the question of cost increase has not previously been approached from the perspective of an individual player, the only bounds we have on this come directly from the bounds on the sum-cost.

The lower bounds on the max-cost increase in the previous section can of course be applied to this setting. Interestingly we find that in the most general case BRD can result in an arbitrary (with respect to their initial cost) increase to an individual's cost. In the uncapacitated case, we improve on the results from the previous section by showing it is possible for an individual to experience an increase in cost of exactly n .

3.5.1 Uncapacitated Games

We will now show that in the uncapacitated case, the worst case increase to any player's cost is exactly n .

Theorem 6. *The worst case effect of BRD on any player's cost is exactly n .*

	General		Uncapacitated	
symmetric	⋮	⋮	1	
rooted	⋮	⋮	⋮	⋮
asymmetric	$\Omega(n)$ [29]	unbounded	$\Omega(1)$	$\Theta(\mathbf{n})$

Table 3.4: Worst case effect of BRD on the cost of an individual player.

It follows from Lemma 3 that it is not possible for BRD from a profile S to increase the cost to any player by a factor of more than n times the cost of their path in S . We now provide a lower bound which matches this.

Lemma 15. *There exists a rooted game Δ_u for which BRD results in an increase in cost to some player of exactly n .*

Proof. Take the following structure, an example of which is shown in Figure 3.5.

$$\begin{aligned}
 V &= \{s_a, s_i, t\} \\
 E &= \{(s_a, s_i), (s_a, t), (s_i, t)\} \\
 p(e) &= \begin{cases} n & \text{if } e = (s_a, t) \\ \infty & \text{if } e = (s_a, s_i) \\ 0 & \text{otherwise} \end{cases}
 \end{aligned}$$

Let the node s_a be the source of player 1, who we will also refer to as player a , and the node s_i be the source of all other players. All n players have the common destination node t .

Consider the profile S where all players use the edge (s_a, t) , and note that player a will pay 1 while all other players pay infinitely more than this. All players but a have an improving move using the zero cost edge (s_i, t) , and their deviations to this path will leave player a to pay the full cost of the edge (s_a, t) , which is n . Player a cannot improve, as their only alternative would use the edge (s_a, s_i) , so we have a NE where a pays n , reached from an initial state with a cost of 1. \square

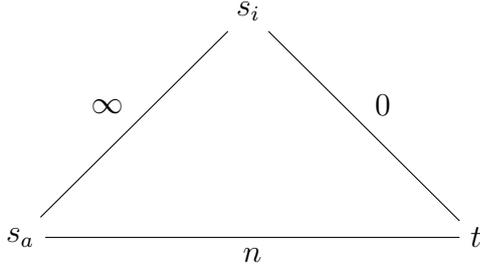


Figure 3.5: Rooted game Δ_u where an individual player experiences an increase in cost of a factor n .

3.5.2 General Games

For the general setting, we get the lower bound of $\Omega(n)$ from the previous lower bound on the worst case increase in max-cost [29]. Our result of a tight bound for max-cost of $\Theta(n \log n)$ implies a lower bound for individual cost increase, improving on what was previously known. In Section 3.5.1 we established an upper bound of $O(n)$ on the increase in cost to any individual in the uncapacitated setting. This upper bound does not hold in the general case, as a player may pay more than the raw edge cost of their cheapest path if it is not a feasible choice i.e. at least one edge in the path is already used by as many players as its capacity allows.

We now show that in the general case, the increase in cost to an individual player cannot be bounded by any function on the number of players, by showing a construction where this is the case.

Theorem 7. *There exists an asymmetric game Δ with 2 players and a start profile for Δ such that BRD increase the cost of a player to an arbitrary factor times the player's cost in the start profile.*

Proof. For $m \in \mathbb{N}$ with $m \geq 3$, consider the general game Δ with $n = 2$ players and underlying undirected graph $G = (V, E)$, defined as follows (see Figure 3.6 for an illustration for $m = 3$):

$$V = \{x_i, z_i \mid 0 \leq i \leq m\} \cup \{y_i \mid 1 \leq i \leq m\}$$

$$E = \{(x_i, x_{i-1}), (z_i, z_{i-1}), (x_i, y_i), (y_i, z_{i-1}), (y_i, z_i) \mid 1 \leq i \leq m\}$$

We denote the two players by a and b . Their source and sink nodes are $s_a = z_0$,

$t_a = x_m$, $s_b = x_0$, and $t_b = z_m$. A horizontal path from x_i to x_j for some $j \geq i$ is denoted by $x_i \rightarrow x_j$, and similarly for $z_i \rightarrow z_j$.

All edges have capacity 1, except those connecting an x node and a y node, which have capacity 2. Only edges incident with y nodes have non-zero cost. For any node y_i , the costs of the edges to z_{i-1} , z_i , x_i are denoted by a_i , b_i , ab_i , respectively. These costs are defined as follows (where $\varepsilon > 0$ is a positive constant satisfying $\varepsilon < 1/m^2$):

$$\begin{aligned} a_1 &= 2^m & a_i &= 2^m - 2^i + 2^{i-2} + 1 - i\varepsilon \quad \text{for } i > 1 \\ b_1 &= 1 + \varepsilon & b_i &= 0 \quad \text{for } i > 1 \\ ab_1 &= 0 & ab_i &= 2^{i-1} + \varepsilon \quad \text{for } i > 1 \end{aligned}$$

Let the start profile be $S = ((z_0, y_1, x_1 \rightarrow x_m), (x_0, x_1, y_1, z_1 \rightarrow z_m))$. Our aim is to enable a sequence of $2m - 2$ best response moves such that the cost of player b increases by an arbitrary factor (depending on m). In the start profile S , player a and b share the edge (x_1, y_1) and their costs are 2^m and $1 + \varepsilon$, respectively. Player a 's best response to player b 's strategy is now the path $z_0, z_1, y_2, x_2 \rightarrow x_m$ with cost $p(a_2) + p(ab_2) = 2^m - 2 - 2\varepsilon + 2 + \varepsilon = 2^m - \varepsilon$, so player a will update to that path. Player b 's best response to a 's new path is now the path $x_0 \rightarrow x_2, y_2, z_2 \rightarrow z_m$ with cost $p(ab_2)/2 + p(b_2) = 1 + \frac{\varepsilon}{2}$, so player b will update to that path. As the edge (x_2, y_2) is now shared, this reduces the cost of player a to $2^m - 1 - 2\varepsilon + \varepsilon/2$.

In the profile reached after $2(i - 1)$ best response moves, for $2 \leq i \leq m$, player a uses path $z_0 \rightarrow z_{i-1}, y_i, x_i \rightarrow x_m$ and player b uses path $x_0 \rightarrow x_i, y_i, z_i \rightarrow z_m$. Player a 's cost is $p(a_i) + p(ab_i)/2 = 2^m - 2^i + 2^{i-2} + 1 - i\varepsilon + 2^{i-2} + \varepsilon/2 = 2^m + 2^{i-1} + 1 - 2^i - i\varepsilon + \varepsilon/2$, and player b 's cost is $2^{i-2} + \varepsilon/2$.

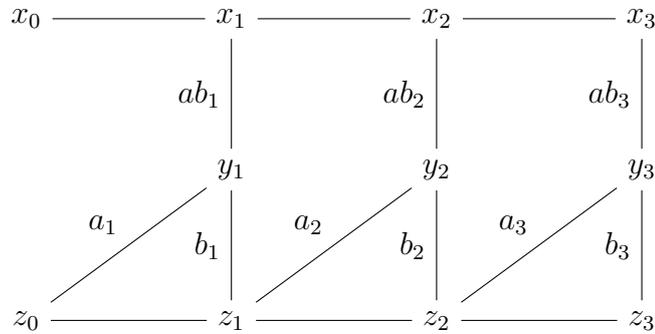


Figure 3.6: Asymmetric game Δ where $m = 3$ and an individuals cost increases by an arbitrary factor w.r.t n .

We prove this by induction on i . For the base case $i = 2$, the claim was shown already. Now assume that the claim holds for i . Player a 's best response is the path $z_0 \rightarrow z_i, y_{i+1}, x_{i+1} \rightarrow x_m$ with cost $p(a_{i+1}) + p(ab_{i+1}) = 2^m - 2^{i+1} + 2^{i-1} + 1 - (i+1)\varepsilon + 2^i + \varepsilon = 2^m + 2^{i-1} - 2^i + 1 - i\varepsilon$, so player a updates to this path. As the edge (x_i, y_i) is no longer shared, player b 's cost increases to $p(ab_i) = 2^{i-1} + \varepsilon$. Player b 's best response is now the path $x_0 \rightarrow x_{i+1}, y_{i+1}, z_{i+1} \rightarrow z_m$ that shares (x_{i+1}, y_{i+1}) with player a and has cost $p(ab_{i+1})/2 + p(b_{i+1}) = 2^{i-1} + \varepsilon/2$. When b updates to this path, this reduces the cost of player a to $p(a_{i+1}) + p(ab_{i+1})/2 = 2^m - 2^{i+1} + 2^{i-1} + 1 - (i+1)\varepsilon + 2^{i-1} + \varepsilon/2 = 2^m + 2^i + 1 - 2^{i+1} - (i+1)\varepsilon + \varepsilon/2$. Hence, the claim also holds for $i+1$.

After $2(m-1)$ best response moves, player a 's path is $(z_0 \rightarrow z_{m-1}, y_m, x_m, t_a)$ with cost $2^{m-1} + 1 - m\varepsilon + \varepsilon/2$ and player b 's path is $(x_0 \rightarrow x_m, y_m, z_m)$ with cost $2^{m-2} + \varepsilon/2$. Denote this strategy profile by S^* . We claim that S^* is a NE. First, note that player a does not have an improving move: As the edges (x_{m-1}, x_m) and (z_m, y_m) have capacity 1 and are used by player b , player a can reach x_m only via the edges (z_{m-1}, y_m) and (y_m, x_m) , and the path that a uses in S^* contains only zero-cost edges in addition to these two edges. Player b 's only alternative paths that are potential improving moves are of the form $(x_0 \rightarrow x_i, y_i, z_i, y_{i+1}, z_{i+1}, \dots, y_{m-1}, z_{m-1}, z_m)$ for some $i < m-1$. Any such path would contain the edge (z_{m-2}, y_{m-1}) with cost $p(a_{m-1}) = 2^m - 2^{m-1} + 2^{m-3} + 1 - (m-1)\varepsilon = 2^{m-1} + 2^{m-3} + 1 - (m-1)\varepsilon > 2^{m-1} + 1 - m\varepsilon + \varepsilon/2$, so it would not be an improving move for b . Therefore, S^* is a NE.

The cost of player b is $1 + \varepsilon$ in the start profile S and $2^{m-2} + \varepsilon/2$ in the NE S^* that is reached by BRD from S . Hence, the cost of player b has increased by a factor arbitrarily close to 2^{m-2} . As m can be chosen arbitrarily large, we have shown that the cost of a player can increase by an arbitrary factor during BRD. \square

3.6 Concluding Remarks

Having shown that, in our setting, the number of steps needed before BRD converge to NE is arbitrary in the number of players, we find it natural to ask whether this fact applies in other cases.

Further to our results showing an arbitrary increase in cost to an individual player, we also propose a closer examination of when such growth can be observed, i.e., which players can experience such growth, and in which settings can such growth be observed.

Chapter 4

Price of Stability With Respect to Max Cost

We now turn our attention to measures related to the loss of efficiency due to selfish behaviour in networks. More specifically, we are interested in comparing stable solutions to the optimum profile. One such measure, the price of stability (PoS), compares the cost of the optimal profile to that of the best NE. We now examine this measure, using the maximum cost of a profile when comparing solutions.

4.1 Introduction

Knowing that sub-optimal stable solutions exist, it is natural to ask how bad these Nash equilibria can be. It is therefore unsurprising that the quantification of inefficiency of Nash equilibria (NE) has received considerable attention in recent years [6, 14, 21, 1, 31].

The concept of the *price of anarchy*, that is, the ratio between the worst case NE and the optimal solution, allows us to bound for a given type of game, the worst possible outcome as a result of individually motivated behaviour, and as such provides a powerful tool in the analysis of the quality of equilibria.

The knowledge that particularly bad NE exist, coupled with the specific configurations which permit such states, is useful from both the perspective of the central authority charged with the construction of multi-agent networks, and for those individuals for whom the networks represents their field of play.

This somewhat pessimistic method of analysis, is of course limited when it comes to questions such as, “Under what conditions do favourable solutions exist?”, and the more optimistic “How bad is the best possible solution?”. The later, which can be more precisely expressed as, what is the worst case ratio between cost of the best NE of a game and that of it’s optimal solution, has given rise to the measure of efficiency we study, suitably called the *price of stability*.

Of course, the quantification of stable solutions requires some metric by which we do so. In our model, where players pay a share of the edges in s - t path, we measure solutions by the cost to the players of the game. The most typical approach is to take the sum of the cost’s of all players in the game. This method of costing, which we will refer to as the *sum-cost* of a game, and is sometimes referred to as the *utilitarian* cost function, has natural economic motivations and has been widely examined in a number of settings.

While useful as an overall measure of it’s quality, an examination of the sum cost of a solution does not give indications as to the fairness of a solutions, in that a solution with a low social cost may come at a particularly high cost to some individual player of the game. To counter this, one can consider the highest cost paid by all players as a measure of a solutions quality. The promotion of solutions which have the lowest possible cost to the player who pays the most may be a crude method for advancing fairness, but has obvious advantages over the sum-cost approach if we are concerned with a measure which promotes individual welfare over social benefits. In this chapter we measure this measure, referring to it as the *max-cost* objective. A somewhat under examined concept in network design games, it is sometimes referred to as the *egalitarian* cost function.

Questions.

We now look to address the question of the PoS related to the max-cost objective (PoS_{mc}), first in the general case, where edges in the underlying network may have capacity less than the total number of players in the game, and then in the special case of games played on networks when edges a capacity at least as great as the number of players in the game (which may be viewed as the uncapacitated version of our model), building on the work of Feldman & Ron [29]. In the general case we as the following:

	Δ_u			Δ		
	symmetric	rooted	asymmetric	symmetric	rooted	asymmetric
PoA _{sc/mc}	$\Theta(n)$ [6]			<i>unbounded</i> [29]		
sc	$\Omega(1)$			$\Theta(\log n)$ [29]		
	$(1 - \Theta(1/k^4))H_k$ [21]					
PoS	1	$\Omega(1)$		$\Omega(n)$ [29]		
mc	$\Theta(n)$ $\Theta(n)$ $\Theta(n \log n)$		
	$O(n \log n)$ [7]		

Table 4.1: Summary of results for the Price of Anarchy and Price of Stability, with regard to the sum-cost and max-cost objectives. Results in bold are presented in this thesis.

1. Are there network topologies where the upper bound of $O(n \log n)$ can be realised?
2. For what types of capacitated networks does this upper bound not apply?

When we come to the uncapacitated case, the same upper bound applies, but we have little in the way of a lower bound (except in the case of symmetric uncapacitated games, where $\text{PoS} = 1$). With the view of closing the gap for the remainder of the setting, we ask:

1. Are there n player games where PoS_{mc} is worse than constant?
2. Does the upper bound of $O(n \log n)$ still apply for this type of games?

In this chapter we will show tight bounds of $\Theta(n)$ for the Price of Stability max-cost measure for all games in our setting, apart from asymmetric general games, which we show to be tightly $\Theta(n \log n)$. A summary of our findings, along with previous results, is shown in Table 4.1.

Limitations.

By measuring the prices of anarchy and stability for a given class of game, one is able to say that, for a specific game of that class, the loss of efficiency as a result of selfish behaviour will fall between the two. Indeed, knowing the specific configurations related to each end of this range, one may use this to promote situations

towards the lower end of the cost spectrum, but to do so may prove difficult in a real world situation due to the limitations of these configurations. Knowing that networks where the strategy sets of all players is symmetric (in the case where each individual has the same source and destination node) is all very well in theory, but less helpful when it comes to constructing more complex multi source networks.

Another limitation of the price of stability measure is the reliance it places on the initial configuration of the game. Many of the results for PoS show the existence of good NE by reasoning about best response dynamics (*BRD*) from the optimal profile. The promotion of these NE would therefore require a coordinating body with both the knowledge of the optimal profile, and the power to dictate it as the starting profile for the game. This approach is further complicated if the players in the game are not fixed, as the addition or removal of a single player may drastically change the optimal profile, and would require a re-run of the process by which the NE is reached. That said, an interesting question is, “If we leave the players of the game to their own devices, how bad can the resultant NE get?”. This question motivates us in Chapter 5 to examine what we will call the *sequential price of anarchy*, which is concerned examine the ratio between the worst case NE which can be reached when no initial configuration is specified and instead players *join* the game one by one, picking their best response to the current configuration.

Outline of the remainder of the chapter.

We start, in Section 4.2, by defining more precisely those measures examined in the remainder of the chapter. We provide a general discussion of what is already known about the Price of Stability, giving some details of what is known for both general and uncapacitated games for both the sum and max cost objective measures.

This is followed, in Section 4.3, by an examination of general games, where edges may have capacity less than the total number of players. For symmetric games with n players, we show that the price of stability with respect to max-cost is $O(n)$. This bound is tight, as implied by the matching lower bound from [29], and hence resolves the open problem posed by Feldman and Ron. A standard proof technique for bounding the price of stability is to bound the increase in social cost during best response dynamics starting from the optimal strategy profile. We show that this technique does not work in our case, as best response dynamics starting from the

optimal strategy profile can actually increase the max-cost by a factor of $\Theta(n \log n)$. Therefore, we use a different approach to bound the price of stability, which may be of independent interest. We show that in any case where BRD increases the max-cost of a profile by more than a factor of n , it is possible to discard this expensive path and generate a new profile in which no player pays more than n times the initial max-cost, and that while this profile may not be an equilibrium this process must eventually terminate with a profile which is both stable and has a max-cost which is at most n times the max-cost in the initial profile. For asymmetric games with n players, we show that the price of stability can be $\Omega(n \log n)$, matching the previously known upper bound.

When considering uncapacitated games, in Section 4.4, we find that in both asymmetric and rooted games the price of stability w.r.t. max-cost is upper bounded by $O(n)$, which improves on the previous upper bound of $O(n \log n)$. We show a rooted construction where the best NE has a max-cost a factor of n times the max-cost of the optimal profile, improving on the previous lower bound (which was constant), and showing that our upper bound of n is tight.

We conclude with a brief summary of the chapter, and discussion of possible directions for future research in Section 4.5.

4.2 Preliminary and Known Results

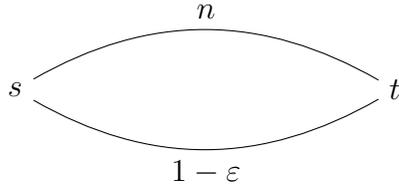
Recall from Section 2.2, that we have two measures of the cost of a solution, namely, the *sum-cost* and *max-cost* objectives, which we denote and calculate as $sc_\Delta(S) =$

$\sum_{i \in [n]} p_i(S)$ and $mc_\Delta(S) = \max_{i \in [n]} p_i(S)$ respectively, for a given profile S .

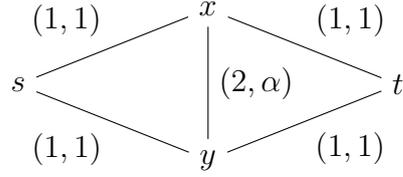
Recall also that the type of games which we study, which are network design games with fair cost sharing, fall into the broader class of *congestion games* as shown by [49], and as such have a potential function which can be used to track the change in cost to an individual when making an update in strategy. We now reproduce the definition for the potential value of a solution S in our games, for ease of reading.

$$\Phi(S) = \sum_{e \in E} \sum_{i=1}^{x_e(S)} \frac{p_e}{i}.$$

Price of Anarchy. The application of Game Theory to the field of Network Design has in the last two decades or so produced numerous insights into the effects of selfish behaviour in the setting of interaction in networks. One of the primary concerns, when examining an interactive situation from a game theoretic perspective, is the analysis of the solutions resulting from individual and selfish actions. The question, in the worst case, how much more expensive does an individually motivated solution become, when compared with the optimal solution for that game, motivated Koutsoupias and Papadimitriou [39] to introduce a measure, later coined the *Price of Anarchy*, which is, for a given class of games, the worst case ratio between the cost of the worst NE of a particular game and the cost of it's optimal solution. The publication of their paper, which showed several bounds for two player games, inspired a flurry of activity, and their measure of network efficiency has been widely adopted and studied in a variety of setting. For a good overview on the type of work on this measure we direct the reader to [15, 17, 54]. It is worth at this point mentioning that the sum-cost measure for the *PoA* has been shown to be tightly $\Theta(n)$ for uncapacitated network design games [6]. Consider a symmetric game with two parallel links connecting the common source and destination nodes. Let the first connection have a cost of 1, and let this be optimal. If the second link were to have a cost of $n - \varepsilon$, where n is the number of players of the game, and ε is some small value approaching 0, the solution where all players share this second link is stable, as each player has a cost of slightly less than 1 and thus has no incentive for a unilateral deviation in strategy to the first link. The price of anarchy has also been studied from the context of *strong equilibria*, where no coalition of players may cooperatively deviate in a way which benefits all it's members, by Andelman et al. [4] and Epstein et al. [22], and by Feldman and Geri [28]. The role of network topology has also been explored in several settings. For some games, such as non-atomic network routing games, the price of anarchy measure are independent from the topology of the underlying network, as shown by Roughgarden and Tardos here [59]. This is not true for all settings. The concept of *Pareto efficiency*, which is a state in which there is no deviation in strategy possible without a negative impact to the cost of at least one other player's profile, has been shown to be strongly dependent on network topology. See the work of Milchtaich [45] and references there. In the



(a) Symmetric game Δ_u where $\text{PoA}_{sc} = n$.



(b) Symmetric game Δ where PoA_{sc} is unbounded in the number of players.

Figure 4.1: Examples showing lower bounds for PoA_{sc} in Δ and Δ_u .

case of symmetric networks, much has been shown for several solution concepts. For the existence of equilibria, both Nash and strong, see Epstein et al. [23], Holzman and Law-Yone [36, 37], and Milchtaich [47], who also studies the uniqueness of equilibrium here [46].

Price of Stability. As a measure of the efficiency, the Price of Anarchy provided an interesting way of measuring the worst case scenario for situations in which no central authority was present. The observation that, while some games do have particularly bad NE, they often have more favourable stable states, motivated the introduction of a new measure by Correa, Schulz, and Moses [16], which examines the ratio between the cost of the best possible NE for a game and that of the optimal solution. This measure, which they refer to as the *optimistic Price of Anarchy*, has also been widely studied in a number of settings, and was later given the name of *Price of Stability* by Anshelevich et al. [6], who examined congestion games with fair cost allocation. A key observation in this paper was that, for games where *best response dynamics* are guaranteed to converge to an NE, one can use the value of the potential function for the game to upper bound the cost of the worst case ratio between the best NE and OPT.

They show logarithmic upper bounds on the Price of Stability with respect to the sum-cost, in the case where fair-cost sharing is used, and matching lower bounds for the directed case. Since its introduction, much has been shown for this measure, and we now outline some of the more important results related to our setting.

For games where players share a single source and destination, which we refer to as *symmetric games*, it is known that the Price of Stability is 1 [7].

In the case of undirected networks, where players may have unique source and destinations, much work has been done on resolving the open problem of the Price of Stability. The upper bound of $O(\log n)$ has been improved on, first by Fiat et al. [31], who show that the price of stability is $O(\log \log n)$ in the case where there is a single common sink and every other vertex is a source vertex. Along with this upper bound, they show a n -player lower bound instance of $12/7$. In the more general case, where agents share a sink but not every other vertex is a source, an upper bound of $O(\log n / \log \log n)$ was shown by Li [41]. The weighted variant of the game, where players pay a share of the cost of an edge proportional to their weight is studied by Chen and Roughgarden [13], and it has been shown by Albers [3] that the Price of Stability in this setting is $\Omega(\log W / \log \log W)$, where W is the sum of players weights.

The Price of Stability for undirected network design is studied by Christodoulou et al. [14], who showed for the first time a separation between the Price of Stability for undirected and directed networks.

4.3 General Games

When considering general games, where edges may have capacity less than the number of players, the actions of a player's opponents may render certain paths infeasible. This suggests that a situation may arise where some player has no option but the most expensive path in their strategy set. When considering the Price of Stability measure of network efficiency, the standard approach is to use the worst case increase through BRD as an upper bound, as any game where OPT is not a NE will have a NE which can be reached by BRD from OPT. We have already seen, in Chapter 3, that in the general setting BRD may induce an increase in max-cost by a factor of up to $n \log n$, even in the symmetric case.

In the following we will show an asymmetric game where the best NE has a max-cost which is $\Theta(n \log n)$ times more expensive than OPT, which gives us a tight bound for PoS_{mc} in general asymmetric games.

We then consider the rooted case and find that this upper bound is not tight. We improve on this upper bound, and show that for rooted games there is always a NE

which has a maximum cost of at most n times the maximum cost of the optimum profile. This, matching the previously known lower bound of n , is tight.

4.3.1 Asymmetric Games

Theorem 8. *There exists an asymmetric game Δ with n players and*

$$\text{PoS}_{mc}(\Delta) = \Theta(n \log n).$$

In the following, we will construct a game Δ with an odd number $n \geq 3$ of players where

$$\min_{S^* \in NE(\Delta)} mc_{\Delta}(S^*) \approx \frac{n}{2} H(\lfloor n/2 \rfloor) \cdot OPT_{mc}(\Delta).$$

When we come to the specifics of the game, the reader may notice that the construction presented in this section bears a striking resemblance to that used in Section 3.4 for the proof related to the worst case effect of BRD on the maximum cost of a profile, and we now outline the differences. When considering the effect of BRD it was necessary to construct a game where a particular sequence of updates was possible, to enable the maximum cost to all players in the game to increase by a certain factor. In Section 3.4 we show that it is possible for max-cost to increase by a factor arbitrarily close to $n \log n$ times the cost of the initial profile, however, the profile arrived at after BRD is not the best NE for the game. As we are now concerned with the Price of Stability w.r.t max-cost, it is necessary to construct a game where BDR from OPT terminates at a solution with a max-cost $O(n \log n)$ times the cost of OPT, where the resultant profile is the only NE (or at the very least the cheapest NE with regard max-cost). To achieve this we present a modified version of the aforementioned construction where each player, apart from the player whose cost increases, is replaced by two players, and each section of the grid has been duplicated, so that the best NE for the game is now a factor $O(n \log n)$ times the cost of OPT. We now give precise details of the construction.

The construction uses parameters $m \in \mathbb{N}$ and $\varepsilon > 0$, where ε is sufficiently small, e.g., $\varepsilon < 0.1$, and m is sufficiently large. It is useful to think of m as approaching infinity. Furthermore, for $2 \leq i \leq n$, let $k(i)$ denote the value $\lfloor (i+2)/2 \rfloor$. Let Δ be the game with underlying graph $G = (V, E)$ defined as follows (see Figure 4.2 for an illustration of the structure of G):

$$\begin{aligned}
& \{s_1, t_1, x_{1,m}, z_{1,m}\} \cup \\
V = & \{s_i, t_i, z_{i,0} \mid 1 < i \leq n\} \cup \\
& \{x_{i,j}, y_{i,j}, z_{i,j} \mid 1 < i \leq n, 1 \leq j \leq m\} \\
& \{(s_1, x_{1,m}), (t_1, z_{n,m}), (x_{1,m}, z_{1,m})\} \cup \\
E = & \left\{ \begin{array}{l} (x_{i,j}, x_{i,j-1}), (x_{i,j}, y_{i,j}), \\ (y_{i,j}, z_{i,j}), (y_{i,j}, z_{i,j-1}), (z_{i,j}, z_{i,j-1}) \end{array} \mid 1 < i \leq n, 1 < j \leq m \right\} \cup \\
& \left\{ \begin{array}{l} (z_{i,0}, z_{i-1,m}), (z_{i,0}, z_{i,1}), (z_{i,0}, y_{i,1}), (x_{i,1}, y_{i,1}), \\ (z_{i,1}, y_{i,1}), (x_{1,m}, x_{i,1}), (z_{i,0}, s_i), (s_i, z_{1,m}), (t_i, x_{i,m}) \end{array} \mid 1 < i \leq n \right\} \\
c(e) = & \begin{cases} n & \text{if } e = (x_{1,m}, z_{1,m}) \\ 2 & \text{if } e = (x_{i,j}, y_{i,j}) : 1 < i \leq n, 1 \leq j \leq m \\ 1 & \text{otherwise} \end{cases} \\
p(e) = & \begin{cases} 2 + 2\varepsilon & \text{if } e = (x_{1,m}, z_{1,m}) \\ \frac{1}{k^{(i)}2^{j-1}} & \text{if } e = (x_{i,j}, y_{i,j}) : 1 < i \leq n, 1 \leq j \leq m \\ H(k^{(i)}) - \frac{3}{k^{(i)}2^j} + \frac{\varepsilon}{k^{(i)m+j}} & \text{if } e = (y_{i,j}, z_{i,j}) : 1 < i \leq n, 1 \leq j \leq m \\ \frac{\varepsilon}{k^{(i)+j}} & \text{if } e = (y_{i,j}, z_{i,j-1}) : 1 < i \leq n, 1 \leq j \leq m \\ 0 & \text{otherwise} \end{cases}
\end{aligned}$$

The source and sink of player i , for $1 \leq i \leq n$, are s_i and t_i , respectively. We refer to the nodes of the form $x_{i,j}$ as the x -row, to the nodes of the form $y_{i,j}$ as the y -row, and to the nodes of the form $z_{i,j}$ as the z -row. We divide the main part of the graph into grids as follows: For any $i \geq 2$, the i -th grid is the induced subgraph of all x, y, z nodes with subscript i, j for any j . Note that the edge costs in pairs of consecutive grids, namely the $(2k-2)$ -th and $(2k-1)$ -th grid, are the same for $2 \leq k \leq \lceil n/2 \rceil$. For fixed i and j , we refer to the subgraph induced by $x_{i,j}, y_{i,j}, z_{i,j}$ and $z_{i,j-1}$ as *column j of the i -th grid*. For simplicity we will refer to the costs of the connections from $y_{i,j}$ to $x_{i,j}, z_{i,j}, z_{i,j-1}$ as $ab_{k^{(i)},j}, a_{k^{(i)},j}, b_{k^{(i)},j}$, respectively.

Lemma 16. $OPT_{mc}(\Delta) \leq \frac{2+2\varepsilon}{n}$.

Proof. Consider the profile S' where player 1 uses the path $(s_1, x_{1,m}, z_{1,m} \rightarrow z_{n,m}, t_1)$ and player i , for $2 \leq i \leq n$, uses the path $(s_i, z_{1,m}, x_{1,m}, x_{i,1} \rightarrow x_{i,m}, t_i)$. All n players share the edge $(x_{1,m}, z_{1,m})$ and use no other edge with non-zero cost. Each player

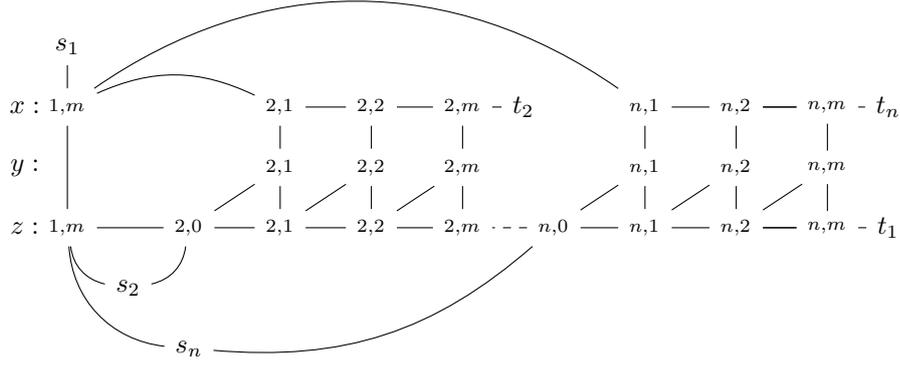


Figure 4.2: Asymmetric game Δ with $\text{PoS}_{mc}(\Delta) = \Theta(n \log n)$.

has cost $(2 + 2\varepsilon)/n$. To show that this profile is optimal w.r.t max-cost, we will show that an alternative to player 1's current path must use some edge $y_{i,j}, z_{i,j}$ for $2 \leq i \leq n, 1 \leq j \leq m$. To reach her destination, player 1 must travel between some $x_{i,j}, z_{i,j}$. If $i \neq 1$, she may use one of the following two sub-paths to do so: $(x_{i,j}, y_{i,j}, z_{i,j-1}, z_{i,j})$, $(x_{i,j}, y_{i,j}, z_{i,j})$. In both cases the chosen subpath will be immediately proceeded by the edge $(x_{i,j-1}, i_{i,j})$, and followed later by the edge $z_{i,m}, z_{i+1,0}$. In the first case, her use of the edges $(x_{i,j-1}, x_{i,j}), (y_{i,j}, z_{i,j-1}), (z_{i,j-1}, z_{i,j}), (z_{i,m}, z_{i+1,0})$ forms a cut of the i^{th} grid, blocking player i from reaching their destination. Any profile where player 1 uses a subpath of the first type is therefore not a feasible solution. Now consider a path containing the second sub-path. Player 1 must pay the full cost of any edge $(y_{i,j}, z_{i,j})$ they use. The cheapest of these, $(y_{2,1}, z_{2,1})$ has a cost of roughly $3/4$.

The optimal max-cost is therefore at most $(2 + 2\varepsilon)/n$. \square

Lemma 17. $\min_{S \in NE(\Delta)} mc_{\Delta}(S) \geq a_{k(n),m} = H(\lfloor \frac{n+2}{2} \rfloor) - \frac{3}{k(n)2^m} + \frac{\varepsilon}{(k(n)+1)m}$.

Proof. We first establish some inequalities that will be required to identify improving moves in the remainder of the proof. Recall that $k(i) = \lfloor (i+2)/2 \rfloor$ and note that this implies $k(i) > i/2$. Since $ab_{k(i),1} = 1/k(i) < 2/i$ and $b_{k(i),1} = \frac{\varepsilon}{k(i)+1} < \frac{2\varepsilon}{i}$ we have

$$\frac{2 + 2\varepsilon}{i} > ab_{k(i),1} + b_{k(i),1} \quad (4.1)$$

for all $2 \leq i \leq n$. For $j < m$ we have $ab_{k(i),j} = 2ab_{k(i),j+1}$. Furthermore, $b_{k(i),j} = \frac{\varepsilon}{k(i)+j}$ and $b_{k(i),j+1} = \frac{\varepsilon}{k(i)+j+1}$, so $b_{k(i),j} > b_{k(i),j+1}$. We get

$$b_{k(i),j} + ab_{k(i),j}/2 > b_{k(i),j+1} + ab_{k(i),j+1} \quad (4.2)$$

for all $2 \leq i \leq n$ and $1 \leq j < m$. Noticing that $a_{k(i),j} + ab_{k(i),j} = H(k(i)) - \frac{3}{k(i)2^j} + \frac{\varepsilon}{k(i)m+j} + \frac{1}{k(i)2^{j-1}} = H(k(i)) - \frac{1}{k(i)2^j} + \frac{\varepsilon}{k(i)m+j}$ and $a_{k(i),j+1} + ab_{k(i),j+1}/2 = H(k(i)) - \frac{3}{k(i)2^{j+1}} + \frac{\varepsilon}{k(i)m+j+1} + \frac{1}{k(i)2^{j+1}} = H(k(i)) - \frac{1}{k(i)2^j} + \frac{\varepsilon}{k(i)m+j+1}$, we get

$$a_{k(i),j} + ab_{k(i),j} > a_{k(i),j+1} + ab_{k(i),j+1}/2 \quad (4.3)$$

for all $2 \leq i \leq n$ and $1 \leq j < m$.

We have $ab_{k(i),m}/2 + a_{k(i),m} = \frac{1}{k(i)2^m} + H(k(i)) - \frac{3}{k(i)2^m} + \frac{\varepsilon}{k(i)m+m} = H(k(i)) - \frac{1}{k(i)2^{m-1}} + \frac{\varepsilon}{(k(i)+1)m}$ and $ab_{k(i)+1,1}/2 + a_{k(i)+1,1} = \frac{1}{2(k(i)+1)} + H(k(i)+1) - \frac{3}{2(k(i)+1)} + \frac{\varepsilon}{(k(i)+1)m+1} = H(k(i)) + \frac{\varepsilon}{(k(i)+1)m+1}$. For the former to be greater than the latter, we need that $\frac{\varepsilon}{(k(i)+1)m} - \frac{1}{k(i)2^{m-1}} > \frac{\varepsilon}{(k(i)+1)m+1}$, which is equivalent to $\frac{\varepsilon}{(m(k(i)+1))^2 + m(k(i)+1)} > \frac{1}{k(i)2^{m-1}}$. If m is large enough, this holds and we have:

$$a_{k(i),m} + ab_{k(i),m}/2 > a_{k(i)+1,1} + ab_{k(i)+1,1}/2 \quad (4.4)$$

Furthermore, we have $ab_{k(i),1}/2 + a_{k(i),1} = \frac{1}{2k(i)} + H(k(i)) - \frac{3}{2k(i)} + \frac{\varepsilon}{k(i)m+1} = H(k(i) - 1) + \frac{\varepsilon}{k(i)m+1}$. For large enough m , we therefore also have:

$$a_{k(i),m} + ab_{k(i),m}/2 > a_{k(i),1} + ab_{k(i),1}/2 \quad (4.5)$$

As $k(i+1), k(i+2) \in \{k(i), k(i)+1\}$, we can combine (4.4) and (4.5) to get:

$$a_{k(i),m} + ab_{k(i),m}/2 > a_{k(i'),1} + ab_{k(i'),1}/2 \text{ for } i' \in \{i+1, i+2\} \quad (4.6)$$

In the rest of the proof we show that player 1 must pass through $x_{n,m}$ or $x_{n-1,m}$ and hence use edges $(x_{i,m}, y_{i,m})$ and $(y_{i,m}, z_{i,m})$ for $i = n-1$ or for $i = n$ in any NE (note that using edges $(y_{i,m}, z_{i,m-1})$ and $(z_{i,m-1}, z_{i,m})$ would block player i from reaching t_i), thus paying at least $a_{k(n-1),m} = a_{k(n),m}$. To show that player 1 must pass through $x_{n,m}$ or $x_{n-1,m}$ in any NE, we show that in all other cases some player has an improving move.

Let S be a NE. Assume that the path S_1 of player 1 does not pass through $x_{n,m}$ or $x_{n-1,m}$. Consider the last x node (i.e., node in the x -row) that the path S_1 of player 1 visits. Let $x_{i,j}$ be that node. Note that $i < n-1$ or $j < m$. If $i \geq 2$, note that $x_{i,j}$ must be followed directly by $y_{i,j}$ and $z_{i,j}$ on S_1 because using the subpath $(y_{i,j}, z_{i,j-1}, z_{i,j})$ would block player i from reaching her destination t_i (by the cut argument described in the proof of Lemma 16). If $i = 1$, the path S_1 must use the edge $(x_{1,m}, z_{1,m})$.

We distinguish three cases for the location of the last x node on S_1 as follows.

Case 1: The last x node on S_1 is $x_{1,m}$. Observe that S_1 must travel from $x_{1,m}$ to $z_{1,m}$ and then reach t_1 by visiting all nodes $z_{i,j}$ from left to right, possibly visiting some y nodes in between adjacent z nodes.

The i -th grid is separated from the rest of the graph by a cut of four edges, namely the edges $(z_{i,0}, z_{i,1})$, $(z_{i,0}, y_{i,1})$, $(x_{1,m}, x_{i,1})$, and $(z_{i,m}, z_{i+1,0})$. (For $i = n$, the former three edges already form such a cut.) As player 1 passes through all grids and uses two edges from each cut (or one edge for $i = n$), and as player i must reach t_i and therefore use at least one edge of the cut of the i -th grid as well, we can see that in S the only players who use edges in the i -th grid are player 1 and player i .

First, we claim that no other player can share the edge $e_1 = (x_{1,m}, z_{1,m})$ with player 1. Assume otherwise. Let i be the largest index of a player who shares edge e_1 with player 1. Observe that player i pays at least $c(e_1)/i = (2 + 2\varepsilon)/i$ for e_1 . If the path $(s_i, z_{i,0}, y_{i,1}, x_{i,1} \rightarrow x_{i,m}, t_i)$ is available to player i , moving to that path of cost $b_{k(i),1} + ab_{k(i),1}$ will be an improving move by (4.1). The only way that path could be blocked is if player 1 were to pass through $(z_{i,0}, y_{i,1}, z_{i,1})$, but then player 1 would have an improving move by replacing $(z_{i,0}, y_{i,1}, z_{i,1})$ by $(z_{i,0}, z_{i,1})$ (which must be available if player i uses e_1).

Hence, player 1 is the only player using e_1 and therefore pays $2 + 2\varepsilon$ to reach $z_{2,0}$. The node after $z_{2,0}$ on S_1 could be $y_{2,1}$ or $z_{2,1}$. In the former case, player 1 has an improving move, namely replacing her path to $y_{2,1}$ of cost at most $2 + 2\varepsilon$ by $(s_1, x_{1,m}, x_{2,1}, y_{2,1})$ of cost at most $1/2$. In the latter case, i.e., if player 1 uses $(z_{2,0}, z_{2,1})$, note that player 2 must be using $(s_2, z_{2,0}, y_{2,1}, x_{2,1} \rightarrow x_{2,m}, t_2)$ of cost $b_{k(2),1} + ab_{k(2),1} = 1/2 + \varepsilon/3$ as any other path would use edge $(y_{2,1}, z_{2,1})$ and cost at least $a_{k(2),1} = H(k(2)) - 3/(2k(2)) + \varepsilon/(k(2)m + 1) = H(2) - 3/4 + \varepsilon/(2m + 1) > 3/4$ (which is larger than $1/2 + \varepsilon/3$ as $\varepsilon < 0.1$). Then player 1 has an improving move by replacing her path to $z_{2,1}$ of cost $2 + 2\varepsilon$ by the path $(s_1, x_{2,1}, y_{2,1}, z_{2,1})$ of cost $ab_{k(2),1}/2 + a_{k(2),1} = \frac{1}{4} + H(2) - \frac{3}{4} + \frac{\varepsilon}{2m+1} = 1 + \frac{\varepsilon}{2m+1}$.

Case 2: The last x node on S_1 is $x_{i,j}$ for some $i \geq 2$, $j < m$. Note again that edges in the i -th grid can only be used by player 1 and player i . Node $x_{i,j}$ must be followed on S_1 by $y_{i,j}$, $z_{i,j}$ because going via $y_{i,j}$, $z_{i,j-1}$, $z_{i,j}$ would cut off player i from her destination.

From $z_{i,j}$ the path S_1 could continue either via $(z_{i,j}, y_{i,j+1}, z_{i,j+1})$ or via $(z_{i,j}, z_{i,j+1})$.

In the former case, player i must be using the edge $(z_{i,j}, z_{i,j+1})$ as otherwise player 1 would have an improving move by replacing $(z_{i,j}, y_{i,j+1}, z_{i,j+1})$ by that edge. Hence, player i does not use edge $(x_{i,j}, y_{i,j})$ and player 1 pays its full price $ab_{k(i),j}$. Thus, player 1 has an improving move by replacing her current path from $x_{i,j}$ to $y_{i,j+1}$ of cost greater than $ab_{k(i),j}$ by the path $(x_{i,j}, x_{i,j+1}, y_{i,j+1})$ of cost at most $ab_{k(i),j+1} < ab_{k(i),j}$.

Now consider the latter case, i.e., the node after $z_{i,j}$ in S_1 is $z_{i,j+1}$. Player i must visit $z_{i,j-1}$ and continue from there via (a) $(z_{i,j-1}, y_{i,j}, x_{i,j} \rightarrow x_{i,m}, t_i)$, (b) $(z_{i,j-1}, z_{i,j}, y_{i,j+1}, x_{i,j+1} \rightarrow x_{i,m}, t_i)$, or (c) a path starting with $(z_{i,j-1}, z_{i,j}, y_{i,j+1}, z_{i,j+1})$. By (4.2), the cost of path (b) to player i is less than that of the path (a). Any path of type (c) has cost at least $a_{k(i),j+1} > H(k(i)) - 3/(k(i)2^{j+1}) \geq H(k(i)) - 3/8 > 1$, while the path (b) has cost

$$b_{k(i),j+1} + ab_{k(i),j+1} = \frac{\varepsilon}{k(i) + j + 1} + \frac{1}{k(i)2^j} < 1.$$

Hence, player i must use path (b), otherwise she will have an improving move. Then player 1 has an improving move by replacing her current path from $x_{i,j}$ to $z_{i,j+1}$ of cost $ab_{k(i),j} + a_{k(i),j}$ by the path $(x_{i,j}, x_{i,j+1}, y_{i,j+1}, z_{i,j+1})$ of cost $ab_{k(i),j+1}/2 + a_{k(i),j+1}$ (note that edge $(x_{i,j+1}, y_{i,j+1})$ is shared with player i) as shown by (4.3).

Case 3: The last x node on S_1 is $x_{i,m}$ for $2 \leq i < n-1$. Player 1 must continue from $x_{i,m}$ via $(x_{i,m}, y_{i,m}, z_{i,m}, z_{i+1,0})$, so her current cost is at least $ab_{k(i),m}/2 + ak(i), m$. To reach t_1 , player 1 must pass through the i' -th grid for all $i \leq i' \leq n$, and hence edges in any such grid can only be used by player 1 and player i' . Note also that player 1 must pass through node $z_{i+2,0}$.

Similar to Case 1, we claim that no player i' for $i' > i$ can use the edge $e_1 = (x_{1,m}, z_{1,m})$. Assume otherwise. Let $i' > i$ be the largest index of a player who uses edge e_1 . Observe that player i' pays at least $c(e_1)/i' = (2 + 2\varepsilon)/i'$ for e_1 . If the path $(s_{i'}, z_{i',0}, y_{i',1}, x_{i',1} \rightarrow x_{i',m}, t_{i'})$ is available to player i' , moving to that path of cost $b_{k(i'),1} + ab_{k(i'),1}$ will be an improving move by (4.1). The only way that path could be blocked is if player 1 were to pass through $(z_{i',0}, y_{i',1}, z_{i',1})$, but then player 1 would have an improving move by replacing $(z_{i',0}, y_{i',1}, z_{i',1})$ by $(z_{i',0}, z_{i',1})$ (which must be available if player i' uses e_1).

Let $i' \in \{i+1, i+2\}$ be such that player i' reaches the i' -th grid via the edge $(s_{i'}, z_{i',0})$. We claim that such an i' exists. Note that the only other possibility for a player $i' \in \{i+1, i+2\}$ to reach the i' -th grid is by a path starting $(s_{i'}, z_{1,m}, z_{2,0})$ and eventually reaching $x_{1,m}$ and arriving at the i' -th grid via $(x_{1,m}, x_{i',1})$. As the edge $(z_{1,m}, z_{2,0})$ has capacity 1, at most one player among $i+1$ and $i+2$ can use such a path, and the other must reach the i' -th grid via node $z_{i',0}$.

As observed above, player 1 must visit $z_{i',0}$. Player 1 has two options from $z_{i',0}$: If player 1 goes directly to $z_{i',1}$, the cheapest way for player i' to continue her path from $z_{i',0}$ is $(z_{i',0}, y_{i',1}, x_{i',1} \rightarrow x_{i',m}, t_{i'})$ of cost $b_{i',1} + ab_{i',1} = \varepsilon/(k(i') + 1) + 1/k(i') \leq 1/2 + \varepsilon/3$. (Any other path would use the edge $(y_{i',1}, z_{i',1})$ of cost $a_{i',1} = H(k(i')) - 3/(2k(i')) + \varepsilon/(k(i')m + 1) \geq 1.5 - 3/4 = 3/4$, which is larger than $1/2 + \varepsilon/3$ since $\varepsilon < 0.1$.) Player i' must use that path, otherwise she has an improving move. By (4.6), player 1 has an improving move by reaching $z_{i',1}$ via $(s_1, x_{1,m}, x_{i',1}, y_{i',1}, z_{i',1})$ instead of her current path. (Note that the cost $ab_{k(i'),1}$ is shared with player i' .) If player 1 uses the second option and continues from $z_{i',0}$ via $(z_{i',0}, y_{i',1}, z_{i',1})$, she has an improving move by reaching $y_{i',1}$ via $(s_1, x_{1,m}, x_{i',1}, y_{i',1})$ of cost at most $ab_{k(i'),1} = 1/k(i') \leq 1/2$ instead of her current path of cost at least $a_{k(i),m} = H(k(i)) - 3/(k(i)2^m) + \varepsilon/((k(i)+1)m) \geq 1.5 - 3/2^{m+1} > 1$ (for $m \geq 2$). No matter whether player 1 uses the first or the second option for continuing from $z_{i',0}$, she has an improving move.

In all three cases for the location of the last x node on S_1 , at least one player has an improving move. Therefore, in any NE S of Δ , S_1 will pass through $x_{n,m}$ or $x_{n-1,m}$ and contain an edge with cost

$$a_{k(n),m} = a_{k(n-1),m} = H(\lfloor \frac{n+2}{2} \rfloor) - \frac{3}{k(n)2^m} + \frac{\varepsilon}{(k(n)+1)m}$$

that is used only by player 1. □

Proof of Theorem 8. We have constructed a general asymmetric game Δ with an optimal strategy profile with a max-cost of at most $\frac{2+2\varepsilon}{n}$ by Lemma 16, while the max-cost of the best NE is arbitrarily close to $H(\lfloor \frac{n+2}{2} \rfloor)$ by Lemma 17. Hence, $PoS_{mc}(\Delta)$ approaches $nH(\lfloor \frac{n+2}{2} \rfloor)/(2 + \varepsilon) = \Theta(n \log n)$ arbitrarily closely. □

4.3.2 Rooted Games

We now consider the case where all players have the same sink, and find that the upper bound of $O(n \log n)$ on the price of stability is not tight. First, we note that BRD starting with a strategy profile with optimal max-cost can increase the max-cost by a factor of $\Omega(n \log n)$ even in the symmetric case.

Hence, in order to bound the price of stability in both the rooted and symmetric case, we cannot use the standard proof technique of analysing BRD starting with the optimal strategy profile. Instead, we use a different approach that may be of independent interest. We iteratively discard a single expensive path from the NE reached by BRD and recombine the remaining $n - 1$ paths with the optimal strategy profile, until a NE with small max-cost is obtained. In this way we are able to show that for every general rooted game Δ there is always a NE where no player pays more than n times $OPT_{mc}(\Delta)$.

Theorem 9. *For any rooted game Δ in directed or undirected networks, $PoS_{mc}(\Delta) \leq n$.*

Proof. We present the proof for directed networks. The result for undirected networks follows using the standard transformation of an undirected network into an equivalent directed network discussed in Section 2.3.

Let Δ be a rooted game with directed graph $G = (V, E)$, n players, each of whom may have a unique source s_i , and a common destination t . Note that G can be modified by the addition of a single source s , connected to each s_i by an edge with capacity 1 and cost 0, to produce a game which is symmetric.

Let S be the optimal strategy profile with respect to max-cost. Without loss of generality, we can scale the edge costs so that the sum-cost of S is n . This implies $mc(S) \geq 1$.

Consider the NE S^* that is obtained from S using BRD. If $mc(S^*) \leq n$, then $PoS(\Delta) \leq n$ and we are done. Otherwise, we have $n < mc(S^*) \leq sc(S^*) \leq \Phi(S^*) < \Phi(S)$. Let $\Phi(S) = n + \alpha$ and $mc(S^*) = n + \beta$ for some $\alpha, \beta > 0$, and let $\Phi(S^*) = \Phi(S) - \delta$ for some $\delta > 0$. Note that $0 < \beta \leq \alpha - \delta$. The following table illustrates

these quantities:

	mc	Φ
S	≥ 1	$n + \alpha$
S^*	$n + \beta$	$n + \alpha - \delta$

Now consider the $(n - 1)$ -player profile S_{-1}^* that consists of the $n - 1$ cheapest (in terms of cost to the respective player) strategies in S^* . For simplicity we will call the player with the most expensive path player 1. As the change in potential function equals the cost to an individual player when making some change in strategy, we have that $\Phi(S_{-1}^*) = \Phi(S^*) - mc(S^*) = n + \alpha - \delta - (n + \beta) = \alpha - \beta - \delta$.

We construct a new n -player strategy profile S' by combining S and S_{-1}^* using an augmentation step in a suitably defined flow network. (We refer the reader to [2] for background on network flow, residual networks, and augmenting paths.) First, define the general network $\bar{G} = (V, \bar{E})$ from $G = (V, E)$ by letting $\bar{E} = \{e \in E \mid x_e(S) > 0 \text{ or } x_e(S_{-1}^*) > 0\}$ and setting the capacity $\bar{c}(e)$ for each $e \in \bar{E}$ to $\bar{c}(e) = \max\{x_e(S), x_e(S_{-1}^*)\}$.

The strategy profile S_{-1}^* induces a flow f of value $n - 1$, from each s_i for $i \in [2..n]$ to t in \bar{G} . The network \bar{G} admits a flow of value n from each s_i for $i \in [1..n]$ to t as the profile S induces such a flow. Hence, the residual network \bar{G}_f of \bar{G} with respect to flow f admits an augmenting path P from s_1 to t . Let f' be the flow of value n obtained by augmenting f with P . Decompose the flow f' into n paths, one from each s_i for $i \in [n]$ to t , and let S' be the strategy profile corresponding to these n paths.

In going from f to f' , the flow on any edge increases by at most 1, and every edge on which the flow increases satisfies $x_e(S) > 0$. Let X be the set of edges on which the flow increases. Observe that $\Phi(S') \leq \Phi(S_{-1}^*) + p(X)$, because increasing the number of players on an edge e by 1 adds at most $p(e)$ to the potential. As $X \subseteq \{e \in E \mid x_e(S) > 0\}$, we have $p(X) \leq sc(S) = n$ and hence $\Phi(S') \leq \alpha - \beta - \delta + n < \Phi(S^*)$.

Let S^{**} be the NE obtained from S' via BRD. Note that $\Phi(S^{**}) \leq \Phi(S') < \Phi(S^*)$. If $mc(S^{**}) \leq n$, we have found a NE with max-cost at most n times the optimal max-cost and we are done. Otherwise, we can repeat the construction that we used to create S^{**} from S' , but starting with S^{**} in place of S' . Each time we

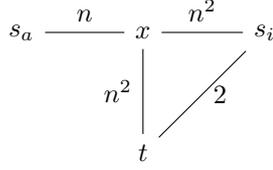


Figure 4.3: Rooted game Δ_u where the best NE is roughly n times the cost of the optimum profile, in terms of max-cost.

repeat the construction and obtain a NE with max-cost greater than n , that NE has strictly smaller potential than the previous NE. As the number of strategy profiles is finite, we must eventually obtain a NE whose max-cost is at most n . This shows $\text{PoS}_{mc}(\Delta) \leq n$. \square

4.4 Uncapacitated Games

For the uncapacitated case, the upper bound of $O(n \log n)$ applies. It is known that there is no loss of efficiency in symmetric networks, as the optimum profile will also be a NE. We now examine the open problem of the rooted and asymmetric cases, and show that the Price of Stability w.r.t. max-cost is $\Theta(n)$ in the worst case. We do this by first showing a rooted construction where the best NE is a factor of n times the cost of the optimum profile, thus giving a lower bound of n for both rooted and asymmetric games. We conclude with a matching upper bound of n for the asymmetric case.

Theorem 10. *For any rooted or asymmetric game Δ_u with n players the worst case PoS_{mc} is $\Theta(n)$.*

Lemma 18. *There exists a rooted game Δ_u where*

$$\min_{S^* \in NE(\Delta_u)} mc_{\Delta_u}(S^*) \approx \frac{n}{2} \cdot OPT_{mc}(\Delta_u).$$

Proof. Let Δ_u be a game for n players with the underlying graph $G = (V, E)$ defined as follows:

$$V = \{s_a, s_i, x, t\}$$

$$E = \{(s_a, x), (s_i, x), (x, t), (s_i, t)\}$$

$$p(e) = \begin{cases} n & \text{if } e = (a, x) \\ 2 & \text{if } e = (s_i, t) \\ n^2 & \text{otherwise} \end{cases},$$

an illustration of which is shown in Figure 4.3. In this game, player 1 has a source of s_a , while all other players have a source of s_i . All players must connect to the common destination node t .

Player a has the choice of paths $A^0 = s_a, x, t$ and $A^1 = s_a, x, s_i, t$, while her opponents have the choice of paths $I^0 = s_i, t$ and $I^1 = s_i, x, t$. No player $i \in \{2..n\}$ can pay more than 2 for the path I^0 , or less than $2n$ for the path I^1 , so the former is the dominant strategy for all opponents of a , to which her best response is the path A^0 for a cost $n^2 + n$. The profile $S = (A^0, I^0, \dots, I^0)$ is therefore the only NE, and $mc(S) = n^2 + n$. Now consider the profile $S^* = (A^0, I^1, \dots, I^1)$: note that $p_a(S^*) = 2n$ and $p_{i \in \{2..n\}}(S^*) = n^2/n + n^2/(n-1) = 2n + 1 + 1/(n-1)$. Any player $i \in \{2..n\}$ can improve with the path I^0 for a cost of 2, after which player a will pay $2n + 1 + 1/(n-1)$ and all players using I^1 will pay more than this. Some player must always pay at least $2n + 1 + 1/(n-1)$, so the profile S^* is optimal w.r.t max-cost, and so $\text{OPT}_{mc}(\Delta_u) = 2n + 1 + 1/(n-1)$. We therefore have a game where

$$\min_{S \in \text{NE}(\Delta_u)} mc_{\Delta_u}(S) / \text{OPT}_{mc}(\Delta_u) = (n^2 + n) / (2n + 1 + 1/(n-1)) \approx \frac{n}{2}.$$

□

Proof of Theorem 10. In Lemma 18 we show a rooted game where the best NE is roughly a factor of n times more expensive (in terms of max-cost) than the optimum. As rooted games are a subset of asymmetric games this gives us a lower bound of $\Omega(n)$ for both cases.

In Lemma 3 we showed that it is not possible for BRD to increase the max-cost of a profile by a factor greater than n times the max-cost in the initial profile, and so it follows that given any game there must always be an equilibrium with a max-cost at most n times the max-cost of the optimal solution. We have therefore shown matching upper and lower bounds for both the asymmetric and rooted case. □

4.5 Concluding Remarks

To prove that the price of stability with respect to max-cost is bounded by n for general symmetric games, we iteratively combined a NE with large max-cost with the optimal strategy profile. It would be interesting to explore whether this method

could be turned into an efficient procedure for constructively finding a good NE. As the approach mainly relies on arguments about the reduction in potential of the strategy profiles constructed, it may be possible to apply it to other potential games.

Chapter 5

Sequential Network Design Games

In the analysis of stable solutions, an often used approach is to examine BRD from an optimal profile. The importance of the starting configuration is the main weakness of this approach, as it does not give a practical method for discovering good NE. We now propose an alternative method, which is to analyse the behaviour of agents from a start profile which they have chosen themselves.

5.1 Introduction

We now examine a variant of our general model in which the game is initialised with an empty strategy set, to which each of the n players must add their strategy. We say that players may join in any order, so long as they do so sequentially, and that their choice of strategy must be the best response to the current configuration of the game, at both the joining stage and for any subsequent updates. We place the following restrictions on players' updates: each player begins with a single update. When joining, a player will choose the best available path given the configuration at the point they join. Once all players have joined the game, we lift all restrictions and allow updates in strategy with any order and frequency.

We will refer to this variant of network design games as *sequential network design games*, and define the price of anarchy of these games to be the ratio of the worst NE to OPT.

The motivation for this method of analysis is that it gives a more realistic model in which to analyse the effects of selfish behaviour, being a closer representation of

the interactions of agents in a dynamic setting. While the traditional approach is for qualitative measure such as the Price of Stability to bound the ratio between the optimal and some Nash, we now examine a subset of stable states which we will call *reachable* Nash Equilibria (rNE), in the sense that it may be arrived at by players making selfish choices. It is worth noting that while a particularly bad, i.e. expensive, Nash Equilibrium may exist, this state may never be reached through selfish behaviour. We now analyse the negative impact of selfishness, with the aim of a more robust measure of a network's efficiency.

In this chapter we first ask the question, “can the existence of reachable equilibria be guaranteed?”, and find that in most general case of games, where edges in the underlying graph may have arbitrary capacity, there are games where no rNE exists, despite the optimal solution being stable. In the uncapacitated case, we show that there is always a reachable Nash Equilibrium.

Having defined a sub-set of stable states (rNE), we then turn to questions related to the efficiency of these states. We define the *reachable Price of Anarchy* (rPoA) to be the worst possible ratio between the worst rNE of a game and its optimal solution, the *reachable Price of Stability* (rPoS) the worst case ratio between the best rNE of a game and its optimal solution.

In the general case, we find that the ratio between the only reachable NE and OPT cannot be bounded by any function on the number of players, meaning that the rPoA and rPoS measure are unbounded w.r.t n for sum-cost and max-cost.

Having examined the general case, we turn our attention to games where the capacity of edges in the underlying graph is at least as great as the number of players. We show that, for games with two players, both the reachable Price of Anarchy and reachable Price of Stability are tightly $8/5$ w.r.t. sum-cost, for both rooted and asymmetric games. For games with an arbitrary number of players, we show that the worst case best rNE of a game is at most $O(\sqrt{n} \log^4 n)$ times the cost of the optimal solution in terms of sum-cost, and show an example where the sum-cost of the best and only rNE is $O(\sqrt{n})$ times more expensive than OPT.

A summary of results presented in this chapter are shown in bold in Table 5.1, along side previously known results.

The class of *congestion games*, into which our games fall, was first defined by Rosenthal [57], and has been widely studied by [36, 44, 49, 61, 62].

The sequential version of network design games, to which we now turn our attention, has been studied under various names. The work of Chekuri et al. [6] was motivated by the observation that the worst case Nash Equilibrium in the game for which the Price of Anarchy is n , cannot be reached if players start from an initial empty profile, and then choose their path one by one. The importance of the start profile cannot be understated; the basis for the logarithmic upper bound [6] on the price of stability relies heavily on the assumption that the start point is an optimal Steiner tree, from which it can be shown that a stable solution at most the cost of the potential value of the optimal tree can be reached through best response updates. This raises the question, what happens if there is no central authority which starts the game with a specific configuration? Another motivation for the analysis of states which can be reached in this way is that, even if there were some central authority able to dictate a favourable and stable initial profile, computing such a state is particularly difficult [18, 33, 19].

Related Work.

In the special case where every node in the network is a terminal, Fiat et al. [31] show an upper bound of $O(\log \log n)$ for games on undirected networks. This was marginally improved on by Agarwal and Charikar [1], with an upper bound of $O(\log n / \log \log n)$ for undirected graphs.

The case where players have an associated weight, and the cost share of a player for a particular edge is proportional to that player's weight, is studied by Chen and Roughgarden [13]. They show bounds of the price of stability for α -approximate Nash equilibria, i.e. configurations where no player can decrease their cost by more than a α multiplicative factor. Mirrokni and Vetta [48] study the *round* model to analyse convergence issues in competitive games.

The reachable price of anarchy and stability, which we now examine, was introduced by Leme, Syrgkanis, and Tardos [40], and has since been studied in various settings. Angelucci et al. [5] consider *isolation games*, de Jong and Uetz [20] *atomic congestion games*.

	General	Uncapacitated
symmetric	unbounded	1
rooted	⋮	$\Omega(\log n), O(\log^4 n)$ [11]
asymmetric	⋮	$\Omega(\sqrt{n}), O(\sqrt{n} \log^4 n)$

Table 5.1: Summary of results pertaining to the Price of Anarchy (sum-cost) for sequential network design games. Results in bold are shown within this chapter.

Mamageishvili and Mihalák examine rooted games where the underlying network is a ring [42], expanding on existing work here [27]. Along with results for the price of stability and the *potential-optimum* price of stability, which they show to be exactly $4/3$ and 2 respectively, they also examine the sequential price of anarchy and stability, under the name *myopic sequential price of stability*. They show that, given the worst possible ordering (with regard to joining) of players, the sequential price of anarchy is exactly 2 , and that the sequential price of stability for these games is at most $\frac{26}{19}$ and at least $4/3$, conjecturing that the lower bound is in fact the worst case.

5.2 Existence of Reachable Nash Equilibria

5.2.1 General Games

Theorem 11. *There are games in which no rNE exist (despite having good NE).*

Proof. Take the following symmetric game, with two players, each travelling from s to t , and an underlying graph defined as

$$\begin{aligned} V &= \{s, x, y, t\} \\ E &= \{(s, x), (s, y), (x, y), (x, t), (y, t)\} , \end{aligned}$$

and associated costs and capacities of

$$p(e) = \begin{cases} 1 & \text{if } e = (s, x) \text{ or } e = (y, t) \\ 0 & \text{otherwise} \end{cases} \quad c(e) = 1 .$$

Figure 5.1 gives an illustration of this game. Note that this game has a NE in the profile where the one player travels through x , the other through y , and that this

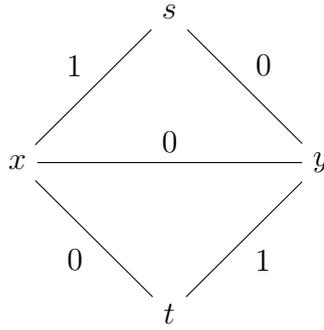


Figure 5.1: Symmetric game Δ for two players where no sequential NE exists. All edges have a capacity of 1, with costs as labelled.

profile is optimal. Now consider the actions of the first player choosing their path: as there is a zero cost s - t path travelling through x and then y , this will be her choice, which means that the second player has no feasible path. \square

5.3 Price of Anarchy for Sequential Games

5.3.1 General Games

We now show that for networks where edges may have capacity less than the total number of players, the cost of the best possible reachable stable state cannot be bounded by any function on n . We do this by showing a symmetric game where this is the case.

Theorem 12. *For any sequential game Δ , the price of anarchy, w.r.t. sum-cost or max-cost, cannot be bounded by any function on the number of players.*

Proof. Consider the symmetric game Δ with the underlying graph

$$\begin{aligned} V &= \{s, x, y, t\} \\ E &= \{(s, t), (s, x), (s, y), (x, y), (x, t), (y, t)\} \end{aligned}$$

with edge costs and capacities

$$p(e) = \begin{cases} \infty & \text{if } e = (s, t) \\ 1 & \text{if } e = (s, x) \text{ or } e = (y, t) \\ 0 & \text{otherwise} \end{cases} \quad c(e) = 1$$

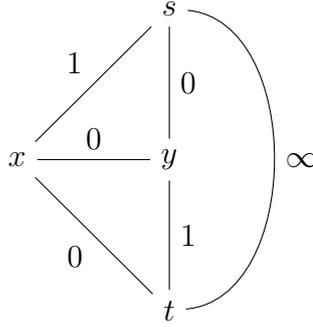


Figure 5.2: Symmetric sequential game Δ where PoA for both the sum-cost and max-cost objectives is unbounded in the number of players. All edges have a capacity of 1, and cost as labelled.

We show this structure in Figure 5.2. Both players must form a connection between the common source destination pair s, t . Observe that there is a path s, x, y, t which has a cost of 0. This being the cheapest option, it will be the choice of the first player to join the game. This leaves one path for the second player, using the direct edge (s, t) , for a cost of ∞ . This gives us the only, and therefore best, rNE for this game, which has an arbitrarily high sum cost, and an arbitrarily high max-cost. Now consider the optimal profile. Note that there are two paths, s, x, t and s, y, t , with a cost of 1. □

5.3.2 Uncapacitated Games

Symmetric Uncapacitated Games

Theorem 13. *For symmetric sequential games Δ_u , both the $rPoA$ and $rPoS$ is 1.*

Proof. Recall from Section 4.2 that the price of anarchy and the price of stability for symmetric uncapacitated games is 1. As all rNE are NE , it follows that in sequential games the $rPoA$ and $rPoS$ measure are also 1 □

Rooted Uncapacitated Games

We now turn our attention to games where all players have a common destination node, but may have unique sources. This case of our model has been studied by Charikar et al. [11] who show that in games with n players, the reachable price of anarchy is $O(\log^4 n)$, and give an example where the best rNE of the game is $\Omega(n)$

times the cost of the optimal solution [11]. With the aim of closing the gap between the upper and lower bound we now examine rooted games for two players.

Theorem 14. *The reachable price of anarchy for rooted Δ_u with two players is tightly $8/5$, w.r.t sum-cost.*

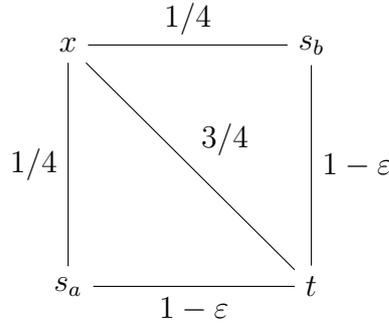


Figure 5.3: Rooted game Δ_u for two players where the sequential Price of Stability is $8/5$. Edges are labelled by cost, with $\lim_{\varepsilon \rightarrow 0}$

Lemma 19. *For two player sequential Δ_u , $PoA_{sc} \geq 8/5$.*

Proof. Consider the network illustrated in Figure 5.3. Both players have the choice of travelling directly to t , of travelling to t from x , or of travelling via their opponent's source, for a cost of at most $1 - \varepsilon$, 1 , and $3/2 - \varepsilon$ respectively. The first to pick will choose the edge from their source to t . When their opponent joins, they will likewise have the choice of a direct edge to t , of travelling to x and then to t , or of travelling to their opponent's source via x , and from here on to t . The cost of these paths with respect to their opponent's current path will be $1 - \varepsilon$, 1 , and $1 - \varepsilon/2$, so they too will choose the direct link to t . This gives us a stable solution with a sum cost of $2(1 - \varepsilon)$. Now consider the optimal solution. It is easy to see that the profile with the lowest sum-cost is the one where both players travel to x , and then share the connection to t , costing $5/4$. We have shown a game where the only stable solution which can be reached from an empty start profile costs $2(1 - \varepsilon)$, and where the optimal profile has a cost of $5/4$. As PoS is the ratio of the best sequential NE and the optimal profile, we have $PoS = \frac{2(1-\varepsilon)}{5/4} = \frac{8}{5}$, and so the worst PoS cannot be a constant smaller than $\frac{8}{5}$. \square

Lemma 20. *In the two player sequential game, $PoA_{sc} \leq 8/5$.*

Proof. Any strategy for two players will consist of edges used by only one player, and edges shared by both. For a profile S we denote the cost of edges which are unique to the strategy of player a as C_a , the cost of those unique to the strategy of b as C_b , and the total cost of those edges used by both players as C_s . Note that the sum cost for a profile is the sum of these three components, and that when two players are present, the cost to a player is the cost of the edges unique to their strategy plus half the cost of the shared edges.

Consider a stable solution S^* which has been reached from an empty strategy profile.

In the sequence of updates leading to S^* , we say that the first player to pick their path was the player with the most expensive initial choice. Let this player be a . Without loss of generality we can scale all costs so that the cost to a of her initial choice is 1.

As no player will ever use a path which is more expensive than the raw edge cost of their most expensive path, we can say that the sum cost of the stable solution will cost at most 2, and the cost of the stable solution to a player may never be more than 1, giving us the following inequalities.

$$\begin{aligned} C_a^* + C_b^* + C_s^* &\leq 2 \\ C_a^* + C_s^*/2 &\leq 1 \\ C_b^* + C_s^*/2 &\leq 1 \end{aligned}$$

As the profile S^* is stable, the cost to player a must be less than the raw edge cost of their path in the optimal solution. The same being true for player b , we have

$$\begin{aligned} C_a^* + C_s^*/2 &\leq \text{OPT}_a + \text{OPT}_s \\ C_b^* + C_s^*/2 &\leq \text{OPT}_b + \text{OPT}_s \end{aligned}$$

In a profile which is stable or optimal, the edges in S will be a contiguous sequence of edges in the graph. We name one end point of this sequence x , the other y , and say that players' sources connect to x , destinations to y . Now consider the path a uses to connect to the point x . She has an alternative to this route, which is to travel along the unique section of her optimal path, and at the point of convergence with the optimal path of b , travelling back to b 's source and from here reaching x by sharing the route currently taken by b . As S^* is stable, this will not

be an improving move. The same being true for the section of a 's path between y and their destination, we can say that

$$\begin{aligned} C_a^* &\leq OPT_a + OPT_b + C_b^*/2 \\ C_b^* &\leq OPT_a + OPT_b + C_a^*/2 \end{aligned}$$

We have shown that these inequalities represent requirements for a stable solution which has cost at most 2.

If we now take these, and with them formulate a minimisation problem for a linear program with an objective value of $OPT_a + OPT_b + OPT_s$, we find that the sum cost for the optimal solution cannot have a cost of less than $5/4$, given the existence of a stable solution with a sum-cost of 2. We therefore have that

$$PoA_{sc} \leq 2/1.25 = \frac{8}{5}$$

as required. □

Proof of Theorem 14. Matching upper and lower bounds. □

5.4 Price of Stability for Sequential Games

5.4.1 General Games

In section 5.3 we show a symmetric game Δ where the only rNE, by any ordering of players, is unbounded in the number of players.

Corollary 1. *For general games Δ , rPoS is unbounded in the number of players.*

5.4.2 Uncapacitated Games

We now examine uncapacitated asymmetric sequential games. We begin with a lower bound of $\Omega(\sqrt{n})$ by constructing a game where the cheapest rNE is roughly a factor \sqrt{n} times more expensive than the optimum profile. We then show an upper bound of $O(\sqrt{n} \log^4 n)$.

Lower Bound.

Theorem 15. *There exists a sequential asymmetric game Δ_u , where the best possible rNE has a sum cost $\Omega(\sqrt{n})$ times that of the optimum.*

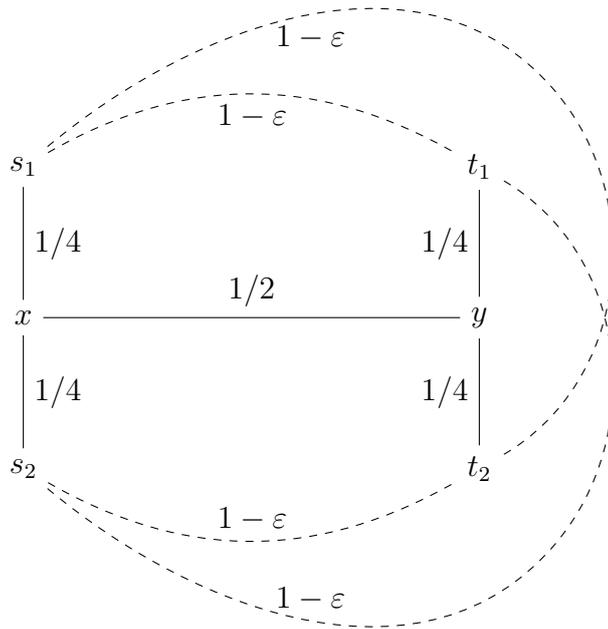


Figure 5.4: Asymmetric sequential game Δ_u for 4 players, where PoS_{sc} is $\Omega(\sqrt{n})$. Each player has a unique source destination pair, namely $(s_1, t_1), (s_1, t_2), (s_2, t_1), (s_2, s_2)$. Solid edges are those used in the optimal profile, dashed edges being those used in the only stable solution. Edges are labelled by cost, and the value of ε can be thought of as approaching 0.

Proof. Let Δ_u be a game for n players, where n is a square number, with the underlying network defined as

$$\begin{aligned} V &= \{s_i, t_i : i \in [\sqrt{n}]\} \cup \{x, y\} \\ E &= \{(s_i, x), (y, t_i), (s_i, t_j) : i, j \in [\sqrt{n}]\} \cup \{(x, y)\} \\ p(e) &= \begin{cases} 1 - \varepsilon & \text{if } e = (s_i, t_j) : i, j \in [\sqrt{n}] \\ 1/2 & \text{if } e = (x, y) \\ 1/4 & \text{otherwise,} \end{cases} \end{aligned}$$

as shown in Figure 5.4. We define the source and destination nodes of each player such that each player has a unique s, t pair. Let the node s_i be the source of players $\{\sqrt{n}(i-1) + j : j \in [\sqrt{n}]\}$, and let t_i be the destinations for players $\{\sqrt{n}(j-1) + i : j \in [\sqrt{n}]\}$. First we will define possible $s-t$ paths for a player i . By D_i we denote the direct connection between the source and destination node of player i . Player i may also connect to their destination by travelling through x, y , which we denote as E_i . Alternatively, a player may reach some other source node via x and from there travel directly to their destination, or travel to some destination other than their own and from there to their destination via y . They also have the option of first reaching some source other than their own via x , travelling from there to some destination other than their own, finally reaching their destination via y . We denote the cheapest of these paths (with respect to the current strategy profile) as F_i .

Now consider the actions of the first player to join. We will refer to this player as a . Note that, as no edges are currently in use, she will pay the raw edge cost of whichever path she chooses.

Noticing that $p(D_a) = 1 - \varepsilon < p(E_a) = 1$, and that any path of the form F_a will use some edge (s_i, t_j) and at least two other edges with a cost of $1/4$ giving $p(F_a) = 3/2 - \varepsilon$, we have that D_a is the best response for a .

The second player to join, who we will call b , will have the choice of D_b , E_b , and F_b . As player a only uses one edge, we need only adjust the cost of F_b . The cheapest path of the form F_b will share player a 's path, paying at least an additional $1/2$ to reach the source of a from their own source, or to reach their destination from the destination node of a . This gives us $p(F_b) = (1 - \varepsilon)/2 + 1/2 = 1 - \varepsilon/2$.

Hence, the second player to join will also have the best response of the direct edge connection between their source and destination. The cost of the paths available to each subsequent player to join will be as they were for player b . After all n players have joined, we have the stable profile S , where each player uses the direct connection between their source destination pair. This gives us $sc(S) = n \cdot (1 - \varepsilon)$, and $mc(S) = 1 - \varepsilon$.

Now consider the optimal profile. Let S^* be the profile in which all players use the path E_i . This profile uses all edges except the direct connections between each s, t pair. Note that there are \sqrt{n} source nodes, each of which is connected to the node x for a cost of $1/4$, and \sqrt{n} destination, each of which connects to y for a cost of $1/4$, and a single connection (x, y) with a cost of $1/2$. The cost of S^* is therefore $(\sqrt{n} + 1)/2$ in total, and $(\sqrt{n} + 1)/2n$ to each player.

In both the optimal and stable profiles outlined above, the maximum cost to any player is exactly $1/n$ times the sum cost of the solution, and so the ratio between the sum-costs of the solutions will also be the ratio between the max-costs of the solutions.

$$\text{PoS} \geq \frac{n(1 - \varepsilon)}{\sqrt{n} + 1} = \Omega(\sqrt{n})$$

□

Upper Bound. We will now show that the price of anarchy for asymmetric network creation games is $O(\sqrt{n} \log^4 n)$.

Our proof is largely based on the unpublished manuscript *Best Response Dynamics in Network Creation with Egalitarian Cost Sharing* [11], by Charikar et al. (a preliminary version of their SPAA paper [12]), which shows an upper bound of $O(\log^4 n)$ in the rooted case. Our proof uses the same approach, which is to introduce a threshold by which we divide players into two categories (by the reduction in their costs as a result of subsequent players joining the game), showing first a method to construct a solution for one group, and then introducing a bound on the second group based relative to the cost to the first. The difference in our proof is that we first use a result on set partitioning to reduce an asymmetric game into at most \sqrt{n} groups, which we can then analyse in a similar way to the rooted case. In this way, we can show that the total cost to players in both categories is bounded above by $O(\sqrt{n} \log^4 n)$ times the cost of the optimal solution to all players.

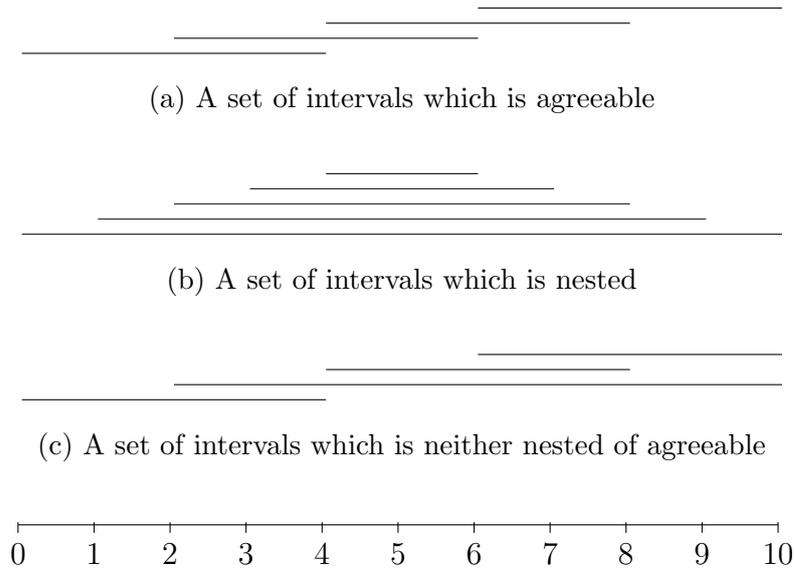


Figure 5.5: An illustration of intervals which are agreeable, nested, and neither.

We now outline the result on set partitioning on which we base our result. This result pertains to the partitioning of an arbitrary set of intervals, into groups which are either *agreeable* or *nested*.

Lemma 21. *A set of k intervals can be partitioned into $O(\sqrt{k})$ sets that are nested or agreeable [24].*

A set of intervals is called nested if it is laminar and there is a point contained within all intervals, i.e. for any two intervals, one is contained within the other. Agreeable intervals are so called if they can be ordered in such a way that the sequence of left endpoints and the sequence of right endpoints are both non-decreasing. See Figure 5.5 for an illustration of sets of intervals which are agreeable and nested. We will not go into the specific method by which Lemma 21 is proved, and direct the reader to [24] for full details.

Theorem 16. *For sequential asymmetric Δ_u the price of stability w.r.t. sum-cost is $O(\sqrt{n} \log^4 n)$.*

Consider an asymmetric sequential game Δ_u with n players.

Recall from Section 2.2 the potential function Φ , which for a given solution is defined as

$$\Phi = \sum_{e \in E} \left(\sum_{i=1}^{n(e)} \frac{c(e)}{i} \right),$$

where $c(e)$ is the cost of an edge, and $n(e)$ is its usage in the given profile, i.e. the number of paths to which it belongs.

We now define the cost function we will use, which we will call the *cost share*. For each edge, order the paths to which it belongs by time. For an edge $e \in E$, its cost share in the i^{th} path is calculated as $c(e)/i$. The cost share for a path p_i , which we denote $c(i)$, is the sum of the cost shares of each edge in this path. The sum of the cost shares of all player's paths is by definition the potential of the solution.

$$\Phi = \sum_{i \in n} c(i) .$$

For each edge $e \in E$ we define its *revised cost* as $c^+(e) = c(e) / (n(e) + 1)$. For a path p , we define the revised cost of this path to be $c^+(p) = \sum_{e \in p} c^+(e)$.

Our aim is to select a set of paths, connecting each player's source to their destination, so as to minimise the total cost of the selected paths. The game is split into two phases.

Phase One. In the first, which we will call the *joining* phase, players join the game one at a time, choosing a path which connects their source to their destination. Assume the *arrival sequence* in this phase is $1, 2, \dots, n$. Upon joining the game, a player $i \in [n]$ chooses a path p_i using a greedy algorithm, which minimises the revised cost of the chosen path ($c^+(i)$) with respect to the solution induced by p_1, p_2, \dots, p_{i-1} . In this phase, each player has exactly one update, which is their initial choice of path.

Phase Two. During the second phase, or *update* phase, players are allowed to make updates in strategy in an arbitrary order, choosing their *best response* with respect to the current solution. As the potential of the solution decreases with each update, the sequence of updates must be finite and so will terminate at some NE.

We will now show an upper bound on the value of the potential on the solution reached at the end of the first phase. Note that $\Phi(S) \geq sc(S)$, so we also have a bound on the sum-cost of the solution after phase one. In the second phase, each update in strategy will decrease the potential of the solution, and so for the stable

solution S^* reached from S , we have that

$$\Phi(S) \geq \Phi(S^*) \geq sc(S^*) ,$$

and so a bound on the potential of the solution reached during the first phase is a bound on the sum cost of the solution reached after the second phase.

We will now introduce some of the notation we will use for the following proof.

Consider the optimal solution to the game. It cannot contain cycles. Assuming it does, there must be two points x and y between which two separate paths are used, one of which must be an improving move for some player, and so the profile could not be stable, and more importantly, cannot be optimal. We therefore have an optimal solution which will consist of some number of trees. Let this number be m . By OPT_i we denote the cost of the i^{th} tree in the optimal solution. The cost of the optimal solution for this game, which we will denote by OPT_{Δ_u} will be the sum of the costs of each of these trees, i.e.

$$\text{OPT}_{\Delta_u} = \sum_{i=1}^m \text{OPT}_{R_i}$$

Take P as the set of all players. Let P_i be the set of players in P whose optimal paths contributes to OPT_i . We therefore have that $P = P_1 \cup P_2 \cup \dots \cup P_m$.

We now introduce a threshold $\gamma \in (0, 1)$, the value of which we will fix towards the end of this proof, and use it to divide players into two categories, γ -good and γ -bad. We define the first group, γ -good, as players whose revised cost falls by a factor of at least γ , i.e. a player i is γ -good if $c^+(i) \leq (1 - \gamma) \cdot c(i)$, with all other players being in the second category of γ -bad. We will denote the set of γ -good players as Q , the set of γ -bad players as R , so that $P = Q \cup R$. When considering players in either of these two sets, we will use the same subscript notation to identify a subset of these players whose optimal paths are contained within the same tree. In this way $P_i = Q_i \cup R_i$, and $Q = Q_1 \cup Q_2 \cup \dots \cup Q_m$.

The proof for Theorem 16 requires us to first bound the cost to players whose choice of path is γ -good relative to the value of γ , after which we fix a value for γ allowing us to bound the cost to all players whose paths are γ -bad.

Theorem 17. *The cost to players whose choice of path is γ -good is bounded as*

$$\sum_{i \in Q} c(i) \leq \sqrt{n} \cdot \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_{\Delta_u} .$$

Consider the subset P^j of players whose optimal paths are contained in the tree OPT_j . For the purposes of this proof, we will view the paths in the tree OPT_j as a set of intervals. This is achieved by generating an Euler cycle from an arbitrary leaf in the tree OPT_j , which can be viewed as a path X along which each player's start and destination nodes appear at least once. First remove all duplicate start and end nodes, so that only the first occurrence of each, from left to right, remain. Each player's path can now be viewed as an interval between their respective s, t points on the path X .

By Lemma 21 we partition the set of players P_j into δ groups, where $\delta \leq \sqrt{|P_j|} \leq \sqrt{n}$, so that the paths in a group can be viewed as either *nested* or *agreeable* on X . Within the set P_j , we now have δ groups, whose paths are either nested or agreeable along X : $P_j = P_j^1 \cup \dots \cup P_j^\delta$. For a group of players P_j^k , we will denote the set of these players who are γ -good as Q_j^k , those who are γ -bad as R_j^k , so that $P_j^k = Q_j^k \cup R_j^k$, and $Q_j = Q_j^1 \cup \dots \cup Q_j^\delta$.

We will now refer to groups of players as either nested or agreeable, if the intervals of their paths in the optimal solution are either nested or agreeable. For any of these groups, we will refer to the section of X in which their optimal paths appear as X' . Note that the cost of the path X is at most twice the cost of the optimal solution OPT_j , and that the cost of any path X' is at most the cost of the path X .

Lemma 22. *For a group of players P_j^k , whose paths are either agreeable or nested in X , the cost to the γ -good players is bounded as*

$$\sum_{i \in Q_j^k} c(i) \leq \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_j.$$

Proof. Let ω be the number of γ -good players in the group P_j^k , i.e. $\omega = |Q_j^k|$. For this proof only, begin by relabelling players in the set Q_j^k as $1, 2, 3, \dots, \omega$ in the order they join the game. Let σ be the permutation of players in the order that their left endpoints appear along X' from left to right, so that $\sigma(1)$ is the first player, $\sigma(\omega)$ the last.

We will now construct a rooted tree T with vertex set $Q_j^k \cup \{r\}$ and root r . We define r as a player whose source and destination is the start node of the player $\sigma(1)$. We will also use $\sigma(0)$ to refer to the root. As we view r as a player $\sigma(0)$, it is

necessary to define the updated cost of this player. As $c(\sigma(0)) = 0$ we can say that $c^+(\sigma(0)) = 0 \leq (1 - \gamma) \cdot c(\sigma(0))$.

The tree T , the specific method of construction for which we will come to shortly, will have the property that for any player i , its parent will be a player of lower index, i.e. a player who joins earlier than i .

We will use the notation $d(i, j)$ to denote the sum of the distances between players i and j 's sources and destinations in the underlying graph G , so that $d(i, j) = d(s_i, s_j) + d(t_i, t_j)$. For each edge in this tree we assign a cost so that the cost of the edge connecting players i and j is $d(i, j)$.

As players will always choose the cheapest available path, we have, for a player j joining after a player i , that

$$c(j) \leq d(t_i, t_j) + d(s_i, s_j) + c^+(i) \leq d(i, j) + (1 - \gamma) \cdot c(i) ,$$

Consider a path $\langle i_k, i_{k-1}, i_{k-2}, \dots, i_2, i_1, i_0 = r \rangle$ in the tree T , from player i_k to the root. It follows from the above inequality that for $1 \leq j \leq k$, $c(i_j) \leq d(i_j, i_{j-1}) + (1 - \gamma) \cdot c(i_{j-1})$. Therefore, we have that for player i_k ,

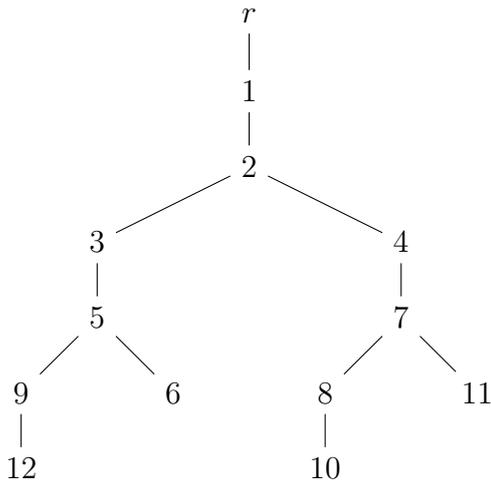
$$c(i_k) \leq \sum_{j=0}^{k-1} (1 - \gamma)^j \cdot d(i_{k-j}, i_{k-j-1}). \quad (5.1)$$

We will now bound the sum of all players' cost shares in the group Q_j^k , i.e. $\sum_{i \in Q_j^k} c(i)$ from above by substituting the upper bound in Inequality (5.1) for each $c(i)$ in the path from i to r (in the tree T).

To do so we introduce the notation $n_e(l)$ to be the number of nodes exactly l levels below the edge $e \in T$, where the level is determined by the number of hops from e so that the lower end point of e is 0 levels below e , its child is 1 level below e and so on. This gives us

$$\sum_{i \in Q_j^k} c(i) \leq \sum_{e \in T} \left(d(e) \cdot \sum_{l \geq 0} (1 - \gamma)^l \cdot n_e(l) \right). \quad (5.2)$$

We will now turn our attention to the construction of the tree, and show a construction with properties which give a propitious upper bound on the right hand side of the Inequality (5.2). For some positive integer τ , let T be the tree with the following properties:



$$n = 12, \quad \tau = 2$$

$$\sigma = \langle 3, 9, 12, 1, 5, 6, 10, 2, 8, 4, 7, 11 \rangle$$

$$S = \{1, 2\}$$

$$\sigma' = \sigma/S = \langle 3, 9, 12, 5, 6, 10, 8, 4, 7, 11 \rangle$$

$$\sigma_1 = \langle 3, 9, 12, 5, 6 \rangle \quad \sigma_2 = \langle 10, 8, 4, 7, 11 \rangle$$

(b) Contents of σ , σ' , S , σ_1 , and σ_2 for first iteration of the algorithm.

(a) Completed tree T with input σ

Figure 5.6: Example of the construction of a tree T for $n = 12$, $\tau = 2$.

1. The root is r and has one child.
2. Each player is a child of a player that preceded it in the arrival sequence (or of the root).
3. All non leaves have one or two children.
4. A non-leaf of T has two children if and only if it has depth divisible by τ and is not the root.
5. The depth of every leaf t in T , $a(t)$, satisfies $2^{a(t)/\tau} \leq n$, so that $a(t) \leq \tau \log n$.
6. The sum of $d(e)$, taken over all edges e in T , is at most $(4(\tau + 1) \log n) \cdot \text{OPT}_j$.

To construct the tree T we will first construct a tree T' , which takes the list σ and constructs a rooted tree using those players as vertices. The algorithm for the construction of T' is recursive, and has the following steps:

1. Take the τ smallest players in σ , by order of arrival, as the set S .
2. Let σ' be the sequence σ less all elements of S . Split σ' into two halves, σ_1 and σ_2 , so that $|\sigma_1| \leq |\sigma_2| \leq |\sigma_1| + 1$, and recursively construct trees T_1 and T_2 on these sequences.

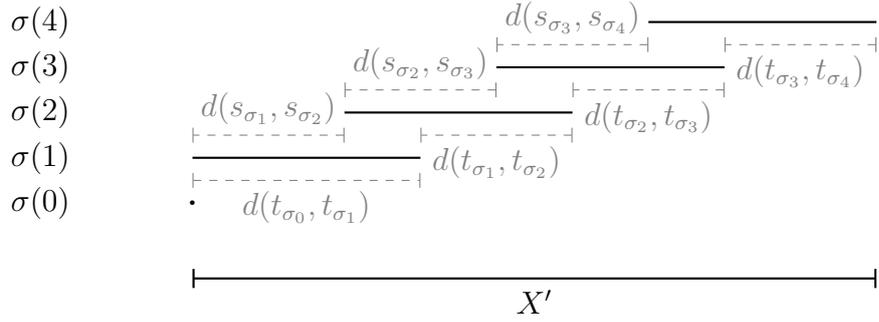


Figure 5.7: Illustration of the sum of distances between end points of a set of players whose paths are agreeable on X' . Note that the player $\sigma(0)$ is a player with a path of cost 0 connected to the start node of the player $\sigma(1)$ by an edge of cost 0, and so appears as a point at the left end of X' .

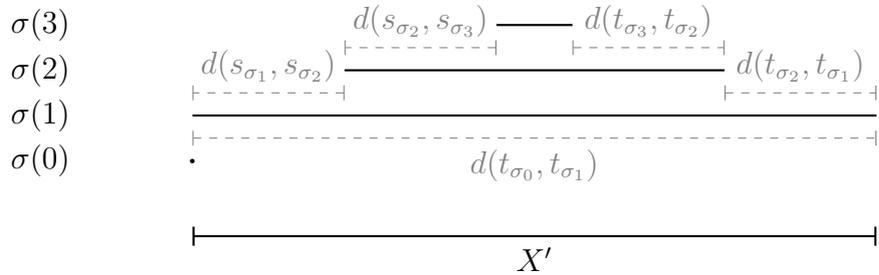


Figure 5.8: Illustration of the sum of distances between end points of a set of players whose paths are nested on X' . Note that the player $\sigma(0)$ is a player with a path of cost 0 connected to the start node of the player $\sigma(1)$ by an edge of cost 0, and so appears as a point at the left end of X' .

3. Build the path π on S by listing elements from smallest (x) to largest (y) by arrival time. The root of T' is x , with the sequence π from there to y , and an edge connecting y to the root of each of T_1 and T_2 .

Having constructed T' , we now join the root of this tree to the node r to form the tree T . Properties 1 – 4 being obvious from construction, we will now show the last two properties of our construction.

Claim 1. *The depth of every leaf t in T , $a(t)$, satisfies $2^{a(t)/\tau} \leq n$, so that $a(t) \leq \tau \log n$.*

Proof. As we branch at most $\log n$ times, from the root to a leaf node, and the number of nodes between each branching point is τ , it follows that no leaf can be at a depth greater than $\tau \log n$. \square

We will now show property 6 by induction.

Claim 2. $\sum_{e \in T} d(e) \leq 4(\tau + 1) \log n \cdot \text{OPT}_{R_j}$.

Proof. First observe that, in the tree T , if e is an edge joining $\sigma(j_1)$ and $\sigma(j_2)$ where $j_1 \leq j_2$, then

$$d(e) \leq \sum_{i=j_1}^{j_2-1} d(\sigma(i), \sigma(i+1)).$$

We will now bound the number of edges e to which $d(\sigma(i), \sigma(i+1))$ can contribute, for all i . We will say that, if $d(\sigma(i), \sigma(i+1))$ is a contributing factor in the cost of e , then e *hits* $\{i, i+1\}$.

First, note that an edge e in T that joins $\sigma(j_1)$ and $\sigma(j_2)$ will only hit $\{i, i+1\}$ if $j_1 \leq i < i+1 \leq j_2$. Let x be a node which is the child of the root, or the child of a node which has two children. Let y be the first descendant of x with two children, the left of which we will call y_1 , the right y_2 .

The edges connecting x to y may all hit $\{i, i+1\}$, as may the two edges connecting y to each of y_1, y_2 , which is $\tau + 1$ in total.

Now consider which other edges hit $\{i, i+1\}$. Noticing that, in the sequence σ , all nodes in the sub-tree rooted at y_1 appear to the left of all nodes in the sub-tree rooted at y_2 , we can say that any other edges $e \in T$ which hit $\{i, i+1\}$ must all have both endpoints in one of these two sub-trees, i.e. one of the two sub-trees rooted at y_1 and y_2 contains no edges which hit $\{i, i+1\}$. As we branch at most $\log n$ times from the root to a leaf, it follows by induction that the number of edges which hit $\{i, i+1\}$ is at most $(\tau + 1) \log n$.

Recall that σ is the sequence of players along X' , which has length at most twice that of OPT_j . In the case where players' optimal paths are nested along X' , all start nodes lie to the left of all end nodes, meaning the sum of the distances between all intervals is at most the length of X' . This, added to the distance between the start and end nodes of $\sigma(1)$ and the root node (which is in fact the distance between $s_{\sigma(1)}, t_{\sigma(1)}$ and therefore at most the length of X') gives us a total of $4 \cdot \text{OPT}_j$ for the sum of $d(\sigma(i), \sigma(i+1))$ for all i in a nested group of players. Figure 5.8 gives

a depiction of a set of intervals which are nested on X' , illustrating the sum of distances between the end points of each interval as they appear along X' . In the case where players' optimal paths are agreeable on X' , the sequence of all start nodes is increasing along X' , and so the sum of distances between all start nodes is at most the length of X' , with the same being true for all end nodes. Thus, the sum of the distances between $\sigma(i), \sigma(i+1)$ for all i , is also at most $4 \cdot \text{OPT}_j$ for groups of agreeable players. Figure 5.7 gives a depiction of a set of intervals which are agreeable on X' , illustrating the sum of distances between the end points of each interval as they appear along X' .

It therefore follows that

$$\sum_{e \in T} d(e) \leq 4 \cdot (\tau + 1) \log n \cdot \text{OPT}_j ,$$

as required. \square

By the properties of T , for any node v , the number $n_e(k)$ of nodes exactly k levels below it is at most $2^{(1+k/\tau)}$. As $\tau \geq 1$, we have that $n_e(k) \leq 2 \cdot (1 + 1/\tau)^k$, which combined with Inequality (5.2) implies that

$$\sum_{i \in Q_j^k} c(i) \leq 2 \sum_{e \in T} \left(d(e) \cdot \sum_{k \geq 0} ((1 - \gamma) (1 + 1/\tau))^k \right) \leq 2 \sum_{e \in T} \left(d(e) \cdot \sum_{k \geq 0} (1 - \gamma + 1/\tau)^k \right)$$

Fixing the value of τ to be $\lceil 2/\gamma \rceil$, the inner sum of the above is at most $2/\gamma$.

This gives us

$$\sum_{i \in Q_j^k} c(i) \leq 2 \sum_{e \in T} d(e) \cdot 2/\gamma ,$$

which, combined with the final property of our tree means that

$$\begin{aligned} \sum_{i \in Q_j^k} c(i) &\leq 8 \cdot \left(\lceil \frac{2}{\gamma} \rceil + 1 \right) \cdot \frac{2}{\gamma} \cdot \log n \cdot \text{OPT}_j \\ &\leq \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_j . \end{aligned}$$

\square

Proof of Theorem 17. Recall that we began by partitioning a set of players P_j into $\delta \leq \sqrt{|P_j|} \leq \sqrt{n}$ groups, where n is the total number of players in the game.

Having shown by Lemma 22 that the cost of the γ -good paths for a group of players P_j^k is bounded as

$$\sum_{i \in Q_j^k} c(i) \leq \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_j ,$$

we now consider the cost of the solution to γ -good players in the set P_j . As $Q_j = Q_j^1 \cup \dots \cup Q_j^\delta$,

$$\sum_{i \in Q_j} c(i) = \sum_{k=1}^{\delta} \sum_{i \in Q_j^k} c(i) \leq \sqrt{n} \cdot \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_j.$$

Recalling that the optimal solution for the game consisted of m disjoint trees, and that $\text{OPT}_{\Delta_u} = \sum_{j=1}^m \text{OPT}_j$, we can conclude our proof by noting that

$$\sum_{i \in Q} c(i) \leq \sum_{j=1}^m \sqrt{n} \cdot \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_j \leq \sqrt{n} \cdot \frac{O(\log n)}{\gamma^2} \cdot \text{OPT}_{\Delta_u}.$$

□

We now turn our attention to bounding the cost of those players whose choice of path is γ -bad. To do so we will analyse the set of all γ -bad players simultaneously, and choose a value for γ so that the cost to these paths is not too high.

Theorem 18. For $\gamma = \frac{1}{4H(n)}$,

$$\sum_{i \in R} c(i) \leq 2H(n) \cdot \sum_{i \in Q} c(i).$$

Proof. Consider the game Δ_u , recalling that those players for whom the revised cost of their path, $c^+(i)$, falls below the threshold of $(1 - \gamma) \cdot c(i)$ are called γ -good, with all other players being called γ -bad, and that the sets Q and R represent those players who are γ -good and γ -bad respectively.

We define F_g as the set of edges in G first used by a player in Q , and F_b as the set of edges in G first used by a player in R .

The cost to the set of γ -good players is the sum of each of their individual costs, i.e. $c(Q) = \sum_{i \in Q} c(i)$, with the revised cost to the same set of players being the sum of each of their revised costs. In this way we also define $c(R) = \sum_{i \in R} c(i)$ and $c^+(R) = \sum_{i \in R} c^+(i)$. The cost of all edges first used by a player who is γ -good, $c(F_g)$, is the sum of the cost of each of these edges, with analogous calculation for the cost of all bad edges, $c(F_b)$.

As the cost of an edge in F_b drops by at least a factor of 2 for subsequent players, we have that

$$c^+(R) \leq c(R) - \frac{c(F_b)}{2},$$

which, added to the fact that, by definition $c^+(R) \geq (1 - \gamma) \cdot c(R)$, gives us

$$c(F_b) \leq 2\gamma \cdot c(R) .$$

Since the cost to all bad players must be less than the potential of the overall solution, i.e. $c(R) \leq H(n) \cdot (c(F_b) + c(F_g))$, we have that

$$c(R) \leq H(n) \cdot (c(F_g) + (2\gamma \cdot c(R))) .$$

As $c(Q) \geq c(F_g)$, we have that $\frac{c(R)}{H(n)} - 2\gamma \cdot c(R) \leq c(Q)$, which with some rearranging is

$$c(Q) \geq \frac{c(R)}{H(n)} \cdot (1 - 2\gamma H(n)) .$$

If we now fix the value of γ to $1/4H(n)$, we have that

$$c(R) \leq 2H(n) \cdot c(Q) .$$

□

Proof of Theorem 16. We conclude this proof by combining the results in Theorems 17 and 18.

$$\begin{aligned} \sum_{i \in P} c(i) &= \sum_{i \in Q} c(i) + \sum_{i \in R} c(i) \\ &\leq (2H(n) + 1) \cdot \sum_{i \in Q} c(i) \\ &\leq \sqrt{n} \cdot \frac{O(\log n)(2H(n)+1)}{\gamma^2} \cdot \text{OPT}_{\Delta_u} \\ &\leq \sqrt{n} \cdot O(\log n)(2H(n) + 1)(4H(n))^2 \cdot \text{OPT}_{\Delta_u} \\ &\leq O(\sqrt{n} \log^4 n) \cdot \text{OPT}_{\Delta_u} \end{aligned}$$

□

5.5 Concluding Remarks

In this chapter we have a number of interesting results related to the measures of the efficiency of the solution which sequential games produce. Our method of analysis for asymmetric games finds an interesting application for a result related to the

partitioning of intervals. One might ask how else we can apply such reasoning to our field. Having shown that the reachable price of anarchy for asymmetric games is $\Omega(\sqrt{n})$ and $O(\sqrt{n} \log^4 n)$, another possible direction for future research would be closing this gap.

For games with two players, we show that both the reachable price of anarchy and the reachable price of stability is exactly $8/5$. The fact that $\text{rPoA} = \text{rPoS}$ is in itself an interesting fact. It is known that, for two players, the price of anarchy is 2 and the price of stability is $4/3$. As our bound for the sequential version of these measures sits between these two, an analysis of games for three players may provide some insight into the price of stability measure, for which no tight bound exists.

Chapter 6

Conclusion and Future Research

We now summarise the results presented in Chapters 3, 4, and 5 of this thesis, and provide a brief discussion of some possible directions for future research. For a more detailed discussion of potential further research questions see the final sections of those chapters.

6.1 Summary of the Contribution

Best Response Dynamics

In Chapter 3 we examine the effects of best response dynamics. We begin by examining the maximum number of steps required before a stable solution of arrived at, and find that in both the general and uncapacitated case, convergence is unbounded in the number of players. This is a departure from previous results on BRD convergence, which had previously been bounded by some function on the number of players. We then turn our attention to the effect of BRD on the quality of solutions. We show that the worst case increase in the maximum cost of a profile is $\Theta(n \log n)$ in general games, and $\Theta(n)$ in uncapacitated games. When considering the worst case increase in cost to an individual player of the game, we find this to be unbounded in general games, and $\Theta(n)$ in uncapacitated games.

Price of Stability w.r.t Maximum Cost

In Chapter 4 we examine the Price of Stability measure, and show that in the general case, it is possible for the best NE to be up to a factor of $\Theta(n \log n)$ times the cost

of the optimal profile, if players have unique start and end points. We also show that this bound does not apply to the general case where players have a common destination (i.e. rooted), by showing that in the case where BRD from the optimal profile lead to an expensive NE (more than n time the cost of OPT), it is always possible to construct a better profile where no player pays more than n times the cost of their path in OPT. The method of proof is of independent interest as it departs from the traditional method of bounding the Price of Stability by examining the effects of BRD from the optimal profile.

We then move on to the uncapacitated case, and show a tight bound of $\Theta(n)$ for the maximum cost Price of Stability, by providing an improved lower and upper bound.

Sequential Network Design Games

Our analysis of the sequential variant of network design games, in Chapter 5, can be split into two parts.

We first consider the existence of stable solutions which can be reached from an initial profile which is empty. We find that in general games, the existence of such stable states is not guaranteed, while in the uncapacitated case there must always be a NE which can be reached by each players greedily choosing their initial path from an empty start profile.

We then examine the quality of the solutions which sequential games permit. In the general case, we show a symmetric game where the best and only reachable state is an arbitrary factor more expensive than the optimal profile, and thus that both the reachable price of anarchy and reachable price of stability measures for these games is unbounded in the number of players. For uncapacitated games, we show that it is possible for the best and only reachable Nash Equilibrium to be a factor $\Omega(\sqrt{n})$ more expensive than the optimal profile, thus showing that the reachable price of stability is at least this. We show an upper bound on this measure, by showing that given a game, there must be a reachable Nash Equilibria which is at most $O(\sqrt{n} \log^4 n)$ times more expensive than the optimal profile w.r.t. sum-cost.

For games with two players, we show that both the reachable price of anarchy and the reachable price of stability are exactly $8/5$.

6.2 Directions for Future Research

We now briefly discuss some of the possible directions for future research. We have split these into three categories, by the chapters to which they relate. For a more detailed discussion of possible future work we direct the reader to the final section of the respective results chapters.

Best Response Dynamics

Our contributions to the effects of BRD are fairly complete for our setting, in that we show tight bounds for the effects it can have in several settings. Of course, the open problem of BRD on sum cost still remains, however, our results do not suggest any avenues for further results on this.

One interesting observation is that the number of updates required for BRD to converge to NE is unbounded in the number of players. One obvious direction could be to see what other settings this applies to. The implications of unbounded convergence is that, while in the increase in cost to some individual is bounded per update, i.e. by the potential of the individuals path, an unlimited number of updates allows changes in cost which go beyond simply realising the potential of a path.

Price of Stability w.r.t Maximum Cost

One open question remaining from our results on the Price of Stability measure for maximum cost would be to improve the lower bound to match exactly the upper bound of $O(n \log n)$. It is with some annoyance that that current upper bound is currently half that of the upper bound, and the author feels that with some modifications to the existing structure this gap can be removed.

A more productive avenue of research might be to examine the proof method used for general rooted games. As a departure from the traditional approach, it may in some way prove useful in similar settings.

Sequential Network Design Games

As an alternative approach to bounding the price of stability, it is the sequential variant of network design games which would seem to have the greatest scope for

further research. One could examine the sequential price of stability with respect to the maximum cost to a player, as well as apply the concept of sequentially joining the game to other settings.

Bibliography

- [1] A. Agarwal and M. Charikar. On the price of stability for undirected multicast. 2006.
- [2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice Hall, New Jersey, 1993.
- [3] S. Albers. On the value of coordination in network design. *SIAM Journal on Computing*, 38(6):2273–2302, 2009.
- [4] N. Andelman, M. Feldman, and Y. Mansour. Strong price of anarchy. In *Proceedings of the eighteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 189–198. Society for Industrial and Applied Mathematics, 2007.
- [5] A. Angelucci, V. Bilò, M. Flammini, and L. Moscardelli. On the sequential price of anarchy of isolation games. *Journal of Combinatorial Optimization*, 29(1):165–181, 2015.
- [6] E. Anshelevich, A. Dasgupta, J. M. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *SIAM J. Comput.*, 38(4):1602–1623, 2008.
- [7] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proceedings of the thirty-fifth annual ACM symposium on Theory of computing*, pages 511–520. ACM, 2003.
- [8] K. J. Arrow. *Social choice and individual values*, volume 12. Yale university press, 2012.
- [9] K. J. Arrow and G. Debreu. Existence of an equilibrium for a competitive economy. *Econometrica: Journal of the Econometric Society*, pages 265–290, 1954.

- [10] A. Asadpour and A. Saberi. An approximation algorithm for max-min fair allocation of indivisible goods. *SIAM Journal on Computing*, 39(7):2970–2989, 2010.
- [11] M. Charikar, H. Karloff, C. Mathieu, J. S. Naor, and M. Saks. Best response dynamics in network creation with egalitarian cost sharing. Unpublished paper, copy available on request to the authors, 2006.
- [12] M. Charikar, H. Karloff, C. Mathieu, J. S. Naor, and M. Saks. Online multicast with egalitarian cost sharing. In *Proceedings of the Twentieth Annual Symposium on Parallelism in Algorithms and Architectures*, SPAA '08, pages 70–76, New York, NY, USA, 2008. ACM.
- [13] H.-L. Chen and T. Roughgarden. Network design with weighted players. *Theor. Comp. Sys.*, 45(2):302–324, June 2009.
- [14] G. Christodoulou, C. Chung, K. Ligett, E. Pyrga, and R. Van Stee. On the price of stability for undirected network design. In *International Workshop on Approximation and Online Algorithms*, pages 86–97. Springer, 2009.
- [15] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the Thirty-seventh Annual ACM Symposium on Theory of Computing*, STOC '05, pages 67–73, New York, NY, USA, 2005. ACM.
- [16] J. R. Correa, A. S. Schulz, and N. E. Stier-Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.
- [17] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *ACM Trans. Algorithms*, 3(1):4:1–4:17, Feb. 2007.
- [18] C. Daskalakis, P. W. Goldberg, and C. H. Papadimitriou. The complexity of computing a nash equilibrium. *SIAM Journal on Computing*, 39(1):195–259, 2009.
- [19] C. Daskalakis and C. H. Papadimitriou. Three-player games are hard. In *Electronic colloquium on computational complexity*, volume 139, pages 81–87, 2005.

- [20] J. de Jong and M. Uetz. The sequential price of anarchy for atomic congestion games. In *International Conference on Web and Internet Economics*, pages 429–434. Springer, 2014.
- [21] Y. Disser, A. E. Feldmann, M. Klimm, and M. Mihalák. Improving the H_k -bound on the price of stability in undirected Shapley network design games. *Theor. Comput. Sci.*, 562:557–564, 2015.
- [22] A. Epstein, M. Feldman, and Y. Mansour. Strong equilibrium in cost sharing connection games. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 84–92. ACM, 2007.
- [23] A. Epstein, M. Feldman, and Y. Mansour. Efficient graph topologies in network routing games. *Games and Economic Behavior*, 66(1):115–125, 2009.
- [24] T. Erlebach and M. Li. Partitioning a set of intervals into nested or agreeable families. Unpublished manuscript, advanced copy available on request to the authors., 2016.
- [25] T. Erlebach and M. Radoja. Further results on capacitated network design games. In *International Symposium on Algorithmic Game Theory*, pages 57–68. Springer, 2015.
- [26] E. Even-Dar, A. Kesselman, and Y. Mansour. Convergence time to nash equilibria. In *International Colloquium on Automata, Languages, and Programming*, pages 502–513. Springer, 2003.
- [27] A. Fanelli, D. Leniowski, G. Monaco, and P. Sankowski. The ring design game with fair cost allocation. *Theoretical Computer Science*, 562:90–100, 2015.
- [28] M. Feldman and O. Geri. Do capacity constraints constrain coalitions? *CoRR*, abs/1411.5712, 2014.
- [29] M. Feldman and T. Ron. Capacitated network design games. In *5th International Symposium on Algorithmic Game Theory (SAGT 2012)*, LNCS 7615, pages 132–143. Springer, 2012.

- [30] M. Feldman and T. Tamir. Convergence rate of best response dynamics in scheduling games with conflicting congestion effects. Technical report, Working paper, 2011.
- [31] A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo. On the price of stability for designing undirected networks with fair cost allocations. In *International Colloquium on Automata, Languages, and Programming*, pages 608–618. Springer, 2006.
- [32] D. Fotakis. Congestion games with linearly independent paths: Convergence time and price of anarchy. In *International Symposium on Algorithmic Game Theory*, pages 33–45. Springer, 2008.
- [33] I. Gilboa and E. Zemel. Nash and correlated equilibria: Some complexity considerations. *Games and Economic Behavior*, 1(1):80–93, 1989.
- [34] P. W. Goldberg and C. H. Papadimitriou. Reducibility among equilibrium problems. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 61–70. ACM, 2006.
- [35] M. D. Hirsch, C. H. Papadimitriou, and S. A. Vavasis. Exponential lower bounds for finding brouwer fix points. *Journal of Complexity*, 5(4):379–416, 1989.
- [36] R. Holzman and N. Law-Yone. Strong equilibrium in congestion games. *Games and economic behavior*, 21(1):85–101, 1997.
- [37] R. Holzman and N. Law-Yone. Network structure and strong equilibrium in route selection games. *Mathematical Social Sciences*, 46(2):193–205, 2003.
- [38] K. Jain and V. V. Vazirani. Equitable cost allocations via primal-dual-type algorithms. In *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 313–321. ACM, 2002.
- [39] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS 99)*, LNCS 1563, pages 404–413. Springer, 1999.

- [40] R. P. Leme, V. Syrgkanis, and É. Tardos. The curse of simultaneity. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pages 60–67. ACM, 2012.
- [41] J. Li. An $o(\log n / \log \log n)$ upper bound on the price of stability for undirected shapley network design games. *Inf. Process. Lett.*, 109(15):876–878, July 2009.
- [42] A. Mamageishvili and M. Mihalák. Multicast network design game on a ring. In *Combinatorial Optimization and Applications*, pages 439–451. Springer, 2015.
- [43] E. Maskin. The theory of implementation in nash equilibrium: A survey. *Social Goals and Social Organization*, pages 173–204, 1985.
- [44] I. Milchtaich. Congestion games with player-specific payoff functions. *Games and economic behavior*, 13(1):111–124, 1996.
- [45] I. Milchtaich. Network topology and the efficiency of equilibrium. In *ICM Millennium Lectures on Games*, pages 233–266. Springer, 2003.
- [46] I. Milchtaich. Topological conditions for uniqueness of equilibrium in networks. *Mathematics of Operations Research*, 30(1):225–244, 2005.
- [47] I. Milchtaich. The equilibrium existence problem in finite network congestion games. In *International Workshop on Internet and Network Economics*, pages 87–98. Springer, 2006.
- [48] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 183–194. Springer, 2004.
- [49] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14:124–143, 1996.
- [50] O. Morgenstern and J. Von Neumann. Theory of games and economic behavior. 1944.
- [51] J. Nash. Non-cooperative games. *Annals of mathematics*, pages 286–295, 1951.

- [52] N. Nisan and A. Ronen. Algorithmic mechanism design. In *Proceedings of the thirty-first annual ACM symposium on Theory of computing*, pages 129–140. ACM, 1999.
- [53] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. Cambridge University Press, New York, NY, USA, 2007.
- [54] C. Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753. ACM, 2001.
- [55] C. H. Papadimitriou and T. Roughgarden. Computing equilibria in multi-player games. In *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 82–91. Society for Industrial and Applied Mathematics, 2005.
- [56] C. H. Papadimitriou and T. Roughgarden. Computing correlated equilibria in multi-player games. *Journal of the ACM (JACM)*, 55(3):14, 2008.
- [57] R. W. Rosenthal. A class of games possessing pure-strategy nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.
- [58] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [59] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [60] A. S. Schulz and N. E. S. Moses. On the performance of user equilibria in traffic networks. In *Proceedings of the Fourteenth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA '03)*, pages 86–87, 2003.
- [61] L. Shapley. A value for n-person games1. *Contributions to the Theory of Games (AM-28)*, 2:307, 2016.
- [62] M. Voorneveld, P. Borm, F. Van Megen, S. Tijs, and G. Facchini. Congestion games and potentials reconsidered. *International Game Theory Review*, 1(03n04):283–299, 1999.