

Selfish Network Creation

On Variants of Network Creation Games

ANDREAS CORD-LANDWEHR

Reviewers:

- Prof. Dr. Friedhelm Meyer auf der Heide,
Paderborn University
- Jun.-Prof. Dr. Alexander Skopalik,
Paderborn University

Contents

Preface	v
1 Introduction	1
1.1 A Model for Selfish Network Creation	3
1.2 Thesis Focus & Overview	4
2 Preliminaries	9
2.1 The Classic Model of Network Creation Games	10
2.2 Notions of Stability, Quality, and Convergence	12
2.2.1 Notions of Stability	12
2.2.2 Quality of Equilibria	14
2.2.3 Convergence of Improving-Response Processes	15
2.3 Known Results	17
2.3.1 Model Variants	18
2.3.2 Relationships of Model Variants	22
2.4 Alternative Models	22
3 Loss and Benefit of Friendships	27
3.1 The Friendship Model & Preliminaries	29
3.2 Related Work & Contribution	31
3.3 Worst-Case Friendships in Swap-Games	33
3.3.1 Private Costs in Max-Swap-Game Tree Equilibria	37
3.3.2 The Price of Anarchy in Max-Swap-Games	45
3.3.3 The Price of Anarchy in Sum-Swap-Games	49
3.4 Process Equilibria	51
3.5 Conclusion & Future Work	56

4	The Impact of Choosing Edge Qualities	57
4.1	Model & Notations	58
4.2	Related Work & Contribution	62
4.3	Existence of Equilibria	63
4.3.1	Equilibria in the Sum-Pricing-Game	63
4.3.2	Equilibria in the Max-Pricing-Game	68
4.4	Quality of Equilibria in the Sum-Pricing-Game	71
4.4.1	Employing Characteristic Price Functions	77
4.5	Quality of Equilibria in the Max-Pricing-Game	78
4.6	Conclusion & Future Work	82
5	Limits of Locality	83
5.1	Model & Notations	85
5.2	Related Work & Contribution	87
5.3	Preliminaries	89
5.4	Approximation Quality of Greedy Probing	92
5.4.1	Approximation of the k-Local Sum-Game	94
5.4.2	Approximation Lower Bound in the Sum-Game	95
5.4.3	Approximation Upper Bounds in the Sum-Game	102
5.5	Efficiency of Probing Locality	107
5.5.1	A Clash of Models	107
5.5.2	The Price of Anarchy	111
5.6	Conclusion & Future Work	116
6	Multilevel Network Games	119
6.1	Model & Preliminaries	121
6.2	Related Work & Contribution	123
6.3	Bidirectional Gateways	124
6.3.1	The Sum-Layer-Game	125
6.3.2	The Max-Layer-Game	136
6.4	Unidirectional Gateways	141
6.4.1	Existence of Equilibria	142
6.4.2	Quality of Equilibria	144
6.5	Conclusion & Future Work	149
	Bibliography	151
	Index	161

Preface

I must admit: Writing this thesis turned out to be more work than I originally anticipated. Writing the following lines, however, turns out to be much more fun than I had expected some months ago. Now, at this point, there are many people I want to thank and for so many things: For exciting discussions, for chats in the coffee room, for critical questions, for encouraging words, for making the time of writing this thesis (and the years before) far from boring, for bearing my little free time during the past months, and for numerous other things.

Foremost, I want to thank my supervisor Friedhelm Meyer auf der Heide. It was a great time working in this research group and I am particularly grateful that I had the opportunity to find my research topic on my own. However, filling the topic with interesting content would not have been possible without all the discussions that helped me identify the important questions. While doing this, I had the chance to travel the world to discuss, present, and bring back new inspirational ideas and for this I owe my gratitude especially to the Collaborative Research Center 901 “On-The-Fly Computing”. The postcards on my office wall are dedicated as a small thank-you for that.

Special thanks go to all of my many co-authors: Alex⁴, Barbara, Bastian, Christoph², Daniel³, Frederik, Friedhelm, Kamil, Manuel, Marcus, Markus, Martina, Matthias, Pascal, Peter², Sebastian, and Sven. Specifically, I want to thank Martina for pushing my nose into the field of network creation games, back in the days, which then directly gripped my interest and shaped all of my future research. For all the long research discussions in front of some white

Preface

board or via phone, I want to specifically thank Sebastian and Pascal. I learned from Peter the art of introduction writing and how to tell even the most boring topics with a fancy motivation – and everyone reading this thesis should be grateful to him.

Working in this research group was a great time and it was made a great time by the people I met there. In particular, this holds for my former and current office-roommates Tim, Peter, Sören and Alex. Last but not least, I have to thank all my friends and my family for enduring the last months in which thesis writing was quite in the focus of my activities.

All of you, specifically those I missed to name here, thank you! And you know, there is always a box of free cookies on my desk, reserved just for you.

Andreas Cord-Landwehr
Paderborn, November 2015

CHAPTER 1

Introduction

“ The Internet is the first thing that humanity has built that humanity doesn't understand – the largest experiment in anarchy that we have ever had. ”

Eric Schmidt, former CEO Google

How are networks formed when their participants are acting selfishly? What is the cost for society by allowing selfish behavior instead of enforcing a central control? And how can we predict the impact of rules that narrow the freedom of decision? – These are the key issues of this thesis, and they are studied for large and dynamic networks.

Networks in the context of this thesis are understood as *overlay networks*. On top of any physical communication layer, they feature connections in the form of logical links among their participants. Specifically, those logical links can be created and adapted almost arbitrarily, while no such modification affects the underlying physical connections. The importance of overlay networks comes from their applications: They are an important technique especially used to create peer-to-peer networks. Popular examples of them are search overlays, like Gnutella or Chord (cf. survey by Androutsellis-Theotokis and Spinellis [AS04]), which provide efficient search mechanisms for large networks by establishing a logical overlay. Still, different examples can be found in other

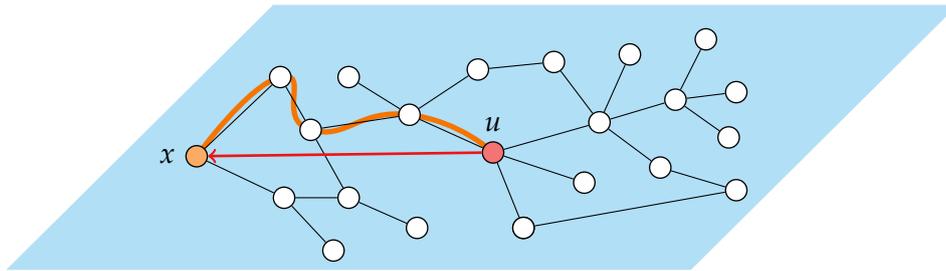


Figure 1.1: Example of a participant in a Gnutella overlay network [AS04]: Participant u creates a logical link to x , who just answered her search request successfully. The expectation of u is that future requests possibly will be served well by x .

fields like overlays providing self-stabilization properties (e.g., [KKS14]).

In such networks, we can be faced with a very dynamic behavior of their participants (e.g., continuous joins/leaves and changing communication patterns). Additionally, networks may have a vast size, literally numbering in billions of participants. These properties combined make it infeasible to install any centralized control that manages the whole network; already obtaining an overview of the complete network would raise immense cost and require a considerable amount of time.

The neat idea usually employed to overcome those problems is to dismiss any role of a centralized organizer in favor of distributed operations: The participants perform local operations in the network to reduce the overall cost. For example, in a search network, the participants locally compute and decide the search paths along which incoming requests will be sent (cf. example of operations in Gnutella search overlay in Figure 1.1). Conceptually, by selecting those paths the participants implicitly form logical links in the network and effectively establish an entire logical overlay network: Paths imply overlay edges and every such edge induces a certain cost for maintaining it.

We encounter a similar behavior when looking at the Internet, where agents in the form of service providers autonomously optimize their connections (cf. Fabrikant et al. [Fab+03]). Specifically visible in their attitude, actions are usually performed in a selfish manner, only with respect to someone's own cost. This brings us to a question valid for overlay networks: Why should participants follow any algorithm and not simply do what is best for them? – And actually, each agent faces her private cost in the form of requests and

link maintenance and by this has a valid incentive to consider her own benefit first, before looking at the overall network. The consequence is that we must consider them as *selfishly acting agents*.

However, by accepting this fact of selfish agents, the classic tools from distributed computing fail to capture the network's behavior. In the following, we use the concept of *network creation games* to tackle this selfish behavior. The model as well as many methods that will be used throughout this thesis originate from the field of *algorithmic game theory*. This field of research is generally interested in issues relating to the result of strategic actions of autonomous agents. Thereby, the main concerns range from the computational complexity of outcome forecasts, over the behavior of dynamic adjustment processes driven by the agents, to the efficiency loss by selfish behavior.

1.1 A Model for Selfish Network Creation

We consider the behavior of a group of network agents who can act arbitrarily, as long as their actions are rational and strategic. This means, every agent is aware of her private cost in the network and only performs actions (e.g., creating or changing connections) that actually improve her cost.

Over the last two decades, various approaches have been introduced to provide a model that captures such strategic behavior of agents. The two most influential models were provided by Jackson and Wolinsky [JW96] and by Fabrikant et al. [Fab+03]. In the first model, the creation of a link is understood as the bilateral decision of both acting agents, whereas both agents share the price of a created link. Thereby, the authors follow the intuition from social and economic networks. Differently, the latter model by Fabrikant et al. is inspired by technical networks and interprets a link creation as a unilateral decision, whereas only the one acting agent has to pay the full price. In any such model, the agents strive for minimizing their trade-offs between the cost for communication to other agents and the cost raised by creating and maintaining links. The communication costs thereby are considered as the sum or the maximum of distances (called the *Sum-Game* or the *Max-Game*, respectively). With its focus on technical networks, in theoretical computer science as well as in this thesis, the model of choice to study the outcome quality of selfish network creation is the *network creation game* model by Fabrikant et

al. [Fab+03].

Our main objective is to study the outcome of the individual selfish actions of the agents. For this, we consider (pure) Nash equilibria as states of the game, where no agent can improve her private cost by a unilateral action. Looking at the set of all Nash equilibrium networks, we can estimate the worst-case loss of selfish behavior by comparing the worst possible overall cost of stable network states to the overall cost of an optimal solution. This gives the so-called *price of anarchy* and is an established worst-case measure in algorithmic game theory (cf. [Nis+07, Chapter 17]).

Evidently, this network creation game model is only an abstraction of reality and cannot capture every effect one can see in practice. Yet, as of today, it is the best model we know to evaluate the cost of selfish behavior in overlay networks. Actually, Chun et al. [Chu+04] show that networks that are created by this model (by adjusting the edge price and applying a maximum degree) lead to overlay networks with very desirable properties with respect to degree distributions, stretch, and resilience. The so-generated networks turn out to behave similarly to the Internet.

1.2 Thesis Focus & Overview

Understanding the impact and outcome of selfish behavior of network participants is crucial when designing network protocols. We want to know if a network protocol is robust against selfish behavior in the sense that selfish actions lead to globally good solutions. Moreover, a proper model for the network behavior allows to make predictions for the outcome of the selfish actions and to take countermeasures like restricting or extending the set of available actions the agents can perform. In addition to all this, we must ensure not to lose focus of aspects and issues governing our practical experience.

Approaching this, my thesis suggests and studies several specializations of the network creation game model by Fabrikant et al. [Fab+03]. Part one (Chapter 3) covers the *impact of non-uniform communication interests*, which is that not every agent is interested in communication with every other agent, i.e., only wants to communicate with her so-called *friends*. In Part two (Chapter 4), I introduce *quality-of-service decisions* into the framework of network creation games, which allows agents to choose not only where to create an edge to but

also of which quality and for which price. Part three (Chapter 5) considers the important topic of the *influence of having only local network information* available to the agents, which specifically limits actions of an agent. Finally, in Part four (Chapter 6), I introduce a new model for studying the *interaction between different network layers*.

The remainder of this section provides a short overview of the main questions, the model extensions, and sketches the respective results of each part.

Loss and Benefit of Friendships. The first part of my thesis considers the impact of non-uniform communication interests on the quality of equilibrium networks. For the game variants derived from the model by Fabrikant et al. [Fab+03] as well as those derived from the model by Jackson and Wolinsky [JW96], the usual assumption is that agents strive to uniformly improve their communication costs to all other agents. However, in real networks we have the observation that agents only want to communicate with a (small) subset of the others, the so-called friends. This raises the question about the impact of non-uniform communication interests on the quality of equilibria. Halevi and Mansour [HM07] were the only ones preceding our work who considered non-uniform communication interests in network creation games, although only for the Sum-Game model by Fabrikant et al. [Fab+03]. For this game variant, they provided an upper bound of $O(\sqrt{n})$ for the price of anarchy for games with n agents.

In this part, first I focus on the overall quality loss by friendships, which is the worst-case impact those non-uniform communication interests can have in Swap-Games (a game variant in which only edges can be changed but no new ones created, [Alo+13]). By discarding the otherwise strong dependency of the edge price to the game's outcome, the Swap-Game variant is particularly well suited to study structural properties of worst-case network instances. In case of non-uniform communication interests, I show that the price of anarchy is worst possible in all considered game variants except for tree equilibria in the Max-Game. For this latter class of tree equilibria, the price of anarchy results in a notably smaller tight worst-case bound of $\Theta(\sqrt{n})$. This is a surprising result since, unlike in most other game variants, here the average distance versions behave worse than the maximum distance versions.

In contrast to this worst-case approach, I further study in the original game

variant by Fabrikant et al. how to exploit that networks typically evolve along the communication interests of their participants and by this derive improved equilibrium quality guarantees. The price of anarchy bounds provided for the newly introduced concept of *process equilibria* are considerably good when the communication interests form distinct groups of agents that do not want to communicate with each other.

The model, analysis, and results presented in this chapter are based on the following publication:

2012 (with M. Hüllmann, P. Kling and A. Setzer). “Basic Network Creation Games with Communication Interests”. In: *Algorithmic Game Theory – 5th International Symposium, SAGT 2012, Barcelona, Spain, October 22–23, 2012. Proceedings*, cf. [Cor+12].

The Impact of Choosing Edge Qualities. The second part of the thesis looks at the impact of providing connections with different qualities and, accordingly, also for different prices. Instead of choosing only where to create an edge to, agents are now able to also specify which quality this edge should have. When considering today’s networks, where connections are offered by several service providers with different bandwidths and latency guarantees, it is natural to assume that agents have several edge quality offers available to choose from. Despite of its significance, this question was not studied before in the context of network creation games.

Given a set of available edge lengths, the price assignment is modeled by a so-called *price function*, which assigns a price to each edge. Such a function must merely be positive and further fulfill that shorter edges are not more expensive than longer ones. Despite this very general model, for both the Sum-Game and the Max-Game and arbitrary combinations of edge lengths and price functions, equilibrium networks always exist. For the Sum-Game, the price of anarchy upper bound utilizes a game characterization of the optimal trade-off between edge price and edge length and is even tight for specific price functions. In particular, for a class of linear price functions the price of anarchy becomes constant. The Max-Game, however, shows a very different behavior and the price of anarchy bounds go in line with the results for the original game with only one available edge length.

The model, analysis, and results presented in this chapter are based on the following publication:

2014 (with A. Mäcker and F. Meyer auf der Heide). “Quality of Service in Network Creation Games”. In: *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings*, cf. [CMM14].

Limits of Locality. Global precise knowledge is a very unrealistic assumption if we consider large and dynamic networks. However, most present models for network creation games assume such knowledge when studying the evolution and outcomes of networks by selfish agents. The only exceptions are provided by Bilò et al. [Bil+14a; Bil+14b]. In their models, for a fixed parameter k the agents only know the network part induced by their k -neighborhoods. Under this locality assumption, one can show that the price of anarchy dramatically increases in terms of lower bounds. This originates from the fact that agents can only partially estimate the outcome of their actions, limited by their local knowledge, and are thus unaware of many possibly profitable actions.

The third part of my thesis studies the limits of locality in the sense: What is possible in terms of strategic decision making and global cost efficiency when facing locality? In contrast to Bilò et al., agents are now equipped with the power to *probe* (i.e., estimate the cost change for a certain action) a certain number of different strategies while still restricted to their k -local knowledge. The main result is that it suffices to probe only a polynomial number of different strategies (namely all one-edge strategy changes) to ensure price of anarchy upper bounds close to those in the global knowledge game and by this to guarantee a good global cost outcome. In particular, this even holds for constant size neighborhoods.

The model, analysis, and results presented in this chapter are based on the following publication:

2015 (with P. Lenzner). “Network Creation Games: Think Global – Act Local”. In: *Mathematical Foundations of Computer Science 2015 – 40th International Symposium, MFCS 2015, Milan, Italy, August 24–28, 2015. Proceedings, Part II*, cf. [CL15].

Multilevel Network Games. The last part of my thesis sheds light on the interaction between different network layers. In particular, how selfish agents of a general purpose network utilize a high-speed layer to improve their communication costs. Such high-speed layers ensure that for every communication path in the general purpose layer there is a shorter path in the high-speed layer. Every agent has the minimal strategy set of either connecting or not connecting to the high-speed layer for a fixed price. Connected agents then act as *gateways* and allow the access to the other layer. Depending on how the high-speed layer is implemented, two different access models for using the high-speed layer are considered: Assuming the high-speed layer is a separated network, it is reasonable that switching between both networks is only possible at specific gateway locations. Otherwise, if the high-speed layer is considered as a logical network layer, then only the access to the high-speed layer is restricted to gateway locations. My model is the first one that analyzes strategic behavior in multilevel networks. The results specifically cover the price of anarchy and show that in general the games are not potential games.

The model, analysis, and results presented in this chapter are based on the following publication:

2014 (with S. Abshoff, D. Jung and A. Skopalik). “Multilevel Network Games”. In: *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings*, cf. [Abs+14].

CHAPTER 2

Preliminaries

IN this chapter, we present an overview of the different concepts and notions used in the research on network creation games. Thereby, our focus lies on variants derived from the model by Fabrikant et al. [Fab+03], which we will call the *classic network creation game models*.¹ We start by formally introducing the Sum-Game and the Max-Game as the major model lines, which either model agents who strive for optimizing their average distances or those who strive for optimizing their maximal distances. After that we discuss different solution concepts, measures for network quality, and convergence characteristics for the agents' operations. Considering these concepts, after that we provide an overview of the results for variants of the classic network creation game.

Most notions introduced and used in this section origin from the field of algorithmic game theory. In particular, solution concepts like for example “Nash equilibria” have a long history in algorithmic game theory, and even predating

¹See Section 2.3 for a broader overview and discussion of models derived from the Sum-Game model by Fabrikant et al. [Fab+03]. As discussed in Section 2.4, there is actually a vast amount of literature in economics and mathematics concerning alternative models that predate the model by Fabrikant et al. [Fab+03]. Yet, we see this model as the most relevant one when considering overlay networks and hence refer to it as the classic model in computer science.

this, in game theory. Instead of discussing them in their full generality, we only introduce and use them in the sense they are required to study the outcome of selfish behavior in network creation games. For a general introduction into algorithmic game theory and discussion of those concepts, we refer to Nisan et al. [Nis+07].

The notions and definitions presented in this chapter form the basis for introducing and discussing the model variants in later chapters. Specifically, the game definition from Section 2.1 and the equilibrium concepts from Section 2.2 are essential for later discussions.

2.1 The Classic Model of Network Creation Games

A network creation game consists of a set of agents $V = \{v_0, \dots, v_{n-1}\}$ (also called *peers* or *players*). These agents are interpreted as network nodes that can create edges to other agents. Thereby, each agent can individually decide about the edges she wants to buy in order to minimize her *private cost*, which is the cost of the bought edges plus the cost for communicating with other agents. In the game introduced by Fabrikant et al., agents strive to minimize the sum of distances to all other agents and are able to perform arbitrary changes of their edges to achieve this goal. In particular, they can exchange any current set of own incident edges with another set of incident edges. Throughout this thesis, we call this game the *Sum-Game*.²

Stating the Sum-Game in the terms of a strategic game, every agent $u \in V$ has a strategy space $S_u := \mathcal{P}(V \setminus \{u\})$, consisting of all possible sets of incident edges as given by the possible edge endpoints. Her current strategy $s_u \in S_u$ specifies the currently selected edges, i.e., the edges *owned* by u . Then, the combination of all agents' strategies $S := (s_{v_0}, \dots, s_{v_{n-1}}) \in S_0 \times \dots \times S_{n-1}$ denotes the *strategy profile*, which we interpret as a graph $G[S] = (V, E)$: Agents are the graph nodes and each strategy $s_u = \{v_1, \dots, v_m\}$ implies the graph edges $\{u, v_1\}, \dots, \{u, v_m\}$. Note that all edges are undirected edges and, moreover, that even though the definition admits multi-edges, throughout this thesis the selfish nature of the agents ensures that no multi-edges will ever be created – creating an already existing edge cannot be a cost-improving operation. Given

²In the literature, the Sum-Game is also called Buy-Game [Len12; KL13], SUMGAME (for example, [MS13]), or sometimes simply network creation game, [Fab+03].

2.1 The Classic Model of Network Creation Games

a strategy profile S and the induced network $G[S]$, we denote the length of the shortest path between two agents u and v as $d_{G[S]}(u, v)$. The length of the longest shortest path, $\text{diam}(G[S]) := \max_{u, v \in V} d_{G[S]}(u, v)$, gives the diameter of the network.

Agents strive for minimizing their private costs, given by a *private cost function* $c_u(S)$. Thereby, the cost of an agent is given only by the current strategy profile. Each edge in an agent's strategy raises a fixed cost value of $\alpha > 0$. There are two variants of private cost functions that yield two different versions of the game:

Sum-Game: If the private cost is given by the sum of distances to all other agents plus α for every bought edge, we name the game the *Sum-Game* (introduced by Fabrikant et al. [Fab+03]). Given a strategy profile S and an agent $u \in V$ with strategy s_u , formally the private cost of u is:

$$c_u(S) = \alpha \cdot |s_u| + \sum_{v \in V} d_{G[S]}(u, v) \quad (2.1)$$

For this cost function, we refer to the first term as *edge* $u(S)$, called *the edge cost of u* , and to the second as $\text{dist}_u(S)$, called *the distance cost of u* .

Max-Game: If the private cost is given by the maximum distance to any other agent plus α for every bought edge, we name the game the *Max-Game* (introduced by Demaine et al. [Dem+07]). Given a strategy profile S and an agent $u \in V$ with strategy s_u , formally the private cost of u is:

$$c_u(S) = \alpha \cdot |s_u| + \max_{v \in V} d_{G[S]}(u, v) \quad (2.2)$$

For this cost function, we refer to the first term as *edge* $u(S)$, called *the edge cost of u* , and to the second as $\text{dist}_u(S)$, called *the distance cost of u* .

Whereas the private cost is a measure of the local cost of an agent, the sum over all agents' private cost values,

$$\text{cost}(S) := \sum_{u \in V} c_u(S), \quad (2.3)$$

estimates the overall quality of a network and we refer to it as the *social cost*.

2.2 Notions of Stability, Quality, and Convergence

We are interested in the outcome of the selfish actions of the agents. Specifically, how do “stable” states look like in those games? – To answer this, there are different approaches employed in the literature. The most commonly used solution concept for network creation games is that of a (pure) Nash equilibrium, which focuses on network states where no agent can improve her private cost by unilateral strategy changes (cf. Fabrikant et al. [Fab+03]). But there are also different concepts of stability, like the stability of bilateral strategy changes, called *pairwise stability* (cf. Jackson and Wolinsky [JW96]), which is widespread in economics literature.

2.2.1 Notions of Stability

The concept of a *Nash equilibrium* was introduced by Nash [Nas51] in his seminal work and since then “has emerged as the central solution concept in game theory” (Nisan et al. [Nis+07, p. 12]). We call a state in a game with a strategy profile S a *Nash equilibrium* (NE) if no agent can improve her private cost by unilaterally changing her current strategy.³ Formally, for every agent v_i and every strategy change $s'_{v_i} \in S_{v_i}$ with the accordingly changed strategy profile $S' := (s_{v_0}, \dots, s_{v_{i-1}}, s'_{v_i}, s_{v_{i+1}}, \dots, s_{v_{n-1}})$, it holds that $c_{v_i}(S) \leq c_{v_i}(S')$.

Depending on the allowed strategy changes of the agents, we distinguish the following three different Nash equilibria. For each of these equilibrium notions, we further introduce a so-called *greedy equilibrium* variant that denotes the stability of single-edge changes.

Buy Equilibrium (BE): A strategy profile S forms a *buy equilibrium* if the agents are allowed to arbitrarily buy, remove, and swap own incident edges and S forms a Nash equilibrium. If no agent can buy one incident edge, remove one own incident edge, or swap one incident edge, S forms a *greedy buy equilibrium*.

Asymmetric Swap Equilibrium: A strategy profile S forms an *asymmetric swap equilibrium* if the agents are only allowed to arbitrarily swap own

³In related literature, this solution concept is often called a “pure” Nash equilibrium to distinguish it from the so-called “mixed” Nash equilibria, where strategies are chosen only with certain probabilities. Since we will focus our analysis only on pure equilibria, we omit the extra term.

incident edges and S forms a Nash equilibrium. If no agent can swap one own incident edge, S forms a *greedy asymmetric swap equilibrium*. Note that for this equilibrium notion, the edge price α will be omitted, since the number of edges does not change.

Swap Equilibrium (SE): A strategy profile S forms a *swap equilibrium* if the agents are allowed to arbitrarily swap any incident edges and S forms a Nash equilibrium. If no agent can swap one arbitrary incident edge, S forms a *greedy swap equilibrium*. Note that for this equilibrium notion, there are no edge ownerships and thus the edge price α can be omitted.

For both, the game variants with sum cost function and with maximum cost function, all these different equilibria exist. Depending on the edge price α , either a clique (Sum-Game for $\alpha \leq 1$, Max-Game for $\alpha \leq 1/(n-1)$) or a star (Sum-Game for $\alpha > 1$, Max-Game for $\alpha > 1/(n-1)$) constitutes a buy equilibrium (cf. Fabrikant et al. [Fab+03] and Demaine et al. [Dem+07]). For the swap equilibrium and the asymmetric swap equilibrium, always a star network forms an equilibrium (cf. Alon et al. [Alo+13] and Mihalák and Schlegel [MS12]). Note that the named equilibrium networks are also stable for the corresponding greedy equilibrium variants.

For the above equilibrium concepts, we consider agents who always want to perform a strategy change when they improve their costs. Yet, this can lead to situations where the gain of an agent is negligibly small but the effort in terms of to be changed edges is very high. Facing this, the notion of an *ε -approximate equilibrium* (e.g., Chien and Sinclair [CS07] and Skopalik and Vöcking [SV08]) captures agents that perform strategy changes only if they reduce their cost by a reasonable fraction. We say for $\varepsilon > 1$ that a strategy profile is an ε -approximate equilibrium if no agent can decrease her private cost by a factor of at least ε by unilaterally changing her strategy, i.e., to be at most $1/\varepsilon$ times the former cost value.⁴ Note that this notion of approximate equilibria applies to all above-mentioned equilibrium variants.

⁴In the literature, there is a similar solution concept with the same name that considers the additive improvement by ε instead of the multiplicative improvement (for example, Daskalakis et al. [DMP07]). Another popular choice for the approximation parameter is α .

2.2.2 Quality of Equilibria

The typical way of evaluating the quality of a network is by estimating its social cost. Our main interest here is the quality of equilibrium networks; in other words, what is the quality of solutions in a network creation game? Specifically, we ask:

- (a) How bad are equilibria in the worst-case?
- (b) How good are equilibria in the best-case?

Since for most variants of network creation games the equilibria are not unique, these two questions constitute the maximal and minimal loss by the selfish acting of the agents, which can possibly be far apart. The maximum loss by selfish behavior was formalized by Koutsoupias and Papadimitriou [KP99] as the *price of anarchy (PoA)* and is defined as the ratio of the highest social cost of any equilibrium network and the optimal social cost. The minimal loss by selfish behavior was first studied by Schulz and Moses [SM03] and nowadays is known as the *price of stability (PoS)*.⁵ Its value is given by the ratio of the smallest social cost of any equilibrium network and a minimum social cost network (not necessarily forming an equilibrium).

Definition 2.1 (Price of Anarchy and Price of Stability). Consider a game with social cost function,

$$\text{cost} : \mathcal{S} \rightarrow \mathbb{R}_{>0},$$

whereas \mathcal{S} is the set of all possible strategy profiles. Let $\mathcal{S}_{NE} \subseteq \mathcal{S}$ be the set of all equilibrium strategy profiles and S_{Opt} be the strategy profile with minimal social cost. Then we define:

- (a) price of anarchy: $\max_{S \in \mathcal{S}_{NE}} \frac{\text{cost}(S)}{\text{cost}(S_{\text{Opt}})}$
- (b) price of stability: $\min_{S \in \mathcal{S}_{NE}} \frac{\text{cost}(S)}{\text{cost}(S_{\text{Opt}})}$

Note that both maximum and minimum are defined over any number of agents.

Given the equilibrium networks from the previous section, it is easy to see that the price of stability is bounded to be at most two and it is even close to one for

⁵The price of stability is sometimes also named the *optimistic price of anarchy*, see [Ans+03]. It was first mentioned under the name “price of stability” by Anshelevich et al. [Ans+04].

several equilibrium concepts and parameters. On the other hand, bounding the price of anarchy is a challenging task that was considered in a remarkable series of papers (see Section 2.3).

If we have a strategy profile S that is an ε -approximate buy equilibrium with a corresponding network $G[S]$, we can derive an upper bound for the price of anarchy by generalizing an argument by Albers et al. [Alb+14, proof of Lemma 3.4].

Theorem 2.2. *For the Sum-Game with $\alpha \geq 2$, let S be a strategy profile that is an ε -approximate buy equilibrium and let S_{Opt} be a strategy profile with minimal social cost. Then the ratio of them is at most:*

$$\frac{\text{cost}(S)}{\text{cost}(S_{\text{Opt}})} \leq \varepsilon(3 + \text{diam}(G[S]))$$

Proof. Let u be an arbitrary fixed agent and consider T to be a shortest path tree rooted at u . (Note that paths to all other agents exist, since S is an ε -approximate buy equilibrium.) For every agent $v \in V$, we consider the strategy change of removing all own edges that do not belong to T and creating one new edge to u . Thereby, let $T_v \subseteq T$ be the set of tree edges owned by v . Since S is an ε -approximate buy equilibrium and no agent v changes $\text{dist}_u(S)$ by this operation, we get $c_v(S) \leq \varepsilon(\alpha \cdot |T_v| + \alpha + (n - 1) + \text{dist}_u(S))$. Hence, for the social cost we get:

$$\begin{aligned} \text{cost}(S) &= \sum_{v \in V} c_v(S) \leq \sum_{v \in V} \varepsilon(\alpha \cdot |T_v| + \alpha + (n - 1) + \delta_u) \\ &\leq \varepsilon(\alpha \cdot |T| + (n - 1)\alpha + (n - 1)^2 + n\delta_u) \\ &\leq \varepsilon(2(n - 1)\alpha + (n - 1)^2 + n(n - 1) \cdot \text{diam}(G[S])) \end{aligned}$$

Since the optimal solution is a star and has social cost of $\alpha(n - 1) + n(n - 1)$, we get as the upper bound for the social cost ratio $\varepsilon(2 + 1 + \text{diam}(G[S]))$. \square

2.2.3 Convergence of Improving-Response Processes

For both games, the Sum-Game and the Max-Game, we know that equilibria exist. In particular, this also holds for all variants of buy and swap equilibrium concepts as introduced above. Yet, if we consider some non-equilibrium

strategy profile as a starting point, it is a valid question whether agents can ever *reach* such an equilibrium state from there. Specifically we ask: Can we find for every initial strategy profile a sequence of improving strategy changes that transforms it into an equilibrium strategy profile? And, if yes, how long is such a sequence?

These sequences of iterative applications of cost-improving operations of the agents are called *improving-response processes*. Here, an *improving response* (IR) denotes any cost-decreasing strategy change of an agent. An improving response is called a *best response* (BR) if this strategy change is optimal regarding the maximum private cost decrease for this agent. We say an improving-response process (or best-response process) *converges to an equilibrium* if the final strategy profile of the process is an equilibrium. If, for a game with a finite number of strategies, there is an infinite long improving-response process, then the process must contain a cycle. We call such a cycle an *improving-response cycle* (or *best response cycle*, respectively).

A game is called a *weakly acyclic game* (WAG) (introduced by Young [You93]) if, starting from any initial strategy profile, there exists *some* finite sequence of improving responses that eventually converges to an equilibrium state. This concept resembles the natural class of games that possibly reach equilibrium states via simple and globally asynchronous strategic actions, independently of their starting states. For this, even very simple dynamics, like randomized improving- or best-response dynamics or regret-based dynamics, suffice (cf. [You93; Mar+09]). Examples for such weakly acyclic games are given by Engelberg and Schapira [ES14] and Milchtaich [Mil96]. In [Mil96], Milchtaich considered a variant of congestion games but with individual payoff functions for every player. Engelberg and Schapira [ES14] introduced a class of routing games that models aspects of Internet-routing algorithms. For both of these weakly acyclic games we have that one can find improving-response cycles and hence, not every sequence of improving strategy changes leads to an equilibrium.

A class of games that was subject to substantially more research interest is the class of *potential games* (cf. Monderer and Shapley [MS96]). This is the subclass of all weakly acyclic games for which it holds that every sequence of improving-response operations terminates in an equilibrium: i.e., every such sequence is finite. In particular, this is known as the *finite improvement*

property (FIP). Monderer and Shapley [MS96] showed that a game has the finite improvement property if and only if there exists a generalized ordinal potential function,

$$\Phi : S_0 \times \dots \times S_{n-1} \rightarrow \mathbb{R}_{\geq 0},$$

that maps strategy profiles to real numbers such that if an agent performs an improving response, then the potential value decreases. One of the most prominent examples for games belonging to this class are *congestion games* (introduced by Rosenthal [Ros73]). Monderer and Shapley [MS96] showed that the class of congestion games is actually isomorphic to potential games. Note that though every sequence of improving responses is finite, these sequences still may be exponentially long (cf. Fabrikant et al. [FPT04]).

The convergence properties in the context of network creation games were studied by Kawald and Lenzner [KL13]. For the Max-Game and the Sum-Game, i.e., with agents who can arbitrarily buy, delete, and swap edges, they showed that improving-response cycles may exist and hence these games cannot be potential games. They further showed that these negative results still hold if agents are only allowed to perform greedy operations as well as if the agents are only allowed to swap edges. The only positive exception, where such a game is known to fulfill the finite improvement property, are swap games where the starting network is a tree (cf. [Len11; KL13]). This means, these game variants are potential games and every sequence of improving response operations converges to an equilibrium state.

For network creation games with bilateral edge operations [CP05], Kawald and Lenzner [KL13] showed that the game is not even weakly acyclic, meaning that there are strategy profiles for which no sequence of best-response operations leads to an equilibrium. The question whether the classic network creation games with unilateral edge operations by Fabrikant et al. are weakly acyclic games or not is still an open question, though.

2.3 Known Results

Starting with the study by Fabrikant et al. [Fab+03], computing the price of anarchy in network creation games attracted a lot of attention. Figure 2.1 summarizes the currently best known price of anarchy results for the Sum-

Game and the Max-Game. In the following, we give an overview of those results. The detailed discussions and comparisons with the models considered in this thesis are postponed to the individual chapters.

2.3.1 Model Variants

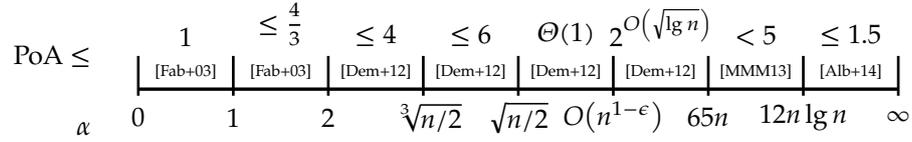
For the Sum-Game, Fabrikant et al. [Fab+03] proved an upper bound of $O(\sqrt{\alpha})$ on the price of anarchy for the case of $\alpha < n^2$, as well as a constant price of anarchy upper bound for tree network equilibria. By using the fact that for $\alpha \geq n^2$ every equilibrium is a tree network, they also obtain a constant price of anarchy for $\alpha \geq n^2$. In their initial paper, they further formulated their famous *tree conjecture*, which states that there is a constant A such that for every $\alpha > A$ all non-transient buy equilibrium networks⁶ are trees [Fab+03, Conjecture 1].

This conjecture was later disproved by Albers et al. [Alb+14] by constructing a family of non-tree buy equilibrium networks that contain cycles of length three and five. In the same paper, Albers et al. further showed that the price of anarchy is constant for $\alpha = O(\sqrt{n})$ and presented the first sublinear price of anarchy upper bound of $O(n^{1/3})$ for general $\alpha > 0$. Demaine et al. [Dem+07] were the first to prove an $o(n^\epsilon)$ bound for α in the range of $\Omega(n)$ and $O(n \log n)$. By Mihalák and Schlegel [MS13] and improved by Mamageishvili et al. [MMM13], it was shown that for $\alpha \geq 65n$ all equilibria are tree networks and thus the price of anarchy is constant. For non-integral constant values of $\alpha > 2$, Graham et al. [Gra+13] showed that the price of anarchy tends to 1 as $n \rightarrow \infty$. Halevi and Mansour [HM07] considered the case that agents are only interested in a subset of the other agents (see detailed discussion in Section 3.2).

Demaine et al. [Dem+07] introduced the Max-Game variant, in which agents desire to minimize the maximum distance to any other agent, rather than the sum of distances. For this model, they showed that the price of anarchy is at most 2 for $\alpha \geq n$, at most $O\left(\min\left\{4\sqrt{\lg n}, (n/\alpha)^{1/3}\right\}\right)$ for α in range of $2\sqrt{\lg n} \leq \alpha \leq n$, and at most $O(n^{2/\alpha})$ for $\alpha < 2\sqrt{\lg n}$. Using a similar technique like for the Sum-Game, Mihalák and Schlegel [MS13] showed that for $\alpha > 129$ all equilibria are tree networks and hence the price of anarchy is constant.

⁶Here, a strategy profile is called a *non-transient equilibrium* if it is a buy equilibrium and also no agent can perform a strategy change that preserves her current private cost.

(a) Sum-Game:



(b) Max-Game:

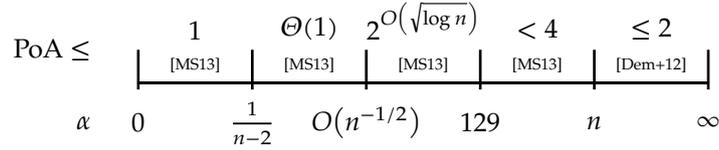


Figure 2.1: Overview of the currently best known upper bounds for the price of anarchy for buy equilibria. In case (a), the parameter ϵ is an arbitrary positive constant.

Swap Equilibria. Considering the price of anarchy results for buy equilibria, we see a strong dependency on the edge price α . To overcome this, Alon et al. [Alo+10] introduced the *basic network creation game* (revisited in [Alo+13]), which avoids the parametrization by an edge price parameter. In their game variant, agents of a given network can arbitrarily swap their incident edges in order to minimize the maximum or sum of distances to the other agents. By the change from edge ownerships to anonymous edges, which do not belong to any specific agent anymore, this game is mostly suited to analyze structural properties of equilibria rather than convergence processes. In the following, we call equilibria of this game *swap equilibria* (cf. Section 2.2), although they are conceptually different to the original Sum-Game and Max-Game, as there are no edge ownerships anymore. Specifically, Mihalák and Schlegel [MS12] showed that the equilibrium concepts of swap equilibrium and buy equilibrium are different in the sense that equilibrium networks from one concept are not always equilibria in the other concept and vice versa (cf. Section 2.3.2).

Restricting the initial networks to trees, Alon et al. [Alo+13] showed that for the Sum-Game the only equilibrium is a star network, while in the Max-Game it can be either a star network or a double-star (i.e., any tree graph with diameter of at most three). Without restrictions to the initial networks, all equilibria in the Sum-Game are proven to have a diameter of at most $2^{O(\sqrt{\lg n})}$, which is also the upper bound for the price of anarchy. For the Max-Game, the authors

only show the existence of an equilibrium of diameter $\Theta(\sqrt{n})$, which especially gives the same lower bound for the price of anarchy. Computationally, the estimation of best-response strategies can still be \mathcal{NP} -hard in general. Yet, the greedy swap game version, which was considered by Alon et al. [Alo+13], makes the best-response computation feasible in polynomial time. In Chapter 3, we will discuss the effects of non-uniform communication interests in swap games and thus omit details here.

Budgeted Buy Equilibria. Ehsani et al. [Ehs+15] introduced a different approach for removing the edge price α from the classic model. Their idea was to build upon a game by Laoutaris et al. [Lao+14], in which every agent has only a limited budget for buying directed edges of different lengths and prices in order to minimize the sum of distances according to her preference list. The model by Ehsani et al. [Ehs+15] applied the budget restriction to both the Sum-Game and the Max-Game. Specifically, they changed the private cost functions to only consist of the distance cost term.

In the changed model, one can observe that the agents' budgets significantly change the outcomes of the games. Given an arbitrary non-negative assignment of budgets to the agents, equilibria always exist for both the Sum-Game and the Max-Game. While Ehsani et al. [Ehs+15] provide an upper bound of $O(2^{\sqrt{\log n}})$ for the price of anarchy in the Sum-Game, achieved by a similar technique as in Alon et al. [Alo+13], in the Max-Game they can show a worst possible price of anarchy of $\Theta(n)$. If all agents have positive budgets, in the Max-Game the price of anarchy result improves to $O(\sqrt{\log n})$.

Asymmetric Swap Equilibria. Mihalák and Schlegel [MS12] proposed a variant of the swap game model in which edges are owned by agents and specifically only own edges can be swapped. The so-called *asymmetric swap game* hereby generalizes both the budget game model and the swap game model, in the sense that any equilibrium in one of these models forms an equilibrium in the asymmetric swap game model.

Their main results are particular structural insights into the equilibria of all of these games. On the one hand, they show that every asymmetric swap equilibrium has at most one (non-trivial) 2-edge-connected component. On the other hand, they show a logarithmic upper bound on the equilibrium network

diameter for the case that the minimum degree of the unique 2-edge-connected component is at least n^ϵ with $\epsilon > \frac{4 \lg 3}{\lg n}$. These structural results were later reused for proving the nowadays best known price of anarchy upper bounds in Mihalák and Schlegel [MS13] and Mamageishvili et al. [MMM13].

Greedy Equilibria. In the models we discussed so far, agents are able to perform radical strategy changes in terms of the number of edges changed at a time. Specifically, for these models we made the assumption that agents are able to compute their best strategy changes – even if this is computationally intractable, as, for example, the computation of a best-response strategy change in the Sum-Game is \mathcal{NP} -hard. Lenzner [Len12] introduced the notion of greedy operations (cf. Section 2.2) by limiting agents to perform only single-edge changes. This notion can be applied to every before-mentioned game.

Restricting agents to greedy operations immediately makes best-response computations feasible in polynomial time. Surprisingly, in the Sum-Game greedy buy equilibria are 3-approximate buy equilibria. Although it is not directly proven in [Len12], by Theorem 2.2 the results for the price of anarchy in the Sum-Game for buy equilibria essentially apply here, too (only impaired by a small multiplicative factor). This is quite different for the Max-Game, where greedy equilibria exist that are not even $\Omega(n)$ -approximate buy equilibria.

Local Knowledge Buy Equilibria. Agents with only local information about the current status of the network were first considered by Bilò et al. [Bil+14a] and then extended by a slightly relaxed locality notion in Bilò et al. [Bil+14b]. In both papers, the local knowledge is modeled in a very pessimistic way: An agent is aware of her neighborhood within a bounded distance only and has to assume a worst-case network structure outside of this view range. In particular, every agent estimates the result of a strategy change by the worst-case private cost over all possible network structures compliant with her current view. We refer to Chapter 5 for a detailed discussion of this model and the results therein. Specifically, there we provide a much more optimistic model for locality that takes into account that agents are able to test the outcome of (some) different strategy changes and finally choose the best of them.

2.3.2 Relationships of Model Variants

For understanding the structural similarities and differences between the various equilibrium concepts, we compare their respective sets of equilibrium networks. Specifically, we are interested in set relationships between them in the following form: Given an equilibrium strategy profile S of a fixed game, for which other games is S also an equilibrium? For example, for every network that is a buy equilibrium for a fixed edge price α , this network is also a greedy buy equilibrium for the same edge price.

Lenzner [Len14, pp. 22–24] discussed this question for various games. In Figure 2.2, we provide an overview of their results for the game variants used in this thesis and refer to [Len14] for the discussion and proofs. Specifically note, as discussed by Mihalák and Schlegel [MS12], that buy equilibria and swap equilibria are very different solution concepts in the regard of equilibrium sets. In particular, there are constellations where both equilibrium sets intersect but none is a subset of the other one. However, asymmetric swap equilibria are a generalizing class of both concepts.

The equilibria in the edge pricing game from Chapter 4, where agents can also choose the length and price for their edges, are evidently a subset of the buy equilibria, assuming that edges of length 1 for price α are available. Supplementary to this diagram, in Chapter 5 we will further discuss how the different versions of local knowledge equilibria fit into this set structure.

2.4 Alternative Models

The research on network formation by selfish agents has a long history in economics research and was considered only later in mathematics and computer science. With a focus on economics literature, Dutta and Jackson [DJ03] and Jackson [Jac08] provide a very good overview of the different approaches and results so far. In the following, we name only some of the most relevant models that help us to provide a better classification of the variants that follow the model by Fabrikant et al. [Fab+03]. Note that especially in economics literature the incentives of agents are modeled by *utilities*. In this light, those models consider agents to utilize their connections to others and by this receive a gain, which is then reduced by the effort spent on creating connections. The more

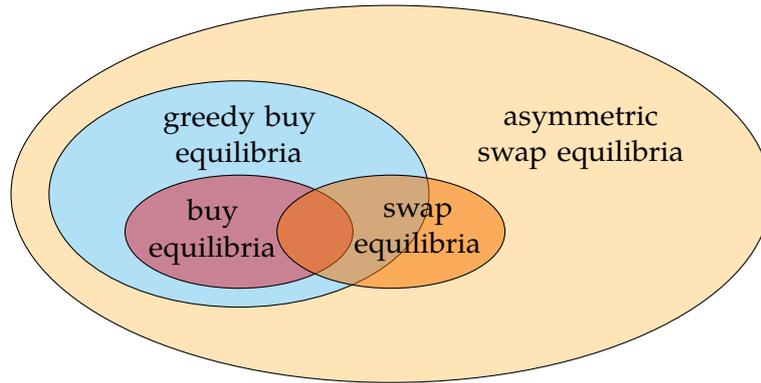


Figure 2.2: Structural relationships between Nash equilibria in different model variants. The areas indicate how equilibria of one game are included in the set of equilibria of another game variant. Specifically, an equilibrium strategy profile for a game in this diagram is also an equilibrium for all game variants indicated as supersets. Cf. Lenzner [Len14, pp. 22–24] for a detailed discussion of the inclusions.

technical view on networks, like it is promoted by Fabrikant et al. [Fab+03], rather focuses on the trade-offs between the cost for utilizing edges and the creation cost as the goal of an agent.

While the early models, like in Myerson [Mye77], were motivated by cooperative game theory and hence only took into account which agents are connected in a subnetwork, Jackson and Wolinsky [JW96] and Bala and Goyal [BG00] considered network distances to estimate the utility of an agent. Hence, different networks connecting the same agents can lead to different utilities of the agents. Jackson and Wolinsky [JW96] introduced the notion of *pairwise stability* for a game where the utility of an agent depends on the distances in the network. A network in this game is called pairwise stable if both agents of every established edge want to keep that connection and no pair of agents wants to create a further edge. Note that here the edge price has to be paid by both incident agents of an edge. Their approach uses a multiplicative utility distance measurement such that for a fixed parameter $\delta \in (0, 1]$, a path of length d gives an utility of δ^d . The utility of an agent then is the sum of utilities to all other agents minus the cost for the established connections. In [JW96], the authors also explicitly compare stable networks to optimal networks. Later, Baumann and Stiller [BS08] continued this research and further analyzed convergence properties and the price of anarchy for different ranges of the edge

price depending on δ . Watts [Wat01] proposed a dynamic process to analyze the outcomes of the selfish decisions of the agents. This process uniformly at random proposes possible edges to the agents, who then can decide whether they want to create the edge or not.

Corbo and Parkes [CP05] considered bilateral edge creation games for agents with the Sum-Game cost function from Fabrikant et al. [Fab+03]. The authors show that the price of anarchy is worse than for the unilateral game by Fabrikant et al. Along their analysis, they showed the equivalence of pairwise stability and a two-player coalition refinement of buy equilibria. Interestingly, Kawald and Lenzner [KL13] showed that best-response dynamics in this game are not even weakly-acyclic in the sum distance variant, and admit best-response cycles in the maximum distance variant.

In contrast to the pairwise stability notions mentioned so far, Bala and Goyal [BG00] considered games with unilateral edge creations. Edges in their game variants are either unidirectional or bidirectional, but there is always only one agent who decides to create and pay for an edge. The utility of agents is given by exponential payoffs like in the game by Jackson and Wolinsky [JW96]. Although the cost function of the agents is different, the game is very close to the game by Fabrikant et al. [Fab+03].

Moscibroda et al. [MSW06] considered the selfish behavior of agents as peers in peer-to-peer networks, which are modeled as metric spaces. Like in the Sum-Game, the agents strive for minimizing their trade-off between the edge cost and the sum of distances to all other agents. In this setting with an underlying metric space, the authors can show a price of anarchy of $O(\min\{\alpha, n\})$. They further provide negative convergence results and moreover, they show that buy equilibria do not always exist; even deciding if such an equilibrium exists is \mathcal{NP} -complete.

A series of research focuses on the formation of social networks. Agents in these networks especially seek for being well connected with agents who have a high influence or centrality in the network. For example, Nikolettseas et al. [Nik+15] introduced a swap-based model where the agents' revenue is based on the sum of degrees of their direct neighbors. With this, the authors aim to provide a model for large distributed systems that are similar to power law or preferential attachment graphs. In this game, there exists an exact potential and hence improving-response processes always converge. Hereby,

the convergence time is polynomially bounded. This still holds even when restricting the agents by a local view such that they can only probe the degrees of a fixed number of other agents; improving-response dynamics still converge in expected polynomial time.

A different approach is provided by Brautbar and Kearns [BK11]. They proposed a model driven by the observation that friendships in social networks are often transitive and thus define the utility of an agent essentially by the number of triangle she is part of. Specifically, using the clustering coefficient of an agent, which is the probability of two uniformly at random selected neighbors being connected, the utility of an agent is the clustering coefficient minus the edge cost (hence, only an edge price of $\alpha \in (0, 1)$ is reasonable). Considering the agents having a high clustering coefficient, we can see which agents are important in the network in terms of being well connected via cliques.

Note that for the remainder of this thesis, we will only consider variants of the classic game by Fabrikant et al. [Fab+03].

Loss and Benefit of Friendships

IN this chapter, we analyze the impact of non-uniform communication interests on the quality of equilibrium networks: Given a large and dynamic network, the agents are usually not interested in communicating with all other agents but only with a subset of them. Our focus lies on the different aspects of influences by such non-uniform communication regarding the negative and the positive effects on the quality of equilibria.

Throughout this chapter, two agents are called *friends* when they want to communicate with each other. In our model, friendships are mutual and an agent is only interested in her direct friends and not necessarily her friends' friends. This means, we do not assume any gain by having many first or second order friends, like it may be in social networks. Rather, we understand the friendships as some given allocation, which simply specifies which agents want to communicate with each other. Our analytical tool for modeling these friendships is a so-called *friendship graph*. Given two nodes in this graph, the respective agents are friends of each other if and only if there is a friendship graph edge between them.

First, we consider the worst-case impact on equilibrium networks by friendship allocations in the Swap-Game [Alo+10]. By discarding the strong dependency of the edge price, which is present in most other models, this model is

particularly well suited to study structural equilibrium properties. On the one hand, we seek for combinations of a friendship graph and a corresponding equilibrium network that maximizes the worst-case social cost ratio when compared to an optimal solution. On the other hand, we aim for upper bounds on the price of anarchy when facing arbitrary friendship allocations. Thereby, we will show a worst-case behavior for almost all considered variants. The only exceptions are tree equilibria for games with agents who strive for minimizing their maximum distances to their friends. In this case, we provide an interesting structural property of equilibrium networks which leads to a surprising bound for the price of anarchy of $\Theta(\sqrt{n})$.

Facing these negative results, we change our focus to the analysis of beneficial effects of friendships. We exploit the properties given by friendship allocations in the Sum-Game and the Max-Game (cf. Section 3.1, [Fab+03]) that ensure best-response processes to lead to equilibria with not too high social costs. Specifically, we introduce a new concept that we name *process equilibrium* and show that equilibria in this natural class, for which connected components of the friendship graph correspond to connected components in the equilibrium, lead to a drastically improved price of anarchy results.

For all such game variants, note that if the friendship graph is a clique, our games with friendship allocations coincide with their original versions in which every agent is interested in every other agent.

Chapter Basis. The model, analysis, and results presented in the remainder of this chapter are based on the following publication:

2012 (with M. Hüllmann, P. Kling and A. Setzer). “Basic Network Creation Games with Communication Interests”. In: *Algorithmic Game Theory – 5th International Symposium, SAGT 2012, Barcelona, Spain, October 22–23, 2012. Proceedings*, cf. [Cor+12].

Chapter Outline. In Section 3.1, we introduce friendship graphs to model a non-uniform communication behavior of agents in network creation games. An overview of our results and a comparison with related work is provided in Section 3.2. The main part of this chapter is given in Section 3.3, which is the analysis of the worst-case behavior of friendship allocations in Swap-Games.

Section 3.4 contrasts these negative results with a more optimistic view on friendship allocations and shows how non-uniform communication interests can have a positive effect on the overall quality of networks. Section 3.5 recaps the results and presents an outlook for future research.

3.1 The Friendship Model & Preliminaries

As usual for network creation games, we consider a set of n selfish agents $V = \{v_1, v_2, \dots, v_n\}$ who unilaterally perform strategy changes in order to improve their private costs. The models considered in this chapter consist of two main ingredients:

- (a) The *friendship model*, which states with respect to whom agents want to reduce their communication costs, and
- (b) the *game model*, which states how agents can act.

Note that we use the notions and notations from Section 2.1 and Section 2.2.3 and name only differences explicitly here.

Friendship Model. Every agent $u \in V$ has a fixed *set of friends* $F(u) \subseteq V$, whereas $F : V \rightarrow \mathcal{P}(V)$ is called a *friendship allocation*. Throughout this chapter, if not specified differently, we only consider friendship allocations that fulfill:

- (a) Friendships are mutual and hence for every $v \in F(u)$ it holds $u \in F(v)$.
- (b) Every agent $u \in V$ has at least one friend: i.e., $|F(u)| \geq 1$.

Considering such a friendship allocation, we define a *friendship graph* $G_F = (V, F)$, whereas the agents V form the graph nodes and there is an edge between two nodes if and only if the respective agents are friends. Edges in this graph are bidirectional.

Game Model. We combine friendship allocations with two different game concepts. On the one hand, we consider Swap-Games (cf. Section 2.2), which are convenient for analyzing structural properties of worst-case equilibrium settings by dismissing the use of an edge price parameter. On the other hand, we study friendship allocations in Buy-Games (cf. Section 2.1) with respect to the positive effects of friendships regarding the social cost.

Swap-Game: In the Swap-Game variant, the agents V are connected by a set of bidirectional edges S . These edges are not owned by anyone and hence any edge can be swapped arbitrarily by any incident agent. Here, the *swap operation* of an agent is the simultaneous removal of an incident edge and replacement by a different incident edge, formally stated as $\{u, v\} \rightarrow \{u, w\}$ for agent u swapping the edge $\{u, v\}$ to edge $\{u, w\}$ (cf. Figure 3.1). An agent's operation can consist of an arbitrary combination of simultaneously executed swaps. The current strategy profile, which is equal to the current set of edges in the network, is called S and in conformity with other models we denote the implied network as $G[S]$. In these games, we only consider connected networks and restrict the agent's actions such that agents must always preserve connectivity.

Any agent strives to minimize her private cost, which is given either by the average distance or by the maximum distance cost function. Namely, in the Sum-Swap-Game,¹ the private cost of an agent is

$$c_u(S) := \frac{1}{|F(u)|} \sum_{v \in F(u)} d_{G[S]}(u, v),$$

and in the Max-Swap-Game, it is

$$c_u(S) := \max_{v \in F(u)} d_{G[S]}(u, v).$$

Here, $d_{G[S]}(u, v)$ denotes the shortest path distance in the network $G[S] = (V, S)$.

Buy-Game: For the Buy-Game variant, we consider buy equilibria of the Sum-Game and the Max-Game (cf. Section 2.1). In this chapter, we name these games Sum-Buy-Game and Max-Buy-Game to avoid confusion with the Swap-Games. In the considered Buy-Games, agents can arbitrarily buy incident edges to other agents, each for a fixed price of $\alpha > 0$. The set of edges of an agent $u \in V$ is given by s_u and S is the joint strategy profile of all individual strategies. For the Sum-Buy-Game, the private cost

¹Note that for the Sum-Swap-Game we normalize the distance cost by the number of friends and thus gain the average distances. This was not necessary in the original games with uniform communication interests, where every agent wanted to communicate with exactly $n - 1$ other agents.

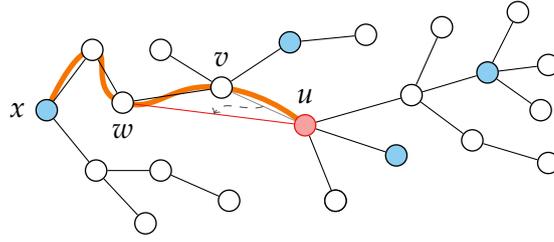


Figure 3.1: Illustration of a swap operation in the Max-Swap-Game. The blue agents denote the friends $F(u)$ of agent u (red), the orange line states the initial longest shortest path from agent u to any of her friends. The swap $\{u, v\} \rightarrow \{u, w\}$ then reduces u 's private cost from 4 to 3.

function is

$$c_u(S) := \alpha \cdot |s_u| + \frac{1}{|F(u)|} \sum_{v \in F(u)} d_{G[S]}(u, v),$$

and in the Max-Buy-Game, it is

$$c_u(S) := \alpha \cdot |s_u| + \max_{v \in F(u)} d_{G[S]}(u, v).$$

The overall quality of a network $G[S]$ is measured by the sum over all private costs and is called the *social cost* $\text{cost}(S) := \sum_{u \in V} c_u(S)$. Using the same terms as in Section 2.2, we denote a strategy profile as equilibrium if no agent can improve her private cost by a unilateral strategy change. To quantify the worst-case loss of selfish behavior, we use the price of anarchy, which is the worst-case ratio of any equilibrium's social cost and the minimal cost of any strategy profile.

3.2 Related Work & Contribution

While network creation games, as introduced by Fabrikant et al. [Fab+03] and their variants, seem to capture the dynamics and evolution caused by the selfish behavior of agents in an accurate way, there is a major drawback: Most of those models assume agents to be interested in communicating with *all* other agents in the network. Given the immense size of communication networks, this seems rather unrealistic. In reality, agents usually communicate in small groups and each only has a small subset of the network participants she is interested in.

The only paper apart from [Cor+12] that considers such non-uniform communication interests in the framework of network creation games is by Halevi and Mansour [HM07]. They also use the above stated concept of friendship allocations to model the non-uniform communication interests of the agents. However, their focus only lies on the Sum-Buy-Game (cf. Section 2.1), for which they proved the existence of equilibria for almost all edge prices of α (in particular, $\alpha \leq 1$ and $\alpha \geq 2$). For general α , they provided an upper bound of $O(\sqrt{n})$ for the price of anarchy. For an average degree d of the friendship graph, i.e., the average number of friends of the agents, in the case of α or d being a constant and $\alpha = O(nd)$, they upper bounded the price of anarchy by a constant. Furthermore, the authors provided a family of problem instances for which the price of anarchy is lower bounded by $\Omega\left(\frac{\log n}{\log \log n}\right)$.

A different approach to introduce non-uniform communication interests was used by Albers et al. [Alb+06]. They apply a so-called weighted traffic matrix to the Sum-Buy-Game such that for the communication cost of an agent every distance is multiplied by a traffic value from the interval $(0, 1)$, which indicates how much traffic should be sent to the target. The special case of having only 0/1-weights results in the friendship graph considered in this chapter. However, their model and applied techniques explicitly require all traffic values to be greater than 0.

Different to explicit a-priori given friendship allocations, there are also models where the utility of an agent is based on how many other agents are in her two-neighborhood, like Nikolettseas et al. [Nik+13], for example. In that game, the utility of an agent is the sum of degrees of her neighbors.

For the results about the uniform interest case in the Buy-Game and the Swap-Game, we refer to Section 2.3.

Contribution. In this chapter, we introduce a generalized class of swap equilibria in network creation games (cf. Section 2.3) by taking the different friends of individual agents into account. For the Swap-Game with friendship allocations, we provide tight price of anarchy results for all interesting model variants: The price of anarchy is worst possible for the Sum-Swap-Game, this even when restricting to the class of tree equilibria. For the Max-Swap-Game it is worst possible for arbitrary equilibria and turns out to be only $\Theta(\sqrt{n})$ for tree

equilibrium networks. The latter result uses an interesting structural insight into equilibrium networks (see Binding-Sequence, Definition 3.7). We show that the price of anarchy for tree equilibrium networks in the Max-Swap-Game can be further characterized by the size M of a maximum independent set in the friendship graph, which gives a price of anarchy of at most $2M$ and hence an improved bound if $M \leq \sqrt{n}$. For example, for a complete friendship graph we have $M = 1$ and hence a constant price of anarchy.

Moreover, we turn our interest to a more optimistic approach of how selfish behavior can deteriorate the social cost. Thereby, we identify a structural property of certain best-response processes, namely that the connected components of the friendship graph are also connected components in respective equilibrium networks. Using this, we introduce the class of *process equilibria* and for this class provide an improved price of anarchy bounds. For the Sum-Buy-Game we provide a so-called process price of anarchy of $O(\log n + \sqrt{N})$; whereas for the Max-Buy-Game it is $O(n^{2/\alpha} + (N/\alpha)^{1/3})$. Here, N is the size of the largest connected component in the friendship graph.

3.3 Worst-Case Friendships in Swap-Games

In this section, we consider the worst-case impact of friendship allocations in Swap-Games. Our focus lies on the existence of equilibria, the convergence of best-response processes, and specifically on bounds for the price of anarchy. For the original Swap-Games with uniform interests and thus with a complete friendship graph, we know from the discussion in Section 2.3 that equilibria always exist, that for tree equilibrium networks the price of anarchy is constant and that for the Sum-Swap-Game with n agents it is at most $2^{O(\sqrt{\lg n})}$, whereas for the Max-Swap-Game only a lower bound of $\Omega(\sqrt{n})$ is known (cf. Alon et al. [Alo+10]).

We start our analysis with games utilizing the maximum-distance price function, for which we show an interestingly different behavior with regard to the price of anarchy, when differentiating between tree equilibrium networks and arbitrary equilibria. Later, the Sum-Swap-Game will show a worst-case behavior also for the class of tree equilibrium networks. This is a remarkable difference to the games with uniform interests, where tree equilibria behave

similarly in both variants and, moreover, the general case actually shows a gap where the Sum-Swap-Game guarantees a much better price of anarchy than the Max-Swap-Game can provide.

To avoid networks becoming disconnected, we restrict the agents to perform only those swaps that preserve the connectivity of the network. This restriction is valid from a practical point of view, where a lost network connectivity is to be avoided, since re-connecting a network causes unpredictable costs – if at all possible. Note that in games with uniform communication interests, agents want to communicate with all other agents and hence a network would never be disconnected.

We start with studying the existence of equilibria when we are given arbitrary friendship allocations. For arbitrary networks this is a trivial question, since having the network equal to the friendship graph raises the cost of 1 for every agent, is an equilibrium, and affirms the price of stability to be 1. On the other hand, for tree network equilibria it is not clear that an equilibrium for a given (possibly cyclic) friendship graph always exists. In the following, we provide an algorithm that computes equilibria for both game variants with social cost that are at most two times those of a socially optimal solution.

Theorem 3.1 (Max-Swap-Game: price of stability). *In the Max-Swap-Game, for any friendship allocation F there exists a tree equilibrium network and the price of stability is at most 2.*

Proof. Let F be a friendship allocation for n agents V . By using Algorithm 1 we construct a tree network, which we claim to be an equilibrium.

In the first loop (lines 4–6), we only add edges between agents v and w who are friends and where one of them has exactly one friend. Thus, the resulting network at line 6 is acyclic. Note that for every agent $v \in B$, who is considered in the second loop (lines 7–12), it holds: (1) v is a friend of at most one agent to whom v is not yet connected and (2) v is connected to an agent $u \in A$ (otherwise, v would have degree one in the friendship graph and therefore would be an element of A instead of B). Thus, adding an edge $\{v, w\}$ with $v \in B$ and $F(v) \setminus \{u \in F(v) \mid \exists u' (u, u') \in S\} = \{w\}$ in line 9 does not create cycles in the network. Finally, in the next loop (lines 14–16) we create a star that connects all formerly created disjoint trees, which contain agents formerly been in set B . By this, the constructed network (V, S) is a set of trees. In the

Algorithm 1: Computation of Equilibrium Networks

```

1  $S \leftarrow \emptyset$ 
2  $A \leftarrow \{v \in V \mid |F(v)| = 1\}$ 
3  $B \leftarrow V \setminus A$ 
4 foreach  $v \in A$  do
5    $S \leftarrow S \cup \{\{v, w\} \mid F(v) = \{w\}\}$  ▷ note that  $w$  is unique
6 end
7 while  $\exists v \in B$  with  $|\{w \in F(v) \mid \{v, w\} \notin S\}| \leq 1$  do
8   if  $|\{w \in F(v) \mid \{v, w\} \notin S\}| = 1$  then
9      $S \leftarrow S \cup \{\{v, w\} \mid w \in F(v) \wedge \{v, w\} \notin S\}$  ▷ note that  $w$  is unique
10  end
11   $B \leftarrow B \setminus \{v\}$ 
12 end
13 select an arbitrary agent as center agent  $x \in B$  and assign  $B \leftarrow B \setminus \{x\}$ 
14 foreach  $w \in B$  do
15    $S \leftarrow S \cup \{\{x, w\}\}$ 
16 end
17 foreach connected component  $C \subset (V, S)$  with  $x \notin C$  do
18   select an arbitrary agent  $y \in C$ 
19    $S \leftarrow S \cup \{x, y\}$ 
20 end
21 return  $S$ ;

```

final loop (lines 17–20), we connect all remaining unconnected components.

It remains to show that (V, S) is an equilibrium. By construction, each agent added during the first two loops (lines 4–12) has a private cost of 1. The center agent x (selected in line 13) of the star also has a private cost of 1, since each friend of x either was connected to x in the previous two loops (lines 4–12) or is connected to x in the last loop (lines 14–16). For each friend u of an agent $w \in B$ who is chosen in line 14 it holds: Either u is chosen in the previous loops (lines 4–12) and the edge $\{u, w\}$ is added to S , or u is connected to x in the last loop (lines 14–16). In both cases, the distance from w to u is at most 2. Thus, w has a private cost of 2. Since $w \notin A$, we have $|F(w)| > 1$ and thus, w cannot perform any improving response. Hence, we get a private cost of at most 2 for every agent.

Finally, we use that every network implies social cost of at least n . Comparing this to the network constructed by the above algorithm, we get a ratio of

$2n/n = 2$, and by this an upper bound on the price of stability. \square

Theorem 3.2 (Sum-Swap-Game: price of stability). *In the Sum-Swap-Game, for any friendship allocation F there exists a tree equilibrium network and the price of stability is at most 2.*

Proof. Let F be a friendship allocation for n agents V . By using Algorithm 1 we construct a tree network, which we claim to be an equilibrium. The arguments from Theorem 3.1 apply unchanged to show that the resulting network is acyclic. Thus, it only remains to show that the computed network $G[S]$ is an equilibrium regarding F .

First, for any agent u with exactly one friend, the first loop (lines 4–12) guarantees that her friend is at distance 1, which is minimally possible. The same holds by loop (lines 7–12) for agents that have exactly two friends, whereas one of them has only one friend. Now consider some agent u with at least two friends $|F(u)| \geq 2$. For every $v \in F(u)$ who was handled in the first two loops, we know that the distance from v to all her friends is 1 and thus minimally possible. If we have $u = x$, i.e., u being the selected center agent, then her distances to all of her friends are exactly 1 and also best possible. Otherwise, u is connected by one edge to the center agent x and has at least two friends to which both shortest paths contain x , since the second loop was not entered for u . Using this, u cannot perform any improving response that swaps her single edge that is used to connect to these agents. \square

For the convergence of best-response processes with a complete friendship graph, we know from Kawald and Lenzner [KL13] that such processes always converge for tree networks but may contain infinite improvement cycles for general networks. In the case of the Max-Swap-Game, we next show that friendship allocations admit best-response cycles also for tree networks.

Proposition 3.3. *The Max-Swap-Game on tree networks with friendships is no potential game.*

Proof. Using the construction provided in Figure 3.2, there are settings that allow cyclic best-response processes over all agents. Hence, the game does not provide the finite-improvement property and cannot be a potential game. \square

3 Loss and Benefit of Friendships

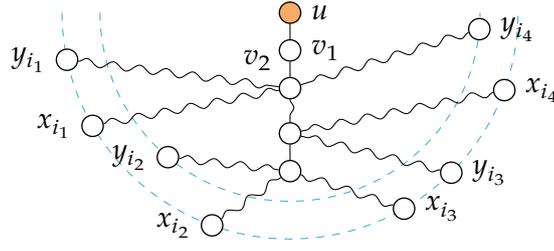


Figure 3.3: Illustration of Lemma 3.6: agent u could perform an improving swap $\{u, v_1\} \rightarrow \{u, v_2\}$ if the T-configuration claim would be wrong.

$u \in V$ with $|F(u)| = 1$ implies $c_u(S) = 1$. Hence, for an agent $u \in V$ with $c_u(S) > 1$ it holds $|F(u)| > 1$.

Lemma 3.6 (T-configuration). *In a Max-Swap-Game, let F be a friendship allocation, $G[S] = (V, S)$ a corresponding tree equilibrium network, and $u \in V$ some agent with $|F(u)| \geq 2$. Then, there exist agents $x, y \in F(u)$ such that*

$$|d_{G[S]}(x, u) - d_{G[S]}(u, y)| \leq 1$$

and u is connected by at most one edge to the shortest path from x to y and $c_u(S) = d_{G[S]}(u, x)$.

Proof. Let $u \in V$ be an agent with $|F(u)| \geq 2$ and $x \in F(u)$ with $d_{G[S]}(u, x) = c_u(S)$. Assume for contradiction that all agents $x' \in F(u) \setminus \{x\}$ are at a distance of $d_{G[S]}(x', u) \leq c_u(S) - 2$ from u . Consider the shortest path (u, v_1, v_2, \dots, x) to agent x . In this case, u can reduce her private cost by the swap $\{u, v_1\} \rightarrow \{u, v_2\}$, since this swap reduces u 's distance to x by 1 but increases the distances to every agent in $F(u) \setminus \{x\}$ by at most 1 each. Using the maximum distance function, we get that u 's private cost decreases and hence, this is a contradiction to S being an equilibrium.

Now we consider all pairs $(x_i, y_i) \in F(u) \times F(u)$ for that it holds $d_{G[S]}(u, x_i) = c_u(S)$ and $d_{G[S]}(u, y_i) \geq c_u(S) - 1$. Let us assume that u is connected to each shortest path from x_i to y_i by at least two edges that do not lie on that path. (cf. Figure 3.3) Thus, u is not located on the shortest path from x_i to y_i . This implies that in the network $G[S] \setminus \{u\}$ for each pair (x_i, y_i) there exists a connected component containing both agents x_i and y_i . Since each two agents at a distance of exactly $c_u(S)$ form such a pair, all agents of $F(u)$ at a distance of exactly $c_u(S)$

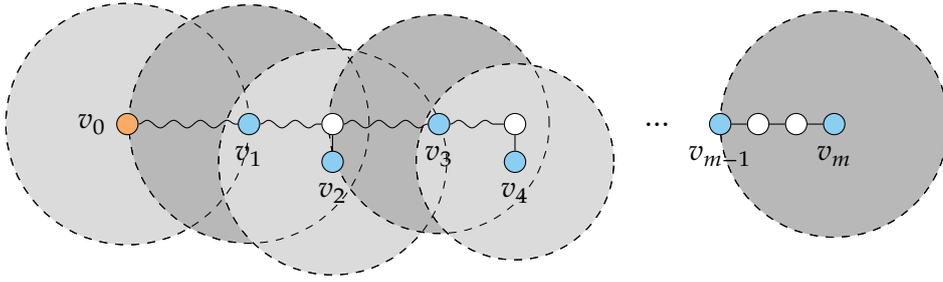


Figure 3.4: Illustration of a Binding-Sequence. The radius of a circle around an agent corresponds to the agent's private cost. Curled lines denote shortest paths.

must be located in the same connected component, which then gives for every pair (x_i, y_i) that both agents are contained in the same connected component. Hence, all agents $x' \in F(u)$ at distance $d_{G[S]}(x', u) \geq c_u(S) - 1$ from u are in the same connected component and by the two edges distance constraint, there must be a path (u, v_1, v_2) that is a subpath of every path from u to every agent x_i and y_i . Hence, u can perform the improving swap $\{u, v_1\} \rightarrow \{u, v_2\}$. This swap decreases her distances to all agents x_i, y_i by 1 each and increases her distances to other agents by at most 1 (i.e., agents $w \in F(u)$ with $d_{G[S]}(w, u) \leq c_u(S) - 2$) and hence contradicts S being an equilibrium. \square

Our main tool for the remainder of the private cost upper bound proof will be so-called *Binding-Sequences*. The definition is as follows:

Definition 3.7 (Binding-Sequence). In a Max-Swap-Game tree equilibrium network $G[S]$ with a friendship allocation F , let $v_0 \in V$ and $v_1 \in F(v_0)$ be friends such that $d_{G[S]}(v_0, v_1) = c_{v_0}(S)$ and further let v_1, \dots, v_m be a sequence of agents such that

- (a) they form a sequence of friends, i.e., $v_i \in F(v_{i-1})$ for $i = 1, \dots, m$,
- (b) all agents have a private cost of $c_{v_i}(S) > 3$, for $i = 0, \dots, m - 1$,
- (c) the last agent has a private cost of $c_{v_m}(S) = 3$, and
- (d) for $i = 2, \dots, m$ it holds:

$$v_i = \arg \max_{v_i \in F(v_{i-1})} \left\{ d_{G[S]}(v_{i-2}, v_i) \mid \begin{array}{l} v_{i-1} \text{ is connected by } \leq 1 \text{ edge to the} \\ \text{shortest path from } v_{i-2} \text{ to } v_i \end{array} \right\}$$

3 Loss and Benefit of Friendships

Then we call this sequence a *Binding-Sequence* (cf. illustration in Figure 3.4).

For such a Binding-Sequence, we will show two key properties that hold in any tree equilibrium network: Given a Binding-Sequence and some agent v_i therein, then

- (a) v_i 's successor v_{i+1} cannot have a much lower private cost than v_i (cf. Lemma 3.8) and
- (b) the shortest path from v_i to v_{i+1} can overlap by at most one edge with the shortest path to v_i 's Binding-Sequence predecessor (cf. Lemma 3.9).

Later, we will show that for any agent there necessarily exists a Binding-Sequence of about the same length as her private cost value. Then, by bounding the maximum length of a Binding-Sequence, we will obtain a private cost upper bound.

Lemma 3.8. *For a friendship allocation F , let S be a Max-Swap-Game tree equilibrium strategy profile and v_0, \dots, v_m a Binding-Sequence. Then, for each two consecutive sequence agents v_i and v_{i+1} , with $0 \leq i < m$, it holds $d_{G[S]}(v_i, v_{i+1}) \geq c_{v_i}(S) - 1$ and $c_{v_{i+1}}(S) \geq c_{v_i}(S) - 1$.*

Proof. For $i \in \{0, \dots, m - 1\}$ consider an agent v_i in the Binding-Sequence. Then, by Lemma 3.6 there exist $x, y \in F(v_i)$ with $d_{G[S]}(v_i, x) = c_{v_i}(S)$ and $c_{v_i}(S) \geq d_{G[S]}(v_i, y) \geq c_{v_i}(S) - 1$ such that v_i is connected by at most one edge to the shortest path from x to y . At least one of these agents is a valid candidate for the next Binding-Sequence agent v_{i+1} . Yet, even if v_{i+1} is neither x nor y , still we gain a lower bound for the maximum distance:

$$d_{G[S]}(v_i, v_{i+1}) \geq \min\{d_{G[S]}(v_i, x), d_{G[S]}(v_i, y)\} \geq c_{v_i}(S) - 1$$

This further gives $c_{v_{i+1}}(S) \geq c_{v_i}(S) - 1$. □

Lemma 3.9 (Increasing Distance). *For a friendship allocation F , let S be a Max-Swap-Game tree equilibrium strategy profile and v_0, \dots, v_m a Binding-Sequence. Then, the distances to v_0 are monotonously increasing, i.e., $d_{G[S]}(v_0, v_i) \leq d_{G[S]}(v_0, v_{i+1})$ for $i = 1, \dots, m - 1$.*

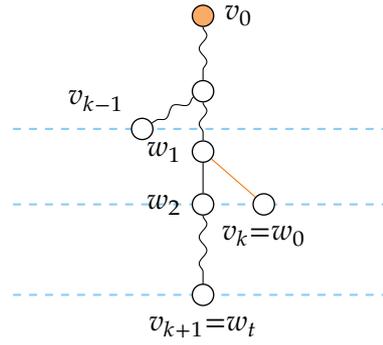


Figure 3.5: Illustration of Lemma 3.10: edge $\{w_0, w_1\}$ is used only two times.

Proof. Using $c_{v_1}(S) \geq 3$ we get with Remark 3.5 that $|F(v_1)| \geq 2$. Hence, by Lemma 3.6 there exists an agent v_2 such that the paths from v_1 to v_0 and from v_1 to v_2 overlap by at most one edge. By construction of the Binding-Sequence, the distance $d_{G[S]}(v_0, v_2)$ is maximal among all distances from v_0 to agents $v \in F(v_1)$ and hence we get $d_{G[S]}(v_0, v_1) \leq d_{G[S]}(v_0, v_2)$.

Now assume that there is an agent v_i with the smallest index $i \geq 2$ in the Binding-Sequence for which the claim does not hold. This is, $d_{G[S]}(v_0, v_{i-1}) \leq d_{G[S]}(v_0, v_i)$ and $d_{G[S]}(v_0, v_i) > d_{G[S]}(v_0, v_{i+1})$. Denote by x the most distant agent from v_0 who is on all shortest paths from v_0 to v_{i-1} , from v_0 to v_i , and from v_0 to v_{i+1} . Such an agent x exists, since especially v_0 fulfills the restrictions. By the choice of i and since all these paths contain agent x , we get:

$$d_{G[S]}(x, v_{i-1}) \leq d_{G[S]}(x, v_i) > d_{G[S]}(x, v_{i+1}) \quad (3.1)$$

By definition of the Binding-Sequence, v_i is connected by at most one edge to the shortest path from v_{i-1} to v_{i+1} . Hence, x must be an agent on the path from v_{i-1} to v_{i+1} . First note that x cannot be v_i or a neighbor of v_i , since for those cases (3.1) yields $d_{G[S]}(x, v_{i+1}) < d_{G[S]}(x, v_i) \leq 1$. Furthermore, x must lie on the shortest path from v_{i-1} to v_i , since otherwise x would lie on the shortest path from v_i to v_{i+1} , which together with $d_{G[S]}(v_{i-1}, v_i) \geq 3$ would imply $d_{G[S]}(x, v_i) < d_{G[S]}(x, v_{i-1})$. But this gives $d_{G[S]}(x, v_i) \leq d_{G[S]}(x, v_{i+1})$ and is a contradiction. \square

Lemma 3.10. *For a friendship allocation F , let S be a Max-Swap-Game tree equilibrium strategy profile with network $G[S] = (V, S)$ and v_0, \dots, v_m a Binding-Sequence.*

3 Loss and Benefit of Friendships

Then, no edge in S is used more than two times by the shortest path visiting the agents v_0, \dots, v_m in the given order.

Proof. We label the agents of $G[S]$ by their distances to v_0 . This is, for every $v \in V$ we define $\text{level}(v) := d_{G[S]}(v_0, v)$ to be the distance to v_0 . For an arbitrary agent v_k with $k \in \{1, \dots, m-1\}$ we consider the corresponding shortest path ($v_k := w_0, w_1, \dots, w_t := v_{k+1}$) to agent v_{k+1} of some length t . By definition, v_k is connected by at most one edge to the shortest path from v_{k-1} to v_{k+1} (for an illustration cf. Figure 3.5). By Lemma 3.9 we have $\text{level}(v_{k-1}) \leq \text{level}(v_k) \leq \text{level}(v_{k+1})$. Hence, for $i = 2, \dots, t-1$ we get $\text{level}(w_i) < \text{level}(w_{i+1})$. This means that at most one edge (specifically edge $\{w_0, w_1\}$) of the shortest path from v_0 to v_k is used a second time by the shortest path traversal from v_k to v_{k+1} . By Lemma 3.8 we have $t \geq c_{v_k}(S) - 1 \geq 3$ and get $\text{level}(v_k) < \text{level}(v_{k+1})$. \square

Finally, we conclude the proof of the private cost upper bound by considering a pair of most distant friends and show that their distance corresponds to a Binding-Sequence of similar length. Using that for the traversal of a Binding-Sequence every tree edge is used at most twice, we get an upper bound on its length and by this an upper bound in the maximal distance.

Theorem 3.11 (Max-Swap-Game: private cost upper bound). *For a friendship allocation F let $G[S] = (V, E)$ be a Max-Swap-Game tree equilibrium network with $n := |V|$ agents. Then, for all $u \in V$ we have $c_u(S) = O(\sqrt{n})$.*

Proof. Let $v_0 \in V$ be an agent with maximal private cost. We can assume that v_0 has at least one friend at a distance of at least 3, since otherwise the claim already holds. Let v_1 be a most distant friend $v_1 \in F(v_0)$ and denote the distance between v_0 and v_1 as $D := d_{G[S]}(v_0, v_1) = c_{v_0}(S)$.

(Existence.) Agents v_0, v_1 obviously fulfill the conditions for a Binding-Sequence. Thus, it suffices to show that given the beginning of a Binding-Sequence v_0, \dots, v_i with $c_{v_j}(S) > 3$, for $j = 0, \dots, i-1$, either we can find a next agent v_{i+1} who suffices the conditions or otherwise $c_{v_i}(S) = 3$ and the sequence terminates. If we assume $c_{v_i}(S) > 3$, then by Lemma 3.6 there exist agents $x, y \in F(v_i)$ with $d_{G[S]}(v_i, x) = c_{v_i}(S)$ and $c_{v_i}(S) \geq d_{G[S]}(v_i, y) \geq c_{v_i}(S) - 1$ such that v_i is connected by at most one edge to the shortest path from x to y . Since $c_{v_i}(S) > 3$, both $c_x(S) \geq 3$ and $c_y(S) \geq 3$ hold. Now, for at least one

agent (x or y) we have that this agent is most distant to v_{i-1} , she is not v_{i-2} , and thus she fulfills the conditions for a Binding-Sequence.

(*Traversal.*) Given the existence, now we can apply the previous lemmas for providing the minimal length of such a Binding-Sequence: Lemma 3.8 states that by construction of the Binding-Sequence we always have $c_{v_{i+1}}(S) \geq c_{v_i}(S) - 1$. Lemma 3.9 implies that no agent can be contained more than once in a Binding-Sequence. By the arguments above we get that we can always find a new agent for the Binding-Sequence until we reach an agent w with $c_w(S) = 3$. Hence, the Binding-Sequence contains at least $c_{v_0}(S) - 2$ agents. Since the distance between two succeeding agents of the Binding-Sequence decreases by at most one per agent, a traversal of this Binding-Sequence consists of at least

$$\sum_{i=3}^{c_{v_0}(S)} i = \frac{c_{v_0}(S)^2 + c_{v_0}(S) - 6}{2}$$

edges. From these edges, by Lemma 3.10, at least $(c_{v_0}(S)^2 + c_{v_0}(S) - 6)/4$ many edges are different.

(*Private cost upper bound.*) Finally, we use that the traversal of the Binding-Sequence uses at least $\frac{D^2+D-6}{4}$ -many different edges. Since the tree has exactly $n-1$ edges, we get $(D^2+D-6)/4 \leq n-1$ as an upper bound for the size of every Binding-Sequence and hence the private cost upper bound is $D = O(\sqrt{n})$. \square

Next we show that this private cost bound is actually tight. This means, there are combinations of a friendship allocation and a tree equilibrium network of n agents such that there is an agent with private cost of $\Omega(\sqrt{n})$. For this, we consider the following ring friendship graph.

Theorem 3.12. *There exists a friendship allocation F and corresponding Max-Swap-Game tree equilibrium network $G[S]$ of n agents V in which some agent has a private cost of $\Omega(\sqrt{n})$.*

Proof. For the agents $V = \{v_1, \dots, v_n\}$, we consider the friendship allocation $F := \{\{v_i, v_{i+1}\} \mid i = 1, \dots, n-1\} \cup \{\{v_n, v_1\}\}$, forming a ring friendship graph, and a corresponding network $G[S] = (V, S)$ as stated in Figure 3.6. We claim that the network is an equilibrium and yields a private cost of $c_{v_i}(S) = \Omega(\sqrt{n})$ for agent $v_i \in V$ (index i will be specified later). Specifically, for the private costs we have

3 Loss and Benefit of Friendships

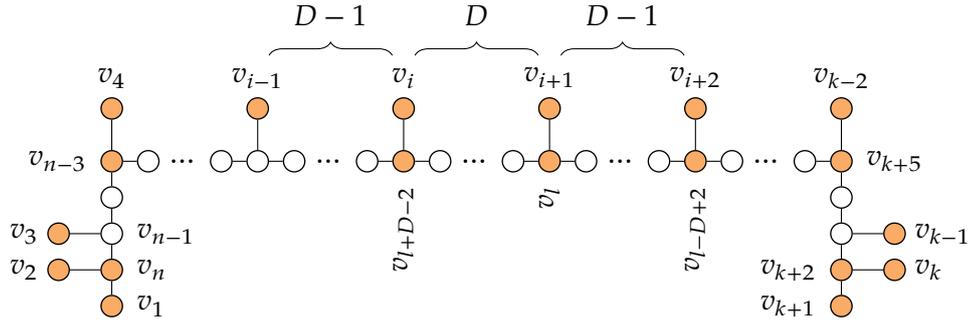


Figure 3.6: Illustration of a Max-Swap-Game tree equilibrium strategy profile $G[S] = (V, S)$ of $n := |V|$ agents with a friendship allocation ring graph such that the private cost of v_i is $\Omega(D)$, with $D := \sqrt{n-2} + 1$, $k := 2D - 3$, and $l = n - \sum_{i=1}^D i$.

- for $j = 1, \dots, i-1$ that $c_{v_{j+1}}(S) = c_{v_j}(S) + 1$ and
- for $j = i+1, \dots, k$ that $c_{v_j}(S) = c_{v_{j+1}}(S) + 1$.

We first compute the exact value for $c_{v_i}(S)$ given by this setting, then we argue why no agent in this network can perform an improving response. Denote the maximal distance from v_i to any of her friends by D . Then D must fulfill

$$n = \sum_{i=1}^{D-2} i + \sum_{i=1}^{D-3} i + 2(D-2) + 3 = D^2 - 2D + 3,$$

and hence, $D = \sqrt{n-2} + 1$. This yields a private cost of $\sqrt{n-2} + 1$ for agent v_i when we fix the parameters as $i := D - 1$ and $k := 2D - 3$.

For each agent with a degree greater than 1 in $G[S]$ we have a private cost of 1 and hence no improving response is possible. Otherwise, consider some agent v_j of degree 1 in $G[S]$. Agent v_j cannot perform any swap if and only if it holds both, $|d_{G[S]}(v_{j-1}, v_j) - d_{G[S]}(v_j, v_{j+1})| \leq 1$ and v_j is connected by one edge to the shortest path from v_{j-1} to v_{j+1} . Since this property is given by construction, v_j cannot perform any improving response. \square

An interesting insight from the last theorem is that the used stability argument for T-configurations (cf. Lemma 3.6) characterizes ring friendship graphs in general: Every agent must be in the center of her two friends.

3.3.2 The Price of Anarchy in Max-Swap-Games

Continuing the analysis of the worst-case behavior of friendships in the Max-Swap-Game, next we consider the price of anarchy. At first, we will provide a lower bound for tree equilibrium networks and then use the private cost upper bound to show that this bound is tight. In the setting of tree equilibria, we will further characterize the price of anarchy by the structure of the friendship graph, namely the size of a maximum independent set therein. We will conclude this section by showing that the price of anarchy is worst possible when considering arbitrary networks.

Lemma 3.13. *There exists a friendship allocation F and a corresponding Max-Swap-Game tree equilibrium network $G[S]$ of n agents such that the social cost is $\Omega(n^{3/2})$.*

Proof. We create a network $G[S] = (V, S)$ of agents v_1, \dots, v_n . For a fixed parameter $D := \frac{\sqrt{2n-7}-3}{2}$, we first connect agents $v_{D+1}, \dots, v_{n/2-D-1}$ as a line and then further connect agents $v_{n/2}, \dots, v_n$ to agent v_l , whereas $l := \frac{n}{2} - D - (\sum_{i=1}^D i + 2)$. The remaining agents are connected as leaves to specific places at the line: For agents v_1, \dots, v_D , first v_1 is connected to v_{l+D} , then v_2 is connected to v_{l+2D-1} , and further up to v_D , the agents are connected such that the distance between each next pair decreases by 1 (cf. Figure 3.8). We make the same construction for agents $v_{n/2-1}$ to $v_{n/2-D}$, whereas $v_{n/2-1}$ is connected to v_{l-D} and the remaining agents are again connected such that the distances decrease by 1 with each pair. Note that by the choices for l and D , we have $n = 2 \sum_{i=1}^D i + \frac{n}{2} + 2D + 4$ and hence the network can actually be constructed as stated above.

The corresponding friendship graph consists of a ring, which connects agents $v_1, \dots, v_{n/2}$, and additionally connects agents $v_{n/2+1}, \dots, v_n$, such that each of them is a friend of both agent $v_{n/2-1}$ as well as agent v_1 (cf. Figure 3.7). Considering the network $G[S]$, this implies a private cost of $(\sqrt{2n-7} + 1)/2$ for all agents $v_{n/2}, \dots, v_n$. Thus, the social cost of S is $\text{cost}(S) = \Omega(n^{3/2})$. The arguments that $G[S]$ is an equilibrium for the given friendship graph apply analogously to Lemma 3.12. \square

Theorem 3.14 (Max-Swap-Game: price of anarchy for tree networks). *In the Max-Swap-Game with friendship allocations, the price of anarchy for tree network equilibria is $\text{PoA} = \Theta(\sqrt{n})$, with n being the number of agents.*

3 Loss and Benefit of Friendships

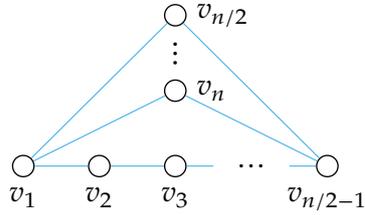


Figure 3.7: The friendship graph for the proof of Lemma 3.13.

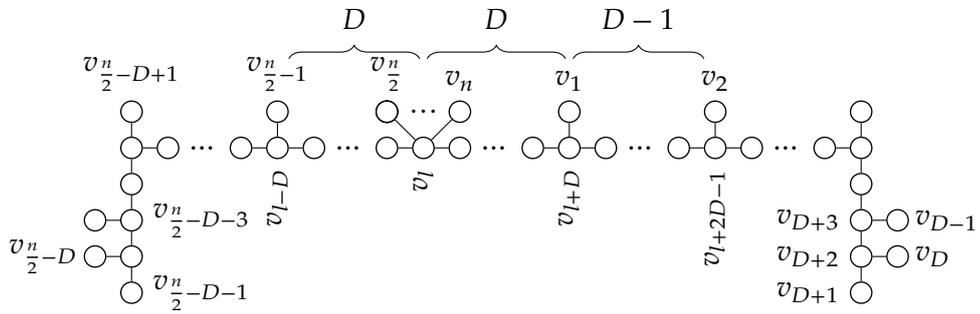


Figure 3.8: The network for the proof of Lemma 3.13. This tree equilibrium network corresponds to the friends as given in Figure 3.7. The parameters are $D = \frac{\sqrt{2n-7}-3}{2}$ and $l = \frac{n}{2} - D - \left(\sum_{i=1}^D i + 2\right)$.

Proof. For the upper bound, we apply Theorem 3.4, which states for tree equilibrium networks that the private cost of every agent is at most $O(\sqrt{n})$. By this, the social cost for every tree network is at most $O(n^{3/2})$. Using Lemma 3.13, we get that this bound is actually tight and the worst-case social cost is $\Theta(n^{3/2})$. On the other hand, for any friendship allocation in an optimal solution the social cost is between n and $2n$. Hence, the worst-case ratio of both is $\text{PoA} = \Theta(\sqrt{n})$. \square

Next, we provide a different characterization of the price of anarchy for tree equilibria, namely by the size of a maximum independent set in the friendship graph. We will get this bound by using the maximum independent set size to bound the maximal length of a Binding-Sequence, similar to the proof of Theorem 3.11. Here, a *maximum independent set* (MIS) of a given graph $G = (V, S)$ is a subset $M \subset V$ of maximum size such that for no two $u, v \in M$ there is an edge connecting them.

Lemma 3.15. *For a friendship allocation F , let $G[S] = (V, S)$ be a Max-Swap-Game tree equilibrium network of $n := |V|$ agents and let $M \subset V$ be a maximum independent set in the friendship graph (V, F) . Then, the length of every Binding-Sequence is at most $2M$.*

Proof. Let v_0, \dots, v_m be a Binding-Sequence with maximal length. We will prove that the agents of this sequence with even index form an independent set in the friendship graph (V, F) . For this, consider an even index i and assume for contradiction that there is an even index $k < i$ such that $v_k \in F(v_i)$. By Lemma 3.9 we get $d_{G[S]}(v_k, v_{k+1}) \leq d_{G[S]}(v_k, v_{k+2})$. If $v_{k+2} \neq v_i$, then by Lemma 3.8 and by $c_{v_j}(S) > 3$, for all v_j in the Binding-Sequence, we get:

$$d_{G[S]}(v_k, v_i) > d_{G[S]}(v_k, v_{k+2}) + 1 \geq c_{v_k}(S)$$

Yet, this is a contradiction.

Thus, consider the case $v_{k+2} = v_i$. Since v_{k+1} is connected by at most one edge to the shortest path from v_k to v_{k+2} and $d_{G[S]}(v_{k+1}, v_{k+2}) \geq 3$ we get that $v_{k+2} \notin F(v_k)$. Otherwise, we either get the same contradiction as before or v_{k+1} would contradict to be the most distant agent in $F(v_k)$ who fulfills the Binding-Sequence conditions.

Hence, the agents with an even index of the Binding-Sequence form an independent set in (V, F) . Since an independent set has at most M agents, we

3 Loss and Benefit of Friendships

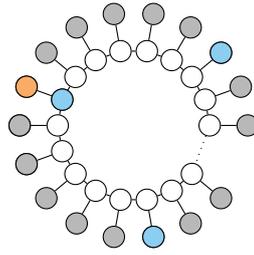


Figure 3.9: Lower bound construction for the price of anarchy in the Max-Swap-Game for general networks. Each of the ring agents is a friend of her three neighbors. Each satellite agent is a friend of her ring neighbor as well of the two agents at a distance of exactly $n/6 + 2$. In this illustration, the three friends of a satellite agent (marked in orange) are marked in blue.

get an upper bound of $2M$. □

Theorem 3.16. *In the Max-Swap-Game with friendship allocations, let n be the number of agents and M the size of a maximum independent set in the friendship graph. Then, the price of anarchy for tree equilibria networks is $\text{PoA} = O(M)$.*

Proof. By using Lemma 3.15, we know that the maximum Binding-Sequence length is $2M$. Now we use the same arguments as in the proof of Theorem 3.11, yet with $2M$ as the maximum length, and get $O(M)$ as the upper bound on the private cost of every agent. With the arguments from Theorem 3.14 we deduce the price of anarchy upper from the private cost upper bound. □

This theorem further shows how the maximum independent set characterization of the friendship graph provides a nice parametrization of tree equilibrium networks for the original Max-Swap-Game with uniform communication interests, as considered by Alon et al. [Alo+10]. Given a game with a complete friendship graph, the maximum independent set has size 1 and hence yields a constant price of anarchy. Then, with increasing size of the independent set, the upper bound for the price of anarchy linearly increases.

Corollary 3.17. *In the Max-Swap-Game with friendship allocations, if the friendship graph forms a clique, then for tree equilibrium networks the price of anarchy is $O(1)$.*

In the following theorem we will show that, in contrast to tree equilibrium networks, the price of anarchy will become worst possible when considering arbitrary equilibrium networks. As a reminder, for this class of arbitrary

equilibrium networks we know from Alon et al. [Alo+10] that for uniform communication interests the price of anarchy is at least $\Omega(\sqrt{n})$, although no non-trivial upper bound is known.

Theorem 3.18 (Max-Swap-Game: price of anarchy for general networks). *In the Max-Swap-Game with friendship allocations, with n being the number of agents, the price of anarchy is $\text{PoA} = \Theta(n)$.*

Proof. First note that the social cost of every strategy profile is upper bounded by $n(n-1)$ and lower bounded by n . Secondly, we provide a friendship allocation for n agents (with n being a multiple of 6) and a corresponding equilibrium network $G[S] = (V, S)$ such that the social cost is $\Omega(n^2)$ (cf. Figure 3.9). For this, we connect $(n/2)$ -many agents as a ring and call them ring agents. For each ring agent, we connect one additional satellite agent to her. Each of the ring agents is a friend of her three adjacent agents in $G[S]$, whereas each satellite agent is a friend of her neighbor at the ring and of both satellite agents at a distance of exactly $n/6 + 2$. This construction is an equilibrium and all $n/2$ satellite agents have a private cost of $n/6 + 2$ each, which gives the claimed price of anarchy of $\Omega(n)$. \square

3.3.3 The Price of Anarchy in Sum-Swap-Games

In the following, we will consider the price of anarchy in Sum-Swap-Games. In comparison to the games with complete friendship graphs, as considered by Alon et al. [Alo+10], we will prove that for tree equilibrium networks as well as for arbitrary networks the price of anarchy will become worst possible. By this, the results are in stark contrast to the non-uniform variant. Specifically, we use very sparse friendship allocations to obtain these worst-case results. Note that the following result specifically applies for general networks, too.

Theorem 3.19 (Sum-Swap-Game: price of anarchy for tree networks). *In the Sum-Swap-Game with friendship allocations, for tree equilibrium networks the price of anarchy is $\text{PoA} = \Theta(n)$.*

Proof. We consider a line network of agents v_1, \dots, v_n and select the biggest integer D such that it holds $3D + 2 \leq n$. All agents on the line are friends of their direct neighbors. Furthermore, for $i = D + 2, \dots, 2D + 2$, we define the friends of agent v_i as $F(v_i) = \{v_{i-1}, v_{i+1}, v_{i-D}, v_{i+D}\}$ (cf. Figure 3.10).

3 Loss and Benefit of Friendships

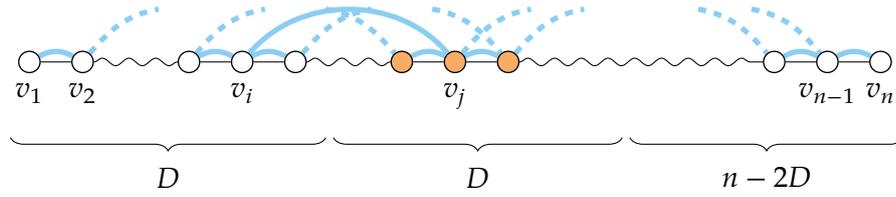


Figure 3.10: Illustration of tree equilibrium network $G[S] = (V, S)$ with social cost $\Omega(n^2)$ as used in Theorem 3.19. Both leaf agents have only their direct neighbors as friends. The agents in the first and last part have 3 friends each, whereas one friend is at distance D each. The orange agents in the center part of the line have 4 friends each, whereas two of them are at distance D .

For the leaf agents of the line, we get that each of them has exactly one friend at distance 1 and hence will not perform an improving response. For any other agent, say v_i , we have a degree of 2 in $G[S]$. Since the network is a tree and v_i has friends in both directions of the line, after any possible improving response still one edge must point in each direction of the line, since otherwise the tree would become disconnected and the private cost of v_i would be unbounded. Considering only one line direction for agent v_i , she has exactly two friends in this direction and is connected to the closest of them. Yet, swapping this edge to any other agent cannot decrease her private cost, since her distance cost to one friend increases by at least the same amount that the distance decreases to the other friend.

For a social cost lower bound of this network, we consider only the agents who have at least 3 friends. For each such agent, the private cost is at least $(2 + D)/3$, which gives for the social cost $3D(2 + D)/3 \geq \frac{(n-4)}{3} \left(2 + \frac{n-4}{3}\right)$. On the other hand, by considering a star network, the private cost of any agent in an optimal solution is at most 2, since an agent has a maximum distance of at most 2 to her friends. Hence, by comparing both social cost bounds we get $\text{PoA} = \Omega(n)$. Since the price of anarchy for connected tree equilibrium networks cannot be higher than n , this bound is tight. \square

This result is actually surprising when compared to the behavior of other network creation games. Usually, for those games one can observe that the average distance version leads to better social costs than the maximum version. Yet, for Swap-Games with friendship allocations this is the contrary – although, in our case, it only holds for tree equilibria and in the case of arbitrary

equilibrium networks, both games show a similar worst-case behavior.

3.4 Process Equilibria

Typically, in real networks equilibrium states result from the best-response strategy changes of the agents and are not instances specifically crafted to show a particular worst-case behavior, as we saw it in the first part of this chapter; although it is the usual approach in the overwhelming body of literature about network creation games. In the remainder of this chapter, we address this issue by taking a more optimistic view on the effects of non-uniform communication interests in the Buy-Game. Specifically, we want to restrict our analysis to only those networks that result from best-response processes when starting with an empty network. By Kawald and Lenzner [KL13] we know that such best-response processes are not guaranteed to converge. Hence, approaches that exploit the behavior of potential functions, as they were studied for several other games (for example Nisan et al. [Nis+07, Chapters 18 and 19]), do not apply here.

Our approach is to identify properties that every best-response process ensures. By using such properties, we can specify a class of equilibria that excludes incompatible equilibrium. Although this class is possibly larger than the class we really want to consider, price of anarchy results therein directly apply for the class of best-response process equilibria we are interested in. We call the so-shaped equilibrium class *process equilibria* and name the price of anarchy restricted to only process equilibria the *process price of anarchy*:

Definition 3.20 (Process Equilibrium). Given agents V and a friendship allocation F , then an equilibrium is called a *process equilibrium* if it can be reached by a sequence of best-response strategy changes of the agents that starts from an network without any edges.

The main property we will use for bounding the process price of anarchy is given by the following observation. Briefly, we use that the size of the largest connected component in the friendship graph also bounds the size of the largest connected component in any process equilibrium network.

Observation 3.21. Let G_F be a friendship graph for an agent set V and N the size of the largest connected component in the friendship graph. Then, for every

3 Loss and Benefit of Friendships

process equilibrium it holds: For an agent $u \in V$ it is never a best response to connect to a connected component that does not contain any of her friends. Thus, by starting from an empty network, an agent will never connect to a connected component not containing any of her friends and consequently in any process equilibrium no connected component of the network is bigger than N .

First, we provide a bound for the Sum-Buy-Game by adapting a technique by Halevi and Mansour [HM07]. For this, we utilize the following lemma by Dutton and Brigham [DB91] to estimate the number of edges in equilibrium networks of limited girth. Note that the following proofs also hold for unidirectional friendships and hence G_F could also be modeled as a directed graph yielding the same results.

Lemma 3.22 (Dutton and Brigham [DB91]). *Given a graph G and a connected component $C \subseteq G$ of n nodes, then the length of a shortest cycle of odd length g is at most $O(n^{1+2/(g-1)})$.*

Theorem 3.23. *In the Sum-Buy-Game with agents V and a friendship graph G_F , let N be the size of the largest connected component of the friendship graph. Then, the process price of anarchy is at most $O(\log n + \sqrt{N})$.*

Proof. In the following, let S be a process equilibrium strategy profile and $G[S] = (V, E)$ the corresponding network. By Observation 3.21, we know that for every process equilibrium no connected component consists of more than N agents.

(Distance cost upper bound.) For an agent u , let T_u be a shortest path tree consisting of the shortest paths from agent u to all her friends $F(u)$. For a parameter $h \in \mathbb{N}$, define $X(h) := \{v \in F(u) \mid h \leq d_{G[S]}(u, v)\}$ to be the set of agents u 's friends who have a distance of at least h to u . Using a fixed parameter h , we can write u 's distance cost as:

$$\begin{aligned} \text{dist}_u(S) &= \sum_{v \in X(h)} d_{G[S]}(u, v) + \sum_{v \in F(u) \setminus X(h)} d_{G[S]}(u, v) \\ &\leq \sum_{v \in X(h)} d_{G[S]}(u, v) + (h-1)(|F(u)| - |X(h)|) \\ &= \sum_{v \in X(h)} (d_{G[S]}(u, v) - h + 1) + (h-1)|F(u)| \end{aligned}$$

We consider the first term of the last estimation: For every $x \in V$, let m_x denote the number of u 's friends whose unique paths to u in the tree T_u contain x . Since the network is an equilibrium, u cannot perform any improving response. By this, we get $m_x(d_{G[S]}(u, x) - 1) \leq \alpha$. For each $v \in X(h)$, we can interpret the value $d_{G[S]}(u, v) - h + 1$ as the number of different agents with distance of at least h to u who are visited on the shortest path from u to v in T_u . Hence, when looking at all agents $X(h)$, we can write:

$$\begin{aligned} \sum_{v \in X(h)} (d_{G[S]}(u, v) - h + 1) &= \sum_{x \in T_u: d_{G[S]}(u, x) \geq h} m_x \\ &\leq \sum_{x \in T_u: d_{G[S]}(u, x) \geq h} \frac{\alpha}{d_{G[S]}(u, x) - 1} \leq \frac{N-2}{h-1} \alpha \end{aligned}$$

Setting $h := 1 + \left\lceil \sqrt{\alpha \frac{N-2}{|F(u)|}} \right\rceil$, we get $\text{dist}_u(S) \leq 3\sqrt{(N-2)\alpha \cdot |F(u)|}$.

(*Edge cost upper bound.*) Next, we estimate an upper bound on the number of edges (utilizing the technique from Halevi and Mansour [HM07], their Theorem 4). For every agent $u \in V$ and any of her edges $v \in s_u$, we define $C(u, v)$ to be u 's friends to which the distance from u increases if the edge $\{u, v\}$ is removed. Accordingly, we assign a weight $w(u, v) := |C(u, v)|$ to this edge and define the total weight to be $W := \sum_{\{u, v\} \in E} w(u, v)$. (Note that in an equilibrium no edge is built twice and hence the weight is well-defined according to the agent who created this edge.) For some parameter $\beta > 0$ (to be specified later), by taking (V, E) and removing all edges $\{u, v\} \in E$ with weight $w(u, v) \geq \beta$, we transform the network into a new network (V, E') and implicitly also gain a new strategy profile S' .

Let m be the number of removed edges, then the total removed weight is at least βm and at most W . For each agent $u \in V$ and any $v \in F(u)$ there is at most one $x \in s_u$ such that the shortest path from u to v uses the edge $\{u, x\}$. Hence, the sum over all weights is upper bounded by twice the number of friendships, i.e., $W \leq 2 \sum_{u \in V} |F(u)|$, and we get:

$$|E'| \geq |E| - 2 \sum_{u \in V} \frac{|F(u)|}{\beta}$$

Since the edge $\{u, v\}$ exists in the equilibrium network $G[S]$, the edge price α

3 Loss and Benefit of Friendships

is at most the distance cost increase for u that we get by removing the edge. Each such distance cost decrease is upper bounded by the distances in $G[S']$ and we get:

$$\alpha \leq w(u, v)(d_{G[S']}(u, v) - 1) \quad (3.2)$$

The new network $G[S']$ is either cycle-free, and the number of edges is limited by at most $n - 1$, or there exists a cycle of a minimal length g . Let $\{u, v\}$ be an arbitrary edge of this cycle and let it be owned by u . By using (3.2), we get:

$$\alpha \leq w(u, v)(g - 2) \leq \beta(g - 2) \quad \Rightarrow \quad g \geq 2 + \alpha/\beta$$

Now, we can choose $\beta := \frac{\alpha}{2 \log n}$ and apply Lemma 3.22. This gives an upper bound of $O(n^{1+1/\log n}) = O(n)$ for $|E'|$. Applying this to $G[S]$, we have $|E| = O\left(n + \log n \sum_{u \in V} \frac{|F(u)|}{\alpha}\right)$.

(Price of anarchy.) Finally, we compare both bounds for distance cost and for edge cost with the social cost lower bound of $\alpha n/2 + \sum_{u \in V} |F(u)|$. The total edge cost easily gives us a ratio upper bound of $O(\log n)$. For estimating the ratio for the distance cost, we first define $\bar{f} := \sum_{u \in V} |F(u)|/n$ to be the average number of friends of an agent and then separately consider if $\alpha \geq \bar{f}$ or not:

$$O\left(\frac{\sum_{u \in V} 3\sqrt{(N-2)\alpha|F(u)|}}{\alpha n/2 + \sum_{u \in V} |F(u)|}\right) = O\left(\frac{n\sqrt{N\alpha\bar{f}}}{\alpha n + n\bar{f}}\right) = O(\sqrt{N})$$

Summed up, the price of anarchy is at most $O(\log n + \sqrt{N})$. \square

Compared to the results by Halevi and Mansour [HM07], the process price of anarchy for the Sum-Buy-Game gives much better results when the friendship graph is sparse enough. Specifically, if the maximal size of a connected component is at most $(\log n)^2$, the process price of anarchy becomes $O(\log n)$.

For process equilibria in the Max-Buy-Game, we can bound the process price of anarchy in a similar way by incorporating the size of the largest connected component. Thereby, we use a technique by Demaine et al. [Dem+12].

Theorem 3.24. *In the Max-Buy-Game with agents V and a friendship graph G_F , let N be the size of the largest connected component of the friendship graph. Then, the process price of anarchy is at most $O(n^{2/\alpha} + (N/\alpha)^{1/3})$.*

Proof. In the following, let S be a process equilibrium strategy profile and $G[S] = (V, E)$ the corresponding network. By Observation 3.21, we know that for every process equilibrium no connected component consists of more than N agents.

(*Distance cost upper bound.*) For any agent $u \in V$, we define the set of all agents within a distance of at most k to agent u as $B_k(u) := \{v \in V \mid d_{G[S]}(u, v) \leq k\}$ and the set of all agents with a distance of exactly k to agent u as $\overline{B_k(u)} := \{v \in V \mid d_{G[S]}(u, v) = k\}$. Given these sets, we claim that for any parameter k with $k \leq \text{diam}(G[S])/2$ it holds:

$$|B_{k+1}(u)| \geq k^2/(2\alpha)$$

Now consider the possible strategy change of u that consists of buying the edges to all agents $v \in \overline{B_k(u)}$, this would decrease u 's distance cost by at least k , while increasing her edge cost by $\alpha \cdot |\overline{B_k(u)}|$. Since S is an equilibrium, we get $|\overline{B_{k+1}(u)}| \geq k/\alpha$. By summing over all ranges this yields:

$$|B_{k+1}(u)| \geq \sum_{i=1}^k k \frac{i}{\alpha} > \frac{k^2}{2\alpha}$$

Next, fix the parameter k such that $k := \frac{\text{diam}(G[S])}{4} - 1$ and consider an agent u such that $\text{dist}_u(S) = \text{diam}(G[S])$. We select a set of cluster centers C by the following iterative algorithm: First mark all agents within a distance of $2k$ from u , then iteratively select an arbitrary unmarked agent c , add c to C , mark all agents within distance of at most $2k$ from c , and continue this procedure with the next iteration until all agents are marked.

Now assume a strategy change of u that creates edges to all agents in C . This would raise an additional edge cost of $\alpha \cdot |C|$ for u . Since for any $c_i, c_j \in C$ with $c_i \neq c_j$ it holds $B_k(c_i) \cap B_k(c_j) = \emptyset$, we get $N \geq |C| \cdot \frac{(k-1)^2}{2\alpha}$ and hence $|C| \leq \frac{2n\alpha}{(k-1)^2}$. Using that S is an equilibrium and thus $\alpha \cdot |C| \geq 2k$, we get $2k \leq \frac{2n\alpha^2}{(k-1)^2}$, which yields: $\text{diam}(G[S]) = O((N\alpha^2)^{1/3})$.

(*Edge cost upper bound.*) The minimal length of a cycle is $\alpha + 1$, since otherwise an agent owning such a cycle edge could improve her costs by removing it. By this, we can apply Lemma 3.22 and get an upper bound on the edges of

$O(n^{1+2/\alpha})$.

(*Price of anarchy.*) For an optimal solution, we know that every agent is connected to at least one other agent, which gives a simple social cost lower bound of $\alpha n/2$. Comparing this to the above upper bounds gives for the process price of anarchy:

$$O\left(n^{2/\alpha} + \left(\frac{N}{\alpha}\right)^{1/3}\right)$$

□

3.5 Conclusion & Future Work

This chapter provided two different approaches to study the impact of friendships on equilibria in network creation games. First, driven by the commonly used worst-case approach, we saw that in Swap-Games the social cost can become worst possible. The only exception is the quality of tree equilibria in the Max-Swap-Game case, which states a remarkably different behavior. In particular, the Binding-Sequence gives a very interesting insight into the structure of worst-case equilibria. Secondly, looking from a much more optimistic view angle, we showed that using only some simple structural insights of the best-response processes suffices to drastically improve the results. Specifically, in the process price of anarchy, we tie the upper bound to the structure of the friendship graph and α .

Throughout this chapter, we only considered static friendship graphs: The set of friends never changes. Yet, in practice, friends of network participants might change over time. Introducing a time model and considering (possibly restricted) changes of the friendship graph seems to be a natural way to generalize our model, yielding an interesting online problem. In particular, the combination of dynamic friendship graphs and more problem tailored price of anarchy concepts, like the stated process price of anarchy, seems to be an interesting further direction.

The Impact of Choosing Edge Qualities

NETWORK creation games try to capture the behavior of Internet-like networks, which are created by the autonomous decisions of multiple strategic agents. Specifically, the game by Fabrikant et al. [Fab+03] was introduced to study the outcome of such interactions with respect to the impact of the agents' selfish behavior to the overall quality. For this challenging task, their classic model stays very simple and only provides one parameter, namely the edge price α , which has major influence on the outcome. In this chapter, we extend their model by enabling agents to select edges of different qualities for different prices.

When considering today's networks, where connections are offered by several service providers with different bandwidths and latency guarantees, choosing both the target and the quality of a connection seems to be a very natural extension. We are specifically interested in latency costs, which can be modeled as the shortest path lengths in a weighted network. Our model extension introduces a set of available edge lengths, from which the agents can choose when creating or changing an edge, and a price function, which assigns an individual price to every available edge length. For this generalized model, we show that equilibrium networks exist for any combination of available edge lengths and price functions. Considering the quality loss by the selfish behavior of the

agents, we analyze the price of stability and the price of anarchy.

Chapter Basis. The model, analysis, and results presented in the remainder of this chapter are based on the following publication:

2014 (with A. Mäcker and F. Meyer auf der Heide). “Quality of Service in Network Creation Games”. In: *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14-17, 2014. Proceedings*, cf. [CMM14].

Chapter Outline. This chapter is organized as follows. In Section 4.1, we introduce our model extensions for the Sum-Game and the Max-Game variants of the classic model by Fabrikant et al. [Fab+03], in particular the notion of edge lengths and price functions, and discuss several important properties of price functions that are needed for the later analysis. A comparison to other price functions, as typically used in economics literature, is provided in Section 4.2. For the introduced game variants, in Section 4.3 we first analyze the existence and structure of equilibrium networks. Supplementing this, in Section 4.4 and Section 4.5 we provide answers on the minimal and maximal quality loss by selfish behavior of agents.

4.1 Model & Notations

The considered model variants are extensions of the Sum-Game and the Max-Game models as introduced in Section 2.2. In each game, there is a set of n selfish agents V and a set $L \subseteq [\check{\beta}, \hat{\beta}]$ of available edge lengths with $0 < \check{\beta} \leq \hat{\beta}$. For convenience, throughout this chapter we assume $\check{\beta} = \min\{x \in L\}$ and $\hat{\beta} = \max\{x \in L\}$, which also gives that a specific minimum and maximum edge length always exists in L . Every agent $u \in V$ can create edges to other agents of any available edge length $x \in L$. The individual price of an edge of a length x is given by a monotonously decreasing function $p : L \rightarrow \mathbb{R}_{\geq 0}$, which is called a *price function*.

Every agent $u \in V$ aims to minimize her private cost by selfishly selecting a *strategy* $s_u \subset V \times L$. Hereby, each $(v, x) \in s_u$ represents an undirected weighted edge $(\{u, v\}, x)$ from u to v of length x , which is created by u and has a price

of $p(x)$. For a strategy profile $S = (s_{v_1}, \dots, s_{v_n})$ of agents $V = \{v_1, \dots, v_n\}$, the resulting weighted graph $G[S]$ consists of the vertices V and the weighted edges $\bigcup_{u \in V} \{(u, v), x \mid (v, x) \in s_u\}$.

Game Variants. We consider the two natural network creation game variants as discussed in Section 2.1. On the one hand, these are games in which agents want to minimize the sum of distances to all other agents, and on other hand these are games with agents who aim for minimizing their maximal distances. The private cost of an agent u in the *Sum-Pricing-Game* with strategy profile S is given by:

$$c_u(S) = \sum_{(v,x) \in s_u} p(x) + \sum_{v \in V} d_{G[S]}(u, v)$$

Here, $d_{G[S]}(u, v)$ denotes the shortest weighted path distance from u to v in the weighted graph $G[S]$. For the *Max-Pricing-Game*, the private cost function is:

$$c_u(S) = \sum_{(v,x) \in s_u} p(x) + \max_{v \in V} d_{G[S]}(u, v)$$

The social cost in both games is estimated as:

$$\text{cost}(S) = \sum_{u \in V} c_u(S)$$

We refer to the first term of a cost function as edge $c_u(S) = \sum_{(v,x) \in s_u} p(x)$, called the *edge cost*, and to the second term as $\text{dist}_u(S)$, called the *distance cost*.

Price Functions. In this chapter, for a game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ a monotonically decreasing function

$$p : L \rightarrow \mathbb{R}_{\geq 0} \tag{4.1}$$

is called a *price function*. Considering only monotonously decreasing functions means that we only consider price functions for which shorter (better) edges are more expensive than longer (inferior) ones. As noted previously, we assume $\check{\beta} = \min\{x \in L\}$ and $\hat{\beta} = \max\{x \in L\}$ and by this know that L contains a specific minimum and maximum value.

Most of the analysis in this chapter makes use of some characteristic values

4 The Impact of Choosing Edge Qualities

of a price function. Given a domain of available edge lengths L and a price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, we consider the edge lengths that minimize the following functions (cf. Figure 4.1):

- (a) $x \mapsto p(x) + x$,
- (b) $x \mapsto p(x) + (n - 1)x$, and
- (c) $x \mapsto (n - 1)p(x) + x$.

The minimizing values can be understood in the following way: If we consider the Sum-Pricing-Game, where agents aim to minimize the sum of distances to all other agents, $x \mapsto p(x) + x$ is the trade-off function between an edge length and its price for an edge that is used only for one shortest path and complementary, $x \mapsto p(x) + (n - 1)x$ illustrates the trade-off between an edge length and its price if the edge is used for $n - 1$ shortest paths. Different for the Max-Pricing-Game, $x \mapsto (n - 1)p(x) + x$ illustrates the trade-off between an edge length and its price for an edge that is used only for one shortest path, while $x \mapsto p(x) + x$ now illustrates the trade-off between an edge length and its price, if the edge is used for $n - 1$ shortest paths. The following lemma gives an overview of these functions and their relations, as they are needed in the later analysis.

Lemma 4.1. *Let $L \subseteq [\check{\beta}, \hat{\beta}]$ be a set of edge lengths and $p : L \rightarrow \mathbb{R}_{\geq 0}$ a price function. Then for the values*

- $x^* := \arg \min_{x \in L} p(x) + x$,
- $\bar{x} := \arg \min_{x \in L} p(x) + (n - 1)x$,
- $\tilde{x} := \arg \min_{x \in L} (n - 1)p(x) + x$,
- $\chi^* := \arg \min_{x \in L} \frac{p(x)}{2} + x$, and
- $\tilde{\chi} := \arg \min_{x \in L} p(x) + 2(n - 1)x$,

it holds:

- (a) $\bar{x} \leq x^* \leq \tilde{x}$ and $p(\tilde{x}) \leq p(x^*) \leq p(\bar{x})$,
- (b) $p(x^*) + x^* \leq (n - 1)p(\tilde{x}) + \tilde{x}$ and $p(x^*) + x^* \leq p(\bar{x}) + (n - 1)\bar{x}$, and

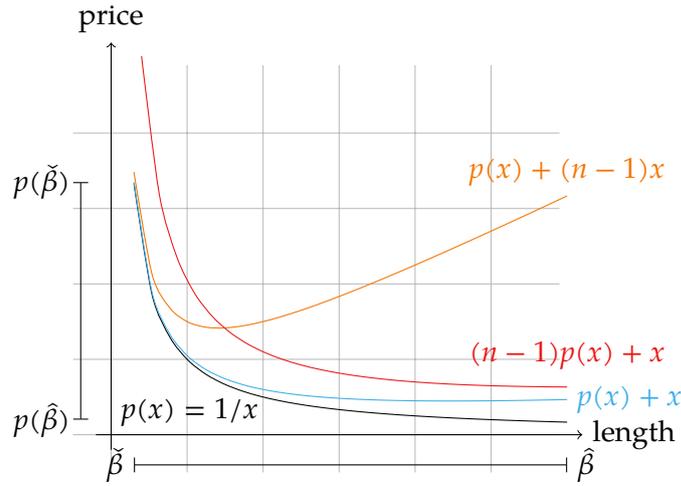


Figure 4.1: For some example price function $p(x)$, the figure illustrates the differences between the functions $p(x) + x$, $p(x) + (n-1)x$, and $(n-1)p(x) + x$. The minimal values of these functions in the domain of available edge lengths are characteristic for the later discussed prices of stability and anarchy.

$$(c) \chi^* + \frac{p(\chi^*)}{2} \geq \frac{x^* + p(x^*)}{2} \text{ and } p(\bar{\chi}) + 2(n-1)\bar{\chi} \geq p(\bar{x}) + (n-1)\bar{x}.$$

Proof. The first set of inequalities directly follows from the fact that p is a monotonically decreasing positive function. For the second set of inequalities, we only need that x^* minimizes the term $p(x) + x$, which cannot have a higher value than the compared values. And finally, for the last set of inequalities, using the definition of the values gives that the respective terms at the right hand side are minimized by the used parameters. \square

Solution Concepts. Using the terms of a buy equilibrium from Section 2.2, we call a strategy profile $S = (s_{v_1}, \dots, s_{v_n})$ a buy equilibrium, if for every agent v_i and every strategy $s'_{v_i} \neq s_{v_i}$ it holds that the strategy profile $S' := (s_{v_1}, \dots, s_{v_{i-1}}, s'_{v_i}, s_{v_{i+1}}, \dots, s_{v_n})$ does not have a lower private cost for v_i . If a strategy profile is not a buy equilibrium, then there exists at least one agent who can perform a strategy change that decreases her private cost. Such a strategy change is called an *improving response*. If the strategy change is the best possible for the agent in terms of reducing her private cost, it is called a *best response*.

4.2 Related Work & Contribution

As discussed in Section 2.2, both the Sum-Game and the Max-Game were studied extensively by various authors, in particular with respect to the question of the price of anarchy (cf. Section 2.2.2 and Section 2.3). This includes the study of different solution concepts, which usually restrict the available strategies or strategy changes of the agents. Yet, there is not much work on extending the capabilities of the agents, specifically not for the question of different edge qualities.

Still, when we consider network formation problems in general, edges of different qualities are quite common. For example, in congestion games (introduced by Rosenthal [Ros73]) for every edge there is a function that denotes its quality (latency) depending on the number of agents using the edge. Also in network design games (e.g., Augustine et al. [Aug+15]), edges usually have prices (weights) that are shared evenly among the agents that use them.

When it comes to modeling how a price function assigns a price to a good of a certain quality, there is a great deal of microeconomics literature (e.g., Jehle and Reny [JR11, pp. 135–145] and Mas-Colell et al. [MWG95, pp. 144–147]). As a reference, we refer to convex and linear price functions (cf. Mas-Colell et al. [MWG95, p. 144]), when benchmarking our results. In the related problem of provider competition in an Infrastructure-as-a-Service market, where providers offer access to computing resources and the resource prices change with the current load, Künsemöller et al. [Kün+14] considered a similar set of price functions by using piecewise linear functions.

Contribution. For every set of available edge lengths and every price function, we show that in the Sum-Pricing-Game and the Max-Pricing-Game buy equilibrium networks exist. Specifically, our constructions yield a constant price of stability for both games.

In the Sum-Pricing-Game, we can show that the price of anarchy is upper bounded by at most $O(\min\{n, (p(x^*) + x^*)/\check{\beta}\})$, with $x^* \in L$ being the edge length that minimizes $p(x) + x$. This emphasizes the importance of the trade-off between edge price and quality. In particular, we can show that the price of anarchy bound is nearly tight for a class of linear price functions, given by $p : [1, \alpha - 2\varepsilon] \rightarrow \mathbb{R}_{\geq 0}$ with $p(x) = \alpha - (1 + \varepsilon)x$, for $\alpha > 0$ and $\varepsilon \in (0, 1/2)$. This

is in considerable contrast to the classic Sum-Game for which no non-constant lower bound is known.

For the Max-Pricing-Game we provide a price of anarchy upper bound of $O(\sqrt[3]{n})$. Here, we note that unlike in the Sum-Pricing-Game, introducing price functions has no major effect to the game.

Note that in both games, by setting the available edge lengths to $L := \{1\}$ and the price function to $p(1) := \alpha$, we obtain the original Max- and Sum-Games.

4.3 Existence of Equilibria

Compared to the classic network creation games by Fabrikant et al. [Fab+03] and Demaine et al. [Dem+07], being able to select edge lengths and hence edge prices equips agents with much more freedom than before. For example, any interval $L \subseteq \mathbb{R}_{\geq 0}$ of positive length gives an infinite number of available edge lengths and hence an infinite number of possible strategy choices. Since a larger strategy space can make equilibria from the classical games unstable, in this section we start by asking whether equilibria always exist.

For this, given an arbitrary price function we make use of the optimal trade-offs between edge length and edge price for edges that are used only for one shortest path and edges that are used for $n - 1$ shortest paths. Using these edge lengths, we can construct equilibrium networks that look similar to those for the Max-Game and the Sum-Game, i.e., being either star or clique networks. In particular, the structure of the equilibrium networks depends on the characteristic price function values as introduced in Lemma 4.1.

4.3.1 Equilibria in the Sum-Pricing-Game

In the following, for any combination of a given set of edge lengths and a price function, we first compute the optimal solutions regarding the social cost and secondly show that always a buy equilibrium network exists. These results will be used in later sections to estimate bounds for the prices of stability and anarchy.

Lemma 4.2. *For the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, let S be a strategy profile such that $G[S]$ is connected and no edge can be removed without increasing the social cost. Denote by \tilde{x} the minimal*

4 The Impact of Choosing Edge Qualities

length of any edge in $G[S]$ and by m the total number of all edges. Then, for $x^* := \arg \min_{x \in L} p(x) + x$ it holds:

$$\text{cost}(S) \geq 2\check{x}n(n-1) + m(p(x^*) + x^* - 4\check{x})$$

Proof. Let $E_x := \bigcup_{u \in V} \{\{u, v\} \mid (v, x) \in s_u\}$ denote the edges in $G[S]$ of length x . Using that not directly connected agents have a (weighted) distance of at least $2\check{x}$, we can make the following estimation:

$$\begin{aligned} \text{cost}(S) &\geq \sum_{x \in L: |E_x| \neq 0} |E_x| \cdot (p(x) + 2x) + 2\check{x}(n(n-1) - 2m) \\ &\geq 2\check{x}n(n-1) + \sum_{x \in L: |E_x| \neq 0} |E_x| \cdot (p(x) + x) - 4m\check{x} \\ &\geq 2\check{x}n(n-1) + m(p(x^*) + x^* - 4\check{x}) \end{aligned}$$

□

Lemma 4.3. *For the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, define $\chi^* := \arg \min_{x \in L} p(x) + 2x$ and $\bar{\chi} := \arg \min_{x \in L} p(x) + 2(n-1)x$. Then, the network with optimal social cost is either given by a star network with all edges having length $\bar{\chi}$ or by a clique network with all edges having length χ^* .*

Proof. First, we argue that it is enough to only consider networks with a hop-diameter of at most 2 by converting any network G into a network G' with a not higher social cost but a hop-diameter of either one or two. For this, let $x_1 \leq \dots \leq x_m$ be the lengths of all edges of G (same lengths are listed multiple times), ordered by increasing length. We create G' with one center agent c , $n-1$ satellite agents v_1, \dots, v_{n-1} , and an initial edge set E_c by connecting each v_i to c by an edge of length x_i .

In the following, we create a one-to-one mapping from agent pairs in G to agent pairs in G' such that no distance compared to G is increased. First, for every directly connected agent pair $\{u, v\}$ such that their edge is associated with an edge in E_c , we map u and v to the end points of this edge. Secondly, for every agent pair $\{u, v\}$ that is not directly connected, but where both the first and the last edge of a shortest path in G are associated with edges in E_c , we map the agents to the satellite agents adjacent to the corresponding edges in G' . Thirdly, for every agent pair $\{u, v\}$ that is not directly connected

in G , but where the first and not the last edge of a shortest path from u to v in G is associated with an edge in E_c , we map to the satellite agent adjacent to the corresponding edge and to an arbitrary other agent to which not more than $n - 2$ agents are already mapped (including the mappings done by edge associations during the current step). Fourthly, for all remaining not directly connected agent pairs in G , we map them to arbitrary pairs of agents to which no mapping is performed yet. Finally, for all directly connected edges in G , which are not mapped yet, we map them to an arbitrary pair of agents. Here, if the distance of the mapped pair in G' is at most the distance as in G , then do nothing. Otherwise, create an edge between them of same length as in G .

By construction, this mapping is bijective as all agent pairs are mapped and to every agent pair in G' only one pair is mapped. In particular, no distance of a mapped pair of agents is bigger than the corresponding distance in G . This is obvious for the first and last step. For the other steps, this holds by the fact that x_1, \dots, x_{n-1} are the minimal lengths of all edges in G . Finally, the edge cost in G' is at most the edge cost in G and hence $\text{cost}(G') \leq \text{cost}(G)$.

We claim that the so constructed network is either a star or a complete network. By construction, a shortest distance path between any two not directly connected agents $u, v \in V \setminus \{c\}$ must contain c . Hence, any edge connecting two satellite agents $u, v \in V \setminus \{c\}$ is used exclusively for the shortest paths u to v and v to u and thus has length χ^* . Since the social cost is optimal, all satellite agents must have the same degree and be connected to c by edges having the same length x . For m being the total number of edges that connect any two satellites, the social cost is:

$$m(p(\chi^*) + 2\chi^* - 4x) + (n - 1)(p(x) + 2(n - 1)x) \quad (4.2)$$

We see that for any fixed x this term is minimized either with $m = 0$ (lower bound for m) or $m = (n - 1)n/2 - (n - 1)$ (upper bound for m). Hence, the optimal solution is either a star or a complete network. For a complete network, all lengths are χ^* and the social cost is $n(n - 1)(\chi^* + p(\chi^*)/2)$. Otherwise, for a star network the edge length $\bar{\chi}$ minimizes the social cost given by:

$$2(n - 1)x + (n - 1)(n - 2)2x + (n - 1)p(x) = (n - 1)(2(n - 1)x + p(x))$$

4 The Impact of Choosing Edge Qualities

Hence, the star and clique networks are the only two solutions with optimal social cost. \square

Next, we show for the Sum-Pricing-Game that for any set of available edge lengths L and any price function $p : L \rightarrow \mathbb{R}_{\geq 0}$ a buy equilibrium network exists. Depending on x^* and \bar{x} for the considered combination of price function and edge lengths (cf. definition in Theorem 4.4), this will either be a star or a clique network.

Theorem 4.4 (Sum-Pricing-Game: existence of equilibria). *For the Sum-Pricing-Game with a set of edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and arbitrary price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, a buy equilibrium strategy profile exists. Depending on the values $x^* := \arg \min_{x \in L} p(x) + x$ and $\bar{x} := \arg \min_{x \in L} p(x) + (n - 1)x$, one of the following networks is a buy equilibrium:*

$$\begin{cases} \text{clique, all lengths } x^* & \text{for } p(x^*) \leq \bar{x}, \\ \text{star, all lengths } \bar{x} & \text{otherwise.} \end{cases}$$

Proof. The following proof is about distinguishing when it is cheaper for a single agent to unilaterally replace an expensive short edge by many cheap long ones, or the other way around. In particular, we consider the cases when $p(x^*) > \bar{x}$ holds and then, for the contrary case of $p(x^*) \leq \bar{x}$, whether $p(\bar{x}) < x^*$ holds or not.

(*Stability of star, first case.*) For $\bar{x} < p(x^*)$, we consider a spanning star network consisting of a center agent u and $n - 1$ satellite agents v_1, \dots, v_{n-1} . Every satellite agent v_i owns one edge towards u of length \bar{x} . Considering the center agent, u will not create any edge since she is directly connected to all other agents and by Lemma 4.1 we know $\bar{x} \leq x^*$: i.e., every edge length is already shorter than the length agent u would choose when connecting to exactly one agent. But also no satellite agent v_i can perform an improving response, since on the one hand the length of v_i 's only edge is optimal for being the only connection to $n - 1$ agents. And on the other hand, for the optimal cost $p(x^*)$ to improve the distance to exactly one other satellite agent, the gain is $2\bar{x} - x^* - p(x^*) \leq 0$. Hence, the star forms a buy equilibrium network.

(*Stability of clique.*) For $p(x^*) \leq \bar{x}$, we check when a clique with all edges having length x^* and arbitrarily assigned edge ownerships forms a buy equilibrium

network. First, we see that decreasing the length of any edge $\{u, v\}$ exclusively decreases the distance between u and v . Since the edge length x^* is the optimal length for an edge used for exactly one shortest path, no agent will change any edge length as well as create a new edge. Hence, we only have to consider the unilateral strategy change of an agent of removing some edges and creating one new edge of a length $x \leq x^*$. Since all strategy changes are unilateral, after this change the hop-diameter is two. This means that other agents are either directly connected or at distance $x + x^*$. We consider the best possible strategy change for an agent u , consisting of removing (and by this gaining the edge cost reduction of) at most $n - 1$ edges of length x^* and creating one new edge of length \bar{x} to an arbitrary agent. This changes the cost of u by:

$$\begin{aligned} -(n-1)p(x^*) + p(\bar{x}) + (n-2)\bar{x} - (x^* - \bar{x}) &= \\ p(\bar{x}) - x^* + (n-1)(\bar{x} - p(x^*)) &\geq \\ p(\bar{x}) - x^* + (n-1)(p(x^*) - p(x^*)) &= p(\bar{x}) - x^* \end{aligned}$$

In particular, for $p(\bar{x}) \geq x^*$ the private cost of u increases and hence for this case the clique network forms a buy equilibrium.

(*Stability of star, second case.*) It remains to consider the last case, which is $p(x^*) \leq \bar{x}$, $p(\bar{x}) < x^*$, and $p(\bar{x}) - x^* + (n-1)(\bar{x} - p(x^*)) < 0$. Again, like in the first case, we claim that a star network with all edges of length \bar{x} and owned by the satellite agents is a buy equilibrium. By choosing the edge lengths to be \bar{x} , we only have to show that creating a new edge gives no gain. The optimal cost change by creating a new edge is $p(x^*) - 2\bar{x} + x^*$. This operation is an improving response if and only if $x^* < 2\bar{x} - p(x^*)$. But combining all constraints gives:

$$\begin{aligned} 0 &> p(\bar{x}) - x^* + (n-1)(\bar{x} - p(x^*)) \\ &> p(\bar{x}) - (2\bar{x} - p(x^*)) + (n-1)(\bar{x} - p(x^*)) \\ &\geq p(\bar{x}) - 2\bar{x} + p(x^*) + (n-1)\bar{x} - (n-1)p(x^*) \\ &\geq (n-3)\bar{x} - (n-3)p(x^*) \\ &= (n-3)(\bar{x} - p(x^*)) \geq 0 \end{aligned}$$

This is a contradiction and hence the star network is a buy equilibrium, since

no improving response exists. \square

Comparing this existence proof to the equilibrium analysis by Fabrikant et al. [Fab+03] for the classic Sum-Game, where only edges of length 1 for price α are available, one can see an interesting relation. If we set $L := \{1\}$ and $p(1) := \alpha$ in our proof for the Sum-Pricing-Game, the case distinction simplifies to the first two cases: i.e., whether $\alpha > 1$ or not. Specifically, this is the same case distinction as in the original game and gives the same existence result.

4.3.2 Equilibria in the Max-Pricing-Game

In the following, we show that also in the Max-Pricing-Game, for every combination of available edge lengths and price function, there exists a buy equilibrium network.

Lemma 4.5. *Let $L \subseteq [\check{\beta}, \hat{\beta}]$ be a set of edge lengths, $p : L \rightarrow \mathbb{R}_{\geq 0}$ be a price function, $\chi^* := \arg \min_{x \in L} x + p(x)/2$, and S a strategy profile. Then, in the Max-Pricing-Game the social cost is at least:*

$$\text{cost}(S) \geq \left(\chi^* + \frac{p(\chi^*)}{2} \right) n$$

Proof. For every agent $u \in V$, we consider an arbitrary fixed, longest shortest path and denote the length of the first edge of this path of agent u by x_u . Since an edge has only two incident agents, it can only be considered twice by the above procedure, whereas one agent has to pay for the edge. Hence, by summing over all private costs, we get:

$$\text{cost}(S) = \sum_{u \in V} c_u(S) \geq \sum_{u \in V} \left(x_u + \frac{p(x_u)}{2} \right) \geq \left(\chi^* + \frac{p(\chi^*)}{2} \right) n$$

\square

Note that this lower bound is actually tight up to a constant. This can be seen by considering a star network with all edges having length of χ^* and thus giving a social cost of $\Theta(\chi^* n + p(\chi^*) n)$.

Theorem 4.6 (Max-Pricing-Game: existence of equilibria). *For the Max-Pricing-Game with a set of edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and arbitrary price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, a*

buy equilibrium strategy profile exists. Depending on $x^* := \arg \min_{x \in L} p(x) + x$ and $\tilde{x} := \arg \min_{x \in L} (n-1)p(x) + x$, one of the following networks is a buy equilibrium:

$$\begin{cases} \text{star, all lengths } x^* & \text{for } (n-1)p(\tilde{x}) + \tilde{x} \geq p(x^*) + 2x^*, \\ \text{star, all lengths } \tilde{x} & \text{for } (n-1)p(\tilde{x}) + \tilde{x} < p(x^*) + 2x^* \wedge \tilde{x} \leq (n-2)p(\tilde{x}), \\ \text{clique, all lengths } \tilde{x} & \text{otherwise.} \end{cases}$$

Proof. In the following, we use the relations from Lemma 4.1 and perform a case distinction on the properties of L and p .

(*Stability of star network with satellites owning all edges.*) If it holds $(n-1)p(\tilde{x}) + \tilde{x} \geq p(x^*) + 2x^*$, we consider a star network with all edges having length x^* and being owned by the satellite agents. Then, the center agent can only improve her private cost by creating $n-1$ edges of some length $x < x^*$, leading to a gain of:

$$x^* - ((n-1)p(x) + x) \leq x^* - ((n-1)p(\tilde{x}) + \tilde{x}) \leq x^* - (p(x^*) + 2x^*) < 0$$

For any satellite agent v , there are two kinds of possible improving responses. First, if v only creates edges of some length $x \leq x^*$ to satellite agents, then the gain is at most:

$$p(x^*) + 2x^* - ((n-2)p(x) + x^* + x) \leq 0$$

Secondly, for a strategy change that connects v to every other agent, the gain is at most:

$$p(x^*) + 2x^* - ((n-1)p(\tilde{x}) + \tilde{x}) \leq 0$$

Hence, no improving response exists and the star network is a buy equilibrium. In the following, we consider the remaining case $(n-1)p(\tilde{x}) + \tilde{x} < p(x^*) + 2x^*$, for which we make a further case distinction.

(*Stability of star with center owning all edges.*) Now, let $(n-1)p(\tilde{x}) + \tilde{x} < p(x^*) + 2x^*$ and further consider the case of $\tilde{x} \leq (n-2)p(\tilde{x})$. We want to show that then a star network consisting only of edges of length \tilde{x} , whereas all edges are owned by a center agent u , is a buy equilibrium. By construction, the center agent cannot perform any improving response and it suffices to consider the possible improving responses of the satellite agents. For this, let v be a satellite agent.

4 The Impact of Choosing Edge Qualities

First, if v creates edges of length x to all other $n - 2$ satellites (note that $x \leq 2\tilde{x}$), the gain is:

$$2\tilde{x} - (\max\{\tilde{x}, x\} + (n - 2)p(x))$$

If $x < \tilde{x}$, this value is negative. But still for $x \geq \tilde{x}$, the gain is at most

$$2\tilde{x} - (x + (n - 2)p(x)) \leq \tilde{x} + (n - 2)p(\tilde{x}) - (x + (n - 2)p(x)) \leq 0$$

since the value minimizing $x + (n - 2)p(x)$ lies in the interval $[x^*, \tilde{x}]$ and thus the only possibly improving choice for v is $x = \tilde{x}$. Secondly, if v creates edges of length x to all other agents, the gain is:

$$2\tilde{x} - ((n - 1)p(x) + x) \leq (n - 2)p(\tilde{x}) + \tilde{x} - (n - 1)p(\tilde{x}) - \tilde{x} \leq 0$$

Thirdly, if v creates only one edge to the center agent of length x , the gain is:

$$\begin{aligned} 2\tilde{x} - (p(x) + x + \tilde{x}) &\leq \tilde{x} - p(x^*) - x^* \\ &\leq \tilde{x} - ((n - 1)p(\tilde{x}) + \tilde{x} - x^*) \\ &\leq (n - 2)p(\tilde{x}) - (n - 1)p(\tilde{x}) - \tilde{x} + x^* < 0 \end{aligned}$$

Hence, this star network is a buy equilibrium.

(*Stability of clique with one agent owning $n - 1$ edges.*) As the final case, we have to consider $(n - 1)p(\tilde{x}) + \tilde{x} < p(x^*) + 2x^*$ and $\tilde{x} > (n - 2)p(\tilde{x})$. Here, we construct a star with one agent u owning $n - 1$ edges of length \tilde{x} and complete this star to a clique with all edges having length \tilde{x} , yet arbitrary edge ownerships. We claim that this network is a buy equilibrium. At first we note that by construction, u has optimal lengths for all of her edges. Also every other agent has optimal lengths for her edges, since by unilaterally changing her edge lengths the diameter stays at least \tilde{x} . We further show that no agent will change her edge set by considering the following kinds of possible improving responses. First, for any agent replacing her current set of edges by edges to all other agents, the optimal length is \tilde{x} and hence, doing so cannot improve her private cost. Secondly, by simply removing all own edges the gain is at most $(n - 2)p(\tilde{x}) - \tilde{x} < 0$. Thirdly, by removing all own edges and creating

4.4 Quality of Equilibria in the Sum-Pricing-Game

one edge to u of length x , the gain is at most:

$$(n-1)p(\tilde{x}) + \tilde{x} - (\tilde{x} + x) - p(x) \leq (n-1)p(\tilde{x}) - x - p(x) \leq \tilde{x} + p(\tilde{x}) - (x^* + p(x^*)) \leq 0$$

Concluding, no improving response exists for any agent and hence the network is a buy equilibrium. \square

4.4 Quality of Equilibria in the Sum-Pricing-Game

In this section, we consider the quality of equilibria in the Sum-Pricing-Game. In particular, we provide bounds for the prices of stability and anarchy.

Corollary 4.7 (Sum-Pricing-Game: price of stability). *For the Sum-Pricing-Game with edge lengths $L \subseteq [\tilde{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, the price of stability is at most 4.*

Proof. In the following, we compute the social cost ratio when comparing the equilibrium networks from Theorem 4.4 with the optimal solutions from Lemma 4.3. For this, define the characteristic values $x^* := \arg \min_{x \in L} p(x) + x$, $\bar{x} := \arg \min_{x \in L} p(x) + (n-1)x$, $\chi^* := \arg \min_{x \in L} p(x) + 2x$, and $\bar{\chi} := \arg \min_{x \in L} p(x) + 2(n-1)x$. If the equilibrium network and the socially optimal network have the same topology, i.e., both being star networks or both being clique networks, the price of stability is at most 2. This directly follows from the relations of Lemma 4.1, when comparing the social costs.

Now consider the case when the equilibrium network is a star with all edges having length \bar{x} , but the optimal solution being a clique. In this case we get:

$$\begin{aligned} \text{PoS} &\leq \frac{(n-1)(2(n-1)\bar{x} + p(\bar{x}))}{n(n-1)\left(\chi^* + \frac{p(\chi^*)}{2}\right)} \leq 4 \frac{(n-1)\bar{x} + p(\bar{x})}{n(x^* + p(x^*))} \\ &\leq 4 \frac{(n-1)x^* + p(x^*)}{n(x^* + p(x^*))} \leq 4 \end{aligned}$$

For the second-last estimation note that by definition \bar{x} is the argument in L for which the function $x \mapsto (n-1)x + p(x)$ is minimized.

Finally, consider the case when the buy equilibrium network is a clique with all edges of length x^* and the optimal solution being a star with all edges of length $\bar{\chi}$. Considering equation (4.2) from Lemma 4.3, for the optimal solution

4 The Impact of Choosing Edge Qualities

to be a star it must hold that $p(\chi^*) + 2\chi^* - 4\bar{\chi} \leq 0$ and hence $p(\chi^*) + 2\chi^* \leq 4\bar{\chi}$. This gives:

$$\begin{aligned} \text{PoS} &\leq \frac{n(n-1)\left(x^* + \frac{p(x^*)}{2}\right)}{(n-1)(2(n-1)\bar{\chi} + p(\bar{\chi}))} \leq \frac{n\left(x^* + \frac{p(x^*)}{2}\right)}{2(n-1)\bar{\chi} + p(\bar{\chi})} \\ &\leq \frac{n(\chi^* + p(\chi^*))}{2(n-1)\bar{\chi} + p(\bar{\chi})} \leq \frac{4n\bar{\chi}}{2(n-1)\bar{\chi}} \leq 4 \end{aligned}$$

□

Similar to Albers et al. [Alb+14], we start our analysis of the price of anarchy by bounding the social cost of a buy equilibrium network by the diameter of the network but now incorporate arguments about maximum edge lengths and edge prices. This will yield the upper bound for the price of anarchy as stated in Theorem 4.10.

Lemma 4.8. *For the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, let S be a buy equilibrium strategy profile and define $x^* := \arg \min_{x \in L} p(x) + x$. Then, for any agent $u \in V$ it holds:*

$$\text{cost}(S) \leq n \cdot \text{dist}_u(S) + x^*(n-1)^2 + 2(p(x^*) + x^*)n(n-1)$$

Proof. First, we claim that in an equilibrium network all edges have a price of at most $n(p(x^*) + x^*)$. For this, assume there is an edge of price $p(x) > (p(x^*) + x^*)n$ and consider replacing it by a new edge of length x^* . This would decrease the owner's edge cost by $p(x) - p(x^*) > nx^* + (n-1)p(x^*)$, while increasing the distance cost by at most $(x^* - x)(n-1)$. Since $(x^* - x)(n-1) < nx^* + (n-1)p(x^*)$, this is an improving response and hence contradicts S forming a buy equilibrium.

Next, fix an arbitrary agent $u \in V$ and consider a shortest path tree T rooted at u in $G[S]$. For every $v \in V$, define $m_v := |\{\{v, w\} \mid (w, x) \in s_v \wedge \{v, w\} \in T\}|$ to be the number of tree edges maintained by v . Then, for any agent $v \neq u$ we argue that it must hold

$$c_v(S) \leq (p(x^*) + x^*)n(m_v + 1) + \text{dist}_u(S) + x^*(n-1),$$

which we can see as follows: Since S forms an equilibrium, deviating from her

4.4 Quality of Equilibria in the Sum-Pricing-Game

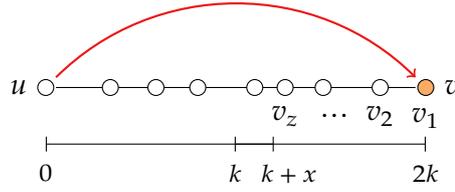


Figure 4.2: Illustration of the diameter argument of Lemma 4.9: Agent u creates an edge to agent v and improves her distance cost to agents v_1, \dots, v_L .

current strategy cannot decrease v 's private cost. In particular, the resulting private cost when removing all own edges, except those belonging to T , and additionally creating one new edge of length x^* to u cannot be less than $c_v(S)$. Since this strategy change would not modify any edges of the shortest path tree T , after the strategy change v 's distance cost would be at most $\text{dist}_u(S) + (n-1)x^*$, while her edge cost would be $(p(x^*) + x^*)n(m_v + 1) + p(x^*)$.

Using this bound for every agent $v \neq u$ and the fact that u only owns edges belonging to T (removing a non-tree edge would reduce u 's cost), we get:

$$\begin{aligned}
 \text{cost}(S) &\leq \text{dist}_u(S) + (p(x^*) + x^*)nm_u \\
 &\quad + \sum_{v \neq u} ((p(x^*) + x^*)n(m_v + 1) + \text{dist}_u(S) + x^*(n-1)) \\
 &= n \cdot \text{dist}_u(S) + x^*(n-1)^2 + (p(x^*) + x^*)nm_u \\
 &\quad + \sum_{v \neq u} (p(x^*) + x^*)n(m_v + 1) \\
 &= n \cdot \text{dist}_u(S) + x^*(n-1)^2 + 2(p(x^*) + x^*)n(n-1)
 \end{aligned}$$

For the last equality we use that a tree with n agents has $n-1$ edges. \square

Lemma 4.9. *In the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, define $x^* := \arg \min_{x \in L} p(x) + x$, and let S be a buy equilibrium strategy profile. Then, the diameter of $G[S]$ is at most $O(p(x^*) + x^*)$.*

Proof. First, we show that no edge can be longer than $p(x^*) + x^*$. Assuming there is an agent u who owns an edge $(v, x) \in s_u$ of length $x > p(x^*) + x^*$, which connects u to an agent v , we consider the replacement of this edge by an edge of length x^* . Such a new strategy $s'_u := (s_u \setminus \{(v, x)\}) \cup \{(v, x^*)\}$ decreases u 's distance cost by at least $x - x^* > p(x^*)$, but increases u 's edge cost by at

4 The Impact of Choosing Edge Qualities

most $p(x^*) - p(x)$. Since an improving response contradicts S being a buy equilibrium, we get the upper bound on the edge length.

Next, we consider the length of a longest shortest path in $G[S]$, of which we call the incident agents u and v . If this path only consists of one edge, the edge would have a length of at most $p(x^*) + x^*$ and the claim holds. Otherwise, the path consists of at least two edges and we define a parameter $k \in \mathbb{R}_{\geq 0}$ such that $2k = d_{G[S]}(u, v)$ and consider the strategy change $s'_u := s_u \cup \{(v, x)\}$ of agent u that consists of creating an edge $\{u, v\}$ of some length $x \in L$. This strategy change (cf. Figure 4.2) decreases u 's distance cost to agents on the path that have a distance of at least $k + x$ to u . Let $v =: v_1, v_2, \dots, v_Z$ denote these agents, ordered by increasing distance to v . Since each edge has a length of at most $\min\{p(x^*) + x^*, \hat{\beta}\}$, we get $Z \geq \left\lceil \frac{k-x}{\min\{p(x^*)+x^*, \hat{\beta}\}} \right\rceil$. With the strategy change s'_u , each distance from u to any v_i decreases from $2k - d_{G[S]}(v, v_i)$ to be at most $x + d_{G[S]}(v, v_i)$, resulting in a distance cost decrease of at least:

$$\begin{aligned} \sum_{i=1}^Z (2k - d_{G[S]}(v, v_i)) - \sum_{i=1}^Z (x + d_{G[S]}(v, v_i)) &= Z(2k - x) - 2 \sum_{i=1}^Z d_{G[S]}(v, v_i) \\ &\geq Z(2k - x) - 2Z(k - x) \\ &= Z(2k - x - 2k + 2x) = Zx \end{aligned}$$

Since S is a buy equilibrium, this cannot be an improving response and hence we get $Zx \leq p(x)$. This gives $p(x) \geq \frac{k-x}{\min\{p(x^*)+x^*, \hat{\beta}\}}x$ and hence:

$$k \leq \min\{p(x^*) + x^*, \hat{\beta}\} \frac{p(x)}{x} + x$$

If $\min\{p(x^*) + x^*, \hat{\beta}\} = \hat{\beta}$, then the diameter of $G[S]$ is at most $2(p(\hat{\beta}) + \hat{\beta}) \leq 2(p(\hat{\beta}) + p(x^*) + x^*) = O(p(x^*) + x^*)$. Otherwise, if $\min\{p(x^*) + x^*, \hat{\beta}\} = p(x^*) + x^*$, then the diameter is at most $(p(x^*) + x^*) \frac{p(x)}{x} + x$. For $p(x^*) \leq x^*$ the lemma follows by setting $x := x^*$. In case $p(x^*) > x^*$, by setting $x := p(x^*)$ the diameter is at most $O\left((p(x^*) + x^*) \frac{p(p(x^*))}{p(x^*)}\right)$. Using the monotonicity of p , it holds $p(p(x^*)) \leq p(x^*)$ and we get $O(p(x^*) + x^*)$. \square

Theorem 4.10 (Sum-Pricing-Game: price of anarchy upper bound). *In the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$, price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, and*

4.4 Quality of Equilibria in the Sum-Pricing-Game

$x^* := \arg \min_{x \in L} p(x) + x$, the price of anarchy is at most:

$$\text{PoA} = O\left(\min\left\{n, \frac{p(x^*) + x^*}{\check{\beta}}\right\}\right)$$

Proof. Let S be a buy equilibrium strategy profile. Then by Lemma 4.9, the diameter of $G[S]$ is at most $O(p(x^*) + x^*)$. Applying this to Lemma 4.8, the social cost of S is at most:

$$\text{cost}(S) = O(n(n-1)(p(x^*) + x^*))$$

Moreover, by Lemma 4.2 the social cost of an optimal solution is at least

$$2\check{\beta}n(n-1) + m(p(x^*) + x^* - 4\check{\beta}),$$

whereas m denotes the number of edges. When comparing both bounds, we get for $p(x^*) + x^* \leq 4\check{\beta}$ that the lower bound is minimized with $m = n(n-1)/2$ and then becomes $n(n-1)(p(x^*) + x^*)/2$, which gives a price of anarchy of $O(1)$. Otherwise, for $p(x^*) + x^* > 4\check{\beta}$ the lower bound is minimized with $m = n-1$ and we get $\text{PoA} = O\left(\frac{n(p(x^*) + x^*)}{\check{\beta}(2n-4+(p(x^*) + x^*)/\check{\beta})}\right)$. When separately considering whether $n < \frac{p(x^*) + x^*}{\check{\beta}}$ holds or not, we get the claimed price of anarchy upper bound. \square

Applying the price and length value ranges, we can deduce a price of anarchy upper bound, which is independent of the price function, but depends only on the range limits.

Corollary 4.11. *In the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$, for every price function $p : L \rightarrow \mathbb{R}_{\geq 0}$ it holds:*

$$\text{PoA} = O\left(\min\left\{1 + \frac{p(\check{\beta})}{\check{\beta}}, \frac{p(\hat{\beta}) + \hat{\beta}}{\check{\beta}}, n\right\}\right)$$

In the following, we will see that the price of anarchy upper bound is even tight for a broad class of price functions, including for all price functions that decrease faster than the linear function $x \mapsto -x$ and where both $p(\hat{\beta}) \leq \check{\beta}$ and $p(\check{\beta}) \leq \hat{\beta}$ hold. Examples of such functions are provided in the following section.

4 The Impact of Choosing Edge Qualities

However, note that the bound cannot be tight for every price function. To see this, consider $p : [1, 1] \rightarrow [\alpha, \alpha]$, which constitutes the original game by Fabrikant et al. [Fab+03] and for which it is known that for most ranges of α the price of anarchy is constant (cf. Section 2.2.2).

Theorem 4.12 (Sum-Pricing-Game: price of anarchy lower bound). *In the Sum-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$, let $p : L \rightarrow \mathbb{R}_{\geq 0}$ be a price function with $p(\hat{\beta}) \leq \check{\beta}$, $p(\check{\beta}) \leq \hat{\beta}$, and $\hat{\beta} = x^* := \arg \min_{x \in L} p(x) + x$, then:*

$$\text{PoA} = \Omega\left(\min\left\{n, \frac{p(x^*) + x^*}{\check{\beta}}\right\}\right)$$

Proof. Using the given constraints, we get $p(x^*) = p(\hat{\beta}) \leq \check{\beta} \leq x$ for every $x \in L$. In particular, this gives $p(x^*) \leq \bar{x}$, whereas $\bar{x} := \arg \min_{x \in L} p(x) + (n-1)x$, and hence by Lemma 4.4 we know that a clique network with all edges having length $\hat{\beta}$ is a buy equilibrium.

We compare the social cost of this clique network to the social cost of a star network with all edges having length $\check{\beta}$. This gives a price of anarchy lower bound of:

$$\text{PoA} \geq \frac{p(\hat{\beta})n(n-1)/2 + \hat{\beta}n(n-1)}{p(\check{\beta})(n-1) + 2\check{\beta}(n-1)(n-2) + 2(n-1)\check{\beta}} = \frac{n(p(\hat{\beta})/2 + \hat{\beta})}{p(\check{\beta}) + 2\check{\beta}(n-1)}$$

Next, we separately consider the cases of $n \geq (p(\hat{\beta}) + \hat{\beta})/\check{\beta}$ and $n < (p(\hat{\beta}) + \hat{\beta})/\check{\beta}$. For $n \geq (p(\hat{\beta}) + \hat{\beta})/\check{\beta}$, we get:

$$\frac{n(p(\hat{\beta})/2 + \hat{\beta})}{p(\check{\beta}) + 2\check{\beta}(n-1)} \geq \frac{n(p(\hat{\beta})/2 + \hat{\beta})}{\hat{\beta} + p(\hat{\beta}) + 2\check{\beta}(n-1)} \geq \frac{n(p(\hat{\beta})/2 + \hat{\beta})}{n\check{\beta} + 2\check{\beta}(n-1)} = \Omega\left(\frac{p(\hat{\beta}) + \hat{\beta}}{\check{\beta}}\right)$$

Otherwise, for $n < (p(\hat{\beta}) + \hat{\beta})/\check{\beta}$ we get:

$$\frac{n(p(\hat{\beta})/2 + \hat{\beta})}{p(\check{\beta}) + 2\check{\beta}(n-1)} \geq \frac{n(p(\hat{\beta})/2 + \hat{\beta})}{p(\check{\beta}) + 2\check{\beta}((p(\hat{\beta}) + \hat{\beta})/\check{\beta}) - 1)} \geq \frac{n(p(\hat{\beta})/2 + \hat{\beta})}{3\hat{\beta}} = \Omega(n)$$

Combining both bounds gives the claim. \square

4.4.1 Employing Characteristic Price Functions

Concluding the analysis of the Sum-Pricing-Game, we apply our price of anarchy results to some typical price functions (cf. Mas-Colell et al. [MWG95, pp. 143–147]). These are firstly the price function $x \mapsto \alpha/x$, whereas $\alpha > 0$, as an example for a convex function, and secondly a class of linear functions. Recalling Corollary 4.7, we know that the price of stability is constant for every price function.

The convex function $x \mapsto \alpha/x$ (cf. Figure 4.1) illustrates the scenario where edge prices increase very fast for good connections but do not vary much for the tail of slow connections. Note that we only provide a price of anarchy upper bound but no lower bound, since the lower bound from Theorem 4.12 does not apply. This is due to the fact that the only interval of edge lengths that fulfills all constraints of the theorem is $L = \{\sqrt{\alpha}\}$.

Corollary 4.13. *Given an interval $L := [1, \beta]$ of available edge lengths for some parameter $\beta > 1$, then for the price function $p : L \rightarrow \mathbb{R}_{\geq 0}, x \mapsto \alpha/x$ with $\alpha \in [1, \beta^2]$, the price of anarchy is at most $\text{PoA} = O(\sqrt{\alpha})$.*

Proof. For the upper bound, we consider Theorem 4.10 and have to compute $\min_{x \in L} p(x) + x$, which is given by $x^* := \sqrt{\alpha}$ and yields the claim. \square

Next, we consider a linear function and show that actually a very high price of anarchy lower bound is possible, in particular higher than anyone known for the classic Sum-Game without edge price functions. For our construction, we choose a set of linear functions that decrease just slowly enough such that a clique network is a buy equilibrium, while the optimal solution is a star.

Corollary 4.14. *Given an interval $L := [1, \alpha - (1 + \varepsilon/2)]$ of available edge lengths, for some value $\alpha > 2$ and a positive value $\varepsilon < \frac{1}{n-1}$, we consider the Sum-Pricing-Game of n agents. Then, for the price function $p : L \rightarrow \mathbb{R}_{\geq 0}, x \mapsto \alpha - (1 + \varepsilon)x$ the price of anarchy is $\text{PoA} = \Theta(\alpha(1 - \varepsilon))$.*

Proof. First we see that $x^* := \arg \min_{x \in L} p(x) + x$, which is $\arg \min_{x \in L} \alpha - \varepsilon x$ in this case, has the value $x^* := \alpha - (1 + \varepsilon/2)$. Using this, by applying Theorem 4.10 we get $\text{PoA} = O((1 - \varepsilon)\alpha)$ as upper bound for the price of anarchy. Considering the ranges of L , we see that now both constraints of Theorem 4.12 are fulfilled, i.e., $p(1) = \alpha - (1 + \varepsilon) \leq \alpha - (1 + \varepsilon/2)$ and $p(\alpha - (1 + \varepsilon/2)) = \alpha - (1 + \varepsilon)(\alpha - 1 - \varepsilon/2) < 1 - \varepsilon \leq 1$. Thus we get the corresponding lower bound. \square

4.5 Quality of Equilibria in the Max-Pricing-Game

In this section, we consider the quality of equilibria in the Max-Pricing-Game. In particular, these are the prices of stability and anarchy.

Corollary 4.15 (Max-Pricing-Game: price of stability). *In the Max-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, the price of stability is at most 4.*

Proof. We compare the social cost of the three equilibrium networks from Theorem 4.6 with the social cost lower bound from Lemma 4.5. For this, define the values $x^* := \arg \min_{x \in L} p(x) + x$ and $\tilde{x} := \arg \min_{x \in L} (n-1)p(x) + x$.

For $(n-1)p(\tilde{x}) + \tilde{x} \geq p(x^*) + 2x^*$, by Theorem 4.6 a star network with all edges having length of x^* is a buy equilibrium. Compared to the social cost lower bound, the social cost ratio is:

$$\text{PoS} \leq \frac{(n-1)p(x^*) + n2x^*}{(x^* + p(x^*)/2)n} \leq \frac{p(x^*)}{x^* + p(x^*)/2} + \frac{2x^*}{x^* + p(x^*)/2} \leq 4$$

For $(n-1)p(\tilde{x}) + \tilde{x} < p(x^*) + 2x^* \wedge \tilde{x} \leq (n-2)p(\tilde{x})$, by Theorem 4.6 a star network with all edges having length of \tilde{x} is a buy equilibrium. Compared to the social cost lower bound, we get the following social cost ratio by applying the first constraint:

$$\text{PoS} \leq \frac{(n-1)p(\tilde{x}) + n2\tilde{x}}{(x^* + p(x^*)/2)n} \leq 2 \frac{p(x^*) + 2x^*}{x^* + p(x^*)/2} \leq 4$$

For the remaining case, by Theorem 4.6 a clique network with all edges having length \tilde{x} is a buy equilibrium. Compared to the social cost lower bound, the social cost ratio is:

$$\text{PoS} \leq \frac{p(\tilde{x})(n-1)n/2 + n\tilde{x}}{(x^* + p(x^*)/2)n} \leq \frac{(n-1)p(\tilde{x}) + \tilde{x}}{(x^* + p(x^*)/2)} \leq 2 \frac{p(x^*) + 2x^*}{2x^* + p(x^*)} \leq 2$$

□

Lemma 4.16. *In the Max-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$, price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, and $x^* := \arg \min_{x \in L} p(x) + x$, let S be a buy equilibrium strategy profile. Then, for any agent $u \in V$ it holds:*

$$\text{cost}(S) \leq n \cdot \text{dist}_u(S) + 2(n-1)(p(x^*) + x^*)$$

4.5 Quality of Equilibria in the Max-Pricing-Game

Proof. First, we show that in a buy equilibrium network no edge costs more than $p(x^*) + x^*$. For this, consider a strategy profile S' such that some agent v owns an edge of price $p(x) > p(x^*) + x^*$, which has length $x < x^*$. If v replaces this edge by one of length x^* , her edge cost will decrease by $p(x) - p(x^*) > x^*$, while her distance cost increases by at most $x^* - x$. Since this contradicts S' being an equilibrium, we get the upper bound on the edge prices.

Next, take an arbitrary agent u and consider a shortest path tree T rooted at u in $G[S]$. For every $v \in V$, define $m_v := |\{\{v, w\} \mid (w, x) \in s_v \wedge \{v, w\} \in T\}|$ to be the number of tree edges maintained by v . Then, for any agent $v \neq u$ we argue that it must hold that

$$c_v(S) \leq (p(x^*) + x^*)(m_v + 1) + \text{dist}_u(S),$$

which can be seen as follows: Since S forms an equilibrium, deviating from her current strategy cannot decrease v 's cost. In particular, the resulting private cost when removing all own edges, except those belonging to T , and additionally creating one new edge of length x^* to u cannot be less than $c_v(S)$. Since this strategy change does not modify any edges of the shortest path tree T , with the changed strategy v 's distance cost would be at most $\text{dist}_u(S) + x^*$, while her edge cost would be at most $(p(x^*) + x^*)(m_v + 1) + p(x^*)$. Summing over all agents' costs and using the fact that agent u only owns edges belonging to T (otherwise removing a non-tree edge would improve u 's cost), we get:

$$\begin{aligned} \text{cost}(S) &\leq \text{dist}_u(S) + (p(x^*) + x^*)m_u \\ &\quad + \sum_{v \in V: v \neq u} ((p(x^*) + x^*)(m_v + 1) + \text{dist}_u(S)) \\ &= n \cdot \text{dist}_u(S) + (n - 1)(p(x^*) + x^*) + \sum_{v \in V} (p(x^*) + x^*)m_v \\ &= n \cdot \text{dist}_u(S) + 2(n - 1)(p(x^*) + x^*) \end{aligned}$$

For the last equality we use that the number of edges in a tree of n agents is $n - 1$. \square

Using a similar approach like Demaine et al. [Dem+07], we derive a bound for the diameter and hence for the social cost of every buy equilibrium network in the Max-Pricing-Game.

4 The Impact of Choosing Edge Qualities

Lemma 4.17. *In the Max-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, let S be a buy equilibrium strategy profile. Then, the diameter of $G[S]$ is at most $O\left(\sqrt[3]{p(x)^2xn + x}\right)$, for $x \in L$ arbitrary.*

Proof. Let $u \in V$ be an agent with maximal distance cost, i.e., with $\text{dist}_u(S) = \text{diam}(G[S])$. In the following, we first show for every agent a lower bound on the number of other agents within a specific range and then use this lower bound to derive an upper bound on the private cost of u .

For an edge length $x \in L$ and a parameter $k \geq 0$, we define $N_k(v)$ to be the set of agents within a distance of at most kx to agent $v \in V$. We claim that for any $k \leq \text{dist}_u(S)/(2x)$ it holds:

$$|N_k(v)| \geq \frac{x}{2p(x)}(k^2 - 3k + 2) \quad (4.3)$$

Let P_1, P_2, \dots, P_{n-1} be the shortest paths from v to all agents $V \setminus \{v\}$ and for every $k \in \mathbb{N}$ define Q_k to be the set containing the first agent of every path P_i that has a distance of at least kx and at most $(k+1)x$ to v . (Note that such agents do not necessarily exist and Q_k might even be empty.) Then, for an arbitrary fixed $k \leq \frac{\text{dist}_u(S)}{2x}$ consider the strategy change where v creates edges of length x to every agent in Q_k . This strategy change decreases v 's distances to all agents in Q_k by at least $(k-1)x$ each, while increasing the edge cost by $p(x) \cdot |Q_k|$. Since S is a buy equilibrium, it must hold that $|Q_k| \geq \frac{(k-1)x}{p(x)}$, which gives us the following lower bound:

$$|N_k(v)| \geq \sum_{i=1}^{k-1} |Q_i| \geq \sum_{i=1}^{k-1} (i-1) \frac{x}{p(x)} = \frac{x}{2p(x)}(k^2 - 3k + 2)$$

Let $k := \frac{\text{diam}(G[S]) - x}{4x}$ be a fixed parameter, then we compute a set of cluster centers C as follows: Starting with $C = \{u\}$, we mark u as well as all agents within a distance of less than $2kx$ to u . Iteratively, while there is still an unmarked agent v , we further mark v and all agents within a distance of less than $2kx$ and add v to C . By construction, the resulting set C meets the following two conditions:

- (a) For any agent $v \in V$ the maximal distance to a center is $d_{G[S]}(v, C) \leq 2kx$.
- (b) The distance between any two centers $c, c' \in C, c \neq c'$ is $d_{G[S]}(c, c') > 2kx$.

4.5 Quality of Equilibria in the Max-Pricing-Game

For the selected k , we observe that it holds $k \leq \frac{\text{dist}_c(S)}{2x}$ for all centers $c \in C$, since otherwise there would be a center $c' \in C$ having $\text{dist}_{c'}(S) < 2kx$, which gives $\text{dist}_{c'}(S) < \text{dist}_u(S) = \text{diam}(G[S])$ and hence contradicts the choice of u . By construction of C , we have:

$$n \geq \sum_{c \in C} |N_k(c)| \geq \frac{|C| \cdot x}{2p(x)} (k^2 - 3k + 2)$$

Considering a strategy change where u buys edges of length x to all $c \in C$, this would decrease u 's distance cost by at least $\text{dist}_u(S) - (2k + 1)x \geq 2kx$ while increasing her edge cost by $|C| \cdot p(x)$. Since S is a buy equilibrium, it must hold $|C| \geq 2k \frac{x}{p(x)}$ and we obtain:

$$n \geq |C| \frac{x}{p(x)} \frac{k^2 - 3k + 2}{2} \geq \frac{2kx^2(k^2 - 3k + 2)}{2p(x)^2}$$

This further gives $p(x)^2 n / x^2 \geq k^3 - 3k^2 + 2k = k(k-1)(k-2) \geq (k-2)^3$ and hence $k \leq (p(x)^2 n / x^2)^{1/3} + 2$ must hold. By this, we get an upper bound for the diameter of: $\text{diam}(G[S]) = 4kx - x \leq 4(xp(x)^2 n)^{1/3} + 7x$. \square

Theorem 4.18 (Max-Pricing-Game: price of anarchy upper bound). *In the Max-Pricing-Game with edge lengths $L \subseteq [\check{\beta}, \hat{\beta}]$ and price function $p : L \rightarrow \mathbb{R}_{\geq 0}$, the price of anarchy is at most $O(\sqrt[3]{n})$.*

Proof. Define the values $x^* := \arg \min_{x \in L} p(x) + x$ and $\chi^* := \arg \min_{x \in L} p(x) + x$. Then, by Lemma 4.5, every strategy profile incurs a social cost of at least $(\chi^* + p(\chi^*)/2)n$. On the other hand, for every buy equilibrium strategy profile S , by Lemma 4.16 and Lemma 4.17, we obtain the social cost upper bound of $\text{cost}(S) \leq n(p(x^*)^2 x^* n)^{1/3} + 2(n-1)(p(x^*) + x^*)$. Comparing both, we get:

$$\begin{aligned} \text{PoA} &\leq \frac{n(p(x^*)^2 x^* n)^{1/3} + 2(n-1)(p(x^*) + x^*)}{(\chi^* + p(\chi^*)/2)n} \\ &\leq 2 \frac{(p(x^*)^2 x^* n)^{1/3} + 2(p(x^*) + x^*)}{x^* + p(x^*)} \\ &\leq 4 + 2 \frac{(p(x^*)^2 x^* n)^{1/3}}{x^* + p(x^*)} \end{aligned}$$

This gives the claim by case distinction of whether $x^*/p(x^*) < 1$ holds. \square

Note that the price of anarchy result in the Max-Game, unlike in the Sum-Game, is not related to the chosen price function. Intuitively, we can explain this observation by the different impact of edge prices in both game variants. While in the Sum-Pricing-Game, an edge is possibly used for several unique shortest paths and hence its distance is weighted by the number of them, in the Max-Pricing-Game, an edge is considered at most once for the distance cost. Thus, the strategic behavior of the agents in the Max-Pricing-Game is close to their behavior in the original Max-Game, since the major strategic choice of the agents is still where to create edges to, and edge prices change this little.

4.6 Conclusion & Future Work

In this chapter, we proposed a model extension that introduces quality-of-service agreements into the framework of network creation games. Despite of the considerably increased freedom in the strategic decisions of the agents (e.g., for any continuous interval of edge lengths the strategy set is unbounded), equilibria always exist. In the Sum-Pricing-Game, we discovered the optimal price and length trade-off for an edge that is used for exactly one shortest path to characterize the price of anarchy. By exploiting properties of specific price functions, we could further provide a non-constant lower bound for the price of anarchy in this game. Contrary to the behavior of the Sum-Pricing-Game, our results for the Max-Pricing-Game indicate that edge prices have only a minor impact on the price of anarchy when agents strive for minimizing their maximum distances.

Building on the models and results of this chapter, it would be interesting to see how the games change when there is not an unlimited provision of edges. For example, inspired by the bounded budgeted equilibria by Ehsani et al. [Ehs+15], there could be a limit on the number of available edges for each individual edge length. Faced with such limitations, it would also be natural to assume that edge prices are not static anymore but being affected by the demands of the agents. Moreover, one could imagine that the agents are not only buying single connections but high-speed access to certain areas of the network. The special case of buying high-speed access to the whole network will be considered in Chapter 6.

CHAPTER 5

Limits of Locality

THE original motivation to study the Sum-Game model was the rising interest in understanding the structure and evolution of Internet-like networks.¹ As discussed in Section 2.3, a remarkable series of further research and model variants followed this publication. Still, except for a recent approach [Bil+14a], all of these network models have in common that they neglect the fact that agents usually do not have access to global knowledge about the network structure. This is caused by the mere network size as well as by the dynamics of the participating agents, a combination that makes it infeasible for an agent to always maintain a correct overview of the current network. Instead, we have to acknowledge that agents know only certain parts of the network and must perform actions to gather additional information about the unknown network parts.

In this chapter, we want to understand what is possible in terms of strategic decision making and social cost efficiency, if agents are restricted to possess only local knowledge. Our approach is to extend the Sum-Game by Fabrikant et al. [Fab+03] by introducing a locality notion that on the one hand restricts agents to know only the subnetwork within a specific fixed range k , but on the

¹As stated in the introduction of Fabrikant et al. [Fab+03].

other hand allows agents to also probe a certain number of different strategy changes and finally select the best. We benchmark our model according to the following three questions:

- (a) For which locality parameters and which probing techniques are equilibria equal or close to those in the global knowledge Sum-Game?
- (b) Which impact do the locality parameter and the probing technique have on the price of anarchy?
- (c) How costly is the locality restriction for individual agents, i.e., by how much could agents improve their costs if they were not restricted?

In regard of applicability to real networks, both having a fixed (possibly constant or logarithmic) neighborhood parameter and being able to estimate the result of a specific strategy change are realistic assumptions: Whereas restricting the view range of an agent is a typical approach to understand locality in distributed computing (cf. Peleg [Pel00, Section 2]), also the latter assumption is reasonable as discussed by Bilò et al. [Bil+14b]. In particular, obtaining distance information, as it is required for estimating the current cost of an agent, is possible via standard traceroute techniques (cf. Dall’Asta et al. [Dal+06]).

The results of this chapter shed light on what is possible under locality constraints. We will see that probing is an efficient way to overcome locality: Having agents with limited viewing ranges, it is enough to allow them a quadratic number of strategy probes to ensure price of anarchy results close to those for games with global view. Still, we can provide a non-constant lower bound on the price of anarchy and by this show that there is an actually overall quality loss by having locality restrictions – even in the most optimistic case. The individual deterioration of an agent by having a view restriction will be quantified in terms of the approximation ratio of local operations versus global operations. Considering the other extreme of forbidding any strategy probing, our model generalizes the worst-case locality model by Bilò et al. [Bil+14a].

Chapter Basis. The model, analysis, and results presented in the remainder of this chapter are based on the following publication:

2015 (with P. Lenzner). “Network Creation Games: Think Global – Act Local”. In: *Mathematical Foundations of Computer Science 2015* –

40th International Symposium, MFCS 2015, Milan, Italy, August 24–28, 2015. Proceedings, Part II, cf. [CL15].

Chapter Outline. In Section 5.1, we start by introducing our locality model and in Section 5.2 compare it to the existing locality approaches. Section 5.3 contains some simple results about the hardness of best-response computations and the non-convergence of best-response processes. In the main part of this chapter, we first show in Section 5.4 how an agent’s efficiency in terms of approximation ratio is affected, when she is limited by certain locality or probing constraints. Finally in Section 5.5, we provide our results about the price of anarchy and explore the parameter ranges for which equilibria in the Sum-Game coincide with those in our locality model.

5.1 Model & Notations

In the following, we introduce the *k-local Sum-Game*, which is a variant of the Sum-Game model by Fabrikant et al. [Fab+03] and incorporates the same notation as given in Section 2.1. An instance of the *k-local Sum-Game* consists of a set of n selfish agents V , an edge price parameter $\alpha > 0$, a view radius $k \in \mathbb{N}$, and a probing technique. Every agent $u \in V$ strives for minimizing her private cost by selecting a strategy $s_u \subseteq V \setminus \{u\}$. The private cost of an agent is given by:

$$c_u(S) = \alpha \cdot |s_u| + \sum_{v \in V} d_{G[S]}(u, v) \quad (5.1)$$

For this cost function, we refer to the first term as *edge* $c_u(S)$, called *the edge cost of u* , and to the second term as *dist* $c_u(S)$, called *the distance cost of u* . Considering a strategy profile S that is formed by the joint strategies of all agents, we obtain a network $G[S]$. The social cost of a strategy profile is $\text{cost}(S) = \sum_{u \in V} c_u(S)$. Throughout this chapter, we will consider connected networks only as they are the only ones that induce finite private and social cost.

Probing Locality. In our *k-local Sum-Game*, the selfish agents are only aware of their k -neighborhood and hence are restricted to perform operations therein. Given a locality parameter $k \in \mathbb{N}$ and an agent $u \in V$, then for a strategy profile S the set $N_k(u) \subseteq V$ denotes the agents with a distance of at most k to u

in $G[S]$. The subgraph of $G[S]$ that is induced by the set $N_k(u)$ is called the k -neighborhood of u .

In this game, agents are only allowed to perform k -local operations, which is the simultaneous application of any combination of the actions of (1) removing an own edge, (2) swapping an own edge to an agent in the k -neighborhood,² and (3) creating an edge to an agent in the k -neighborhood. In particular, the set of actions that realize a k -local operation must not contradict each other. Since together they form a single operation, they are also performed simultaneously and hence due to the same k -neighborhood. If a k -local operation consists of only exactly one of the actions (1)–(3), it is called a k -local greedy operation.

A *probing technique* limits which and how many different strategies can be tested by an agent before selecting her best operation that she wants to perform. In this chapter, we consider the two probing techniques named *unrestricted probing* and *greedy probing*. Yet, for completeness and to emphasize the relation to the model by Bilò et al. [Bil+14a], we also define *0-probing*, which resembles their worst-case model.

Unrestricted Probing: An agent is enabled to test all possible k -local operations and to select a best-response strategy among them.

Greedy Probing: An agent is enabled to test all possible k -local greedy operations and to select a best-response strategy among them.

0 Probing: An agent is not able to test any k -local operation and thus estimates the result of a strategy change by considering the worst-case of all possible networks that comply with her current k -neighborhood (cf. Bilò et al. [Bil+14a] and the discussion in Section 5.2).

Solution Concepts. Restating the notions of Section 2.2, we call a strategy profile to be a *buy equilibrium* if no agent can unilaterally change her strategy to decrease her cost. The strategy profile is a *greedy buy equilibrium* if no agent can unilaterally change her strategy by any greedy operation. When restricting agents to perform only k -local operations with a respective probing technique, we obtain their k -local counterparts. We say that a strategy profile is a k -local

²A swap-action of an agent is the simultaneous deletion of an own edge and creation of a new incident own edge.

buy equilibrium and call the corresponding network to be k -local stable if no agent can unilaterally decrease her cost by a k -local operation. If no agent can decrease her cost by a k -local greedy operation, the strategy profile is a k -local *greedy buy equilibrium* and the network is called k -local greedy stable.

Supplementing the above notions, we also consider ε -approximate equilibria (cf. Section 2.2). We call a strategy profile to be a ε -approximate *buy equilibrium* if no strategy change of an agent can decrease her cost by more than an ε -fraction of her current cost. Similarly, we say a strategy profile is a ε -approximate *greedy buy equilibrium* if no agent can decrease her cost by more than an ε -fraction of her current cost by a greedy operation. Note that in both approximate equilibria notions arbitrary operations are allowed, not only k -local ones.

For a fixed price parameter α , we define classes of equilibria due to the different solution concepts: BE is the class of all networks that are a buy equilibrium, GBE is the class of all networks a greedy buy equilibrium, k -BE is the class of all k -local buy equilibrium networks, and k -GBE is the class of all k -local greedy buy equilibrium networks.

Our notion of social efficiency of a strategy profile is the ratio of its induced social cost and the optimal social cost. Specifically, we are interested in the worst-case ratio of any equilibrium's social cost and the optimal social cost, which we recall as the price of anarchy (cf. Definition 2.1).

5.2 Related Work & Contribution

The only models in the realm of network creation games that consider games without global knowledge are by Bilò et al. [Bil+14a; Bil+14b]. In their games, the agents have limited viewing ranges and thus can access only a certain subset of information about the network. Hence, they are forced to base their decisions upon such incomplete information. Bilò et al. model this by considering conservatively acting agents, i.e., agents who perform only those actions which they know for certain to reduce their private costs. By this, the private cost of an agent can still depend on the total network, even if an agent cannot estimate it.

In Bilò et al. [Bil+14a], the authors incorporate a most pessimistic locality view, which limits agents to know exactly their k -neighborhoods. For any operation, an agent then estimates her private cost change by making a worst-

case assumption about the unknown network part. Specifically, she computes an operation's private cost change by taking the worst-case over all networks of arbitrary size that comply with her current k -local view. Not surprisingly, the authors can provide several non-constant lower bounds for the price of anarchy in both the Sum-Game and the Max-Game. In particular, for the Sum-Game the price of anarchy is at least $\Omega(n/k)$, when $k = o(\alpha^{1/3})$, and for the Max-Game it is at least $\Omega(n/(1 + \alpha))$. They further show for the Max-Game that their lower bound is still $\Omega(n^{1-\varepsilon})$ for every $\varepsilon > 0$, even if k is poly-logarithmic and $\alpha = O(\log n)$. For none of the considered games, the authors provide upper bounds on the price of anarchy. In the context of our probing models, these games can be understood as the k -local Sum-Game and the k -local Max-Game with 0-probing.

In a follow-up paper, Bilò et al. [Bil+14b] considered a variant where agents have access to certain traceroute-based information, in addition to their k -local views. Specifically, they look at how much the agents gain by having access to (1) a distance vector, (2) a minimum spanning tree, or (3) the set of all minimum spanning trees. Using these traceroute-based information, they provide the first known upper bounds on the price of anarchy for games with local view restricted agents. However, for all considered variants, the price of anarchy bounds are much worse than for the classic Sum-Game and Max-Game. For all versions of the Sum-Game, the price of anarchy is $\Theta(\min\{1 + \alpha, n\})$; while for versions of the Max-Game, it is $\Theta(n)$ for $\alpha > 1$. Notably, their price of anarchy proofs only require agents to have access to distance vector information.

There is a lot of literature about network exploration and in particular about using traceroute strategies. For example, Beerliova et al. [Bee+06] consider the complexity of discovering the topology of a whole network in an online-setting and provide an $O(\sqrt{n \log n})$ -competitive online algorithm for a network of n agents. A good overview of applications of different traceroute protocols for network discovery is provided by Dall'Asta et al. [Dal+06]. Specifically for retrieving simple network information like distance vectors, we refer to the discussion in Bilò et al. [Bil+14b] about how to utilize traceroute protocols for this purpose.

Apart from network creation games, the influence of locality has also been studied for other game-theoretic settings, e.g., in the local matching model by Hoefer [Hoe13], where agents know only their 2-neighborhood and have to

choose their matching partners from this set.

Contribution. Our main contribution is a new model for locality in network creation games. By taking the natural observation into account that agents want to test the outcomes of their strategic changes, we create not only a more optimistic but also a more realistic model. In our model, agents are still limited in their knowledge and actions, but now can probe different operations and choose the best one. Yet, our probing locality (even for viewing ranges as small as 2) has no impact on the hardness of best-response computations or the convergence of best-response processes.

For the Sum-Game with unrestricted probing, we show upper bounds for the price of anarchy that are close to those in the classic Sum-Game. Hence, the results are in stark contrast to the worst-case model by Bilò et al. [Bil+14a], where myopic agents select their strategies without knowing the exact results of their strategy changes. Looking at the Sum-Game with greedy probing, which limits the probes to the quadratic number of greedy operations, we show the surprising insight that n^2 probes suffice to gain the same results for the price of anarchy as with 2^n probes. Moreover, considering the behavior of the individual agents, we discuss how well k -local greedy operations approximate arbitrary greedy operations. For tree networks, we specifically show that k -local operations with greedy probing approximate arbitrary operations by $\Theta\left(\frac{\log n}{k}\right)$.

5.3 Preliminaries

Starting our analysis, we first provide some observations about the structure and relations of equilibria and further show that results about the hardness of computing best responses and about non-convergence of best-response processes still apply for the k -local Sum-Game. In particular, the hardness results hold for any k , while the non-convergence results hold for any $k \geq 2$.

Observation 5.1. Agents in the k -local Sum-Game and the k -local greedy Sum-Game can be characterized as follows:

- (a) Agents with unrestricted probing are equivalent to agents who are aware of the whole network, but whose operations are restricted to be only

k -local operations.

- (b) Agents with greedy probing are equivalent to agents who are aware of the whole network, but whose operations are restricted to be only k -local greedy operations.

Using these characterizations, we can directly derive some set relationships for the equilibrium classes.

Observation 5.2. For a fixed edge price $\alpha > 0$ and any fixed locality parameter $k \in \mathbb{N}$, the following relations between the different equilibrium classes hold:

- (a) $\text{BE} \subseteq k\text{-BE} \subseteq k\text{-GBE}$
- (b) $\text{BE} \subseteq \text{GBE} \subseteq k\text{-GBE}$

Since there exist equilibria for the classic Sum-Game (cf. Fabrikant et al. [Fab+03, Section 2]), note that there are also equilibria for both the k -local Sum-Game and the k -local greedy Sum-Game.

Theorem 5.3. *For the k -local Sum-Game with $k \geq 1$, in general it is \mathcal{NP} -hard to compute a k -local best-response operation.*

Proof. We follow the hardness proof by Fabrikant et al. [Fab+03, Proposition 1], which reduces the MINIMUM DOMINATING SET problem [GJ02] to the computation of an optimal strategy change in the Sum-Game. For a given network $G = (V, E)$, the MINIMUM DOMINATING SET problem is the task to compute a dominating set $D \subseteq V$ of minimal size. Here, D is called a *dominating set* if every agent of G belongs to D or has a neighbor in D .

Let $G = (V, E)$ be an instance of MINIMUM DOMINATING SET, then we obtain an instance of the k -local Sum-Game as follows: Let V be a set of agents and define a strategy profile such that for every edge $\{u, v\}$ in E , there is a corresponding strategy s_u with $v \in s_u$. Thereby, edge ownerships are assigned arbitrarily. Furthermore, we add an additional agent z to the network and set her strategy to $s_z := \{V\}$: i.e., z owns edges to all other agents. For an edge price of $\alpha \in (1, 2)$ and arbitrary $k \geq 1$, we claim that a minimum cost strategy for z forms a minimum dominating set.

For this, let s'_z be the optimal strategy change for z and S' the changed strategy profile. Since $\alpha < 2$, we get that for every agent $v \in V$ it must hold

$d_{G[S']}(z, v) < 3$, since otherwise z could improve her private cost by creating an edge to v . Hence, the distance is either 1 or 2. For $d_{G[S']}(z, v) = 1$, agent z owns an edge to v and otherwise, for $d_{G[S']}(z, v) = 2$, agent v must own an edge to a direct neighbor of v . By this, s'_z forms a dominating set and it remains to show that its size is minimal. Since the private cost of z is

$$c_z(S') = \alpha \cdot |s'_z| + |s'_z| + 2 \cdot |V \setminus s'_z| = |V| + \alpha \cdot |s'_z| + |V \setminus s'_z|,$$

we get by $\alpha > 1$ that z 's private cost is minimized when $|s'_z|$ is minimized. Thus, an optimal strategy change for z is a minimum dominating set in the constructed instance and directly gives a minimum dominating set in G . \square

In the remainder of this section, we consider the convergence properties of best-response processes, as introduced in Section 2.2.3. Given a fixed game with edge price, locality parameter, and probing technique, we consider some initial strategy profile and analyze sequences of best-response strategy changes of the agents. At each time step, exactly one agent acts and we ask if such best-response processes are guaranteed to converge to an equilibrium state. Specifically, does the game possess the finite improvement property and hence, is it a potential game? Or otherwise, can we show a cyclic sequence of best-response strategy changes, i.e., the existence of a best-response cycle that can prevent these processes from terminating?

Theorem 5.4. *In both the 1-local greedy Sum-Game and the 1-local Sum-Game, every sequence of $(n - 1)^2$ -many improving operations converges to an equilibrium. For $k \geq 2$, for both the k -local greedy Sum-Game and the k -local Sum-Game, there are strategy profiles and best-response sequences that result in best-response cycles.*

Proof. For $k = 1$, neither in the Sum-Game nor in the greedy Sum-Game there is an agent who can create or swap an edge. Hence, the number of edges is strictly monotonically decreasing with every operation. We know for the initial strategy profile that for $n = |V|$ agents, there are at most $n(n - 1)$ edges. Since no agent disconnects the network by any operation, there never can be less than $n - 1$ edges. Hence, after at most $(n - 1)^2$ -many improving strategy changes, in both models the network is an equilibrium.

For $k = 2$ and $\alpha \in (2, 3)$, Figure 5.1 provides a best-response cycle in the 2-local Sum-Game: In (1) c swaps edge $\{c, a\} \rightarrow \{c, b\}$, in (2) a buys edge $\{a, e\}$,

in (3) b deletes edge $\{a, b\}$, in (4) a buys edge $\{a, b\}$, in (5) b deletes edge $\{b, e\}$, in (6) d swaps edge $\{d, b\} \rightarrow \{d, a\}$, in (7) c swaps edge $\{c, b\} \rightarrow \{c, a\}$, in (8) b buys edge $\{b, e\}$, in (9) a deletes edge $\{a, b\}$, in (10) b buys edge $\{a, b\}$, in (11) a deletes edge $\{a, e\}$, and in (12) d swaps edge $\{d, a\} \rightarrow \{d, b\}$. This gives the original network from (1) and hence a best-response cycle exists. It is easy to check that in every step of the cycle, the active agent performs a best-response operation. Since every operation is a greedy operation, the best-response cycle also holds for the 2-local greedy Sum-Game. Note that in strategy change (6), agent d could perform a 3-local greedy operation (swapping edge $\{d, b\} \rightarrow \{d, e\}$), if the game was 3-local, and hence this construction cannot be used for $k = 3$.

For $k = 3$ and $\alpha \in (3, 4)$, Figure 5.2 provides a best-response cycle in the 3-local Sum-Game that works as follows: In (1) b buys edge $\{b, h\}$, in (2) d swaps edge $\{d, c\} \rightarrow \{d, b\}$, in (3) a swaps edge $\{a, c\} \rightarrow \{a, b\}$, in (4) b deletes edge $\{b, h\}$, in (5) c buys edge $\{c, h\}$, in (6) d swaps edge $\{d, b\} \rightarrow \{d, c\}$, in (7) a swaps edge $\{a, b\} \rightarrow \{a, c\}$, and in (8) c deletes edge $\{c, h\}$. This again gives the original network from (1). It is easy to check that in every step of the cycle, the active agent performs a best-response operation. Since every operation is a greedy operation, the best-response cycle also holds for the 3-local greedy Sum-Game.

For any $k \geq 4$, we refer to the construction in Kawald and Lenzner [KL13, Theorem 7], which only requires agents to perform 4-local greedy operations and hence provides a best-response cycle for both the k -local Sum-Game and the k -local greedy Sum-Game. \square

5.4 Approximation Quality of Greedy Probing

In this section, we investigate the agents' perspectives in the Sum-Game in terms of how close k -local operations approximate arbitrary strategies. Specifically, for a given k -local game we ask by how much agents could improve their costs if they were allowed to perform arbitrary operations. First, we show that a k -local greedy best-response operation is a 3-approximation of a k -local best-response operation. Then, we shift our focus to the approximation quality of k -local greedy operations versus arbitrary greedy operations. For tree networks, we provide a tight approximation bound of $\Theta\left(\frac{\log n}{k}\right)$ and for any

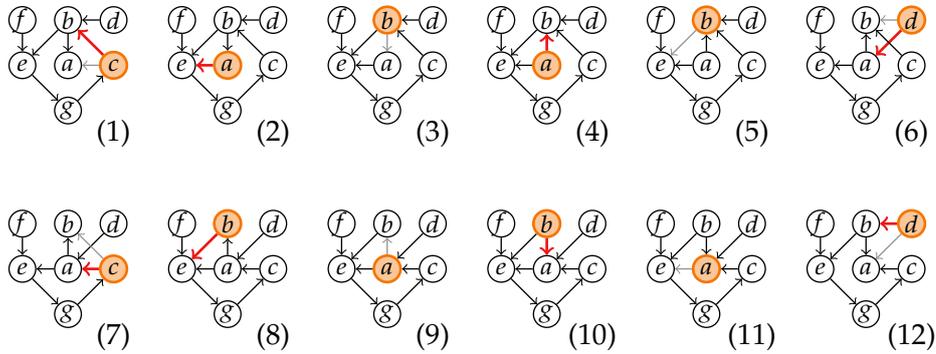


Figure 5.1: The 2-local Sum-Game with a best-response cycle for edge price $\alpha \in (2, 3)$. The orange agent performs a best-response operation: gray edges are removed, red edges are created.

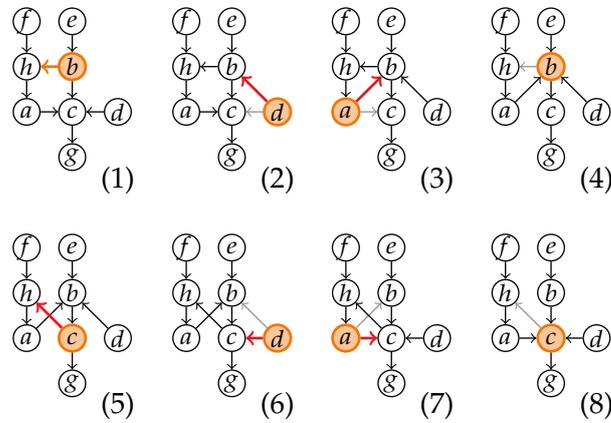


Figure 5.2: The 3-local Sum-Game with a best-response cycle for edge price $\alpha \in (3, 4)$. The orange agent performs a best-response operation: gray edges are removed, red edges are created.

general k -local greedy buy equilibrium network G , we get an approximation upper bound of $O(\text{diam}(G))$.

5.4.1 Approximation of the k -Local Sum-Game

First, we consider the approximation quality of k -local greedy operations regarding arbitrary k -local operations. For this, we show that k -local greedy best-response operations are 3-approximations of arbitrary k -local operations.

Theorem 5.5. *In the k -local Sum-Game with $k \geq 1$, every strategy profile in k -local greedy buy equilibrium is a 3-approximate k -local buy equilibrium.*

Proof. We show that if an agent cannot improve her cost by a k -local greedy operation, then this agent also cannot perform an arbitrary k -local operation that reduces her private cost to a value less than $1/3$ of her current cost.

For this, similar to [Len12], we reduce the best-response computation of any agent to the solution of a corresponding UNCAPACITATED METRIC FACILITY LOCATION instance (UMFL, cf. Williamson and Shmoys [WS11]). UMFL is the problem of selecting a subset $X \subseteq \mathcal{F}$ of facilities with the objective to minimize the term $\sum_{v \in X} f_v + \sum_{x \in \mathcal{C}} \min_{v \in X} d(x, v)$ for a given set of clients \mathcal{C} , individual opening costs $f_v \geq 0$ for every facility $v \in \mathcal{F}$, and a metric distance function $d : \mathcal{F} \times \mathcal{C} \rightarrow \mathbb{R}_{\geq 0}$. Arya et al. [Ary+04, Theorem 4.3] provide a locality gap result that (beside other implications) states: When starting with an arbitrary facility set and performing only the operations of closing a single facility, opening a single facility, or swapping a single facility (i.e., simultaneously closing one facility and opening another one) until no further improvement is possible, this greedy local search heuristic results in a 3-approximation of the optimal solution.

Given a k -local greedy buy equilibrium strategy profile S for a k -local Sum-Game with agents V and edge price α , let $u \in V$ denote an arbitrary agent. For agent u let s_u denote the set of agents to which she owns an edge and let \bar{s}_u be the set of agents who own edges to agent u . Using this, we define an instance $I = (\mathcal{F}, \mathcal{C}, \{f_v\}_{v \in V}, d)$ of the UMFL problem as follows:

- The set of facilities \mathcal{F} is given by $\mathcal{F} := N_k(u) \setminus \{u\}$.
- The set of clients \mathcal{C} is given by $\mathcal{C} := V \setminus \{u\}$.

- For every facility $v \in \mathcal{F} \cap \bar{s}_u$, we define the opening cost as $f_v := 0$ and for all other facilities we set $f_v := \alpha$.
- For a facility $v \in \mathcal{F}$ and a client $x \in \mathcal{C}$, we set the distance as $d(v, x) := d_{G[S]}(v, x) + 1$; the distance is ∞ if there is no path from v to x in $G[S]$.

Note that by using the shortest path metric to define the distances in I , we ensure that the distances are metric. It is easy to see that $c_u(S) = \text{cost}(I) = \sum_{v \in S_u} f_v + \sum_{x \in \mathcal{C}} \min_{v \in S_u} d(x, v)$. Since we assume that agent u cannot perform any improving k -local greedy operation, the locality gap for UMFL [Ary+04] yields that the cost of agent u in S is at most 3 times her cost it would be by performing a k -local best-response operation. \square

The construction by Lenzner [Len12, Theorem 3] further yields an approximation lower bound for $k \geq 2$ such that there exist k -local greedy buy equilibria that are in $(3/2)$ -approximate k -local buy equilibrium. This lower bound also applies here.

5.4.2 Approximation Lower Bound in the Sum-Game

In the following, we prove a lower bound on the approximation ratio for k -local greedy operations versus arbitrary greedy operations. For this, we use the following constructed d - l -Tree-Star network (cf. Figure 5.3). It consists of a complete binary tree subgraph and a star subgraph, both connected by one additional agent.

Complete Binary Tree T_d : For $d \in \mathbb{N}$, define T_d to be a complete balanced binary tree of depth d with root agent r such that every edge is owned by the agent who is closer to r . Let u denote a fixed leaf agent (i.e., an agent with maximal distance to r).

Tree-Star $G_{d,l}$: We consider a combination of a complete binary tree T_d , with root agent r and of even depth d , and a star network consisting of a center agent z and l -many leaves (cf. Figure 5.3). Both subgraphs are connected by one additional agent y who owns one edge to the root agent r and one edge to the center agent z . The tree subgraph contains one (arbitrary) leaf marked as u . In the following, we will also consider the networks:

5 Limits of Locality

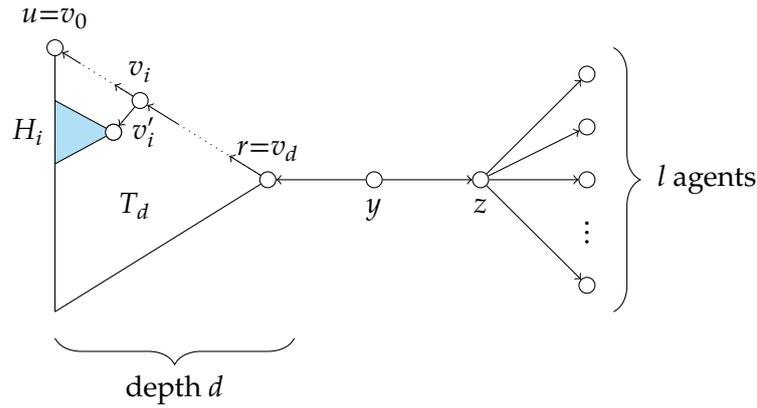


Figure 5.3: Illustration of the d - l -Tree-Star $G_{d,l}$ with depicted path v_0, \dots, v_d and the subtree H_i for an agent v_i . Note that edges are undirected but depicted with directions to indicate edge ownerships.

- $G_{d,l}^y$: the network that is obtained from $G_{d,l}$ when agent u buys the edge $\{u, y\}$.
- $G_{d,l}^z$: the network that is obtained from $G_{d,l}$ when agent u buys the edge $\{u, z\}$.

For a d - l -Tree-Star network with suitable chosen d and l , we will prove that agent u cannot improve her private cost by a k -local greedy operation, yet by an arbitrary greedy operation. We start by computing the maximal distance cost improvement of agent u when performing a greedy operation and then show that u 's operations upper-bound the distance cost improvement of any agent. By comparing the private cost of u in $G_{d,l}$ to her private cost after she performed a (non-local) best-response greedy operation, eventually we will obtain the lower bound.

Different to the notation in the previous chapters, for convenience, in the following we sometimes use $c_u(G)$ to denote the private cost of some agent u with respect to the strategy profile implied by network the G and only implicitly consider the respective strategy profile. Moreover, we denote the agents belonging to a subgraph G' of some network G by $V(G')$ and use

$$\text{dist}_u(G') := \sum_{v \in V(G')} d_G(u, v)$$

to denote the distance cost of agent u only regarding the agents in G' .

Lemma 5.6. *Let d be even and the networks defined as above. Then, the following distance cost estimations hold:*

$$(a) \text{ dist}_u(G_{d,l}) = 9 + l(d + 3) + 3d + 2^{d+1}(2d - 3)$$

$$(b) \text{ dist}_u(G_{d,l}^y) = 9 + 3l + d + 2^{d+1}(d + 1) - 2^{d/2+3}$$

$$(c) \text{ dist}_u(G_{d,l}^z) = 9 + 2l + d + 2^{d+1}(d + 2) - 3 \cdot 2^{d/2+2}$$

Proof. For the following estimations, we consider the shortest path from u to r in T_d and denote the agents on this path as $u =: v_0, \dots, v_d := r$ (cf. Figure 5.3). For every v_i , we call the subtree rooted at the (unique) non-path neighbor of v_i as H_i . Then, the distance cost of v_i for the corresponding subtree is $\text{dist}_{v_i}(H_i) = \sum_{j=1}^i 2^{j-1}j = 2^i i - 2^i + 1$. Note that there are $|H_i| = 2^i - 1$ many agents in the subtree H_i and hence u has distance cost for this tree of:

$$\text{dist}_u(H_i \cup \{v_i\}) = \text{dist}_{v_i}(H_i \cup \{v_i\}) + 2^i i = \text{dist}_{v_i}(H_i) + 2^i i$$

Claim (a): For network $G_{d,l}$ we get the following distance cost for u :

$$\begin{aligned} \text{dist}_u(G_{d,l}) &= 3 + 2d + l(d + 3) + \sum_{i=1}^d (2^i i - 2^i + 1 + 2^i i) \\ &= 3 + 2d + l(d + 3) + 2(2^{d+1}d - 3 \cdot 2^d) + 6 + d \\ &= 9 + 3d + l(d + 3) + 2^{d+1}(2d - 3) \end{aligned}$$

Claim (b): For network $G_{d,l}^y$ we get the following distance cost for u :

$$\begin{aligned} \text{dist}_u(G_{d,l}^y) &= 3 + 3l + \sum_{i=1}^d \text{dist}_{v_i}(H_i) + \sum_{i=1}^{d/2+1} 2^i i + \sum_{i=d/2+2}^d 2^i (d - i + 2) \\ &= 3 + 3l + \sum_{i=1}^d \text{dist}_{v_i}(H_i) + \sum_{i=1}^{d/2+1} 2^i i + \sum_{i=1}^{d/2-1} 2^{d/2+1+i} (d/2 - i + 1) \\ &= 3 + 3l + (4 + d + 2^{d+1}(d - 2)) + (2^{d/2+1}d + 2) \\ &\quad + (-2^{d/2+1}d + 3 \cdot 2^{d+1} - 2^{d/2+3}) \\ &= 9 + 3l + d + 2^{d+1}(d + 1) - 2^{d/2+3} \end{aligned}$$

5 Limits of Locality

Claim (c): For network $G_{d,l}^z$ we get the following distance cost for u :

$$\begin{aligned}
\text{dist}_u(G_{d,l}^z) &= 3 + 2l + \sum_{i=1}^d \text{dist}_{v_i}(H_i) + \sum_{i=1}^{d/2+1} 2^i + \sum_{i=d/2+2}^d 2^i(d-i+3) \\
&= 3 + 2l + \sum_{i=1}^d \text{dist}_{v_i}(H_i) + \sum_{i=1}^{d/2+1} 2^i + \sum_{i=1}^{d/2-1} 2^{d/2+1+i}(d/2-i+2) \\
&= 3 + 2l + (4 + d + 2^{d+1}(d-2)) + (2^{d/2+1}d + 2) \\
&\quad + (-2^{d/2+1}d + 2^{d+3} - 3 \cdot 2^{d/2+2}) \\
&= 9 + 2l + d + 2^{d+1}(d+2) - 3 \cdot 2^{d/2+2} \quad \square
\end{aligned}$$

Considering agent u in a d - l -Tree-Star network $G_{d,l}$, we first compute a parameter l such that the best possible greedy edge creation of u is to buy the edge $\{u, z\}$. Then, we show that u is the agent with the maximum possible private cost improvement among all agents, but that u cannot perform an improving response if l and k are selected appropriately. By this, we get that the network is a k -local greedy buy equilibrium and we can compare u 's private cost with the result of an arbitrary greedy best-response operation.

Lemma 5.7. *For d even and $l \geq 2^{d+1}$, let $G_{d,l}$ be a d - l -Tree-Star. Then, creating the edge $\{u, z\}$ is the best possible greedy edge creation operation of agent u .*

Proof. We compute the minimal parameter l such that $\text{dist}_u(G_{d,l}^z) \leq \text{dist}_u(G_{d,l}^y)$: i.e., that creating an edge to some agent in T_d cannot be a best response. Using Lemma 5.6, we compute:

$$0 \leq \text{dist}_u(G_{d,l}^y) - \text{dist}_u(G_{d,l}^z) = l - 2^{d+1} + 2^{d/2+2}$$

Hence, for $l \geq 2^{d+1}$ creating an edge to z is the best possible greedy edge creation operation for u . \square

Lemma 5.8. *For d even and $l \geq 2^{d+1}$, let $G_{d,l}$ be a d - l -Tree-Star and T_d the contained binary tree of depth d . Consider a $v \in V(T_d)$ who can perform an improving response and let $v =: v_0, v_1, \dots, v_m := z$ denote the shortest path from v to z . Then, for $k \geq 2$:*

$$v\text{'s } k\text{-local greedy best response is } \begin{cases} \text{to create an edge to } v_k & \text{for } k \leq m, \\ \text{to create an edge to } z & \text{otherwise.} \end{cases}$$

Proof. For any agent $x \in T_d$, we use the notation $S(x) \subseteq T_d$ to denote the set of agents in the subtree rooted at agent x in T_d . Furthermore, we use $\text{height}(w)$ to denote the height of agent $x \in T_d$, i.e., the number of edges on the longest downward simple path from x to a leaf of T_d .

First, we consider v creating an edge to an agent $w \in T_d$ with $\text{height}(w) \leq \text{height}(v)$ and assume that this edge creation gives a maximum possible distance cost decrease for v . Let w' be the direct predecessor of w on the shortest path from v to w in $G_{d,l}$. If $d_{G_{d,l}}(v, w) > 2$, we compare the distance cost improvements of creating the edge $\{v, w\}$ and of creating the edge $\{v, w'\}$. Note that the distance cost reduce differs only by the number of agents in the corresponding subtrees. Since T_d is a complete binary tree, creating the edge to w' gives an additional decrease of $2^{\text{height}(w')} - 2^{\text{height}(w)} > 0$. Hence, w must be at distance 2 to v and accordingly gives a distance cost improvement of $2^{\text{height}(w)} - 1 \leq 2^{d-1} - 1 < l$. That is, creating an edge to v_2 would result in a bigger gain and therefore, creating an edge to some agent in $S(v)$ cannot be a best response. By the same arguments, this also holds for edge swaps and thus no improving edge swap to an agent in $S(v)$ is a best response.

Secondly, we assume that creating an edge to some agent $w \in T_d$ with $\text{height}(w) > \text{height}(v)$ is v 's best response. If there exists an agent w' on the shortest path from w to z who is within v 's k -neighborhood, we consider creating the edge $\{v, w'\}$. Compared to the creation of $\{v, w\}$, this edge gives an additional distance cost reduction to all l -many leaves of the star subgraph of $G_{d,l}$ of 1 each, while increasing the distance to at most 2^{d+1} agents in T_d by 1 each. Hence, the target for the edge creation must be closest possible to z .

Finally, note that agent v can neither delete any edge that connects her to an agent in $S(v)$, nor swap such an edge to an agent with a height bigger than $\text{height}(v)$. \square

Lemma 5.9. *For d even and $l \geq 2^{d+1}$, let $G_{d,l}$ be a d - l -Tree-Star and T_d the contained binary tree of depth d . Then, the agent with maximum gain by a k -local greedy best-response operation is agent u .*

Proof. First note that the maximum possible distance cost decrease for all leaves of T_d is equal and it suffices to consider only agent u .

Assume that there is some non-leaf agent $v \in V(T_d)$, who can achieve a bigger distance cost decrease than u by creating an edge $\{v, w\}$ with $w \in N_k(v)$.

By Lemma 5.8, $d_{G_{d,l}}(w, z) = d_{G_{d,l}}(v, z) - k$ must hold. Since v is no leaf, there is a neighbor v' of v with a larger distance to r than v . Moreover, there is a neighbor w' of w who lies on the shortest path from v to w . We claim that agent v' can achieve a strictly larger cost decrease when creating the edge $\{v', w'\}$ than agent v when creating $\{v, w\}$. Let $A_i \subset V(G_{d,l})$ be the set of agents to which agent v decreases her distances by exactly i when creating the edge $\{v, w\}$. Moreover, let $A'_i \subset V(G_{d,l})$ be the set of agents to which agent v' decreases her distances by i each when creating the edge $\{v', w'\}$. Thus, agent v has a distance cost decrease of $\sum_{i=1}^{k-1} i \cdot |A_i|$ when creating edge $\{v, w\}$ and agent v' has a distance cost decrease of $\sum_{i=1}^{k-1} i \cdot |A'_i|$ when creating edge $\{v', w'\}$. Since w' lies on the shortest path from agent v to w and, additionally, w lies on the shortest path from v to z , we have that $A_i \subseteq A'_i$ holds for all $1 \leq i \leq k-1$. Thus, the claim follows.

It remains to consider the maximum possible distance cost decrease for agents $x \in V(G_{d,l}) \setminus V(T_d)$. First note that without loss of generality it suffices to consider only the case of x being a leaf agent of the star subgraph, in which case an agent has a maximum possible distance cost decrease in the set. Thus, consider a leaf agent x creating the edge $\{x, r\}$, which decreases her distance cost by $2 \cdot |V(T_d)| = 2^{d+2} - 2$. On the contrary, we compare the cost reduction of agent x to the distance cost reduction that agent u can achieve by creating an edge to w , where w is on u 's shortest path to z and w has distance 2 to u . This is, u 's distances to all but 3 agents decrease by 1, which gives a distance cost decrease of $|V(T_d)| + 2 + l - 3 = 2^{d+1} - 2 + 2^{d+1} = 2^{d+2} - 2$. Hence, x cannot reduce her distance cost by more than u . \square

Lemma 5.10. *For d even and $l \geq 2^{d+1}$, let $G_{d,l}$ be a d - l -Tree-Star. Then, for*

$$\alpha \geq (k-1)(2^{d+1} + l + 2)$$

and any $k \leq d$, the network is a k -local greedy buy equilibrium.

Proof. By Lemma 5.9, it suffices to consider only agent u 's k -local greedy best-response operation, which is by Lemma 5.8 the creation of an edge to an agent $w \in N_k(u)$ closest to z . We define $\Delta_{d,k,l} := \text{dist}_u(G_{d,l}) - \text{dist}_u(G_{d,l}^w)$ to denote

the distance cost reduction by this operation. For $k \in \{2, \dots, d\}$, this value is:

$$\begin{aligned} \Delta_{d,k,l} &= (k-1) \left(2 + l + \sum_{i=k}^d 2^i \right) + \sum_{i=1}^{\lfloor k/2 \rfloor - 1} (k-2i-1) 2^{k-i} \\ &\leq (k-1)(2^{d+1} - 2^k + l + 2) + 2^k(k-5) + 3 \cdot 2^{k/2+1} \\ &\leq (k-1)(2^{d+1} + l + 2) \end{aligned}$$

Since this is the maximum possible distance cost decrease of any single edge creation by u , for α being at least this value, no agent can perform an improving response. \square

Theorem 5.11 (*k*-local greedy approximation lower bound). *For any $2 \leq k \leq \log_3(n/2)$ and suitably chosen d and l , the d -l-Tree-Star network $G_{d,l}$ is a k -local greedy buy equilibrium, but only an $\Omega\left(\frac{\log n}{k}\right)$ -approximate greedy buy equilibrium.*

Proof. Using Lemma 5.6, the private cost of u is given by $\text{dist}_u(G_{d,l}) = 9 + 4d + l(d+3) + 2^{d+1}(2d-3)$ for u 's private cost in $G_{d,l}$. Comparing this to the result of an arbitrary greedy best-response operation gives us by applying Lemma 5.6 and Lemma 5.10:

$$\lim_{l \rightarrow \infty} \frac{\text{dist}_u(G_{d,l})}{\alpha + \text{dist}_u(G_{d,l}^Z)} = \lim_{l \rightarrow \infty} \frac{l(d+3)}{l(k+1)} = \frac{d+3}{k+1}$$

Hence, the approximation ratio for any constant neighborhood range k may exceed any constant ratio C by choosing $d > C(k+1) - 3$ and l large enough. In particular, it suffices that l grows fast enough to dominate the other terms in the numerator and in the denominator, as l tends to infinity. We choose $d := \log_3 l$, which then yields:

$$\lim_{l \rightarrow \infty} \frac{\text{dist}_u(G_{\log_3 l, l})}{\Delta_{\log_3 l, k, l} + \text{dist}_u(G_{\log_3 l, l}^Z)} = \Omega\left(\frac{l \log_3 l}{kl}\right) = \Omega\left(\frac{\log l}{k}\right)$$

There are $n = 2^{d+1} + 1 + l$ many agents in the network. Choosing $d \geq 2$, for this value it holds $2^{d+1} + 1 \leq 3^d$ and gives $n \geq l \geq n/2$. Hence, we get $\log l = \Theta(\log n)$ and thus a lower bound of $\Omega\left(\frac{\log_3(n/2)}{k}\right) = \Omega\left(\frac{\log n}{k}\right)$ for $2 \leq k \leq \log_3(n/2)$ and $\Omega(\log n)$ for any constant $k \geq 2$. Note that the diameter of $G_{d,l}$ is $d+3 = \Theta(\log l)$. \square

5.4.3 Approximation Upper Bounds in the Sum-Game

In the following, we will show that the approximation lower bound is tight for every k -local greedy buy equilibrium tree network. Our main insight for this result (formalized in the following lemma) is that whenever an agent can perform a swap in a tree network, then there is also a 2-local greedy improving-response swap available for this agent. Since this property does not hold for general networks, we later provide another approximation upper bound that holds for arbitrary networks.

Lemma 5.12. *Let u be an agent in a tree network T . If u can perform an arbitrary greedy edge swap in T , then there exists an improving 2-local greedy edge swap operation for u .*

Proof. Let $\{u, v\} \rightarrow \{u, v_m\}$ be a best-response edge swap of u and assume $m = d_T(u, v_m) > 2$. We define $P := (u, v = v_1, v_2, \dots, v_{m-1}, v_m)$ to be the shortest path from u to v_m (cf. Figure 5.4). For this path, we obtain $v = v_1$ since T is a tree and a swap must preserve connectivity. Thus, the swap only changes distances to agents in the subtree T_{v_1} of agent v_1 , which is rooted at u . For all agents v_i on the path P , let V_{v_i} denote the set of agents who have agent v_i on their shortest path to any neighbor of v_i on P . Let T_i be the tree that results from u performing the edge swap $\{u, v\} \rightarrow \{u, v_i\}$. Since the swap $\{u, v\} \rightarrow \{u, v_m\}$ is a best-response edge swap, we have $c_u(T_i) \geq c_u(T_m)$, for $2 \leq i \leq m - 1$. Using this together with:

$$c_u(T_m) = \sum_{i=1}^m (m - i + 1) \cdot |V_{v_i}| + \sum_{z \in V(T) \setminus V(T_v)} d_T(u, z) + \text{edge}_u(T) \text{ and}$$

$$c_u(T_{m-1}) = \sum_{i=1}^{m-1} (m - i) \cdot |V_{v_i}| + 2 \cdot |V_{v_1}| + \sum_{z \in V(T) \setminus V(T_v)} d_T(u, z) + \text{edge}_u(T),$$

we get:

$$\begin{aligned} 0 &\leq c_u(T_{m-1}) - c_u(T_m) \\ &= \sum_{i=1}^{m-1} (m - i) \cdot |V_{v_i}| + 2 \cdot |V_{v_1}| - \left(\sum_{i=1}^m (m - i + 1) \cdot |V_{v_i}| \right) \\ &= - \sum_{i=1}^{m-1} |V_{v_i}| + |V_{v_1}| \end{aligned}$$

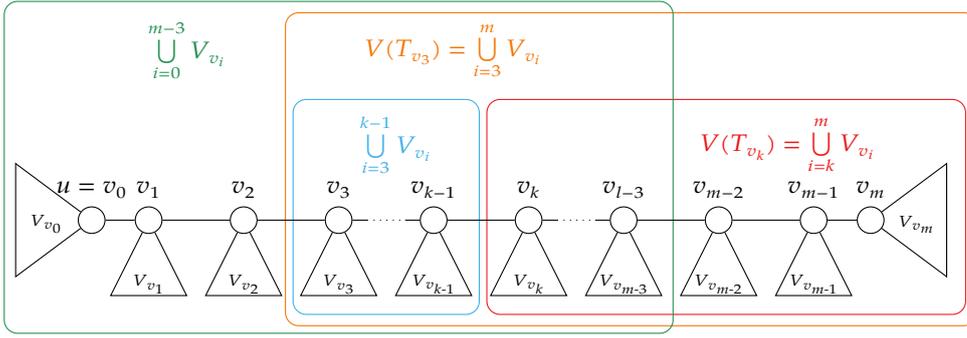


Figure 5.4: Illustration of the sets of agents as used in the proof of Theorem 5.13.

Hence, it must hold:

$$|V_{v_1}| \leq |V_{v_l}| - \sum_{i=2}^{m-1} |V_{v_i}| < |V_{v_l}| + \sum_{i=2}^{m-1} |V_{v_i}| < \sum_{i=2}^m |V_{v_i}|$$

The last estimation gives us that the 2-local edge swap $\{u, v\} \rightarrow \{u, v_2\}$ is an improving response for u , since it decreases u 's distances to exactly $(\sum_{i=2}^m |V_{v_i}|)$ -many agents, yet only increases u 's distances to $|V_{v_1}|$ -many agents by 1. \square

We can now apply this lemma to prove that the approximation lower bound from Theorem 5.11 is tight for tree networks.

Theorem 5.13 (approximation upper bound for tree networks). *Any tree network in k -local greedy buy equilibrium is an $O\left(\frac{\log n}{k}\right)$ -approximate greedy buy equilibrium.*

Proof. Let T be a k -local greedy buy equilibrium tree network. In the following, we show that T is a $O\left(\frac{\text{diam}(T)}{k}\right)$ -approximate greedy buy equilibrium. From this, we then can deduce the claim, since Lemma 5.12 implies that every tree equilibrium is a asymmetric swap equilibrium (cf. Section 2.2), for which we know from Mihalák and Schlegel [MS12] that the equilibrium network diameter is at most $O(\log n)$, where n is the number of agents.

Let v_0 be some agent in T who can buy an edge to decrease her cost. We assume that v_0 buys the edge $\{v_0, v_m\}$ to an agent v_m at distance m and that this is her best-response greedy edge creation. Let the shortest path from v_0 to v_m be given by $P = (v_0, v_1, v_2, \dots, v_{k-1}, v_k, v_{k+1}, \dots, v_{m-1}, v_m)$. We then denote by T_{v_i} the subtree of some agent v_i that is rooted at v_0 and let the sets V_{v_i} for all

5 Limits of Locality

$v_i \in V(P)$ be defined like in the previous proof of Lemma 5.12 (cf. Figure 5.4 for an illustration of these sets).

We assume that agent v_0 cannot decrease her private cost by creating an edge to any agent in her k -neighborhood. Hence, it must hold $\text{dist}_T(v_0, v_m) = m > k \geq 2$. Since we assume that T is a k -local greedy buy equilibrium, v_0 cannot decrease her cost by creating an edge to v_k . But since this operation would decrease v_0 's distances to all agents in $V(T_{v_k})$ by $k - 1$ each, we get:

$$\alpha \geq (k - 1) \cdot |V(T_{v_k})| \quad (5.2)$$

Next, we consider the ratio of agent v_0 's private cost before and after creating edge $\{v_0, v_m\}$. For this, let T' be the network after v_0 has bought the edge and let δ_{v_0} denote the distance cost decrease of agent v_0 . We get:

$$\begin{aligned} \frac{c_{v_0}(T)}{c_{v_0}(T')} &= \frac{c_{v_0}(T)}{c_{v_0}(T) - \delta_{v_0} + \alpha} = \frac{\text{edge}_{v_0}(T) + \sum_{v \in V(T)} d_T(v_0, v)}{\text{edge}_{v_0}(T) + \sum_{v \in V(T)} d_T(v_0, v) - \delta_{v_0} + \alpha} \\ &\leq \frac{\sum_{v \in V(T_{v_3})} d_T(v_0, v)}{\sum_{v \in V(T_{v_3})} d_T(v_0, v) - \delta_{v_0} + \alpha} \end{aligned}$$

The last inequality holds, since all agents to that v_0 decreases her distances by creating $\{v_0, v_m\}$ are in T_{v_3} . We can upper bound the nominator by assuming that all agents in $V(T_{v_3})$ are at maximum distance to v_0 , thus:

$$\sum_{v \in V(T_{v_3})} d_T(v_0, v) \leq \text{diam}(T) \cdot |V(T_{v_3})|$$

Since v_0 has at least distance 1 to all agents in $V(T_{v_3})$ after creating the edge $\{v_0, v_m\}$, we have that $\sum_{v \in V(T_{v_3})} d_T(v_0, v) - \delta_{v_0} > 0$ must hold. Thus, we can lower bound the denominator by $\sum_{v \in V(T_{v_3})} d_T(v_0, v) - \delta_{v_0} + \alpha > \alpha$. Hence, we have:

$$\begin{aligned} \frac{c_{v_0}(T)}{c_{v_0}(T')} &\leq \frac{\sum_{v \in V(T_{v_3})} d_T(v_0, v)}{\sum_{v \in V(T_{v_3})} d_T(v_0, v) - \delta_{v_0} + \alpha} \\ &\leq \frac{\text{diam}(T) \cdot |V(T_{v_3})|}{\alpha} \stackrel{(5.2)}{\leq} \frac{\text{diam}(T) \cdot |V(T_{v_3})|}{(k - 1) \cdot |V(T_{v_k})|} \end{aligned}$$

5.4 Approximation Quality of Greedy Probing

For $k \leq 3$, this already yields $\frac{c_{v_0}(T)}{c_{v_0}(T')} = O(\text{diam}(T))$, since for $i \leq 3$ it holds $|V(T_{v_i})| \geq |V(T_{v_3})|$.

It remains to show that $|V(T_{v_k})| = \Omega(|V(T_{v_3})|)$ for $k > 3$. Since creating edge $\{v_0, v_{m-1}\}$ is a best-response operation for v_0 , its gain must be bigger than when creating edge $\{v_0, v_m\}$. Hence, swapping from $\{v_0, v_m\}$ to $\{v_0, v_{m-1}\}$ would increase agent v_0 's distances to all agents in V_{v_m} by one, as well as decrease agent v_0 's distances to all agents in the set $\bigcup_{i=\lfloor \frac{m}{2} \rfloor + 1}^{m-1} V_{v_i}$ by one. Since all sets V_{v_i} are pairwise disjoint, we get:

$$\sum_{i=\lfloor \frac{m}{2} \rfloor + 1}^{m-1} |V_{v_i}| \leq |V_{v_m}| \quad (5.3)$$

First consider the case $m = 5$, i.e., $m - 2 = 3$ and hence $k = m - 1 = 4$. We use that creating $\{v_0, v_m\}$ strictly decreases agent v_0 's private cost, whereas creating edge $\{v_0, v_k\} = \{v_0, v_{m-1}\}$ does not:

$$|V_{v_m}| > \sum_{i=\lfloor \frac{m}{2} \rfloor + 1}^{m-1} |V_{v_i}| \quad (5.4)$$

Thus, we have $\lfloor \frac{m}{2} \rfloor + 1 = 3$ and we have that $|V(T_{v_3})| < 2 \cdot |V_{v_m}|$, which implies that $|V(T_{v_3})| < 2 \cdot |V(T_{v_k})|$, yielding $|V(T_{v_k})| = \Omega(|V(T_{v_3})|)$.

Next, consider the case $k > 4$. For $m - 2 > 3$ we claim that the edge $\{v_{m-2}, v_{m-1}\}$ must be owned by agent v_{m-1} . This holds, since otherwise agent v_{m-2} could perform the swap $\{v_{m-2}, v_{m-1}\} \rightarrow \{v_{m-2}, v_m\}$ and thereby strictly decrease her cost. This can be seen as follows: If $m - 1 = k$, then by (5.4) we have $|V_{v_m}| > |V_{v_{m-1}}| = |V_{v_k}|$. On the other hand, if $m - 1 > k$, then (5.3) implies $|V_{v_{m-1}}| < |V_{v_m}|$, since the sum on the left has at least one additional non-zero summand. In both cases, we have that the swap $\{v_{m-2}, v_{m-1}\} \rightarrow \{v_{m-2}, v_m\}$ must be improving for agent v_{m-2} . This proves the claim.

Having established that the edge $\{v_{m-2}, v_{m-1}\}$ is owned by agent v_{m-1} and using the assumption that no agent in T can swap an edge in her k -neighborhood to strictly decrease her private cost, the swap $\{v_{m-1}, v_{m-2}\} \rightarrow \{v_{m-1}, v_{m-3}\}$ cannot be an improving response for agent v_{m-1} , which yields $|V_{v_{m-2}}| \geq \sum_{i=0}^{m-3} |V_{v_i}|$. Since $m > 5$, we have that $\lfloor \frac{m}{2} \rfloor + 1 \leq m - 2$. By (5.3) this

5 Limits of Locality

implies $\sum_{i=0}^{m-3} |V_{v_i}| \leq |V_{v_{m-2}}| < \sum_{i=\lfloor \frac{m}{2} \rfloor + 1}^{m-1} |V_{v_i}| \leq |V_{v_m}|$. Thus, if $k = m - 1$ we have that

$$\sum_{i=3}^{k-1} |V_{v_i}| \leq 2 \cdot |V_{v_{m-2}}| \leq 2 \cdot |V_{v_m}| \leq 2 \cdot |V(T_{v_k})|,$$

which implies that $|V(T_{v_3})| \leq 3 \cdot |V(T_{v_k})|$. If $k \leq m - 2$, it follows that

$$\sum_{i=3}^{k-1} |V_{v_i}| < |V(T_{v_k})|$$

and hence $|V(T_{v_3})| < 2 \cdot |V(T_{v_k})|$.

In both cases this yields $|V(T_{v_k})| = \Omega(|V(T_{v_3})|)$. \square

Finally, we prove a general upper bound on the approximation ratio, which is tight for constant k and almost tight in general. As discussed in [CL15, Lemma 2], for general networks the property of Lemma 5.12 does not hold and thus we have to analyze edge swaps and edge creations separately. For both cases, we get the same upper bound on the approximation ratio, which is independent of k .

Theorem 5.14 (general approximation upper bound). *Any k -local greedy buy equilibrium S is an $O(\text{diam}(G[S]))$ -approximate greedy buy equilibrium.*

Proof. Let u be an agent and consider her best-response greedy operation in S . We denote the resulting strategy profile as S' and in the following consider only the cases when this greedy operation is an edge swap or an edge creation, since improving-response edge deletions would contradict $G[S]$ to be a k -local greedy buy equilibrium. We call the agents to which u decreases her distance X^- and the agents to which u increases her distance X^+ . Then for u 's distance cost decrease δ_u we get:

$$\begin{aligned} \delta_u &= \sum_{x \in X^-} (d_{G[S]}(u, x) - d_{G[S']}(u, x)) - \sum_{x \in X^+} (d_{G[S']}(u, x) - d_{G[S]}(u, x)) \\ &\leq \sum_{x \in X^-} (d_{G[S]}(u, x) - d_{G[S']}(u, x)) \end{aligned}$$

Comparing u 's private cost in both networks, when the operation is an edge

swap we get for the approximation ratio:

$$\begin{aligned}
 \frac{c_u(S)}{c_u(S')} &= \frac{c_u(S)}{c_u(S) - \delta_u} = \frac{\text{edge}_u(S) + \text{dist}_u(S)}{\text{edge}_u(S) + \text{dist}_u(S) - \delta_u} < \frac{\text{dist}_u(S)}{\text{dist}_u(S) - \delta_u} \\
 &= \frac{\sum_{v \in V} d_{G[S]}(u, v)}{\sum_{v \in V} d_{G[S]}(u, v) - \delta_u} \leq \frac{\sum_{v \in X^-} d_{G[S]}(u, v)}{\sum_{v \in X^-} d_{G[S]}(u, v) - \delta_u} \\
 &\leq \frac{\sum_{v \in X^-} d_{G[S]}(u, v)}{\sum_{v \in X^-} d_{G[S]}(u, v) - \left(\sum_{v \in X^-} (d_{G[S]}(u, v) - d_{G[S']}(u, v)) \right)} \\
 &\leq \frac{\text{diam}(G[S]) \cdot |X^-|}{\sum_{v \in X^-} d_{G[S']}(u, v)} \leq \frac{\text{diam}(G[S]) \cdot |X^-|}{|X^-|} = \text{diam}(G[S])
 \end{aligned}$$

Considering greedy edge creations, we get for the approximation ratio (note, for this case it holds $X^+ = \emptyset$):

$$\begin{aligned}
 \frac{c_u(S)}{c_u(S')} &= \frac{c_u(S)}{c_u(S) - \delta_u + \alpha} \leq \frac{\sum_{x \in X^-} d_{G[S]}(u, x)}{\sum_{x \in X^-} d_{G[S]}(u, x) - \delta_u + \alpha} \\
 &\leq \frac{\sum_{x \in X^-} d_{G[S]}(u, x)}{|X^-| + \alpha} < \frac{\text{diam}(G[S]) \cdot |X^-|}{|X^-|} = \text{diam}(G[S])
 \end{aligned}$$

For this estimation, the second inequality holds since agent u must have at least distance 1 to all agents in X^- in $G[S']$. \square

5.5 Efficiency of Probing Locality

In this section, we consider the price of anarchy in the k -local Sum-Game, concerning the unrestricted probing and the greedy probing strategies. At first, we analyze for which choices of k the equilibria in the k -local model with unrestricted probing coincide with equilibria in the original Sum-Game. The results of this section are summarized in Figure 5.5. Moreover, in Section 5.5.2 we provide several specific bounds for the price of anarchy.

5.5.1 A Clash of Models

We stated in Observation 5.2 that $k\text{-BE} \subseteq \text{BE}$ holds. In the following, we will discuss the limits of several proof techniques to identify the parameters for which $k\text{-BE} = \text{BE}$, i.e., both equilibria concepts coincide. Specifically, we ask

5 Limits of Locality

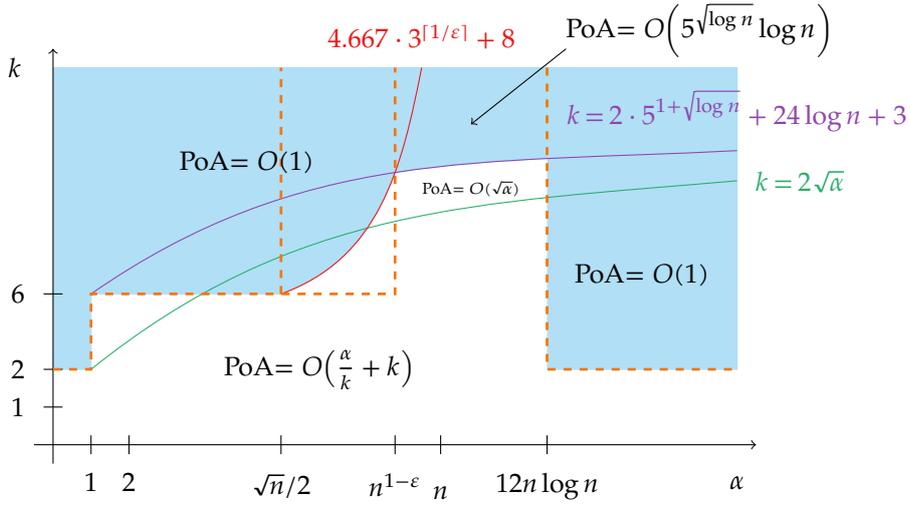


Figure 5.5: Overview of our results from Theorem 5.15 and Theorem 5.28. The light blue area indicates where buy equilibria and k -local buy equilibria coincide and the orange lines mark the ranges that are covered by different proofs.

for which combinations of k and α any k -local buy equilibrium diameter is smaller than k . If this is true, then a k -local operation can achieve the same result as an arbitrary operation. Theorem 5.15 combines the results, which we will prove below.

Theorem 5.15. *The equilibrium concepts k -local buy equilibrium and buy equilibrium coincide for the following parameter combinations and yield the respective price of anarchy results (cf. Figure 5.5): We have k -BE = BE for*

$$\begin{cases} \alpha \in (0, 1) \wedge k \geq 2 & \Rightarrow \text{PoA} = O(1), \\ \alpha \in [1, \sqrt{n/2}] \wedge k \geq 6 & \Rightarrow \text{PoA} = O(1), \\ \alpha \in [1, n^{1-\epsilon}] \wedge \epsilon \geq \frac{1}{\log(n)} \wedge k \geq 4.667 \cdot 3^{\lceil 1/\epsilon \rceil} + 8 & \Rightarrow \text{PoA} = O(3^{\lceil 1/\epsilon \rceil}), \\ \alpha \in [1, 12n \lg n] \wedge k \geq 2 \cdot 5^{1+\sqrt{\lg n}} + 24 \lg(n) + 3 & \Rightarrow \text{PoA} = O\left(5^{\sqrt{\lg n}} \lg n\right), \\ \alpha \geq 12n \log n \wedge k \geq 2 & \Rightarrow \text{PoA} = O(1). \end{cases}$$

Lemma 5.16. *For parameters $0 < \alpha < 1$ and $k \geq 2$, it holds k -BE = BE and the price of anarchy is 1.*

Proof. Given any strategy profile S and $\alpha < 1$, assume there are two closest

agents $u, v \in G[S]$ that are not connected by one edge: i.e., $d_{G[S]}(u, v) = 2$. In this case, creating an edge $\{u, v\}$ is an improving response for u . Hence, the only equilibrium graph for $\alpha < 1$ is a clique, which is also the optimal solution (cf. [Fab+03]). \square

Lemma 5.17 ([Dem+07], Theorem 4). *For parameters $1 \leq \alpha \leq \sqrt{n/2}$ and $k \geq 6$, it holds $k\text{-BE} = \text{BE}$ and the price of anarchy is at most 6.*

Proof. In [Dem+07], the authors show that every shortest path tree rooted at some agent u has a height of at most 5. For this, they assume the contrary and show the existence of an improving response where an agent at a distance of at least 6 buys an edge towards u . This operation is available with $k \geq 6$, hence every k -local equilibrium has a diameter of at most 5. In this case, we get $k\text{-BE} = \text{BE}$ and the price of anarchy bound of [Dem+07] applies. \square

Lemma 5.18 ([Dem+07], Theorem 10). *For parameters $1 \leq \alpha < n^{1-\varepsilon}$, $\varepsilon \geq 1/\lg(n)$ and $k \geq 4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8$, it holds $k\text{-BE} = \text{BE}$ and the price of anarchy is at most $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8$.*

Proof. In Theorem 10 of [Dem+07], the authors use an inductive argument to find some agent u and a radius d such that the d -neighborhood of u contains more than $(n/2)$ -many agents. For this, they start with their Lemma 3 (for which only $k \geq 2$ must hold) and apply their Lemma 9 iteratively. They show that the maximal radius d , for which their Lemma 9 must be applied, is at most $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8$, which gives a first lower bound for k . Using this result, they apply their Corollary 7 to show that actually all agents are contained in a ball of radius $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 7$, for which they need the operation of creating an edge to an agent at distance $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8$, which is the second lower bound for k .

Using both results, they show that the diameter of every equilibrium is at most $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8$. By the choice of k , the same holds for k -local buy equilibria. We get $k\text{-BE} = \text{BE}$ and thus the price of anarchy is at most $4.667 \cdot 3^{\lceil 1/\varepsilon \rceil} + 8$. \square

Lemma 5.19 ([Dem+07], Theorem 12). *For parameters $1 \leq \alpha \leq 12n \log n$ and $k \geq 2 \cdot 5^{1+\sqrt{\lg n}} + 24 \lg(n) + 3$, it holds $k\text{-BE} = \text{BE}$ and the price of anarchy is at most $O\left(5^{\sqrt{\lg n}} \lg n\right)$.*

Proof. Similar to the proof of their Theorem 10 in [Dem+07], the authors provide a price of anarchy upper bound for a larger range of α : Again, they use an inductive argument to find an agent u and a radius d such that the d -neighborhood of u contains more than $(n/2)$ -many agents. For this, they start with looking at the number of agents in any radius $(12 \lg n)$ -neighborhood and then apply their Lemma 11 iteratively. They show that the maximal radius d , for which their Lemma 11 must be applied, is at most $5^{1+\sqrt{\lg n}}$, which gives a first lower bound for k . Using this result, they apply their Corollary 8 to show that actually all agents are contained in a specific ball, for which they need the operation of creating an edge to an agent at distance $2 \cdot 5^{1+\sqrt{\lg n}} + 24 \lg(n) + 3$, which is the second lower bound for k .

Using both, they show that in every equilibrium network there is an agent who contains all others in a ball of radius $(8 \cdot 5^{1+\sqrt{\lg n}} + 24 \lg(n) + 2)$. With the choice of k , the same holds for k -local buy equilibria and we get $k\text{-BE} = \text{BE}$ as well as a price of anarchy upper bound of $O(5^{\sqrt{\lg n}} \lg n)$. \square

Lemma 5.20 ([Alb+14], Theorem 3.6). *For parameters $12n \log n \leq \alpha$ and $2 \leq k$, it holds $k\text{-BE} = \text{BE}$ and the price of anarchy is $O(1)$.*

Proof. In [Alb+14], the authors provide a technical proof that characterizes equilibria for $\alpha \geq 12n \log n$. The main insight that is used for their bound is that there are different types of agents (see their Lemma 3.4, which uses their Lemma 3.2 and Lemma 3.3) with which they characterize equilibria and show that any buy equilibrium network with girth of at least $12 \cdot \lceil \log n \rceil$ has a diameter of less than $6 \cdot \lceil \log(n) \rceil$ and hence is a tree. In their Lemma 3.5, they prove that the considered big α values ensure a girth of at least $12 \lceil \log n \rceil$. The result of their Theorem 3.6 then comes from a comparison to the social optimum and gives a price of anarchy upper bound of at most 1.5.

Interestingly, in all used statements, there are only two statements concerning the creation or deletion of edges. For their Lemma 3.3, the operation of creating an edge to an agent in distance 2 is considered, and for their Lemma 3.5, the operation of deleting an edge is considered. Both operations are available with $k \geq 2$. Hence, for any $k \geq 2$, we have $k\text{-BE} = \text{BE}$ and the price of anarchy bound of 1.5 from [Alb+14] applies. \square

5.5.2 The Price of Anarchy

Our analysis of the price of anarchy focuses on diameter bounds for equilibrium networks. The reason for this is that using the following theorem that translates any diameter bound into an upper bound for the price of anarchy. This correspondence was first shown by Albers et al. [Alb+06] and again formulated by [Nis+07, Lemma 19.4], yet with a different proof. Specifically, the latter proof requires only the availability of 1-local edge deletion operations and hence applies without changes for both the greedy probing and the unrestricted probing k -local games.

Theorem 5.21 ([Nis+07], Lemma 19.4). *For any $k \geq 1$ and any edge price $\alpha \geq 2$, if a k -local greedy buy equilibrium network G has diameter D , then its social cost is at most $O(D)$ times the optimal social cost.*

In the following, we present several upper bounds on the diameter of equilibrium networks and then conclude the price of anarchy results in Theorem 5.28 by using Theorem 5.21. Note that most of the diameter bounds will be given for k -local greedy buy equilibria and, since $k\text{-BE} \subseteq k\text{-GBE}$, also apply directly to k -local buy equilibria, in which arbitrary operations are allowed. We start by providing a general result for tree network equilibria and then proceed with different bounds for respective ranges of the edge price α that concerns general networks. Finally, at the end of this section, we provide a notably non-constant lower bound for the price of anarchy in k -local buy games.

Corollary 5.22 (price of anarchy for tree networks). *For k -local greedy buy equilibrium tree networks with $2 \leq k \leq \log n$: $\text{PoA} = O(\log n)$.*

Proof. By Lemma 5.12, every tree network that is a k -local greedy buy equilibrium also is an asymmetric swap equilibrium (cf. Section 2.3). Therefore, we can apply the diameter upper bounds by Ehsani et al. [Ehs+15] and Mihalák and Schlegel [MS12], and get that every tree network equilibrium has a diameter of at most $O(\log n)$. Combining this with Theorem 5.21, the price of anarchy is at most $\text{PoA} = O(\log n)$. \square

For general networks, we next provide two different network diameter upper bounds. The first bound holds for any $k \geq 2$ and the second one gives improved results when the edge price is smaller than $n^{1-\varepsilon}$ for any constant $\varepsilon \geq \frac{1}{\log n}$.

Theorem 5.23. *Given a k -local greedy buy equilibrium network G with edge price α for $k \geq 2$, then it holds:*

$$\text{diam}(G) \leq \begin{cases} \frac{\alpha}{(k-1)} + k\frac{3}{2} + 1 & \text{for } k < 2\sqrt{\alpha}, \\ 2\sqrt{\alpha} & \text{for } k \geq 2\sqrt{\alpha}. \end{cases}$$

Proof. If $k \geq 2\sqrt{\alpha}$, then the classes of k -local greedy buy equilibria and buy equilibria coincide. This follows by Fabrikant et al. [Fab+03], who showed that no two agents can have a distance of more than $2\sqrt{\alpha}$, since otherwise one of these agents could buy an edge to the other one and decrease her private cost. Since k is large enough to allow any k -local greedy operation, for $k \geq 2\sqrt{\alpha}$ we get $\text{diam}(G) \leq 2\sqrt{\alpha}$.

Otherwise, if $k < 2\sqrt{\alpha}$, let u be an agent with maximal distance to any agent in G and let v be a most distant agent to u . Define $D := \text{diam}(G)$ and consider the distance cost improvement of u by creating an edge to the agent x at distance k on the shortest path from u to v . When creating this edge, u reduces her distance cost by $k - 1$ to each of the $D - k$ last agents on the path. Specifically, her total distance cost decrease is $\sum_{i=1}^{\lfloor k/2 \rfloor} (k - 2i + 1)$, given by the distance cost decrease to the $\lfloor k/2 \rfloor$ -many last agents on the same path from u to x , including x . Since G forms a k -local buy equilibrium, we have

$$\alpha \geq (k - 1)(D - k) + \sum_{i=1}^{\lfloor k/2 \rfloor} (k - 2i + 1),$$

which yields $D \leq \frac{\alpha}{k-1} + \frac{3k}{2} + 1$. □

In the following, we provide a more involved upper bound when the edge price is in range $1 \leq \alpha < n^{1-\varepsilon}$. For this, we modify an approach by Demaine et al. [Dem+07] and start with providing three lemmas that lower bound the number of agents in specific sized neighborhoods for k -local greedy buy equilibrium networks. Then, considering the locality parameter k , we look at the maximal neighborhoods for which these lower bounds still apply and present an estimation on the network's diameter.

First we restate Lemma 3 from [Dem+07], which holds for any $k \geq 2$, since only 2-local operations are considered.

Lemma 5.24 ([Dem+07], Lemma 3). *For any $k \geq 2$ and any k -local greedy buy equilibrium network G with $\alpha \geq 0$, it holds $|N_2(u)| > \frac{n}{2\alpha}$ for every agent $u \in V$.*

Lemma 5.25. *For $k \geq 6$, let G be a k -local greedy buy equilibrium network and $d \leq \frac{k}{3} - 1$ an integer. If there is a constant $\lambda > 0$ such that $|N_d(u)| > \lambda$ holds for every $u \in V$, then*

- (1) either $|N_{2d+3}(u)| > \frac{n}{2}$, for some agent $u \in V$,
- (2) or $|N_{3d+3}(v)| > \lambda \frac{n}{\alpha}$, for every agent $v \in V$.

Proof. The claim directly holds if there is a $u \in V$ with $|N_{2d+3}(u)| > n/2$. Hence, we assume the contrary and fix an arbitrary $u \in V$. Denote u 's $(2d+3)$ -neighborhood as $B := N_{2d+3}(u)$ and name the agents at a distance of exactly $(2d+3)$ as $\partial B := \{v \in V \mid d(u, v) = 2d+3\}$. We now greedily select a maximal subset $X \subseteq \partial B$ by the following iterative algorithm: (1) mark all agents of ∂B as unassigned, (2) while there is an unassigned agent x in ∂B , add x to X and create a new set ∂C_x containing x and all unassigned agents of ∂B within a distance of at most $2d$ to x , and mark these agents as assigned. Note that for the so-computed set X it holds that for any two $x, y \in X$ with $x \neq y$, we have $d_G(x, y) > 2d$.

Next, we lower bound the size of X by $|X| \geq n/\alpha$. For this, enumerate the elements of X with $x_1, \dots, x_{|X|}$ and define cluster sets C_{x_i} such that every C_{x_i} contains all elements of the corresponding ∂C_{x_i} . Further, for every agent $v \in V \setminus B$, we select an arbitrary shortest path from u to v and assign v to the cluster C_{x_i} that contains the (unique) agent on the path belonging to ∂B . By construction, we have $|\bigcup_{i=1}^{|X|} C_{x_i}| \geq n/2$. Now assume that u buys an edge to some $x \in X$, say to x_i . After this operation, the distance from u to every $v \in \partial C_{x_i}$ is at most $2d+1$ and thus the distance to any $w \in C_{x_i}$ decreases by at least 2. Since G forms an equilibrium and creating an edge to x_i is a k -local operation, we get $\alpha \geq 2 \cdot |C_{x_i}|$ for every $x_i \in X$. Hence, $|X| \cdot \alpha \geq 2 \sum_{i=1}^{|X|} |C_i| \geq 2n/2 = n$, i.e., $|X| \geq n/\alpha$.

By construction, for any $x, y \in X$ with $x \neq y$ we have $N_d(x) \cap N_d(y) = \emptyset$. With $|N_d(x)| > \lambda$, this gives $|\bigcup_{x \in X} N_d(x)| > |X| \cdot \lambda$. For every $x \in X$ we have $d_G(u, x) = 2d+3$ and hence, the maximal distance from u to any $v \in N_d(x)$ is at most $3d+3$. This gives $|N_{3d+3}(u)| \geq |\bigcup_{x \in X} N_d(x)| > |X| \cdot \lambda \geq \lambda n/\alpha$. \square

5 Limits of Locality

Lemma 5.26. *For $k \geq 4$, let G be a k -local greedy buy equilibrium network with $\alpha < n/2$ and $d \leq k/2 - 1$ an integer. If there is an agent $u \in V$ with $|N_d(u)| \geq n/2$, then $|N_{2d+1}(u)| \geq n$.*

Proof. We prove the contra-positive: Assume $|N_{2d+1}| < n$, then there is a $v \in V$ such that $d_G(u, v) = 2d + 2$. Since for all $x \in N_d(u)$ it holds $d_G(u, x) \leq d$, by the triangle inequality we get that $d_G(v, x) \geq d + 2$ for all $x \in N_d(u)$. Now consider v creating an edge to u , which reduces $\text{dist}_v(G)$ by at least $|N_d(u)|$. Since G forms an equilibrium, we get $n/2 > \alpha \geq |N_d(u)|$, which gives the claim. \square

Combining the three previous lemmas, we can now estimate the maximal network diameter essentially by using how fast the number of agents increase if we look at increasingly bigger neighborhood ranges around a fixed agent.

Theorem 5.27. *For $k \geq 6$, $n \geq 4$, and $1 \leq \alpha \leq n^{1-\varepsilon}$ with $\varepsilon \geq 1/\log n$, the maximal diameter of any k -local greedy buy equilibrium network is $O(n^{1-\varepsilon(\log(k-3)-1)})$.*

Proof. Let G be a k -local greedy buy equilibrium network. We define a sequence $(a_i)_{i \in \mathbb{N}}$ by $a_1 := 2$ and for any $i \geq 2$ with $a_i := 3a_{i-1} + 3$. We want to apply Lemma 5.25 iteratively with $\lambda_i := (n/\alpha)^i/2$. Lemma 5.24 ensures that with $|N_2(v)| > n/(2\alpha) = \lambda_1$ for all $v \in V$, we have a start for this.

Let m be the highest sequence index with $a_m \leq k/4 - 3$. If there is a $j \leq m$ such that case (1) of Lemma 5.25 applies, then there is an agent $u \in V$ with $|N_{2a_m+3}(u)| > n/2$. For $\alpha \leq n^{1-\varepsilon} < n/2$ we get with Lemma 5.26 that $|N_{4a_m+7}| \geq n$ holds and hence the diameter is at most $a_m < k$. Otherwise, case (2) applies for all $i \leq m$ and we know that for every $v \in V$ it holds $|N_{3a_m+3}(v)| > (n/\alpha)^{m-1}/2$. Using $a_i = \frac{7}{6}3^i - 3/2$ and $a_m \leq k/4 - 3$, we get $m \geq \log(k-3) - 1$.

Let D be the diameter of G and P a longest shortest path. We define a set C by selecting the first agent of P as c_1 and then along the path selecting every further agent with a distance of $2k$ to the last previously selected agent. Now consider the operation of c_1 creating an edge to an agent at distance k in the direction of c_2 . Using $k \geq 3a_m + 3$, $|N_{3a_m+3}(c)| > \frac{(n/\alpha)^{m-1}}{2}$ for all $c \in C$ and that G is an equilibrium:

$$\alpha \geq (k-1)(|C|-1)(n/\alpha)^{\log(k-3)-2}/2 \geq \frac{k-1}{2} \left(\frac{D}{2k} - 1 \right) (n/\alpha)^{\log(k-3)-2}$$

This gives:

$$D \leq \frac{4k}{k-1} \alpha \left(\frac{\alpha}{n}\right)^{\log(k-3)-2} + 2k \leq 5 \frac{\alpha^{\log(k-3)-1}}{n^{\log(k-3)-2}} + 2k \leq 5n^{1-\varepsilon(\log(k-3)-1)} + 2k$$

By using Lemma 5.18, we get the claimed diameter upper bound for any $k \geq 6$. \square

The following theorem summarizes the results of this section, which we derive by applying Theorem 5.21 to the diameter upper bounds.

Theorem 5.28. *For k -local greedy buy equilibria, the price of anarchy is:*

$$\text{PoA} = \begin{cases} \Theta(n) & \text{for } k = 1, \\ O(\min\{\alpha/k + k, \log n\}) & \text{for } k \geq 2, \\ O(n^{1-\varepsilon(\log(k-3)-1)}) & \text{for } k \geq 6 \wedge n \geq 4 \wedge \alpha \in [1, n^{1-\varepsilon}], \\ O(1) & \text{for } k \geq 6 \wedge n \geq 4 \wedge \alpha \in [1, n^{1-\frac{1}{\log(k-3)-1}}], \\ O(\sqrt{\alpha}) & \text{for } k \geq 2\sqrt{\alpha}, \\ O(\log n) & \text{for any } k \geq 1, \text{ if tree network.} \end{cases}$$

Finally, we show a non-constant lower bound for the price of anarchy when the edge price is at least $\alpha \geq (k-1)n$. Note that for the classic buy game no such bound is known and with respect to the known results there, it can only exist – if at all – for α close to n . Hence, this lower bound states that even with the most optimistic model of locality, there still exists a fundamental difference in the efficiency of the local versus the non-local buy games. Our lower bound is now tight for tree networks when $k = O(1)$ and tight in general when $k = \Omega(\log n)$. The latter can be seen immediately, since the lower bound simplifies to $\Omega(1)$, which matches the results of Fabrikant et al. [Fab+03] for $\alpha \geq n \log n$, where they provide an upper bound of $O(1)$.

Theorem 5.29 (price of anarchy lower bound). *For k -local buy equilibrium networks with $2 \leq k \leq \log n$ and edge price $\alpha \geq (k-1)n$ it holds: $\text{PoA} = \Omega\left(\frac{n \log n}{\alpha}\right)$. Specifically for $\alpha = kn$ the price of anarchy is:*

$$\text{PoA} = \Omega\left(\frac{\log n}{k}\right)$$

Proof. For any $2 \leq k \leq d$, we claim that the strategy profile S of a complete binary tree network T_d of depth d , where every edge is owned by the incident agent who is closer to the root, is in k -local buy equilibrium. Independent of the edge price α , no agent in T_d can delete or swap edges to improve her private cost. Thus, we only have to choose α high enough such that no agent can improve her private cost by creating any number of edges in her k -neighborhood. The creation of an edge within a k -neighborhood cannot decrease any distance by more than $k - 1$. Thus, if $\alpha \geq (k - 1)n$, then no agent can create one single edge to decrease her private cost. Moreover, by creating more than one edge, no agent u can decrease her distance cost by more than $(k - 1)n$, which implies that S is a k -local buy equilibrium for $\alpha \geq (k - 1)n$.

We consider the social cost ratio of S and the spanning star strategy profile S_{Opt} on $n = 2^{d+1} - 1$ agents. Since $\alpha \geq 2$, S_{Opt} is the social cost minimizing network. It holds $\text{cost}(S) \geq \alpha(n-1) + n \cdot \text{dist}_r(S)$, since the root r has a minimum distance cost among all agents in T_d . Further note that $\text{dist}_r(S) > \frac{n}{4} \log n$, since r has at least $(n/2)$ -many agents at distance $\frac{\log n}{2}$. In total we get:

$$\text{PoA} \geq \frac{\text{cost}(S)}{\text{cost}(S_{\text{Opt}})} \geq \frac{\frac{n^2}{4} \log n + (n-1)\alpha}{(n-1)\alpha + 2(n-1)^2} > \frac{n^2 \log n}{4(n-1)\alpha} = \Omega\left(\frac{n \log n}{\alpha}\right).$$

□

5.6 Conclusion & Future Work

In this chapter, we introduced a new model that provides a realistic locality notion for network creation games. Our results show that strategy probing, even when restricted to only greedy strategy changes, suffices to overcome most quality regressions caused by the knowledge limitations of a k -neighborhood locality. We can see this by comparing our price of anarchy upper bounds to the respective (much higher) lower bounds that are provided by Bilò et al. [Bil+14a] in their worst-case locality model. Specifically, we see that our upper bounds for the price of anarchy are close to the non-local model and hence agents can still create socially efficient networks. However, our non-constant lower bound on the price of anarchy and our negative results concerning the approximation quality of non-local strategies by local strategies show that

the locality constraints still do have a significant impact on the game. Facing our negative results about the hardness of computing best responses and the non-convergence of best-response strategy changes, we provide further evidence that those questions are intrinsically hard and effectively independent of locality assumptions.

Obviously, the proposed locality model is interesting not only to the studied network creation game variant, but also to various other games. An interesting area for further research seems to be social network games. Like for the game discussed in this chapter, probing only a few strategies in only one's circle of friends is a very natural behavior. Closely linked to this is the interesting question of how many probes suffice to achieve reasonable good social quality results: We know that we can obtain networks with a very high social cost if we prohibit any probing and we know that only with n^2 probes we always achieve networks with a reasonable low social cost. It remains an interesting question what the actual threshold for this behavior is.

Multilevel Network Games

LAYERS provide the architectural basis of most computer networks. Already the OSI reference model (cf. Zimmermann [Zim80]) specified how layers should be used to gain a modular structure for the Internet and by this laid the architectural foundations of many modern computer networks. Consequently, today this architecture is present all over in the design of networks and their communication protocols. The general idea of a layered system is to provide service-specific protocol layers that stack onto each other. Each layer can access the layer below, in some architectures also several layers below, and provides services for the layers on top. At the bottom layer, we have the physical network, at which every operation of a higher layer must be reflected eventually.

In this chapter, we study the interaction of two communication layers in such a layered system: One layer provides general purpose connections, the other one is a high-speed layer that allows agents to improve their communication distances. Unlike in most previous research, we take a game theoretical view on the availability of such a high-speed layer and ask about its influence to the network's total efficiency when faced with selfishly acting agents. Specifically, we consider agents as rational actors who individually decide if they want to connect to the high-speed layer for a fixed price of α or not, depending only on

their private costs.

The availability of such high-speed networks is motivated by various observations. Foremost, techniques as discussed in Chapter 4 for individual connections also allow offering access to a whole high-speed network and not only for single point-to-point connections. A technical different, yet from a theoretical standpoint still similar scenario, is the use of an additional logical overlay network. Similar to the way overlays are used for search overlays (cf. survey by Androutsellis-Theotokis and Spinellis [AS04]), they can provide better routing information (e.g., larger routing tables or addresses to more likely communication partners in case of non-uniform communication interests) for the shortest path communications to other agents. This means, a logical network can also drastically reduce communication costs of the individual agents by providing such routing information.

For using the high-speed layer, we consider two different access models in accordance with the two named motivations. On the one hand, we see the high-speed network as an additional network to which connections have to be created in order to enter or leave it. This is the same concept as one can find in physical networks, for example, which are connected via hardware routers. On the other hand, when looking at multilevel games that originate from quality-of-service agreements like in Chapter 4, it is reasonable that only the access to the high-speed layer raises cost, but switching back to the general purpose layer is allowed everywhere. In our games, we call the first connection model *bidirectional* and the second one *unidirectional* and further denote the connection points between the layers as *gateways*.

Chapter Basis. The model, analysis, and results presented in the remainder of this chapter are based on the following publication:

2014 (with S. Abshoff, D. Jung and A. Skopalik). “Multilevel Network Games”. In: *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings*, cf. [Abs+14].

Chapter Outline. In Section 6.1, we introduce a basic model for multilevel networks. This model splits into two variants: a game variant in which layers

can be switched only at gateway agents and a variant in which just entering the high-speed layer requires gateways. The more specific model descriptions for both variants are then provided in Section 6.3 and Section 6.4, alongside the respective analysis. An overview of our results and a comparison with other models is given in Section 6.2. The chapter concludes with an outlook and a summary of open questions.

6.1 Model & Preliminaries

A multilevel network game (V, L_1, L_2) consists of n agents V who are connected via two network layers L_1 and L_2 . In layer L_1 , the agents form a bidirectional connected graph (V, L_1) and each edge has a length of 1. The second layer L_2 is a supporting high-speed layer, which can be used to improve the agents' communication costs. Thus, the agents are present in both network layers, however, the access to the second layer must be enabled specifically.

We denote the *distance* between two agents $u, v \in V$ in layer L_1 by $d_1(u, v)$, which indicates the shortest path distance in graph (V, L_1) . Likewise, $d_2(u, v)$ denotes the shortest path distance in (V, L_2) . The maximal distance for any pair of agents in (V, L_1) is given by $\text{diam}(L_1) := \max_{u, v \in V} d_1(u, v)$, respectively by $\text{diam}(L_2)$ for (V, L_2) . Agents are able to use the high-speed layer only at *gateway agents*, which means that a path may switch from layer L_1 to layer L_2 . Hence, gateways function as connections between the two layers. In the following, we will study games with two fundamentally different variants of gateways:

Multilevel games with bidirectional gateways: A gateway at agent u forms a bidirectional edge of length 0 between agent u in (V, L_1) and agent u in (V, L_2) . There are no other connections between the layers other than the gateways and thus switching between the layers is only possible at gateway agents.

Multilevel games with unidirectional gateways: At every agent u , there is a unidirectional edge of length 0 from u in (V, L_2) to u in (V, L_1) and thus, switching from layer L_2 to layer L_1 is allowed at every agent. If an agent is a gateway, then this connection from L_2 to L_1 is bidirectional and thus, gateways allow switching from L_1 to L_2 .

Considering a gateway set S , the *communication distance* $\delta_S(u, v)$ constitutes the actual distance between two agents u and v by making use of both layers. Although the shortest path is measured by using both layers, we use the convention that the end points of the path must be the respective agents in layer (V, L_1) . Note that we will omit the index S if it is clear from context. Given an agent u and a range $k > 0$, then $B_k(u)$ denotes the set of all agents within a communication distance of at most k to u .

In our multilevel network game, agents can decide selfishly if they want to become a gateway or not. Being a gateway means that the agent pays a fixed price $\alpha > 0$ and establishes the above mentioned connection between the two network layers. We call the set of gateways S and identify it with the current strategy profile. Agents in $V \setminus S$ are called non-gateways. Analog to network creation games, the decision of becoming a gateway or not is based on the private cost function of an agent. In the *Sum-Layer-Game*, the private cost of an agent u is:

$$c_u(S) := \alpha \cdot |S \cap \{u\}| + \sum_{v \in V} \delta_S(u, v)$$

For the *Max-Layer-Game*, the private cost function is:

$$c_u(S) := \alpha \cdot |S \cap \{u\}| + \max_{v \in V} \delta_S(u, v)$$

For both games, the social costs are given by $\text{cost}(S) := \sum_{u \in V} c_u(S)$.

If an agent improves her private cost by changing her strategy from non-gateway to gateway or vice versa, we call this an *improving response*. For an improving response where an agent u changes her strategy to be a gateway, we say that u *opens*. Analogously, we say u *closes* if she changes her strategy from gateway to non-gateway. We call a strategy profile S a (pure) Nash equilibrium, or simply an *equilibrium*, if no agent can perform an improving response. For the convergence analysis of improving-response processes¹, we ask whether the games provide the *finite improvement property* or (lesser) whether they are *weakly acyclic* (cf. Section 2.2.3).

¹Since the strategy space of any agent contains only two possible choices, in this game every improving response is also a best response.

6.2 Related Work & Contribution

The outcome of the individual strategic connect and disconnect decisions of the network's participants is a key issue in network creation games, as it was discussed in the previous chapters: How good can such an outcome be? How bad is it at most? And is it likely that the agents will ever reach an equilibrium state despite their uncoordinated behavior? – Considering our multilevel network games, these questions still apply in order to understand and quantify the effects of individual strategic decision making.

Regarding the agents' behaviors, multilevel games are actually very similar to network creation games. Specifically, both have the property in common that strategy changes result in changes of the network's topology. However, the substantial difference is the size of the agents' strategy spaces: For multilevel network games, a single agent has only two possible choices, compared to 2^{n-1} options previously. This means that the decision whether to use an improving network layer or not is much more drastic than before in the classic network creation games as agents cannot make any fine-grained decision like connecting for a smaller cost to only specific areas.

Despite its importance, the question of strategic decision making in multilevel networks is barely studied so far. When leaving out the strategic behavior of agents but using random processes to model their actions, the effects of network interactions in complex multilevel networks received various considerations, for example the interaction between a physical layer and a congestion flow by Kurant and Thiran [KT06]. However, such an approach misses the effects of strategic behavior, which is already present when one agent decides against being a gateway in favor of free-riding via her neighbor's high-speed connection.

With a focus on network formation, Shahrivar and Sundaram [SS13] considered a multilevel network game of centralized, strategically acting designers. Together with their follow-up paper [SS15], they provide the only contributions in this field with a game theoretic view. In their games, multiple network designers simultaneously construct networks and the overall efficiency of a designer also depends on the network layers provided by the other designers. Yet, compared to our model, the individual decision making of the agents was not considered.

Contribution. In this chapter, we introduce a new model for analyzing the effects of strategic decision making in multilevel networks. Our model is the first one that captures the effects of individual agents being strategic actors in a multilevel context, namely agents of a general purpose network who can utilize a high-speed layer. Depending on how the general purpose and the high-speed layers interact with each other, we gain two qualitatively different networks games for which we apply the classic sum and maximum private cost functions.

Considering the game with bidirectional gateways, we show that computing the optimal placement of gateways is \mathcal{NP} -hard for both variants of private cost functions. For the Sum-Layer-Game, we show that for $\alpha \leq n - 1$ and $\alpha > n(n - 1)$ equilibria always exist and that then the price of anarchy is $\Theta(1 + n/\sqrt{\alpha})$; for $\alpha \in (n - 1, n(n - 1))$, we upper bound the price of anarchy by $O(\sqrt{\alpha})$. For the Max-Layer-Game, we show that equilibria always exist if the networks are trees or if the girth is not too small. We further provide a price of anarchy bound of 1, for $\alpha < 1$, and otherwise the tight bound of $\Theta(1 + n/\sqrt{\alpha})$. Concerning the dynamics, both the Sum-Layer-Game and the Max-Layer-Game are no potential games, whereas the Sum-Layer-Game is not even weakly acyclic.

Regarding the game with unidirectional gateways, in the Sum-Layer-Game the price of anarchy is at most $O\left(\frac{1}{1-\mu} + \frac{\alpha}{n(1-\mu)^2}\right)$, whereas $\mu \in (0, 1)$ is the improvement factor of the high-speed layer. In the Max-Layer-Game, we provide an algorithm to compute equilibria for tree networks, when the L_2 -layer provides some exact improvement property. For the general case, we show that in this game the price of anarchy is at most $O(\alpha/(1 - \mu)^2)$. Complementing this upper bound, we also provide a high lower bound of $\Omega(\sqrt{n})$ for certain parameters of α .

6.3 Bidirectional Gateways

In this section, we analyze the multilevel network game (V, L_1, L_2) with bidirectional gateways. The high-speed layer L_2 is assumed to provide negligible short connections between all agents. In our sense, this means that every distance is shorter than 1 divided by the number of agents. Without loss of

generality, we can assume then that all distances in layer L_2 have a length of 0. Consequently, for the remainder of this section, we will omit specifications of the L_2 -layer network and use only (V, L_1) to state a game instance. Since the distance between any two gateways is 0, the communication distance for $u, v \in V$ resolves to:

$$\delta_S(u, v) = \min\{d_1(u, v), d_1(u, S) + d_1(S, v)\}$$

Here, $d_1(u, S)$ denotes the shortest path distance from agent u to any gateway. Throughout this section, we further require that one gateway must always be left in the game. Thus, a last gateway is not allowed to close even if that would be an improving response for her. It is easy to see that otherwise $S = \emptyset$ would form an equilibrium for any game instance, since then no agent could improve her cost by a unilateral strategy change.

6.3.1 The Sum-Layer-Game

We start our study with the Sum-Layer-Game. First, we ask the difficulty of computing a gateway set that minimizes the social cost. Note that this set is not required to be an equilibrium.

Theorem 6.1. *For the Sum-Layer-Game with bidirectional gateways, the computation of a gateway set that minimizes the social cost is \mathcal{NP} -hard.*

Proof. Let (V, L_1) be an instance of the Sum-Layer-Game. For two parameters $n, m > 4$, let there be a set of m elements $X := \{x_1, \dots, x_m\}$ and further n subsets $S_1, \dots, S_n \subseteq X$ of this element set. Then, the \mathcal{NP} -complete SET-COVER problem (cf. Karp [Kar72]) is the task to compute a minimal number of subsets that together contain all elements of X . Given such a SET-COVER instance, we construct an instance (V, L_1) of the Sum-Layer-Game as follows (cf. Figure 6.1): First, we create a clique C of k agents and mark one of its agents as c . For every set S_i , we create a corresponding agent S_i and connect her to c . For every element $x_i \in X$, we create w -many agents x_i^1, \dots, x_i^w and connect all x_i^j , for $i = 1, \dots, m$ and $j = 1, \dots, w$, to all set agents S_l with $x_i \in S_l$. Using the parameters $w := n, k := m - 1$, and $\alpha := 4n(m - 1)$, in the following we show that an optimal placement of gateways corresponds to a solution of the SET-COVER problem.

For now, assume that c is a gateway agent in the optimal solution S_{Opt} (we will prove this claim later). We claim that then no other clique agent $v \in C \setminus \{c\}$ is a gateway. For this, assume that l further clique agents are open and compute the social cost decrease by closing all clique agents except agent c . The decrease is at least $l\alpha - 2l(\omega m + n) - l(l + 1) > 0$ and hence c is the only agent in $C \cap S_{\text{Opt}}$.

Next, for an element x_i consider the corresponding element agents x_i^1, \dots, x_i^w and a set S_j such that $x_i \in S_j$. If there is any $x_i^l \in S$ and $S_j \notin S$, closing x_i^l and opening S_j does not increase the social cost. Hence, we can assume that in S_{Opt} there is no closed set agent with an open element agent. Now, let S_j be an open set agent and assume that for $x_i \in S_j$ there are l open element agents. Closing all of these element agents reduces the social cost by at least $l\alpha - 2l(k + n + 2(l - 1) + (\omega - l) + (m - 1)\omega) = l\alpha - 2l(\omega m + n + k + l - 2) > 0$ and hence in S_{Opt} all are closed. Given a set of closed element agents x_i^1, \dots, x_i^w such that for all S_j with $x_i \in S_j$ the set agents are closed, opening S_j reduces the social cost by at least $2(k\omega + (m - 1)\omega + (n - 1)) - \alpha > 0$. Contrarily, opening a set agent whose element agents are already completely covered increases the social cost by at least $\alpha - 2(k + m\omega + n - 1) > 0$.

Finally, we can see that c actually has to be a gateway in S_{Opt} . For this, consider an arbitrary optimal setting with all clique agents closed (if one clique agent is open, we can close it and open c without increasing the social cost). When opening c , we know that without increasing the social cost we can close all element agents and open corresponding set agents. Hence, when opening c we can assume that all element agents are closed and that for each element agent a corresponding set is open. This gives a social cost decrease by opening c of at least $2km\omega - \alpha > 0$.

Hence, the socially optimal solution S_{Opt} is given by a gateway agent c and a minimal number of set agents such that all element agents are covered. \square

We now study the existence of equilibrium networks. Given a Sum-Layer-Game with a moderately small or alternatively very high connection price, we show that equilibria always exist.

Proposition 6.2. *Given a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways, connection price $\alpha \leq n - 1$ or $\alpha > n \cdot \text{diam}(L_1)$, then an equilibrium setting exists.*

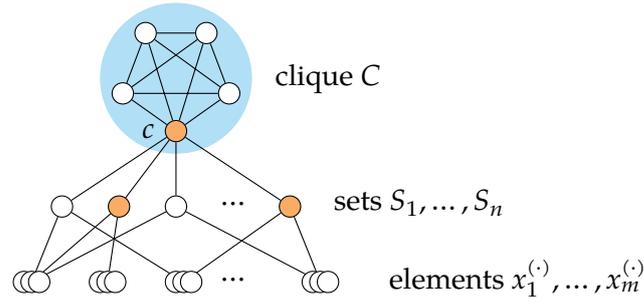


Figure 6.1: Illustration of the \mathcal{NP} -hardness reduction from SET-COVER to optimal gateway placement.

Proof. For $\alpha \leq n - 1$, consider the strategy profile $S := V$ in which every agent has a private cost of α . If any gateway closes in this setting, her distance cost would become at least $n - 1$. This cannot be an improving response and hence $S = V$ is an equilibrium.

For $\alpha > n \cdot \text{diam}(L_1)$, consider an arbitrary setting with $|S| = 1$. Assuming a second agent would open, then her distance cost decreased by not more than $n \cdot \text{diam}(L_1) < \alpha$ and hence this cannot be an improving response. \square

Proposition 6.3. *Given a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways and a connection price $\alpha \leq n - 1$. Then, $S = V$ minimizes the social cost and the price of stability is 1.*

Proof. Let S be a socially optimal solution and assume that there are m closed agents. When opening all of them, then for any gateway $v \in S$ the distances to all these m agents reduce by at least 1 each, while for $u \in V \setminus S$ the distances reduce by at least $n - 1$ each. Hence, setting the strategy profile to $S = V$ changes the social cost by $m\alpha - ((n - m) + m(n - 1)) < 0$. This holds for any setting with fewer than n gateways and thus, it is the socially optimal solution. Since $S = V$ is also in equilibrium, the price of stability is 1. \square

Note that Proposition 6.3 does not contradict the \mathcal{NP} -hardness proof of Theorem 6.1, since in that proof the connection price α was chosen to be bigger than the number of agents.

Convergence Properties

In the following, we want to understand how combinations of the L_1 -layer and the connection price influence the convergence of improving-response dynamics. We start with several negative convergence results, which state that for a wide range of connection prices the Sum-Layer-Game is no potential game, since it does not have the finite improvement property. This holds for $\alpha \in (4, n - 1)$ as well as for $\alpha \in \left(\frac{3}{32}n^2 + n, \frac{5}{32}n^2\right)$. Surprisingly, for specific game instances we can further show that the game is not even weakly acyclic, which means that there are improving-response cycles that can never terminate.

Proposition 6.4. *In general, the Sum-Layer-Game of $n > 7$ agents with bidirectional gateways and a connection price $\alpha \in (4, n - 1)$ has not the finite improvement property.*

Proof. We construct a game instance (V, L_1) as depicted in Figure 6.2 (with $c := 1$): First, we create a path (u, v, w) of three agents and then connect additional r -many agents to agent w , as well as additional $(n - r - 3)$ -many agents to agent u ; here the parameter r will be computed below. Starting with only w being a gateway, we specify the constraints under which u and v form an improving-response cycle:

- I: u opens if $\alpha < 2r + 2$.
- II: v opens if $(n - 3 - r) + r + 2 = n - 1 > \alpha$.
- III: u closes if $\alpha > r + 2$.
- IV: v closes if $\alpha > r + 1$.

Combining these conditions, we get $r + 2 < \alpha < \min\{n - 1, 2r + 2\}$. For $2 \leq r \leq n - 3$, the interval $(r + 2, \min\{n - 1, 2r + 2\})$ is non-empty and thus for $4 < \alpha < n - 1$ the game admits an infinite improving-response cycle. \square

Proposition 6.5. *In general, the Sum-Layer-Game of $n > 16$ agents with bidirectional gateways and a connection price $\alpha \in \left(\frac{3}{32}n^2 + n, \frac{5}{32}n^2\right)$ has not the finite improvement property.*

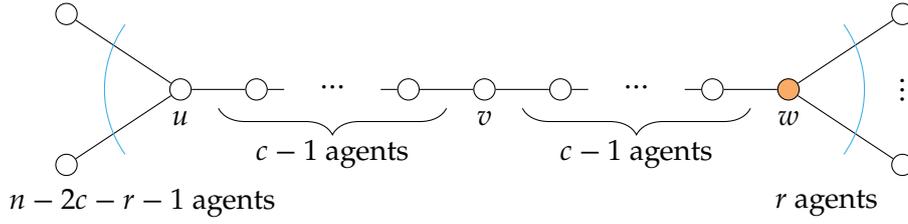


Figure 6.2: Improving-response cycle where agents u and v perform improving responses in turn.

Proof. We construct a game instance (V, L_1) as depicted in Figure 6.2. First, we create a path (u, \dots, v, \dots, w) with $(c - 1)$ -many agents between agents u and v , as well as the same number of agents between v and w . Then, we connect additional r -many agents to w and additional $(n - 2c - r - 1)$ -many agents to u ; here, the parameter r will be computed below and we set $c := n/4$. Starting with only agent w being a gateway, under the following constraints u and v form an improving-response cycle:

$$\text{I: } u \text{ opens if } \alpha < \sum_{i=1}^c 2i + 2rc.$$

$$\text{II: } v \text{ opens if } \alpha < 2 \sum_{i=1}^{\lfloor c/2 \rfloor} 2i + (n - 2c - 1)c.$$

$$\text{III: } u \text{ closes if } \alpha > \sum_{i=1}^{\lfloor c/2 \rfloor} 2i + (r + c + 1)c.$$

$$\text{IV: } v \text{ closes if } \alpha > \sum_{i=1}^{\lfloor c/2 \rfloor} 2i + (r + 1)c.$$

To simplify calculations, we assume n to be a multiple of 4. Since constraint III implies constraint IV, it suffices to consider:

$$\alpha < c^2 + (2r + 1)c \tag{6.1}$$

$$\alpha < -\frac{3}{2}c^2 + nc \tag{6.2}$$

$$\alpha > \frac{5}{4}c^2 + \left(r + \frac{3}{2}\right)c \tag{6.3}$$

Combining (6.1) and (6.3) gives $r \in \left(\frac{1}{2}\left(\frac{\alpha}{c} - c - 1\right), \frac{\alpha}{c} - \frac{5}{4}c - \frac{3}{2}\right)$ as a valid range for r . Plugging in $c = n/4$ gives $r \in \left(\frac{2\alpha}{n} - \frac{n}{8} - \frac{1}{2}, \frac{4\alpha}{n} - \frac{5n}{16} - \frac{3}{2}\right)$: i.e., the interval of valid values for r has a length of $2\alpha/n - 3n/16 - 1$. To ensure that there exist integral solutions for r , we require the interval to have a length of

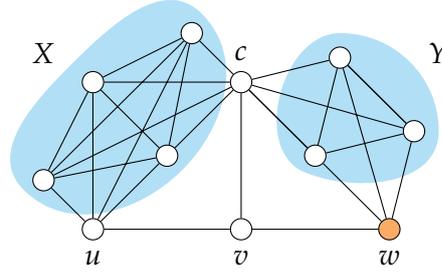


Figure 6.3: Improving-response cycle for the Sum-Layer-Game with $\alpha := 7$. Starting with w being a gateway, agents u and v change their strategies in turn and are the only agents who can perform an improving response.

at least 1, i.e., $2\alpha/n - 3n/16 - 1 \geq 1$, which gives $\alpha \geq n + 3n^2/32$. Considering (6.2), which is $\alpha \leq \frac{5}{32}n^2$, we get $\alpha \in \left(\frac{3}{32}n^2 + n, \frac{5}{32}n^2\right)$ as the permitted range. For $n > 16$, this interval is non-empty and thus agents u and v form the above mentioned infinite improving-response cycle. \square

Theorem 6.6. *The Sum-Layer-Game with bidirectional gateways is not a weakly acyclic game.*

Proof. For $\alpha := 7$ we consider the L_1 -layer as depicted in Figure 6.3. The layer consists of three agents u , v , and w , which are connected as a line. Additionally, we create a clique X of $\lfloor \alpha/2 \rfloor$ agents, a clique Y of $\lfloor \alpha/2 \rfloor$ agents, and a center agent c . All agents of X are connected to agents c and u , all agents of Y are connected to agents c and w , and furthermore agent c is connected to v .

We consider the initial strategy profile $S = \{w\}$ and argue that there exists a unique sequence of improving responses, such that u and v change their strategies in turn. Table 6.1 states that there is always exactly one of these two agents, who can improve her private cost. Note that we explicitly use $\alpha = 7$. \square

With these negative convergence results in mind, we next study some specific game properties that still guarantee convergence despite the general results.

Proposition 6.7. *The Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways is a potential game if the connection price is $\alpha < 1$ or $\alpha > n \cdot \text{diam}(L_1)$.*

Proof. If $\alpha < 1$, then for $S \neq V$ there is a non-gateway $v \in V \setminus S$ who can perform an improving response by opening. Also, no gateway $u \in S$ will deviate from

Table 6.1: Calculation of improving responses in Theorem 6.6. At each time only one improving response is possible, resulting again in the initial strategy profile after four operations.

(a) agent u opens:

	Cost if opened	Cost if closed	State after
$x \in X$	$2\alpha + 2$	$\alpha + \lfloor \alpha/2 \rfloor + 6$	closed
$y \in Y$	$2\alpha + \lceil \alpha/2 \rceil + 1$	$\alpha + \lceil \alpha/2 \rceil + 6$	closed
u	$2\alpha + 1$	$\alpha + 2\lfloor \alpha/2 \rfloor + 5$	opening
v	$2\alpha + \lceil \alpha/2 \rceil + 2$	$2\alpha + 3$	closed
w	$2\alpha + 2\lceil \alpha/2 \rceil + 5$	$\alpha + 2\lceil \alpha/2 \rceil + 5$	opened
c	$2\alpha + 5$	$\alpha + 5$	closed

(b) agent v opens:

	Cost if opened	Cost if closed	State after
$x \in X$	$2\alpha + 1$	$\alpha + \lfloor \alpha/2 \rfloor + 1$	closed
$y \in Y$	$2\alpha + 1$	$\alpha + \lceil \alpha/2 \rceil + 4$	closed
u	$2\alpha + 1$	$\alpha + 2\lfloor \alpha/2 \rfloor + 5$	opened
v	$2\alpha + 1$	$2\alpha + 3$	opening
w	$2\alpha + 3$	$\alpha + 2\lceil \alpha/2 \rceil + 5$	opened
c	$2\alpha + 1$	$2\alpha + 5$	closed

(c) agent u closes:

	Cost if opened	Cost if closed	State after
$x \in X$	2α	$\alpha + 3$	closed
$y \in Y$	2α	$\alpha + 3$	closed
u	$2\alpha + 1$	$\alpha + \lfloor \alpha/2 \rfloor + 4$	closing
v	$2\alpha + 1$	$2\alpha + 3$	opened
w	$2\alpha + 1$	$\alpha + \lceil \alpha/2 \rceil + 4$	opened
c	2α	$\alpha + 3$	closed

(d) agent v closes:

	Cost if opened	Cost if closed	State after
$x \in X$	$2\alpha + \lfloor \alpha/2 \rfloor + 1$	$\alpha + \lfloor \alpha/2 \rfloor + 2$	closed
$y \in Y$	$2\alpha + \lceil \alpha/2 \rceil + 1$	$\alpha + \lceil \alpha/2 \rceil + 4$	closed
u	$2\alpha + 1$	$\alpha + \lfloor \alpha/2 \rfloor + 4$	closed
v	$2\alpha + \lceil \alpha/2 \rceil + 2$	$2\alpha + 3$	closing
w	$2\alpha + \lceil \alpha/2 \rceil + 2$	$\alpha + 2\lceil \alpha/2 \rceil + 5$	opened
c	$2\alpha + 1$	$\alpha + 4$	closed

her current strategy and close. Hence, after at most $n - 1$ improving responses, the strategy profile is $S = V$. Otherwise, if $\alpha > n(n - 1)$, no non-gateway will open and for every gateway it is an improving response to close. \square

Proposition 6.8. *Given a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways. If $\text{diam}(L_1) > 2\alpha + 1$ with $\alpha \in [4, n - 1]$ and initially only one gateway is open, then there exists a sequence of improving responses such that the game converges to an equilibrium.*

Proof. Let $x \in S$ be the initial gateway and consider u and v being two agents with $d_1(u, v) > 2\alpha + 1$. One of these agents (say v) must have a distance greater than $\alpha + 1$ to agent x . By opening, v reduces her distances to at least half of the agents on the shortest path to x . This means, her distance cost decreases by:

$$\sum_{i=1}^{\lceil \alpha/2 \rceil} (2i - 1) = \left\lceil \frac{\alpha}{2} \right\rceil \left(\left\lceil \frac{\alpha}{2} \right\rceil + 1 \right) - \left\lceil \frac{\alpha}{2} \right\rceil > \alpha$$

Next, with $S = \{x, v\}$, also agent u wants to open, since opening reduces her distances to at least half of the agents on a shortest path from u to v , i.e., to $\lceil \alpha \rceil$ -many agents. Considering the agents on the shortest path from u to v , for each of them it is an improving response to open, since opening improves the distances to at least $\lceil \alpha \rceil$ -many agents. Therefore, starting from one end of the path we can open them iteratively and each time it is an improving response for the respective agent. Finally, with $|S| > \alpha$, all other agents also want to open and we reach $S = V$, which is an equilibrium. \square

Price of Anarchy

In the following, we consider the price of anarchy in the Sum-Layer-Game with bidirectional gateways. The next theorem combines all the results that will be proven in this section.

Theorem 6.9. *In a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional*

gateways, the price of anarchy is:

$$\text{PoA} = \begin{cases} 1 & \text{for } \alpha \in (0, 1), \\ \Theta(n/\sqrt{\alpha}) & \text{for } \alpha \in [1, n-1], \\ O(\sqrt{\alpha}) & \text{for } \alpha \in (n-1, n(n-1)), \\ 1 & \text{for } \alpha \geq n(n-1). \end{cases}$$

This theorem directly follows from the following lemmas.

Lemma 6.10. *In the Sum-Layer-Game with bidirectional gateways, for $0 < \alpha < 1$ the price of anarchy is 1.*

Proof. Given a game instance (V, L_1) , then by Proposition 6.3 we know that the social optimum is $V = S$. Since for $\alpha < 1$ opening is an improving response for every non-gateway, this is also the only equilibrium. \square

Lemma 6.11. *In a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways, for $1 \leq \alpha < 2$ the price of anarchy is $\Theta(n/\sqrt{\alpha})$.*

Proof. If $\text{diam}(L_1) \geq 2$, then all agents will open and constitute a socially optimal solution. Otherwise, with $\text{diam}(L_1) < 2$ the network (V, L_1) forms a clique and the only possible equilibria are a setting with all agents being gateways or a setting with exactly one gateway. The first one is again the socially optimal solution and in the latter case we get $\alpha + n(n-1)$ as social cost, which yields a price of anarchy of $\Theta(n/\sqrt{\alpha})$. (Note that here we use $\alpha \in [1, 2)$.) \square

Lemma 6.12. *In a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways, for $2 \leq \alpha \leq n-1$ the price of anarchy is at least $\Omega(n/\sqrt{\alpha})$.*

Proof. First, consider $\alpha \in [2, 4)$ and an L_1 -layer constituting a star graph with one center agent u and $n-1$ satellite agents. If exactly one satellite agent is a gateway, this graph forms an equilibrium with a social cost of $2(n-1)n$. Comparing this to the social optimum of αn , we get:

$$\text{PoA} \geq \frac{2(n-1)}{\alpha} \geq \frac{(n-1)}{\sqrt{\alpha}}$$

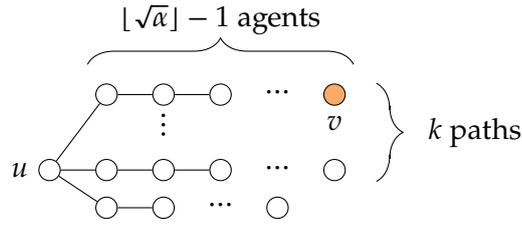


Figure 6.4: Equilibrium construction for the Sum-Layer-Game that gives a lower bound on the price of anarchy with $\alpha \geq 4, k := \lfloor \frac{n-1}{\lfloor \sqrt{\alpha} \rfloor - 1} \rfloor$, and v being the only gateway.

For the remainder of the proof, consider $\alpha \geq 4$. In this case, we construct a star-like L_1 -layer (cf. Figure 6.4) consisting of one center agent u , $k := \lfloor \frac{n-1}{\lfloor \sqrt{\alpha} \rfloor - 1} \rfloor$ -many disjoint paths P_1, \dots, P_k , each consisting of $(\lfloor \sqrt{\alpha} \rfloor - 1)$ -many agents, and possibly an additional path P_{k+1} consisting of the remaining agents. The first agent on each path is connected to u . We select one leaf agent v at a distance of exactly $\lfloor \sqrt{\alpha} \rfloor - 1$ to u to be a gateway. Then, no agent can perform an improving response, since the maximal distance cost decrease by opening is $\sum_{i=1}^{\lfloor \sqrt{\alpha} \rfloor - 1} 2i < \alpha$. We estimate a social cost lower bound by considering the private cost of u , which is minimal for all agents:

$$c_u(S) \geq k \sum_{i=1}^{\lfloor \sqrt{\alpha} \rfloor - 1} i = \frac{k}{2} (\lfloor \sqrt{\alpha} \rfloor - 1) \lfloor \sqrt{\alpha} \rfloor$$

This gives for the social cost:

$$\text{cost}(S) \geq \frac{n}{2} \left\lfloor \frac{n-1}{\lfloor \sqrt{\alpha} \rfloor - 1} \right\rfloor (\lfloor \sqrt{\alpha} \rfloor - 1) \lfloor \sqrt{\alpha} \rfloor$$

Comparing this to the social cost αn of the optimal solution, we get as the result $\text{PoA} = \Omega(n/\sqrt{\alpha})$. \square

Lemma 6.13. *In a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways, for $2 \leq \alpha \leq n - 1$ the price of anarchy is $\mathcal{O}(n/\sqrt{\alpha})$.*

Proof. Let $S \subseteq V$ be an arbitrary equilibrium strategy profile. Using Proposition 6.3, we know that $S = V$ is the socially optimal solution.

If $S \neq V$, then it must hold $|S| \leq \lceil \alpha \rceil$, since otherwise a non-gateway could reduce her distance cost by more than α by opening. Further, for every non-

gateway $v \in V \setminus S$, we get that $d_1(v, S) \leq 2\lceil\sqrt{\alpha}\rceil$, since otherwise opening v would reduce her private cost by at least:

$$\sum_{i=1}^{\lceil\sqrt{\alpha}\rceil} 2i = \lceil\sqrt{\alpha}\rceil(\lceil\sqrt{\alpha}\rceil + 1) > \alpha$$

Thus, for all gateways $v \in S$ it holds $c_v(S) \leq \alpha + |V \setminus S| \cdot 2\lceil\sqrt{\alpha}\rceil$. Since a non-gateway cannot have a higher private cost than a gateway, we get:

$$\text{cost}(S) \leq n\alpha + n \cdot |V \setminus S| \cdot 2\lceil\sqrt{\alpha}\rceil \leq n\alpha + 2n^2\lceil\sqrt{\alpha}\rceil$$

Comparing this to the social optimum yields:

$$\text{PoA} \leq \frac{n\alpha + 2n^2\lceil\sqrt{\alpha}\rceil}{\alpha n} \leq 1 + \frac{2n}{\lceil\sqrt{\alpha}\rceil} = O\left(\frac{n}{\sqrt{\alpha}}\right)$$

□

Lemma 6.14. *In a Sum-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways and a connection price $\alpha > n - 1$, the price of anarchy is:*

$$\text{PoA} = \begin{cases} O(\sqrt{\alpha}) & \text{for } \alpha \in (n - 1, n(n - 1)), \\ 1 & \text{for } \alpha \geq n(n - 1). \end{cases}$$

Proof. First, we show that for an arbitrary strategy profile $S' \subseteq V$ it holds $\text{cost}(S') > \alpha \cdot |S'| + n \cdot |V \setminus S'|$. We define $k := |V \setminus S'|$ to be the number of non-gateways and can use $k(k + 1) < nk$, since $|S'| \geq 1$. This gives:

$$\begin{aligned} \text{cost}(S') &\geq k(n - 1) + |S'| \cdot (\alpha + k) \\ &= kn - k + n\alpha + nk - \alpha k - k^2 \\ &= 2kn - k(k + 1) + \alpha(n - k) \\ &> \alpha(n - k) + kn \end{aligned}$$

Now we consider an equilibrium strategy profile S . If $S = V$, then the social cost is αn . For the case $\alpha > n(n - 1)$, no agent wants to open and hence exactly one gateway exists, which gives $\alpha + n(n - 1)$ for the social cost. Since the social cost lower bound is minimized when having exactly one gateway, we

get $\text{PoA} \leq \frac{\alpha+n(n-1)}{\alpha+(n-1)n} = 1$.

For $n(n-1) \geq \alpha \geq n$, let m be the number of gateways in an equilibrium S . Since S is an equilibrium, the maximal distance from a non-gateway to a gateway is $2\sqrt{\alpha}$. This gives for any gateway $u \in S$ that $c_u(S) \leq \alpha + (n-m)2\sqrt{\alpha}$ and for any non-gateway $v \in V \setminus S$ that $c_v(S) \leq (n-1)4\sqrt{\alpha}$. The social cost can be upper bounded by:

$$\begin{aligned} \text{cost}(S) &\leq m\alpha - m(n-m)2\sqrt{\alpha} + (n-m)4\sqrt{\alpha}(n-1) \\ &\leq m\alpha - m^22\alpha + 4\sqrt{\alpha}n(n-1) \end{aligned}$$

The global maximum of this upper bound is at $\sqrt{\alpha}/4$, which has the value of $\frac{\alpha\sqrt{\alpha}}{8} + 4\sqrt{\alpha}n(n-1)$. Comparing this to the social cost lower bound of $\alpha + n(n-1)$, we get $\text{PoA} = O(\sqrt{\alpha})$. \square

6.3.2 The Max-Layer-Game

Similar to the Sum-Layer-Game, we start our analysis of the Max-Layer-Game by studying the hardness of computing a socially optimal solution, followed by a discussion of the convergence properties of improving-response processes and the price of anarchy.

Theorem 6.15. *For the Max-Layer-Game with bidirectional gateways, the computation of a gateway set that minimizes the social cost is \mathcal{NP} -hard.*

Proof. For two parameters n and m with $m = 2n$, let there be a set of m elements $X := \{x_1, \dots, x_m\}$ and further n subsets $S_1, \dots, S_n \subseteq X$ of this element set. Then the \mathcal{NP} -complete SET-COVER problem (cf. Karp [Kar72]) is the task to compute a minimal number of subsets that together contain all elements of X . Given such a SET-COVER instance (V, L_1) , we construct an instance of the Max-Layer-Game as follows (cf. Figure 6.1). First, we create a clique C of k agents and mark one of them as c . For every set S_i , we create a corresponding agent S_i and connect her to c . For every element $x_i \in X$, we create an agent x_i and connect her to all set agents S_l with $x_i \in S_l$. Using the parameters $\alpha := 3$ and $k := \alpha n = 3n$, in the following we show that an optimal placement of gateways corresponds to a solution of the SET-COVER problem.

For now, assume that c is a gateway agent in the optimal solution S_{Opt} (we will prove this claim later). We claim that then no other clique agent $v \in C$ with

$v \neq c$ is a gateway. For this, assume that l further clique agents are open in S_{Opt} and compute the social cost decrease gained by closing all of these clique agents except c . If $l < k - 1$, then at most the distances of these l agents are increased by one each, which gives a social cost decrease of $l\alpha - l > 0$. Otherwise, the social cost decrease is at least $(k - 1)\alpha - (k - 1) - m - n = 2(3n - 1) - 3n > 0$. Hence, there can be at most one gateway agent c contained in the clique.

Next, assume that there are l open element agents in S_{Opt} . If $l < m$ and if at the same time there are open set agents who form a set cover, then by closing all element agents, only the maximal distances of these element agents increase and the social cost decreases by at least $\alpha l - l > 0$. If there are not yet set agents open that form a set cover, we have to open at most n set agents to form a set cover. By opening them and simultaneously closing all element agents, the maximum distances for all clique agents decrease by one each, which gives a social cost decrease of at least $\alpha l + k - \alpha n - l = 3l + 3n - 3n - l > 0$. Finally, if $l = m$, by closing all element agents and opening a set cover, the social cost decreases by at least $\alpha m - \alpha n - m - k = 6n - 3n - 2n > 0$.

Finally, we can see that c actually has to be a gateway in S_{Opt} . For this, consider an arbitrary optimal setting with all clique agents closed (if one clique agent is open, we can close it and open c without increasing the social cost). When opening c , we know that without increasing the social cost we can close all element agents and open corresponding set agents. Hence, when opening c we can assume that all element agents are closed and that for each element agent a corresponding set is open.

Hence, the socially optimal solution S_{Opt} is given by a gateway agent c and a minimal number of set agents such that all element agents are covered. \square

Equilibria and Convergence Properties

Given a Max-Layer-Game (V, L_1) with bidirectional gateways, next we study the existence of equilibria and the convergence of improving-response processes. For the simple cases when the connection price is very small or very big, we can provide positive convergence results and by this implicitly show the existence of equilibria. Likewise, for the class of tree networks and networks with big girth, whereas the *girth* is the length of a shortest cycle in the network, we can compute equilibrium settings in polynomial time. Yet for the general

case, it will turn out that the Max-Layer-Game is not necessarily a potential game.

If the connection price is at most $\alpha < 1$, then for any non-gateway it is an improving response to open and also no gateway will ever close. Hence, no improving-response process can be longer than $n - 1$ steps and such a process always converges to the equilibrium state $S = V$. Moreover, for $\alpha > \text{diam}(L_1)$ no non-gateway will ever open and every gateway wants to close. Hence, also here we have the same convergence properties.

Next, we consider the non-trivial case of arbitrary networks with big girth.

Proposition 6.16. *Given a Max-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways such that the girth² is at least $\text{girth}((V, L_1)) \geq 4\alpha$, then for $\alpha \in [1, \text{diam}(L_1))$ a Max-Layer-Game equilibrium exists.*

Proof. Let x_1, x_2 be two maximal distant agents in (V, L_1) . If $d_1(x_1, x_2) < 2\alpha$, we get by $\text{girth}((V, L_1)) \geq 4\alpha$ that (V, L_1) is a tree and there exists an agent v who has a maximal distance of less than α to every other agent. In this case, opening v yields an equilibrium.

Otherwise, define $R := \lfloor \min\{\alpha - 1, (d_1(x_1, x_2) - \alpha)/2\} \rfloor$. Since agents x_1 and x_2 are at maximal distance, none of them can be connected to a leaf agent. For both of these agents, we do the following (cf. Figure 6.5): We consider the breadth-first-search trees up to level R , rooted at x_1 and x_2 , respectively. From the agents at level R , we open a maximal set of gateways such that no two gateways are at a distance less than R .

Now we claim that for every agent x in such a tree, there exists a gateway within a distance of at most R . For this, consider a shortest path to an agent u at level R . If u is not a gateway, then there must also be another agent u' at level R who is a gateway. Since the girth is at least 4α and $R < \alpha < \text{diam}(L_1)$, the shortest path from u to u' can only consist of agents of the tree and hence $d_1(x, u') < R$.

Next, iteratively open a maximal set of further agents such that each new agent has a minimal distance of exactly $\lceil \alpha \rceil$ to a gateway. By construction, since every non-gateway has a maximal distance of $\lfloor \alpha \rfloor$ to a gateway, a non-gateway can improve her maximal distance by at most $\lfloor \alpha \rfloor$ and hence cannot perform any improving response. For every gateway v , it holds that her private cost

²Note that for an acyclic graph the girth is infinity.

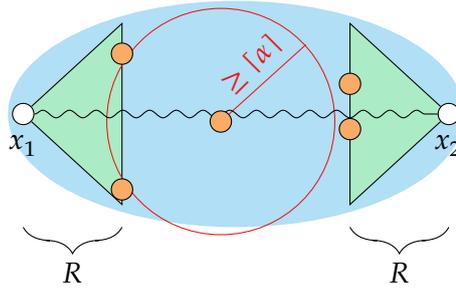


Figure 6.5: Illustration of the equilibrium construction in the proof of Proposition 6.16 for the Max-Layer-Game with bidirectional gateways and girth of at least 4α . Orange agents denote gateways.

is $c_v(S) = \alpha + R$ (with both x_1 and x_2 at maximal distance, since otherwise we get a contradiction to the maximal distance of x_1 and x_2 .) Considering the private cost change of closing v , her maximal distance increases by exactly $\lceil\alpha\rceil$ and hence this is not an improving response. \square

Theorem 6.17. *The Max-Layer-Game with bidirectional gateways and a connection price $\alpha > 1$ is not a potential game.*

Proof. Consider an L_1 -layer consisting of $n := 3\lceil\alpha\rceil + 4$ agents that are connected as a line. We denote the first agent of the line as u , the agent at distance $\lceil\alpha\rceil + 1$ to u as v , and the agent at distance $2\lceil\alpha\rceil + 2$ to u as w . Initially, only u is a gateway. Then, v and w form an improving-response cycle:

- I: w opens since $2\lceil\alpha\rceil + 2 > \alpha + \lceil\alpha\rceil + 1$.
- II: v opens since $2\lceil\alpha\rceil + 2 > \alpha + \lceil\alpha\rceil + 1$.
- III: w closes since $\alpha + \lceil\alpha\rceil + 1 > \lceil\alpha\rceil + 1$.
- IV: v closes since $\alpha + 2\lceil\alpha\rceil + 2 > 2\lceil\alpha\rceil + 2$.

Hence, the game does not provide the finite improvement property. \square

Price of Anarchy

Previously, we already argued that for $\alpha < 1$ the only equilibrium is $S = V$. Since this is also the socially optimal solution, both the price of anarchy and

the price of stability are 1 then. For the remaining connection prices of $\alpha \geq 1$, next we provide a tight price of anarchy result.

Theorem 6.18. *Given a Max-Layer-Game (V, L_1) of $n := |V|$ agents with bidirectional gateways, then for $\alpha \geq 1$ the price of anarchy is $\Theta(1 + n/\sqrt{\alpha})$.*

Proof. (Upper bound.) We start with an upper bound on the price of anarchy. For this, let (V, L_1) be a game instance and $S \subseteq V$ an arbitrary equilibrium strategy profile. With $D := \text{diam}(L_1)$, it trivially holds that $\text{cost}(S) \leq nD$.

Now we want to consider the minimal social cost when placing exactly k gateways on a longest shortest path P . Having only these k gateways, the total cost of the agents on P is:

$$\alpha k + 2k \sum_{i=1}^{\lfloor D/(2k) \rfloor} \left(i + \left\lfloor \frac{D}{2k} \right\rfloor \right) \geq \alpha k + \frac{3}{4k} D^2$$

The total cost of all agents not on P is at least $(n - D) \frac{D}{2k}$, which gives a social cost lower bound of:

$$\alpha k + \frac{3}{4k} D^2 + (n - D) \frac{D}{2k} = \alpha k + \frac{D^2 + 2nD}{4k}$$

This term is minimized by $k = \sqrt{\frac{D^2 + 2nD}{4\alpha}}$, which corresponds to a social cost of at least $\sqrt{\alpha(D^2 + 2nD)}$. Comparing this value to the previously computed social cost upper bound for any equilibrium setting gives

$$\frac{nD}{\sqrt{\alpha(D^2 + 2nD)}} \leq \frac{n}{\sqrt{\alpha}},$$

which is the claimed upper bound for the price of anarchy.

(Lower bound.) For $n \in \mathbb{N}$, $k := \lfloor (n - 1)/3 \rfloor$, we consider the following L_1 -layer: We select one agent c as a center agent, connect two disjoint paths of each k -many agents to c , and finally connect one path consisting of $(n - 2k - 1)$ -many agents to c . When opening the leaf agent of the last connected path, we obtain an equilibrium strategy profile since no agent can improve her maximum

distance by opening. The social cost of this equilibrium is at least:

$$3 \sum_{i=1}^k (i+k) = 3k^2 + \frac{3}{2}(k+1)k = \Omega(n^2)$$

Next, consider the socially optimal solution. For $\sqrt{\alpha} \geq n$, the optimal solution coincides with the equilibrium. Otherwise, we get the optimal solution by opening c as well as a maximal set of agents on each path such that between each two neighboring gateways the distance is $\lfloor \alpha \rfloor$. The resulting social cost of this solution is at most $\alpha \frac{n}{\lfloor \sqrt{\alpha} \rfloor} + n \frac{\alpha}{\lfloor \sqrt{\alpha} \rfloor}$, which gives the price of anarchy lower bound of $\Omega(n/\sqrt{\alpha})$. \square

6.4 Unidirectional Gateways

In this section, we consider multilevel games with the property that switching from layer L_2 to layer L_1 is permitted at any agent, but access to the L_2 layer is restricted to gateway agents. As introduced in Section 6.1, these games are called multilevel games with *unidirectional* gateways. Different to the games with bidirectional gateways, we now consider arbitrary second layer networks, yet still keep our high-speed assumption. Applied to the new model, this means that we restrict our analysis to games (V, L_1, L_2) that ensure for every distance in the L_1 -layer that the corresponding distance in L_2 is at most a μ -fraction of the original one. Formally, we call this property μ -improving and it states that for every $u, v \in V$ it must hold:

$$\mu \cdot d_1(u, v) \geq d_2(u, v)$$

The following lemma states a first consequence of this property: Any shortest path between two agents switches at most once from L_1 to L_2 and then exclusively uses edges from L_2 , until it reaches the target agent.

Lemma 6.19. *Let (V, L_1, L_2) be a μ -improving multilevel game with unidirectional gateways, $S \subseteq V$ a set of gateways, and $u, v \in V$ two arbitrary agents. Then, for the edges (e_1, \dots, e_m) of any shortest path from u to v it holds: there is a $k \in \{0, \dots, m\}$ such that $e_i \in L_1$, for all $i \leq k$, and $e_i \in L_2$, for all $i > k$.*

Proof. Assume there is an edge $e_i \in L_2$ such that the succeeding edge of the

shortest path belongs to $e_{i+1} \in L_1$. Let $\{x, y\} = e_{i+1}$ denote the end points of this edge. Since the L_2 -layer is μ -improving, we know that there also exists a path of length at most μ from x to y purely consisting of L_2 edges. Yet, this would contradict e_{i+1} belonging to a shortest path. \square

6.4.1 Existence of Equilibria

Compared to the multilevel games with bidirectional gateways, the computation of equilibria in the unidirectional model seems to be even harder than before, given that now the L_2 -layer can have an arbitrary structure. In the following, we show the existence of equilibria in the Max-Layer-Game for the case when the L_1 -layer is a tree and furthermore the L_2 -layer provides the so-called *exact- μ -improving* property. For a game (V, L_1, L_2) we say that it is exact- μ -improving if it fulfills $\mu \cdot d_1(u, v) = d_2(u, v)$ for every pair of agents u and v . Specifically, the following theorem provides a polynomial time algorithm for computing an equilibrium setting.

Theorem 6.20. *Let (V, L_1, L_2) be an exact- μ -improving Max-Layer-Game instance with unidirectional gateways such that (V, L_1) is a tree. Then, there exists a set of gateways $S \subseteq V$ forming an equilibrium, which can be computed in polynomial time.*

Proof. We start with an empty gateway set S and compute a solution as follows:

- (a) If $\text{diam}(L_1) \leq \frac{\alpha}{1-\mu}$, then output $S = \emptyset$ as the solution.
- (b) If $\frac{2\alpha}{1-\mu} > \text{diam}(L_1) > \frac{\alpha}{1-\mu}$, then select an arbitrary agent z such that it holds $d_1(z, v) \leq \frac{\alpha}{1-\mu}$ for all $v \in V$ and furthermore, there is some agent $x \in V$ with $d_1(z, x) \geq \frac{\alpha}{1-\mu}$. Then output $S = \{z\}$ as solution.
- (c) If $\text{diam}(L_1) \geq \frac{2\alpha}{1-\mu}$ then:
 - (i) Considering only the first layer L_1 , select an agent with the smallest maximal distance to all other agents, name her r and open r .
 - (ii) Next, iteratively consider the other tree agents in a sequence such that the first layer distance to r is increasing. If for such an agent v it holds that v would reduce the distance to r by at least α by opening, then we open this agent. Otherwise v stays closed.

We claim that the so-computed solution S forms an equilibrium setting. First, if $\text{diam}(L_1) \leq \frac{\alpha}{1-\mu}$, then no agent can improve her maximum distance by more than α and hence $S = \emptyset$ is an equilibrium.

In case $\frac{2\alpha}{1-\mu} > \text{diam}(L_1) > \frac{\alpha}{1-\mu}$, then there exists an agent z with the specified properties. Specifically, every agent has a distance of at most $\frac{\alpha}{1-\mu}$ to z and hence no agent can improve her maximum distance cost of more than α by opening. Since z would increase her maximum distance by closing, it follows that $S = \{z\}$ is an equilibrium.

Now we consider the interesting case of $\text{diam}(L_1) \geq \frac{2\alpha}{1-\mu}$. For this, we first show that no gateway wants to deviate from her strategy and, furthermore, that also no non-gateway wants to open. We use $h(L_1)$ to denote the maximal distance in the L_1 -layer from r to any agent.

Gateways: By construction, for gateway r it holds that the closest other gateway is at a distance of at least $\frac{\alpha}{1-\mu}$. Hence, r would increase her longest shortest path distances by at least α when closing, which cannot be an improving response.

We denote the gateways as r, v_1, \dots, v_k , ordered in the sequence they were opened. For the i -th opened gateway v_i , the shortest path distance from v_i to r was improved by at least α . We further know for v_i that for both strategies, v_i being a gateway or being a non-gateway, there is a longest shortest path containing r . For v_i being a gateway, this directly holds by choice of r . But also if $v_i \notin S$, since r is a gateway, $\delta(v_i, r) + \mu h(L_1)$ is an upper bound on every distance and by choice of r there must be a shortest path over r to some agent x to which this is the distance. Hence, the maximum distance to any other agent, and by this the private cost of v_i , is given by the distance to r .

Thus, agent v_i would close only if this distance increased by less than α . By construction, only the opening of a gateway v_j with $j > i$ can cause a strategy change of v_i . We denote the closest common predecessor of v_i and v_j in the rooted tree by z and get $d_1(r, z) \leq d_1(r, v_i) \leq d_1(r, v_j)$. Hence, if v_i closed, this would incur additional distance cost for v_i of at least $\min\{\alpha, d_1(v_i, z) + d_1(z, v_j) + \mu d_1(v_j, r) - \mu d_1(v_i, r)\}$. Yet, the same also holds for v_j and since v_j was opened after v_i , it must

hold $d_1(v_i, z) + d_1(z, v_j) + \mu d_1(v_j, r) - \mu d_1(v_i, r) \geq d_1(v_i, z) + d_1(z, v_j) + \mu d_1(v_i, r) - \mu d_1(v_j, r) > \alpha$. This means, no gateway in S wants to close.

Non-gateways: Let x be a non-gateway, then, as in the previous discussion, we see for x that $\delta(x, r) + \mu h(L_1)$ is an upper bound on every distance. Hence, the distance improvement from x to r must be at least α such that x wants to open. Yet, if that was the case, x would have been opened by the above described process.

Hence, neither gateways nor non-gateways can improve their private costs by changing their strategies. \square

Although we do not present existence results for all game instances, with respect the following price of anarchy estimations still we can argue in the following way that the price of anarchy is well-defined: For any number of agents n , there exists a network of n agents enabling equilibria. In particular, for a clique network of n agents trivially an equilibrium always exists, even for arbitrary L_2 -layers.

6.4.2 Quality of Equilibria

We now study upper bounds for the Sum-Layer-Game and subsequently give results for the Max-Layer-Game.

Sum-Layer-Game

For the Sum-Layer-Game, we start with characterizing equilibria by the maximal distance that agents can have to the closest gateway. We derive this distance bound by an expansion argument.

Lemma 6.21. *In a μ -improving Sum-Layer-Game (V, L_1, L_2) of $n := |V|$ agents with unidirectional gateways and a connection price α , let S be an equilibrium strategy profile. Then, for any agent $u \notin S$ and any $k \in \mathbb{N}$ with $B_k(u) \cap S = \emptyset$, it either holds*

- (a) $|B_k(u)| \geq \frac{n}{2}$ or
- (b) $\alpha \geq \frac{kn(1-\mu)}{2}$.

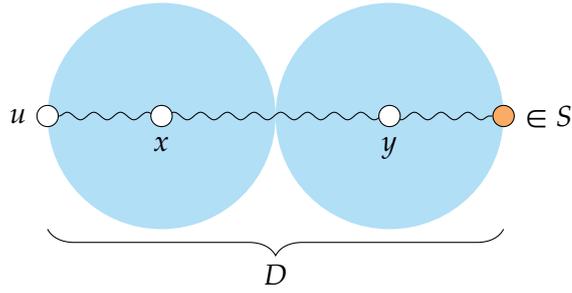


Figure 6.6: Illustration of main argument used in Lemma 6.22.

Proof. If $|B_k(u)| < n/2$, then there is a set $X := V \setminus B_k(u)$ of a size of at least $|X| > n/2$. For any agent $x \in X$, we know that the shortest path from u to x passes an agent y who has a distance of exactly k to u . By the claim, the shortest path to y does not contain any gateways and hence opening u reduces the distance to y by at least $(1 - \mu)d_1(u, y) = (1 - \mu)k$. Since this holds for all agents in X , the total distance cost improvement is at least $|X| \cdot (1 - \mu)k$. With S being an equilibrium, u cannot perform any improving response and thus it must hold

$$\alpha \geq |X| \cdot (1 - \mu)k \geq \frac{kn(1 - \mu)}{2},$$

which concludes the claim. \square

Lemma 6.22. *In a μ -improving Sum-Layer-Game (V, L_1, L_2) of $n := |V|$ agents with unidirectional gateways and a connection price α , the maximal distance of any agent to a gateway is $\frac{16\alpha}{(1-\mu)n}$.*

Proof. Let $u \in V$ be an agent with a maximal minimal distance to any gateway in S and define $D := \min_{v \in S} d_1(u, v)$ (cf. Figure 6.6). Then, there exist two agents $x, y \in B_D(u)$ such that $B_{D/4}(x) \cap B_{D/4}(y) = \emptyset$ and further $B_{D/4}(x) \subseteq B_D(u)$ as well as $B_{D/4}(y) \subseteq B_D(u)$.³

If at least one of these balls contains fewer than $(n/2)$ -many agents, then by Lemma 6.21 we have $\frac{D}{4} \leq \frac{2\alpha}{n(1-\mu)}$ and thus the claim. Otherwise, we know that both balls $B_{D/4}(x)$ and $B_{D/4}(y)$ combined contain n -many agents. In this case, we consider a ball around u of radius $D/2$ and again can apply the same

³Note that if D is odd, then we can assign the agent contained in both balls arbitrarily and the following arguments still hold.

construction as above. Since now at least one of the balls around x or y must contain fewer than $(n/2)$ -agents, we get $D \leq 2 \frac{8\alpha}{(1-\mu)n}$. \square

Theorem 6.23. *In a μ -improving Sum-Layer-Game (V, L_1, L_2) of $n := |V|$ agents with unidirectional gateways and a connection price α , the price of anarchy is:*

$$\text{PoA} = \begin{cases} 1 & \text{for } \alpha < 1 - \mu, \\ O\left(\frac{1}{1-\mu}\right) & \text{for } \alpha \in [1 - \mu, n(1 - \mu)], \\ O\left(\frac{\alpha}{n(1-\mu)^2}\right) & \text{for } \alpha \in (n(1 - \mu), n(n - 1)], \\ 1 & \text{else.} \end{cases}$$

Proof. We derive the price of anarchy upper bounds by individually comparing the private costs of agents in an arbitrary equilibrium S to the private costs in a socially optimal solution S_{Opt} . First note that for $\alpha < 1 - \mu$, in both settings S and S_{Opt} , every agent will open and the social costs are equal. Similarly, for $\alpha > n(n - 1)$ no agent will open and thus the price of anarchy is 1, too.

For the general case, if we have some agent $u \in V$ with $u \in S$ and $u \in S_{\text{Opt}}$, then the private cost of this agent is the same in both solutions. Next, consider some agent $u \in S$ who is a gateway in S but not in S_{Opt} . Note that for some solution S' with $u \notin S'$ it holds $\delta_{S'}(u, v) \geq d_2(u, v) + (1 - \mu)$ for all $v \in V$. Comparing the private costs, we get:

$$\frac{c_S(u)}{c_{S_{\text{Opt}}}(u)} \leq \frac{\alpha + \sum_{v \in V} d_2(u, v)}{(1 - \mu)(n - 1) + \sum_{v \in V} d_2(u, v)} \leq 1 + \frac{\alpha}{(n - 1)(1 - \mu)}$$

If agent u is not a gateway in S but a gateway in S_{Opt} , we get for the private cost comparison:

$$\begin{aligned} \frac{c_S(u)}{c_{S_{\text{Opt}}}(u)} &= \frac{\sum_{v \in V} \delta_S(u, v)}{\alpha + \sum_{v \in V} d_2(u, v)} \\ &\leq \frac{\sum_{v \in V} d_2(u, v)}{\alpha + \sum_{v \in V} d_2(u, v)} + \frac{(n - 1) \frac{16\alpha}{(1-\mu)n}}{\alpha + \sum_{v \in V} d_2(u, v)} \leq 1 + \frac{16}{1 - \mu} \end{aligned}$$

Finally, if agent u is not a gateway in both solutions S and S_{Opt} , the private cost

comparison gives:

$$\begin{aligned}
\frac{c_S(u)}{c_{S_{\text{Opt}}}(u)} &\leq \frac{\sum_{v \in V} \delta_S(u, v)}{(1 - \mu)(n - 1) + \sum_{v \in V} d_2(u, v)} \\
&\leq \frac{\sum_{v \in V} d_2(u, v)}{(1 - \mu)(n - 1) + \sum_{v \in V} d_2(u, v)} + \frac{(n - 1) \frac{16\alpha}{(1 - \mu)n}}{(1 - \mu)(n - 1) + \sum_{v \in V} d_2(u, v)} \\
&\leq 1 + \frac{16\alpha}{n(1 - \mu)^2}
\end{aligned}$$

Combining all private cost ratio upper bounds, the price of anarchy is at most $1 + \max\left\{\frac{\alpha}{(n-1)(1-\mu)}, \frac{16\alpha}{n(1-\mu)^2}, \frac{16}{1-\mu}\right\}$ and yields the respective bounds with case distinction for $\alpha \leq n(1 - \mu)$. \square

For comparing these results to the Sum-Layer-Game with bidirectional gateways, the improvement factor μ has to be close to $\mu \approx 0$: i.e., L_2 ensures very short connections. This either results in a constant or in an $O(1 + \alpha/n)$ bound for the price of anarchy, whereas the latter tends to 1 with growing n . Consequently, the results are usually much better than in the game with bidirectional gateways; and they become worst when the L_2 -layer is improving only marginally.

Max-Layer-Game

For the Max-Layer-Game, we derive our price of anarchy results from characterizing equilibria by the maximal distance that any agent can have to a gateway and the minimal distance any two gateways must have to each other. Complementing this upper bound, we further provide a high lower bound, which will also be a lower bound on the price of stability.

Theorem 6.24. *In a μ -improving Max-Layer-Game (V, L_1, L_2) of $n := |V|$ agents with unidirectional gateways and a connection price α , the price of anarchy is:*

$$\text{PoA} = \begin{cases} 1 & \text{for } \alpha < 1 - \mu, \\ O(1) & \text{for } \alpha \in \left[1 - \mu, \mu \frac{(1+\mu)^2}{1-\mu}\right), \\ O\left(\frac{\alpha}{(1-\mu)^2}\right) & \text{else.} \end{cases}$$

Proof. First, we consider the case with very small connection prices. For $\alpha < 1 - \mu$, assume there is an equilibrium setting $S \neq V$ with an agent $u \in V \setminus S$. Then, we know for agent u that every shortest path to any other agent contains at least one edge of layer L_1 and thus by opening, u would decrease her distance cost by at least $(1 - \mu)$. Since this contradicts S being an equilibrium, the only equilibrium is $S = V$. The same argument holds for the optimal setting, which means $S_{\text{Opt}} = V$. Hence, for this connection price, the price of anarchy is 1.

In the following, we consider the contrary case with connection price $\alpha \geq (1 - \mu)$. For the social optimum S_{Opt} , let $m := |S_{\text{Opt}}|$ be the number of gateways and note that for every $u \notin S_{\text{Opt}}$ it holds $\delta_{S_{\text{Opt}}}(u, v) \geq d_2(u, v) + (1 - \mu)$ for all $v \in V$. With this, we lower bound the social cost of S_{Opt} by:

$$\begin{aligned} \text{cost}(S_{\text{Opt}}) &\geq m\alpha + \sum_{u \in S} \max_{v \in V} d_2(u, v) + (n - m)(1 - \mu) + \sum_{u \in V \setminus S} \max_{v \in V} d_2(u, v) \\ &\geq \sum_{u \in V} \max_{v \in V} d_2(u, v) + n(1 - \mu) \end{aligned}$$

Next, we estimate an upper bound on the social cost for any equilibrium strategy profile. For this, let S be an arbitrary equilibrium and we estimate the maximal distance of any agent to a gateway as well as the minimal distance between any pair of agents.

- (a) (*Maximum gateway distance.*) Let $v \notin S$ be an agent and D the minimal distance from v to any gateway. Then, for every shortest path from v to any other agent, there are D -many edges belonging to L_1 and thus opening would reduce v 's distance cost by at least $D(1 - \mu)$. Since S is an equilibrium, we get $\alpha \geq D(1 - \mu)$ and hence $D \leq \frac{\alpha}{1 - \mu}$.
- (b) (*Minimum gateway distance.*) For any gateway $u \in S$, let D be the minimal distance to any other gateway, i.e., $D = \min_{v \in S \setminus \{u\}} d_1(u, v)$. If u would close, her distance cost increased by at most $D + \mu D$. Since S is in equilibrium, we have $0 \leq -\alpha + D + \mu D$ and hence $D \geq \frac{\alpha}{1 + \mu}$. Using the minimal distance between any pair of gateways, we directly gain an upper bound on the number of gateways of any equilibrium setting. That is, there cannot be more than $\left(n \frac{1 + \mu}{\alpha}\right)$ -many gateways in an equilibrium setting.

Using both estimations, we can finally derive an upper bound on the social cost of any equilibrium strategy profile S , hereby $m := |S|$ denotes the number

of gateways in S :

$$\begin{aligned}
\text{cost}(S) &\leq \sum_{u \in V \setminus S} \left(\frac{\alpha}{1-\mu} + \max_{v \in V} d_2(u, v) \right) + \sum_{u \in S} \left(\alpha + \max_{v \in V} d_2(u, v) \right) \\
&= \sum_{u \in V} \max_{v \in V} d_2(u, v) + (n-m) \frac{\alpha}{1-\mu} + m\alpha \\
&\leq \sum_{u \in V} \max_{v \in V} d_2(u, v) + n \left(1 - \mu \frac{1+\mu}{\alpha} \right) \frac{\alpha}{1-\mu}
\end{aligned}$$

Comparing this to the social cost lower bound, we get for $\alpha < \mu \frac{(1+\mu)^2}{1-\mu}$ that the price of anarchy is constant. And otherwise, the ratio is at most $1 + \frac{\alpha}{(1-\mu)^2}$. \square

Theorem 6.25 (price of stability and price of anarchy lower bound). *There are μ -improving multilevel networks (V, L_1, L_2) of $n := |V|$ agents, where the price of stability and the price of anarchy for the unidirectional Max-Layer-Game is $\Omega(\sqrt{n})$, when $\alpha = \sqrt{n}$.*

Proof. Consider a network (V, L_1) consisting of agents $v_1, \dots, v_{\lfloor \sqrt{n} \rfloor}$ who are connected as a line and additional agents $v_{\lfloor \sqrt{n} \rfloor + 1}, \dots, v_n$ who are all connected to agent v_1 . The second layer (V, L_2) forms a clique network with all edges having a length of zero. Note that this instance is μ -improving for any μ .

For this game, the diameter of (V, L_1) is $\lfloor \sqrt{n} \rfloor + 1$ and hence for $\alpha := \sqrt{n} + 1$, no agent will become a gateway in any equilibrium setting. Thus, for the only equilibrium, $S = \emptyset$, the social cost is at least $\lfloor \sqrt{n} \rfloor \sqrt{n} / 2 + (n - \lfloor \sqrt{n} \rfloor) \sqrt{n}$. On the other hand, for $S = \{v_1\}$ the social cost is at most $\alpha + (n - \sqrt{n}) + \sqrt{n} \sqrt{n}$. Hence, the ratio is $\Omega(\sqrt{n})$. \square

Compared to the Max-Layer-Game with bidirectional gateways, the price of anarchy results become much lower when using unidirectional gateways. This is similar to the Sum-Layer-Game, where unidirectional gateways also led to improved worst-case guarantees for the quality of equilibrium networks.

6.5 Conclusion & Future Work

We introduced a new network model to analyze effects of multilevel network interactions that are not captured by the traditional framework of network creation games. The provided price of anarchy results emphasize that for

a very small or big connection price α (i.e., when tending to the number of agents) equilibria are nearly optimal solutions despite the selfish behavior of the agents. Comparing the bidirectional and the unidirectional variants, the unidirectional games outperform the price of anarchy results of the other variant considerably.

In comparison to the classic network creation games, in the multilevel games the existence of equilibria became a much harder question. Although we have partial answers for special combinations of the connection price α and the network topology, providing a general answer seems to be a very challenging open question that deserves further research.

Regarding the convergence of best-response processes, in the game with bidirectional gateways, both variants with average and maximum cost functions do not provide the finite improvement property and remarkably, the Sum-Layer-Game is not even weakly acyclic. Due to the symmetry of the maximum operator, the equilibria and convergence properties for the maximum cost function seem to be more stable though. The same observation holds for the game with unidirectional gateways, however, for this variant we could not provide a definite answer to the question if and how best-response processes convergence.

Bibliography

- [Abs+14] S. Abshoff, A. Cord-Landwehr, D. Jung, and A. Skopalik. “Multilevel Network Games”. In: *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings*. Ed. by T. Liu, Q. Qi, and Y. Ye. Vol. 8877. Lecture Notes in Computer Science. Springer, 2014, pp. 435–440. doi: 10.1007/978-3-319-13129-0_36.
- [Alb+06] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. “On nash equilibria for a network creation game”. In: *17th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22–26, 2006. Proceedings*. ACM, 2006, pp. 89–98. URL: <http://dl.acm.org/citation.cfm?id=1109557.1109568>.
- [Alb+14] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. “On Nash Equilibria for a Network Creation Game”. In: *ACM Transactions on Economics and Computation* 2.1 (2014), 2:1–2:27. doi: 10.1145/2560767.
- [Alo+10] N. Alon, E. D. Demaine, M. Hajiaghayi, and T. Leighton. “Basic network creation games”. In: *22nd Annual ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2010, Thira, Santorini, Greece, June 13–15, 2010. Proceedings*. Ed. by F. Meyer auf der Heide and C. A. Phillips. ACM, 2010, pp. 106–113. doi: 10.1145/1810479.1810502.
- [Alo+13] N. Alon, E. D. Demaine, M. T. Hajiaghayi, and T. Leighton. “Basic Network Creation Games”. In: *SIAM Journal on Discrete Mathematics* 27.2 (2013), pp. 656–668. doi: 10.1137/090771478.
- [Ans+03] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. “Near-optimal network design with selfish agents”. In: *35th Annual ACM Symposium on Theory of Computing, STOC 2003, June 9–11, 2003, San*

Bibliography

- Diego, CA, USA. *Proceedings*. Ed. by L. L. Larmore and M. X. Goemans. ACM, 2003, pp. 511–520. doi: 10.1145/780542.780617.
- [Ans+04] 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”. In: *45th Symposium on Foundations of Computer Science, FOCS 2004, Rome, Italy, October 17–19, 2004. Proceedings*. IEEE Computer Society, 2004, pp. 295–304. doi: 10.1109/FOCS.2004.68.
- [Ary+04] V. Arya, N. Garg, R. Khandekar, A. Meyerson, K. Munagala, and V. Pandit. “Local Search Heuristics for k-Median and Facility Location Problems”. In: *SIAM Journal on Computing* 33.3 (2004), pp. 544–562. doi: 10.1137/S0097539702416402.
- [AS04] S. Androutsellis-Theotokis and D. Spinellis. “A survey of peer-to-peer content distribution technologies”. In: *ACM Journal of Computing Surveys* 36.4 (2004), pp. 335–371. doi: 10.1145/1041680.1041681.
- [Aug+15] J. Augustine, I. Caragiannis, A. Fanelli, and C. Kalaitzis. “Enforcing Efficient Equilibria in Network Design Games via Subsidies”. In: *Algorithmica* 72.1 (2015), pp. 44–82. doi: 10.1007/s00453-013-9845-5.
- [Bee+06] Z. Beerliova, F. Eberhard, T. Erlebach, A. Hall, M. Hoffmann, M. Mihalák, and L. S. Ram. “Network Discovery and Verification”. In: *IEEE Journal on Selected Areas in Communications* 24.12 (2006), pp. 2168–2181. doi: 10.1109/JSAC.2006.884015.
- [BG00] V. Bala and S. Goyal. “A noncooperative model of network formation”. In: *Econometrica* 68.5 (2000), pp. 1181–1229. doi: 10.1111/1468-0262.00155.
- [Bil+14a] D. Bilò, L. Gualà, S. Leucci, and G. Proietti. “Locality-based network creation games”. In: *26th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2014, Prague, Czech Republic, June 23–25, 2014. Proceedings*. Ed. by G. E. Blelloch and P. Sanders. ACM, 2014, pp. 277–286. doi: 10.1145/2612669.2612680.
- [Bil+14b] D. Bilò, L. Gualà, S. Leucci, and G. Proietti. “Network Creation Games with Traceroute-Based Strategies”. In: *Structural Information and Communication Complexity – 21st International Colloquium, SIROCCO 2014, Takayama, Japan, July 23–25, 2014. Proceedings*. Ed. by M. M. Halldórsson. Vol. 8576. Lecture Notes in Computer Science. Springer, 2014, pp. 210–223. doi: 10.1007/978-3-319-09620-9_17.

- [BK11] M. Brautbar and M. Kearns. “A Clustering Coefficient Network Formation Game”. In: *Algorithmic Game Theory – 4th International Symposium, SAGT 2011, Amalfi, Italy, October 17–19, 2011. Proceedings*. Ed. by G. Persiano. Vol. 6982. Lecture Notes in Computer Science. Springer, 2011, pp. 224–235. doi: 10.1007/978-3-642-24829-0_21.
- [BMP13] A. Bonato, M. Mitzenmacher, and P. Pralat, eds. *Algorithms and Models for the Web Graph – 10th International Workshop, WAW 2013, Cambridge, MA, USA, December 14–15, 2013. Proceedings*. Vol. 8305. Lecture Notes in Computer Science. Springer, 2013. doi: 10.1007/978-3-319-03536-9.
- [BS08] N. Baumann and S. Stiller. “The Price of Anarchy of a Network Creation Game with Exponential Payoff”. In: *Algorithmic Game Theory – 1st International Symposium, SAGT 2008, Paderborn, Germany, April 30 – May 2, 2008. Proceedings*. Ed. by B. Monien and U. Schroeder. Vol. 4997. Lecture Notes in Computer Science. Springer, 2008, pp. 218–229. doi: 10.1007/978-3-540-79309-0_20.
- [Chu+04] B. Chun, R. Fonseca, I. Stoica, and J. Kubiatoicz. “Characterizing Selfishly Constructed Overlay Routing Networks”. In: *IEEE INFOCOM 2004, The 23rd Annual Joint Conference of the IEEE Computer and Communications Societies, Hong Kong, China, March 7–11, 2004. Proceedings*. IEEE, 2004. URL: http://www.ieee-infocom.org/2004/Papers/28_4.PDF.
- [CL15] A. Cord-Landwehr and P. Lenzner. “Network Creation Games: Think Global – Act Local”. In: *Mathematical Foundations of Computer Science 2015 – 40th International Symposium, MFCS 2015, Milan, Italy, August 24–28, 2015. Proceedings, Part II*. Ed. by G. F. Italiano, G. Pighizzini, and D. Sannella. Vol. 9235. Lecture Notes in Computer Science. Springer, 2015, pp. 248–260. doi: 10.1007/978-3-662-48054-0_21.
- [CMM14] A. Cord-Landwehr, A. Mäcker, and F. Meyer auf der Heide. “Quality of Service in Network Creation Games”. In: *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings*. Ed. by T. Liu, Q. Qi, and Y. Ye. Vol. 8877. Lecture Notes in Computer Science. Springer, 2014, pp. 423–428. doi: 10.1007/978-3-319-13129-0_34.
- [Cor+12] A. Cord-Landwehr, M. Hüllmann, P. Kling, and A. Setzer. “Basic Network Creation Games with Communication Interests”. In: *Algorithmic Game Theory – 5th International Symposium, SAGT 2012, Barcelona, Spain, October 22–23, 2012. Proceedings*. Ed. by M. J. Serna.

- Vol. 7615. Lecture Notes in Computer Science. Springer, 2012, pp. 72–83. doi: 10.1007/978-3-642-33996-7_7.
- [CP05] J. Corbo and D. C. Parkes. “The price of selfish behavior in bilateral network formation”. In: *24th Annual ACM Symposium on Principles of Distributed Computing, PODC 2005, Las Vegas, NV, USA, July 17–20, 2005. Proceedings*. Ed. by M. K. Aguilera and J. Aspnes. ACM, 2005, pp. 99–107. doi: 10.1145/1073814.1073833.
- [CS07] S. Chien and A. Sinclair. “Convergence to approximate Nash equilibria in congestion games”. In: *18th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2007, New Orleans, Louisiana, USA, January 7–9, 2007. Proceedings*. Ed. by N. Bansal, K. Pruhs, and C. Stein. SIAM, 2007, pp. 169–178.
- [Dal+06] L. Dall’Asta, J. I. Alvarez-Hamelin, A. Barrat, A. Vázquez, and A. Vespignani. “Exploring networks with traceroute-like probes: Theory and simulations”. In: *Theoretical Computer Science* 355.1 (2006), pp. 6–24. doi: 10.1016/j.tcs.2005.12.009.
- [DB91] R. D. Dutton and R. C. Brigham. “Edges in graphs with large girth”. In: *Graphs and Combinatorics* 7.4 (1991), pp. 315–321. doi: 10.1007/BF01787638.
- [Dem+07] E. D. Demaine, M. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. “The price of anarchy in network creation games”. In: *26th Annual ACM Symposium on Principles of Distributed Computing, PODC 2007, Portland, Oregon, USA, August 12–15, 2007. Proceedings*. Ed. by I. Gupta and R. Wattenhofer. ACM, 2007, pp. 292–298. doi: 10.1145/1281100.1281142.
- [Dem+12] E. D. Demaine, M. T. Hajiaghayi, H. Mahini, and M. Zadimoghaddam. “The price of anarchy in network creation games”. In: *ACM Transactions on Algorithms* 8.2 (2012), 13:1–13:13. doi: 10.1145/2151171.2151176.
- [DJ03] B. Dutta and M. O. Jackson. *Networks and groups: Models of strategic formation*. Springer-Verlag, 2003. doi: 10.1007/978-3-540-24790-6.
- [DMP07] C. Daskalakis, A. Mehta, and C. H. Papadimitriou. “Progress in approximate Nash equilibria”. In: *8th ACM Conference on Electronic Commerce, EC 2007, San Diego, California, USA, June 11–15, 2007. Proceedings*. Ed. by J. K. MacKie-Mason, D. C. Parkes, and P. Resnick. ACM, 2007, pp. 355–358. doi: 10.1145/1250910.1250962.

- [Ehs+15] S. Ehsani, S. S. Fadaee, M. Fazli, A. Mehrabian, S. S. Sadeghabad, M. A. Safari, and M. Saghafian. “A Bounded Budget Network Creation Game”. In: *ACM Transactions on Algorithms* 11.4 (2015), 34:1–34:25. doi: 10.1145/2701615.
- [ES14] R. Engelberg and M. Schapira. “Weakly-Acyclic (Internet) Routing Games”. In: *Theory of Computing Systems* 54.3 (2014), pp. 431–452. doi: 10.1007/s00224-013-9474-z.
- [Fab+03] A. Fabrikant, A. Luthra, E. N. Maneva, C. H. Papadimitriou, and S. Shenker. “On a network creation game”. In: *22nd ACM Symposium on Principles of Distributed Computing, PODC 2003, Boston, Massachusetts, USA, July 13–16, 2003. Proceedings*. Ed. by E. Borowsky and S. Rajsbaum. ACM, 2003, pp. 347–351. doi: 10.1145/872035.872088.
- [FPT04] A. Fabrikant, C. H. Papadimitriou, and K. Talwar. “The complexity of pure Nash equilibria”. In: *36th Annual ACM Symposium on Theory of Computing, STOC 2004, Chicago, IL, USA, June 13–16, 2004. Proceedings*. Ed. by L. Babai. ACM, 2004, pp. 604–612. doi: 10.1145/1007352.1007445.
- [GJ02] M. R. Garey and D. S. Johnson. *Computers and intractability*. Vol. 29. W.H. Freeman and Company, 2002.
- [Gra+13] R. L. Graham, L. Hamilton, A. Levavi, and P. Loh. “Anarchy Is Free in Network Creation”. In: *Algorithms and Models for the Web Graph – 10th International Workshop, WAW 2013, Cambridge, MA, USA, December 14–15, 2013. Proceedings*. Ed. by A. Bonato, M. Mitzenmacher, and P. Pralat. Vol. 8305. Lecture Notes in Computer Science. Springer, 2013, pp. 220–231. doi: 10.1007/978-3-319-03536-9_17.
- [HM07] Y. Halevi and Y. Mansour. “A Network Creation Game with Nonuniform Interests”. In: *Internet and Network Economics, Third International Workshop, WINE 2007, San Diego, CA, USA, December 12–14, 2007. Proceedings*. Ed. by X. Deng and F. C. Graham. Vol. 4858. Lecture Notes in Computer Science. Springer, 2007, pp. 287–292. doi: 10.1007/978-3-540-77105-0_28.
- [Hoe13] M. Hoefer. “Local matching dynamics in social networks”. In: *Information and Computation* 222 (2013), pp. 20–35. doi: 10.1016/j.ic.2012.10.005.
- [Jac08] M. O. Jackson. *Social and economic networks*. Vol. 3. Princeton University Press Princeton, 2008.

Bibliography

- [JR11] G. A. Jehle and P. J. Reny. *Advanced Microeconomic Theory*. 3. ed. Harlow, England: Pearson Education Limited, 2011.
- [JW96] M. O. Jackson and A. Wolinsky. "A strategic model of social and economic networks". In: *Journal of economic theory* 71.1 (1996), pp. 44–74.
- [Kar72] R. M. Karp. "Reducibility Among Combinatorial Problems". In: *Complexity of Computer Computations*. Ed. by R. E. Miller, J. W. Thatcher, and J. D. Bohlinger. The IBM Research Symposia Series. New York, NY, USA: Springer, 1972, pp. 85–103. doi: 10.1007/978-1-4684-2001-2_9.
- [KKS14] S. Kniesburges, A. Koutsopoulos, and C. Scheideler. "Re-Chord: A Self-stabilizing Chord Overlay Network". In: *Theory of Computing Systems* 55.3 (2014), pp. 591–612. doi: 10.1007/s00224-012-9431-2.
- [KL13] B. Kawald and P. Lenzner. "On dynamics in selfish network creation". In: *25th ACM Symposium on Parallelism in Algorithms and Architectures, SPAA 2013, Montreal, QC, Canada, July 23–25, 2013. Proceedings*. Ed. by G. E. Blelloch and B. Vöcking. ACM, 2013, pp. 83–92. doi: 10.1145/2486159.2486185.
- [KP99] E. Koutsoupias and C. H. Papadimitriou. "Worst-case Equilibria". In: *16th Annual Symposium on Theoretical Aspects of Computer Science, STACS 1999, Trier, Germany, March 4–6, 1999. Proceedings*. Ed. by C. Meinel and S. Tison. Vol. 1563. Lecture Notes in Computer Science. Springer, 1999, pp. 404–413. doi: 10.1007/3-540-49116-3_38.
- [KT06] M. Kurant and P. Thiran. "Layered complex networks". In: *Physical review letters* 96.13 (2006), 138701:1–138701:4. doi: 10.1103/PhysRevLett.96.138701.
- [Kün+14] J. Künsemöller, S. Brangewitz, H. Karl, and C. Haake. "Provider Competition in Infrastructure-as-a-Service". In: *IEEE International Conference on Services Computing, SCC 2014, Anchorage, AK, USA, June 27 – July 2, 2014. Proceedings*. IEEE, 2014, pp. 203–210. doi: 10.1109/SCC.2014.35.
- [Lao+14] N. Laoutaris, L. J. Poplawski, R. Rajaraman, R. Sundaram, and S. Teng. "Bounded Budget Connection (BBC) games or how to make friends and influence people, on a budget". In: *Journal of Computer and System Sciences* 80.7 (2014), pp. 1266–1284. doi: 10.1016/j.jcss.2014.04.013.

- [Len11] P. Lenzner. “On Dynamics in Basic Network Creation Games”. In: *Algorithmic Game Theory – 4th International Symposium, SAGT 2011, Amalfi, Italy, October 17–19, 2011. Proceedings*. Ed. by G. Persiano. Vol. 6982. Lecture Notes in Computer Science. Springer, 2011, pp. 254–265. doi: 10.1007/978-3-642-24829-0_23.
- [Len12] P. Lenzner. “Greedy Selfish Network Creation”. In: *Internet and Network Economics – 8th International Workshop, WINE 2012, Liverpool, UK, December 10–12, 2012. Proceedings*. Ed. by P. W. Goldberg. Vol. 7695. Lecture Notes in Computer Science. Springer, 2012, pp. 142–155. doi: 10.1007/978-3-642-35311-6_11.
- [Len14] D.-I. P. Lenzner. “On Selfish Network Creation”. PhD thesis. Humboldt-Universität zu Berlin, 2014.
- [LQY14] T. Liu, Q. Qi, and Y. Ye, eds. *Web and Internet Economics – 10th International Conference, WINE 2014, Beijing, China, December 14–17, 2014. Proceedings*. Vol. 8877. Lecture Notes in Computer Science. Springer, 2014. doi: 10.1007/978-3-319-13129-0.
- [Mar+09] J. R. Marden, H. P. Young, G. Arslan, and J. S. Shamma. “Payoff-Based Dynamics for Multiplayer Weakly Acyclic Games”. In: *SIAM Journal on Control and Optimization* 48.1 (2009), pp. 373–396. doi: 10.1137/070680199.
- [Mil96] I. Milchtaich. “Congestion Games with Player-Specific Payoff Functions”. In: *Games and Economic Behavior* 13.1 (1996), pp. 11–124. doi: 10.1006/game.1996.0027.
- [MMM13] A. Mamageishvili, M. Mihalák, and D. Müller. “Tree Nash Equilibria in the Network Creation Game”. In: *Algorithms and Models for the Web Graph – 10th International Workshop, WAW 2013, Cambridge, MA, USA, December 14–15, 2013. Proceedings*. Ed. by A. Bonato, M. Mitzenmacher, and P. Pralat. Vol. 8305. Lecture Notes in Computer Science. Springer, 2013, pp. 118–129. doi: 10.1007/978-3-319-03536-9_10.
- [MS12] M. Mihalák and J. C. Schlegel. “Asymmetric Swap-Equilibrium: A Unifying Equilibrium Concept for Network Creation Games”. In: *Mathematical Foundations of Computer Science 2012 – 37th International Symposium, MFCS 2012, Bratislava, Slovakia, August 27–31, 2012. Proceedings*. Ed. by B. Rovan, V. Sassone, and P. Widmayer. Vol. 7464. Lecture Notes in Computer Science. Springer, 2012, pp. 693–704. doi: 10.1007/978-3-642-32589-2_60.

Bibliography

- [MS13] M. Mihalák and J. C. Schlegel. “The Price of Anarchy in Network Creation Games Is (Mostly) Constant”. In: *Theory of Computing Systems* 53.1 (2013), pp. 53–72. doi: 10.1007/s00224-013-9459-y.
- [MS96] D. Monderer and L. S. Shapley. “Potential Games”. In: *Games and Economic Behavior* 14.1 (1996), pp. 124–143. doi: 10.1006/game.1996.0044.
- [MSW06] T. Moscibroda, S. Schmid, and R. Wattenhofer. “On the topologies formed by selfish peers”. In: *25th Annual ACM Symposium on Principles of Distributed Computing, PODC 2006, Denver, CO, USA, July 23–26, 2006. Proceedings*. Ed. by E. Ruppert and D. Malkhi. ACM, 2006, pp. 133–142. doi: 10.1145/1146381.1146403.
- [MWG95] A. Mas-Colell, M. D. Whinston, and J. R. Green. *Microeconomic theory*. New York, NY, USA: Oxford University Press, Inc., 1995.
- [Mye77] R. B. Myerson. “Graphs and Cooperation in Games”. In: *Mathematics of Operations Research* 2.3 (1977), pp. 225–229. doi: 10.1287/moor.2.3.225.
- [Nas51] J. Nash. “Non-Cooperative Games”. In: *The Annals of Mathematics*. Second Series 54.2 (Sept. 1951), pp. 286–295. doi: 10.2307/1969529.
- [Nik+13] S. E. Nikolettseas, P. N. Panagopoulou, C. Raptopoulos, and P. G. Spirakis. “On the Structure of Equilibria in Basic Network Formation”. In: *Fundamentals of Computation Theory – 19th International Symposium, FCT 2013, Liverpool, UK, August 19–21, 2013. Proceedings*. Ed. by L. Gasieniec and F. Wolter. Vol. 8070. Lecture Notes in Computer Science. Springer, 2013, pp. 259–270. doi: 10.1007/978-3-642-40164-0_25.
- [Nik+15] S. E. Nikolettseas, P. N. Panagopoulou, C. Raptopoulos, and P. G. Spirakis. “On the structure of equilibria in basic network formation”. In: *Theoretical Computer Science* 590 (2015), pp. 96–105. doi: 10.1016/j.tcs.2015.03.029.
- [Nis+07] N. Nisan, T. Roughgarden, E. Tardos, and V. V. Vazirani. *Algorithmic Game Theory*. New York, NY, USA: Cambridge University Press, 2007.
- [Pel00] D. Peleg. “Distributed computing”. In: *SIAM Monographs on discrete mathematics and applications* 5 (2000).

- [Per11] G. Persiano, ed. *Algorithmic Game Theory – 4th International Symposium, SAGT 2011, Amalfi, Italy, October 17–19, 2011. Proceedings*. Vol. 6982. Lecture Notes in Computer Science. Springer, 2011. doi: 10.1007/978-3-642-24829-0.
- [Ros73] R. W. Rosenthal. “A class of games possessing pure-strategy Nash equilibria”. In: *International Journal of Game Theory* 2.1 (1973), pp. 65–67. doi: 10.1007/BF01737559.
- [SM03] A. S. Schulz and N. E. S. Moses. “On the performance of user equilibria in traffic networks”. In: *14th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2003, January 12–14, 2003, Baltimore, Maryland, USA. Proceedings*. ACM/SIAM, 2003, pp. 86–87. URL: <http://dl.acm.org/citation.cfm?id=644108.644121>.
- [SS13] E. M. Shahrivar and S. Sundaram. “Strategic multi-layer network formation”. In: *52nd IEEE Conference on Decision and Control, CDC 2013, Firenze, Italy, December 10–13, 2013. Proceedings*. IEEE, 2013, pp. 582–587. doi: 10.1109/CDC.2013.6759944.
- [SS15] E. M. Shahrivar and S. Sundaram. “The Strategic Formation of Multi-Layer Networks”. In: *CoRR abs/1501.02866* (2015). URL: <http://arxiv.org/abs/1501.02866>.
- [SV08] A. Skopalik and B. Vöcking. “Inapproximability of pure nash equilibria”. In: *40th Annual ACM Symposium on Theory of Computing, STOC 2008, Victoria, British Columbia, Canada, May 17–20, 2008. Proceedings*. Ed. by C. Dwork. ACM, 2008, pp. 355–364. doi: 10.1145/1374376.1374428.
- [Wat01] A. Watts. “A Dynamic Model of Network Formation”. In: *Games and Economic Behavior* 34.2 (2001), pp. 331–341. doi: 10.1006/game.2000.0803.
- [WS11] D. P. Williamson and D. B. Shmoys. *The design of approximation algorithms*. Cambridge University Press, 2011.
- [You93] H. P. Young. “The evolution of conventions”. In: *Econometrica: Journal of the Econometric Society* (1993), pp. 57–84.
- [Zim80] H. Zimmermann. “OSI Reference Model – The ISO Model of Architecture for Open Systems Interconnection”. In: *Communications, IEEE Transactions on* 28.4 (Apr. 1980), pp. 425–432. doi: 10.1109/TCOM.1980.1094702.

- μ -improving, **141**
 - exact, **142**
- ε -approximate equilibrium, **13, 87**
 - k -local Greedy Buy Equilibrium, **94, 101**
 - Greedy Buy Equilibrium, **21**
 - k -local Greedy Buy Equilibrium, **106**
- 0-probing, **86**
- Asymmetric Swap Equilibrium, **12, 103, 111**
- Basic Network Creation Games, *see* Swap-Game
- BE, *see* Buy Equilibrium
- best-response
 - cycle, **16, 91**
 - operation, **16, 61**
 - process, **16**
- Binding-Sequence, **37, 39**
- BR, *see* best-response
- Budged Buy Equilibrium, **20**
- Buy Equilibrium, **12, 30, 61, 86**
 - of the k -local Sum-Game, **90**
 - of the Max-Pricing-Game, **68**
 - of the Sum-Pricing-Game, **66**
- Buy-Game, *see* Buy Equilibrium
- diameter, **11**
- distance cost, **11, 85**
- edge cost, **11, 85**
- edge ownership, **10**
- finite improvement property, **17, 128, 139**
- FIP, *see* finite improvement property
- friends, **29**
- friendship graph, **27, 29**
- gateway, **121**
 - bidirectional, **121**
 - unidirectional, **121, 141**

INDEX

- GBE, *see* Greedy Buy Equilibrium
- girth, 52, 110, **137**
- Greedy Asymmetric Swap
 - Equilibrium, **13**
- Greedy Buy Equilibrium, **12**, 86
- greedy probing, **86**
- Greedy Swap Equilibrium, **13**

- improving-response, **16**, 61, 122
 - cycle, 16
 - process, **16**
- IR, *see* improving-response

- k-BE, *see* k-local Buy Equilibrium
- k-GBE, *see* k-local Greedy Buy Equilibrium
- k-local Buy Equilibrium, 87
- k-local Greedy Buy Equilibrium, 87
- k-local greedy operation, **86**
- k-local operation, **86**
- k-local Sum-Game, 85
- k-neighborhood, 86

- Max-Buy-Game, *see* Max-Game
- Max-Game, 31
- Max-Layer-Game, **122**
- Max-Pricing-Game, **59**
- Max-Swap-Game, **30**
- maximum independent set, 47
- Minimum Dominating Set
 - Problem, 90

- Nash Equilibrium, **12**
- NE, *see* Nash Equilibrium
- network creation game, **10**
- node, *see* agent

- optimistic price of anarchy, *see*
 - price of stability
- overlay network, **1**

- pairwise stability, **12**, **23**
- peer, *see* agent
- player, *see* agent
- PoA, *see* price of anarchy
- PoS, *see* price of stability
- potential game, **16**, 36, 128, 139
- price function, 6, 58, **59**, 62, 77
- price of anarchy, 14
 - of Asymmetric Swap Equilibria, 20
 - of Budgeted Buy Equilibria, 20
 - of Greedy Equilibria, 21
 - of Local Buy Equilibria, 21
 - of Swap Equilibria, 19
 - of the k -local Sum-Game, 108, 115
 - of the Max-Game, 18
 - of the Max-Layer-Game, bidirectional, 140
 - of the Max-Layer-Game, unidirectional, 147
 - of the Max-Pricing-Game, 81
 - of the Max-Swap-Game, general networks, 49
 - of the Max-Swap-Game, tree networks, 45, 48
 - of the Sum-Game, 18
 - of the Sum-Layer-Game, bidirectional, 132
 - of the Sum-Layer-Game, unidirectional, 146, 149

- of the Sum-Pricing-Game, 74, 76
 - of the Sum-Swap-Game, 49
- price of stability, 14
 - of the Max-Pricing-Game, 78
 - of the Max-Swap-Game, 34
 - of the Sum-Layer-Game, unidirectional, 149
 - of the Sum-Pricing-Game, 71
 - of the Sum-Swap-Game, 36
- private cost, **10**, 42
- private cost function, *see* private cost
- probing technique, 86
- process equilibrium, **51**
- process price of anarchy, **51**
 - of the Max-Buy-Game, 54
 - of the Sum-Buy-Game, 52
- SE, *see* Swap Equilibrium
- Set-Cover Problem, **125**, 136
- social cost, **11**, 31
- strategic game, 10
 - strategy, **10**
 - strategy profile, **10**
 - strategy space, **10**
- Sum-Buy-Game, *see* Sum-Game
- Sum-Game, **10**, 30
- Sum-Layer-Game, **122**
- Sum-Pricing-Game, **59**
- Sum-Swap-Game, **30**
- swap, 30
- Swap Equilibrium, **13**, 19, 30
- Swap-Game, *see* Swap Equilibrium
- tracerroute, 84
- tree conjecture, 18
- tree-star, **95**
- Uncapacitated Metric Facility Location Problem, **94**
- unrestricted probing, **86**
- view radius, **85**
- WAG, *see* weakly acyclic game
- weakly acyclic game, **16**, 128, 130