TEL-AVIV UNIVERSITY

RAYMOND AND BEVERLY SACKLER

FACULTY OF EXACT SCIENCES

SCHOOL OF COMPUTER SCIENCE

A Network Creation Game with Nonuniform Interests

This thesis was submitted in partial fulfillment of the requirements for the M.Sc. degree in the School of Computer Science, Tel-Aviv University

by

Yair Halevi

The research work for this thesis has been carried out at Tel-Aviv University under the supervision of Prof. Yishay Mansour

June 2007

Acknowledgements

I wish to thank Professor Yishay Mansour for his guidance during the work on my thesis. His comments, sound advice, good ideas, and diligent reviews of my writing were a great help, and I could not have hoped for a better supervisor. I am also grateful to him for introducing game theory to me and sparking my interest in this exciting research area.

Abstract

In a network creation game, initially proposed by Fabrikant et. al. [FLM+03], selfish players build a network by buying links to each other. Each player pays a fixed price per link $\alpha > 0$, and suffers an additional cost that is the sum of distances to all other players. We study an extension of this game where each player is only interested in its distances to a certain subset of players, called its friends.

We study the social optima and Nash equilibria of our game, and prove upper and lower bounds for the Price of Anarchy, the ratio between the social cost of the worst Nash equilibria and the optimal social cost. Our upper bound on the Price of Anarchy is

$$O\left(1 + \min\left(\alpha, \bar{d}_H, \log n + \sqrt{n\alpha/\bar{d}_H}, \sqrt{n\bar{d}_H/\alpha}\right)\right) = O(\sqrt{n}),$$

where n is the number of players, α is the edge building cost, and $\bar{d}_H$ is the average number of friends per player. We derive a lower bound of $\Omega(\log n/\log\log n)$ on the Price of Anarchy for a specific family of games which we construct. We also bound the Price of Stability, the ratio between the social cost of the best Nash equilibria and the optimal social cost. We provide additional bounds for specific friendship structures, namely a forest and a cycle. Finally, we show equivalent upper bounds on the Price of Anarchy for a weighted extension of our game.

Contents

Chapter 1

Introduction

In many natural settings entities form links to other entities, in a distributed and selfish manner, in order to increase their utility. A few examples include social networks, trade networks and most importantly communication networks such as the Internet. The structure of such networks, and in particular the Internet, has been recently of prime research interest. In this work we will focus primarily on the loss of efficiency due to the distributed and selfish behavior of the entities.

We model the network creation as a game between selfish players and focus on networks that are in equilibrium (no player can benefit by deviating). Our measure of the "loss of efficiency" would be the well known *Price of Anarchy* [KP99], which is the ratio between the social cost of the worst Nash equilibrium and that of the socially optimal solution. The concept of the Price of Anarchy was introduced by Koutsoupias and Papadimitriou in their seminal paper [KP99], and has been successfully studied for a wide range of problems including job scheduling, routing, facility location, and network design (see e.g. [BG00, JV01, CKV02, CV02, FKK$^+$02, RT02, ADTW03, HS03, ADK$^+$04, CSSM04]).

We study an extension of the network creation game proposed by Fabrikant et. al. [FLM$^+$03], where n players form a network by building edges to other players, and bought edges may be used by any player in both directions. The player's cost models both the "infrastructure cost" (of building the edges) and the "communication cost" (of reaching the other players). Formally, each player pays a fixed price $\alpha > 0$ for each edge it builds, and suffers an additional distance cost, which is the sum of distances to all other players in the resulting network. The social cost is the sum of the players' costs. This basic game and several extensions were studied in [FLM$^+$03, CP05, AEED$^+$06, AFM07].

By focusing on the distance to **all** players for the distance cost, the model implicitly assumes that the communication needs of every player are uniformly distributed over all other players. This may be far from true in many real world networks. For instance, Autonomous Systems in the Internet rarely communicate equally with every other Autonomous System. Our model extends the original game by specifying which pairs of players have non-negligible communication needs. We call such pairs of players *friends*, and model this new information using an undirected graph called the *friendship graph*, whose vertices are the players and edges connect pairs of friends. We initially assume that friendship is symmetric, hence our use of an undirected graph. Players, as before, build an edge to any other player (friend or non-friend) at a cost of α, however, their distance cost now includes only the distance to their friends (rather than the distance to all the players). The social cost, as before, is the sum of the players' costs. We call our game a *network creation game with nonuniform interests* (NI-NCG). It is obvious that the original network creation game of [FLM$^+$03] is a special case of our game, and we call it the *complete interest network creation game*. We also study a weighted extension of our game, where each player has a non-negative friendship weight assigned to every other player, and the distance cost for the player is the weighted sum of distances to all other players.

For the complete interest network creation game, [FLM$^+$03] demonstrated that both from a social perspective and from a selfish perspective, it is not always best to build edges to all friends (i.e., a complete graph). In an NI-NCG we may ask also whether there are cases in which players prefer to build edges to non-friends. In other words, is the network resulting from either socially optimal design, or selfish uncoordinated design, always a subgraph of the friendship graph. As it turns out, this is not the case, and there are situations where it is beneficial to build edges to non-friends.

1.1 Our Results

We first prove, by explicit construction, the existence of pure Nash equilibria for most values of α, specifically for $\alpha \leq 1$ and $\alpha \geq 2$, for $\alpha \in (1,2)$ when the friendship graph has girth at least 6, and for $\alpha \in (1, 1.5)$ when the friendship graph has girth 5. This is an important step in our analysis, since the Price of Anarchy is not defined if no Nash equilibrium exists. Note that while the existence of a Nash equilibrium in mixed strategies is always guaranteed [Nas51], this is not the case for pure strategy Nash equilibria, which is our primary focus. We remark that for games where $\alpha \in (1,2)$ and the friendship graph has girth 3 or 4, or where $\alpha \in [1.5, 2)$ and the friendship graph has girth 5, we have not resolved the question whether some pure Nash equilibrium is always guaranteed to exist.

Our main result is an upper bound of $O\left(1 + \min\left(\alpha, \bar{d}_H, \log n + \sqrt{n\alpha/\bar{d}_H}, \sqrt{n\bar{d}_H/\alpha}\right)\right) = O(\sqrt{n})$, on the Price of Anarchy of any NI-NCG, where n is the number of players and $\bar{d}_H$ is the average degree of the friendship graph H, i.e., the average number of friends per player. For either $\alpha = O(1)$, $\bar{d}_H = O(1)$, or $\alpha = \Omega(n\bar{d}_H)$, this upper bound is a constant. By comparison, Albers et. al. [AEED+06] show that for the complete interest game, the Price of Anarchy is a constant for $\alpha = O(\sqrt{n})$ and $\alpha \geq 12n\log n$, and is at most $O\left(1 + \left(\min\left(\frac{\alpha^2}{n}, \frac{n^2}{\alpha}\right)\right)^{1/3}\right) = O(n^{1/3})$ for any α. As part of the proof of our bound we also show that the ratio of the edge building cost to the social optimum cost is at most $O(\log n)$. We show an improved $O\left(\log n / \log\left(4\frac{\bar{d}_H}{\alpha}\log n\right)\right)$ bound on this ratio when $\alpha < 4\bar{d}_H \log n$. (In contrast, [AEED+06] show that for the complete interest game this ratio is only a constant.)

For a lower bound, we construct a family of games, for any large enough number of players n, for which the Price of Anarchy is $\Omega(\log n / \log\log n)$. Our construction relies on well known constructions of regular graphs of a given girth and degree with a minimal number of vertices, also known as cage graphs. The resulting games have $\alpha = \Theta(\log n / \log\log n)$ and $\bar{d}_H = \Omega(\log n)$, and they have the special property that some Nash equilibrium is achieved when players build exactly the friendship graph, i.e., every pair of friends has an edge connecting them. Our lower bound construction also gives a lower bound of $\Omega(\log n / \log\log n)$ on the ratio of the edge building cost to the social optimum. Hence, our upper and lower bounds for this ratio are tight. We remark that no non-constant lower bound for the Price of Anarchy in the complete interest game is known.

We provide bounds on the *Price of Stability* (the social cost ratio between the best Nash equilibrium and the social optimum), and show that it is at most 2. We also study games with specific classes of friendship graphs. For games with an acyclic friendship graph (i.e., a forest), we show that the friendship graph itself is both the unique Nash equilibrium and a social optimum, hence the Price of Anarchy and Price of Stability are both always 1. For games where the friendship graph is a simple cycle, we show upper and lower bounds in the range $[1, 2)$ for the Price of Anarchy. We also use such games to demonstrate that an NI-NCG may have a Price of Stability that is not 1 even for non-constant α, i.e., an NI-NCG may have a social optimum that is not a Nash equilibrium.

For the weighted extension of our game, we show equivalent upper bounds on the Price of Anarchy to those achieved for the un-weighted variant. Our results allow us to deduce upper bounds on the Price of Anarchy for several special cases: an asymmetric un-weighted variant, and (2) a distributional variant where each player has a probability distribution over all other players, and its distance cost is the expected distance to other players under this distribution. A weighted variant of the game was also studied in [AEED$^+$06], but the model used there did not permit zero weights, and the upper bound proven grows to $O(n)$ (a trivial bound) as w_{min}, the minimum weight in the game, approaches 0, making their bounds unapplicable for our game.

1.2 Organization

The remainder of this paper is organized as follows. In Chapter 2 we give some background and briefly review related research. In Chapter 3 we define our model and notation. Chapter 4 presents basic results for an NI-NCG, including Nash equilibria existence and social optimum bounds. In Chapter 5 we prove our main upper bound on the Price of Anarchy, and in Chapter 6 a lower bound is shown. Chapter 7 contains bounds for the Price of Stability. Chapter 8 studies NI-NCGs for specific families of friendship graphs, and Chapter 9 studies the weighted extension of our game.

Chapter 2

Background and Related Work

The research in this paper is part of a growing body of work on network creation games and computational game theory in general. We provide some background on these fields and review related research. Note that formal definitions of concepts used in this paper are provided in Chapter 3, not in this section.

2.1 Non-Cooperative Games and Nash Equilibria

Non-cooperative game theory studies situations where multiple players make decisions aiming at maximizing their returns, and their returns are the outcome of the collective decisions made by all players. In a game with pure strategies, each player may select an action from a certain action space, and the return of a player is a function (per player) of the actions selected by all players. In a game with mixed strategies, each player may select a probability distribution over its action space, and the return of a player is the expected return under the joint probability distribution over the collective action space. The solution concept of a *Nash equilibrium*, introduced by Nash [Nas51], offers a model of an "optimal state" in such games, and is probably the most widely used and well studied solution concept in this field. A Nash equilibrium is optimal in the sense that no player can gain by unilaterally changing only her own strategy. Nash showed in [Nas51] that in a non-cooperative game with mixed strategies and a finite number of players and actions, a Nash equilibrium is guaranteed to exist. This is not generally true, however, for games that allow only pure strategies.

If the players employ a *greedy best response algorithm*, in which they repeatedly modify their strategies to maximize their return if it is not already maximal, given the current strategies of the other players and in an uncoordinated manner, then Nash equilibria are absorbing states of the resulting dynamic. Even though convergence to a single state is not generally guaranteed in such a scenario, Nash equilibria are sometimes reasonable "candidates" for the final state of such a dynamic. Our interest in the "social quality" of Nash equilibria follows from this property.

Additional solution concepts for games have been developed as well. Among these is the concept of a *strong equilibrium*, proposed by Aumann [Aum59], which is a state where no coalition of players can deviate and improve the return of all its members. It isolates the effects of selfishness from the effects of lack of coordination, focusing on the former. Another interesting solution concept is that of a *correlated equilibrium* [Aum74], where a probability distribution "recommends" joint strategies to the players, and no player gains from deviating given that all other players are following the recommendation. Note that every strong equilibrium is a Nash equilibrium, and every Nash equilibrium is a correlated equilibrium. A *sink equilibrium* [GMV05] extends the concept of a Nash equilibrium to a set of states rather than requiring a single equilibrium state. A sink equilibrium is defined as a strongly connected component in the best response state transition graph, that has no outgoing edges. In other words, it is a set of states for which any greedy best response dynamic starting in the set may reach any other state in the set, but never leave it. These concepts are beyond the scope of this paper, and we mention them for completeness only and as suggestions for further research regarding our game.

2.2 Computational Game Theory

Much of the classical research in computer science is focused on solving certain optimization problems. In the classical setting, our task is to devise an efficient algorithm that finds an optimal solution to a specific problem and is to be carried out by a machine (or several machines), or to understand the computational characteristics of the problem in such a setting. Most of this work is based on the underlying assumption that all machines (or agents) involved in the solution of the problem share a single goal - the optimization of a single global (or social) function. In many cases, however, the various agents involved are selfish and are motivated by optimizing their own utilities, rather than reaching a solution that is socially optimal. This is the classical setting of game theory. The intersection of these two paradigms, computer science and game theory, has spawned a field of research known as computational game theory, which has been of much interest recently.

A basic question we may ask when attempting to optimize some social function using distributed and selfish agents is: how good can the social outcome be, compared to what a central authority could achieve? How bad? Such questions usually serve as a prelude to more complex algorithmic and computational questions. A key metric used to answer some of these questions is the *Price of Anarchy* [KP99], which is defined as the ratio between the social cost of the worst Nash equilibrium and the optimal social cost. The Price of Anarchy measures the worst-case "loss of efficiency" due to selfishness and lack of coordination. Closely related is the Price of Stability [ADK$^+$04], the ratio between the social cost of the best Nash equilibrium and the social optimum. The Price of Stability models the best loss of efficiency that can be guaranteed by an initial coordinated effort, followed by selfish, uncoordinated behavior.

Since their introduction, the Price of Anarchy and Price of Stability have been studied for a wide range of problems, including load balancing [KP99, CKV02, CV02, AART03], routing [MS01, Rou02a, Rou02b, RT02], facility location [Vet02], and network formation [BG00, JV01,

FKK+02, ADTW03, FLM+03, HS03, ADK+04, CSSM04, AEED+06]. Similar metrics have been recently researched for other game theoretic solution concepts, such as the *Strong Price of Anarchy* and the *Strong Price of Stability* [AFM07], and the *Price of Sinking* [GMV05], but these are beyond the scope of this paper.

2.3 Network Creation Games

The formation of networks by distributed and selfish agents has been the focus of much research lately. Anshelevich et. al. [ADTW03] study a game where players aiming to connect a certain set of terminals, select edges to purchase in some underlying graph and offer a payment for each such edge. An edge is built in the resulting network if and only if the sum of the payments made for it by all players is at least the edge's cost. It is shown in [ADTW03] that it is not guaranteed that a pure strategy Nash equilibrium exists in such a game, and that even if it does, both the Price of Anarchy and Price of Stability can be quite high (as high as the n, the number of players). They present, however, algorithms for finding approximate Nash equilibria that are close to the social optimum. Anshelevich et. al. [ADK+04] consider a similar game where each player selects only a set of edges she wishes to purchase, and the cost of an edge is shared equally among all players who purchased it. They show that a pure strategy Nash equilibrium always exists, and that the Price of Stability is at most $H(n) = \sum_{i=1}^{n} \frac{1}{i}$ (although the Price of Anarchy may still be as high as n). They also extend these results to games with design costs and latency costs, under certain conditions. Additional research on cost sharing and network design includes [JV01, GST04, JMT06].

The model used by [ADTW03, ADK+04] allows agents to purchase any link in the network. The model we are interested in is different in that each player is a certain network node, and we allow a player to purchase only links incident on it. Bala and Goyal [BG00] study such a model where the payoff of each player decreases with the number of links it builds, and increases with the number of other players in its connected component in the resulting network. They also provide some results for a similar model where the benefit of being connected to some player exhibits exponential decay depending on the distance from the player. Haller and Sarangi [HS03] extend this model by allowing heterogenous link failure probabilities instead of distance-based exponential decay. The models of [BG00, HS03] seem especially suitable to information exchange networks (such as social networks), where players benefit from the collective information made available to them by the other players in their network, but the actual cost of communication over the links is negligible.

The game proposed by Fabrikant et. al. [FLM$^+$03] and further studied in [AEED$^+$06], which is the basis of the game studied in this paper, is similar to [BG00, HS03] in that players are network nodes and may only purchase links incident on themselves. However, in their game, players suffer an additional cost that is the sum of distances to all other players in the resulting network. The game is therefore suitable for scenarios where communications incur a latency or cost proportional to the hop count in the shortest route available. The results obtained in [FLM$^+$03, AEED$^+$06] were reviewed in the introduction to this paper (see Chapter 1). Albers et. al. [AEED$^+$06] also studies a cost sharing variant of this game where each player may pay a fraction of an edge's cost, and the edge is built only if the sum of fractional payments exceeds 1. Another cost-sharing variant is studied by Corbo and Parkes [CP05], where players must both consent to building an edge between them and they share the edge's cost equally. Their work analyzes the game in the face of pairwise coalitions, and not just unilateral deviations. The effects of coordination on the game from [FLM$^+$03] are further studied by Andelman et. al. [AFM07]. They use the concept of a Strong Equilibrium and define its related Strong Price of Anarchy, to measure the cost of selfishness alone, when coordination is permitted. They show that for almost all values of α a strong equilibrium exists, and that the Strong Price of Anarchy is at most 2.

Chapter 3

The Model

A *network creation game with nonuniform interests* (NI-NCG), is a tuple $\mathcal{N} = \langle V, H, \alpha \rangle$ where $V = \{1 \dots n\}$ is the set of players, $H = (V, E_H)$ is an undirected graph whose vertices are V and edges are E_H (H is called the *friendship graph*) and $\alpha \in \mathbb{R}^+$ is an *edge building price*. A *strategy* S_v for a player v is a subset of the other players $S_v \subseteq V - \{v\}$ (these are the players to which v builds an edge). A *joint strategy* $\vec{S}$ is an n-tuple of player strategies, i.e., $\vec{S} = (S_1, S_2, \dots, S_n)$.

Given an NI-NCG, we say that players v and u are *friends* iff $(v, u) \in E_H$, and player v's neighborhood in H, denoted $N_H(v) = \{u : (v, u) \in E_H\}$, is also called the *friend set* of v. Given a joint strategy $\vec{S}$, we associate a *joint strategy graph* $\mathcal{G}^*(\vec{S}) = (V, E)$ with vertices V and directed edges $\{(v, u) : v, u \in V \wedge u \in S_v\}$. The *network* created by $\vec{S}$, denoted $G = \mathcal{G}(\vec{S})$, is the underlying undirected graph of $\mathcal{G}^*(\vec{S})$. We limit our discussion to NI-NCGs for which H has no isolated vertices (and hence $n \geq 2$ and the number of connected components is at most $n/2$). The special case $H = K_n$ (the complete graph over n vertices) is called the *complete interest network creation game*, and was extensively studied in [FLM$^+$03, AEED$^+$06, AFM07].

The *cost* of a player $v \in V$ in a joint strategy $\vec{S}$ is $C^{\mathcal{N}}(\vec{S}, v)$ and it is the sum of two components: an *edge building cost* $B^{\mathcal{N}}(\vec{S}, v) = \alpha|S_v|$, which implies a cost of α per edge bought, and (2) a *distance* *cost* which is $Dist^{\mathcal{N}}(\vec{S}, v) = \sum_{u \in N_H(v)} \delta_G(v, u)$, where $\delta_G(v, u)$ is the distance between v and u in $G = \mathcal{G}(\vec{S})$ (if there is no path in G between v and u then $\delta_G(v, u) = \infty$).

We define the *social cost* of a joint strategy $\vec{S}$ as the sum of the player costs, i.e., $C^{\mathcal{N}}(\vec{S}) = \sum_{v \in V} C^{\mathcal{N}}(\vec{S}, v)$, and the *social edge building cost* and *social distance cost* are defined similarly. The minimal social cost of all joint strategies of an NI-NCG $\mathcal{N}$ is denoted $OPT(\mathcal{N})$, and any joint strategy $\vec{S}$ that yields this cost is a *social optimum strategy* for $\mathcal{N}$. We say that a graph G is a *social optimum network* if there is a social optimum strategy $\vec{S}$ that creates G, i.e., $G = \mathcal{G}(\vec{S})$.

Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ be a joint strategy, and $v \in V$ some player. We denote by $\vec{S}_{-v}$ the tuple of strategies of players $u \neq v$ in $\vec{S}$, i.e., $\vec{S}_{-v} = (S_1, \ldots, S_{v-1}, S_{v+1}, \ldots, S_n)$. Let $A_v \subseteq V - \{v\}$ be some strategy for v. A *unilateral deviation of v to A_v in $\vec{S}$*, denoted $(\vec{S}_{-v}, A_v)$, is the joint strategy obtained by replacing the strategy for v in $\vec{S}$ by A_v, i.e.,$(\vec{S}_{-v}, A_v) = (S_1, \ldots, S_{v-1}, A_v, S_{v+1}, \ldots, S_n)$.

Given a game $\mathcal{N}$ and joint strategy $\vec{S}$, the best response of player v to $\vec{S}_{-v}$ is $BR(\vec{S}, v) = \arg\min_{A_v \subseteq V - \{v\}} C^{\mathcal{N}}((\vec{S}_{-v}, A_v), v)$. A joint strategy $\vec{S}$ is a Nash equilibrium for the game $\mathcal{N}$ if every player v is playing a best response, i.e., $\forall v \in V, S_v \in BR(\vec{S}, v)$. We say that a network G is a *Nash equilibrium network*, if there exists a Nash equilibrium $\vec{S}$ that creates it, i.e., $G = \mathcal{G}(\vec{S})$. We will show that for almost all values of α (the edge building price), an NI-NCG has a Nash equilibrium.

Let $\Phi(\mathcal{N})$ be the set of Nash equilibria of an NI-NCG $\mathcal{N}$ and assume that $\Phi(\mathcal{N}) \neq \emptyset$. The Price of Anarchy (Price of Stability, respectively), denoted $PoA(\mathcal{N})$ $(PoS(\mathcal{N}))$, is the ratio between the maximal (minimal) social cost of a Nash equilibrium and the cost of a social optimum, i.e.,

$$PoA(\mathcal{N}) = \frac{\max_{\vec{S} \in \Phi(\mathcal{N})} C^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \quad \text{and} \quad PoS(\mathcal{N}) = \frac{\min_{\vec{S} \in \Phi(\mathcal{N})} C^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \quad .$$

Note that the Price of Anarchy and Price of Stability are defined only if some Nash equilibrium exists.

We present some additional terminology used throughout this paper. A *friendship edge* is an edge that connects two friends. A strategy A_v for $v \in V$ is called a *friendship strategy* if it builds edges only to v's friends, i.e., $A_v \subseteq N_H(v)$. Given a joint strategy $\vec{S}$, strategy A_v is called a *complete friendship response* for v in $\vec{S}$ if it builds edges exactly to all friends of v that have not built edges to v in $\vec{S}$, i.e., $A_v = \{u \in N_H(v) : v \notin S_u\}$.

A *sensible* joint strategy is a joint strategy in which no two players build edges to each other, i.e., $\vec{S}$ is sensible iff $\forall v, u \in V, u \in S_v \Rightarrow v \notin S_u$. Clearly, if a joint strategy $\vec{S}$ is not sensible, then it is neither a social optimum nor a Nash equilibrium. For any network G, we denote by $\overline{\mathbf{S}}(G)$ the set of sensible joint strategies that create the network G. Notice that the set of joint strategy graphs corresponding to $\overline{\mathbf{S}}(G)$ is exactly the set of all directed orientations of G, and that they all have exactly $|E(G)|$ edges.

We use the following graph-theoretical notations: for any undirected graph $G = (V_G, E_G)$, we denote by $d_G(v) = |N_G(v)|$ the degree of vertex v in G, and by $\bar{d}_G = \sum_{v \in V_G} d_G(v)/|V_G|$ the average vertex degree of G. The *girth* of G, denoted $g(G)$, is the minimal length of a cycle in G (if G is cycle free then $g(G) = \infty$).

Chapter 4

Basic Results

We begin our analysis by providing a tight bound on the social optimum and proving the existence of Nash equilibria. For the complete interest game ($H = K_n$), the complete graph is a social optimum network for $\alpha \leq 2$ and the star graph is a social optimum network for $\alpha \geq 2$ (see [FLM$^+$03]). For a general graph H, we do not provide a complete characterization of the social optimum network of the related NI-NCG, but rather only bound the minimal social cost as follows:

Theorem 4.1. *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG. Then:*

$$n(\alpha/2 + \bar{d}_H) \leq OPT(\mathcal{N}) < 2n(\alpha/2 + \bar{d}_H) \ .$$

Proof: Let G be some social optimum network. Clearly, G must have at least $n/2$ edges, because otherwise it has an isolated vertex, which has an infinite cost, and yields an infinite social cost, contradicting optimality. Therefore, $n\alpha/2$ is a lower bound on the social edge building cost. Clearly, the distance between any pair of friends is at least 1, hence the social distance cost is at least $2|E_H| = n\bar{d}_H$. Summing these two lower bounds yields the lower bound on $OPT(\mathcal{N})$. For the upper bound, let G' be a network that is a star graph on V, and let $\vec{S} \in \overline{\mathbf{S}}(G')$ be some sensible joint strategy that creates G'. Obviously $\vec{S}$ builds exactly $n-1$ edges. Additionally, the distance in G' between any pair of friends is at most 2, hence:

$$C^{\mathcal{N}}(\vec{S}) \leq (n-1)\alpha + 4|E_H| < 2n(\alpha/2 + \bar{d}_H) \ ,$$

which upper bounds the social optimum. □

The above theorem gives a 2-approximation of $OPT(\mathcal{N})$, without specifying the exact social optimum. We also have the following specific characterization for $\alpha \leq 2$, which follows since adding any friendship edge which is missing from the network increases the social edge building cost by $\alpha \leq 2$ and reduces the social distance cost by at least 2.

Theorem 4.2 (Proof in Appendix A). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG with $\alpha \leq 2$. Then H is a social optimum network, and any sensible joint strategy $\vec{S}$ that creates it (i.e., $\vec{S} \in \overline{\mathbf{S}}(H)$) is a social optimum. Furthermore, for $\alpha < 2$, H is the unique social optimum network.*

We also show the following constraint on the number of edges in a social optimum:

Lemma 4.3 (Proof in Appendix A). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a social optimum for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Then $|E(G)| \leq |E_H|$ with equality possible if and only if $G = H$.*

We proceed to the issue of existence of Nash equilibria, which is important for establishing Price of Anarchy. For the complete interest game, the complete graph K_n is a Nash equilibrium network for $\alpha \leq 1$, and it is unique for $\alpha < 1$ (see [FLM+03]). Since for $\alpha \leq 1$ a player cannot lose by adding a friendship edge, we have the following equivalent result for a general NI-NCG,

Theorem 4.4 (Proof in Appendix A). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG with $\alpha \leq 1$. Then H is a Nash equilibrium network, and any joint strategy $\vec{S} \in \overline{\mathbf{S}}(H)$ is a Nash equilibrium. Furthermore, for $\alpha < 1$, H is the unique Nash equilibrium network.*

Theorem 4.4 and Theorem 4.2 imply that for $\alpha < 1$, any Nash equilibrium is a social optimum, hence $PoA(\mathcal{N}) = PoS(\mathcal{N}) = 1$. We therefore focus in the remainder of our discussion on the case

$\alpha \geq 1$. For $\alpha \geq 1$ the star graph is a Nash equilibrium network for the complete interest game [FLM+03]. For a general NI-NCG we have the following result:

Theorem 4.5 (Proof in Appendix A). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG with $\alpha \geq 2$. There exists a Nash equilibrium for $\mathcal{N}$ whose cost is at most $n(\alpha + 2\bar{d}_H)$.*

Our proof gives a specific construction of a Nash equilibrium for $\alpha \geq 2$, where the distance between any two friends is at most 2 and the resulting network is a forest (but not necessarily a star), consisting of a spanning tree for every connected component of H. It also shows that for certain NI-NCGs, a Nash equilibrium may contain non-friendship edges.

For $\alpha \in (1, 2)$, we have not resolved the question whether some Nash equilibrium always exists. In Chapter 6 we show that if the girth of the friendship graph $g(H)$ is at least 6, existence of a Nash equilibrium is guaranteed for $\alpha \in (1, 2)$ as well, and if $g(H) = 5$, it is guaranteed for $\alpha \in (1, 1.5)$ (Lemma 6.3).

Chapter 5

Upper Bounds On The Price Of Anarchy

In this section we present several upper bounds on the Price of Anarchy and conclude by combining them into our main result, an upper bound of:

$$O\left(1 + \min\left(\alpha, \bar{d}_H, \log n + \sqrt{n\alpha/\bar{d}_H}, \sqrt{n\bar{d}_H/\alpha}\right)\right) = O(\sqrt{n}) \ .$$

Our results in this section apply to games for which a Nash equilibrium exists. We showed that this holds for $\alpha \geq 2$ (Theorem 4.5) and $\alpha \leq 1$ (Theorem 4.4) but not necessarily for $\alpha \in (1, 2)$. Our first step is showing a simple upper bound of $O(1 + \min(\alpha, \bar{d}_H))$, which is due to the following lemma:

Lemma 5.1. *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG and $\vec{S}$ a Nash equilibrium of $\mathcal{N}$. Then:*

$$C^{\mathcal{N}}(\vec{S}) \leq (\alpha + 1)n\bar{d}_H \ .$$

Proof: Every player $v \in V$ is playing a best response in $\vec{S}$, hence its strategy S_v is no worse than the complete friendship response for v. The cost of the complete friendship response for v is at most $\alpha d_H(v) + d_H(v) = (\alpha + 1)d_H(v)$, and summing over all players completes the proof. □

Using Lemma 5.1, Theorem 4.1 and Theorem 4.2, we obtain the following upper bound on the Price of Anarchy:

Theorem 5.2 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, and assume that a Nash equilibrium exists. Then:*

$$PoA(\mathcal{N}) \leq \begin{cases} 1 & \text{if} \quad 0 < \alpha < 1 \\ 2 - \frac{2}{\alpha+2} & \text{if} \quad 1 \leq \alpha \leq 2 \\ 1 + \frac{\bar{d}_H \alpha}{\alpha/2 + \bar{d}_H} & \text{if} \quad 2 < \alpha \end{cases}$$

$$= O(1 + \min(\alpha, \bar{d}_H)) \ .$$

This result already provides us with some basic insights. In particular, if either α or $\bar{d}_H$ are constant, then the Price of Anarchy is at most some constant.

5.1 Upper Bound: Edge Building Cost

In this section we show an upper bound for the edge building cost. For the complete interest game, the contribution of the edge building cost to the Price of Anarchy was shown to be at most a constant (see [AEED$^+$06]). We show a similar (but weaker) result for a general NI-NCG, specifically an upper bound of $O\left(1 + \min\left(1, \frac{\bar{d}_H}{\alpha}\right) \log n\right)$, which is at most $O(\log n)$. We begin by defining weights for edges built in a joint strategy:

Definition 5.3. Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a joint strategy of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. We denote by $\Psi_G(v, x)$ the set of shortest paths between vertices v and x in the graph G. For any $v, u \in V$ such that $u \in S_v$ (i.e., v builds an edge to u in $\vec{S}$), we define:

$$\Gamma_G^H(v, u) = \{x \in N_H(v) : \quad \Psi_G(v, x) \neq \emptyset \quad \textbf{and} \quad \forall p \in \Psi_G(v, x), u \in p\} \ ,$$

i.e., $\Gamma_G^H(v,u)$ is the set of friends of v whose distance from v would increase should the edge (v,u) be removed from the network. We define $w(v,u) = |\Gamma_G^H(v,u)|$, and $W = \sum_{v \in V, u \in S_v} w(v,u)$.

Intuitively, the weight $w(v,u)$ measures effectiveness of the edge (v,u). Note that while w and W are dependent on $\mathcal{N}$ and $\vec{S}$, for simplicity of our notation we ignore this dependance. Next, we show a constraint on these weights. The proof of the following lemma is by analyzing a specific deviation of v that removes the edge (v,u).

Lemma 5.4 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. If v builds an edge to u in $\vec{S}$ (i.e., $u \in S_v$), then $\alpha \leq w(v,u)(\delta_{G'}(v,u) - 1)$, where G' is the graph G without the edge (v,u).*

We use the above lemma to bound the average degree in G, based on a parameter β (that we select later).

Lemma 5.5. *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a Nash equilibrium of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. Then for any $\beta > 0$*

$$\bar{d}_G \leq 1 + \frac{2\bar{d}_H}{\beta} + n^{2\beta/\alpha} \quad .$$

Proof: Our first step is to transform G by removing all edges $(v,u) \in E(G)$ for which $w(v,u) \geq \beta$ (where v builds the edge in $\vec{S}$). We denote the resulting graph by G_β, and by m the number of removed edges. The total weight of edges removed is at least βm, and at most the total weight of edges in the original graph, W. Therefore, $\beta m \leq W$. On the other hand, for any pair of friends v, x, there is at most one $u \in S_v$ such that $x \in \Gamma_G^H(v,u)$ (these sets, for a fixed player v, must be pairwise disjoint). Therefore every such pair is counted at most twice in W (at most once in each direction), hence $W \leq 2|E_H|$. Combining these inequalities we get $|E(G_\beta)| \geq |E(G)| - 2|E_H|/\beta$. After dividing by $n/2$, we have

$$\bar{d}_{G_\beta} \geq \bar{d}_G - \frac{2\bar{d}_H}{\beta} \quad . \tag{5.1}$$

For $\bar{d}_{G_\beta} \leq 2$ this yields $\bar{d}_G \leq \frac{2\bar{d}_H}{\beta} + 2$ and the claim holds trivially since $n^{2\beta/\alpha} > 1$. We therefore assume that $\bar{d}_{G_\beta} > 2$. This implies that G_β has a cycle. Let C be a cycle in G_β of minimal length $g(G_\beta)$, and note that C is also a cycle in G. Let (v, u) be an arbitrary edge in C, built by v, and recall that we have $w(v, u) \leq \beta$. Using Lemma 5.4 for this edge we now obtain $\alpha \leq w(v, u)(g(G_\beta) - 2) \leq \beta(g(G_\beta) - 2)$, or $g(G_\beta) \geq 2 + \frac{\alpha}{\beta}$. We now use a bound due to Alon et. al. [AHL02] on the number of vertices in a graph based on its girth and average degree. This bound is an extension of the *Moore bound* (see [BB94], p. 180) for irregular graphs (see Appendix F). Restated as a bound on average degree, it says that the average degree of a graph with n vertices and girth g is at most $1 + n^{2/(g-2)}$. Therefore, for our graph G_β we have:

$$\bar{d}_{G_\beta} \leq 1 + n^{\left(\frac{2}{g(G_\beta)-2}\right)} \leq 1 + n^{2\beta/\alpha} \quad ,$$

and using this with inequality (5.1) completes the proof. □

Finally, by using the above lemma with $\beta = \frac{\alpha}{2\log_2 n}$ and the lower bound on social optimum (Theorem 4.1), we obtain our improved bound:

Theorem 5.6 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, and $\vec{S}$ a Nash equilibrium of $\mathcal{N}$. Then:*

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq 3 + \frac{\bar{d}_H}{\alpha/2 + \bar{d}_H}(2\log_2 n - 3) = O\left(1 + \min\left(1, \frac{\bar{d}_H}{\alpha}\right)\log n\right) = O(\log n) \quad .$$

5.2 Upper Bound: Distance Cost

In this section we show an upper bound on the contribution of the distance cost to the Price of Anarchy. The bound we show is $O\left(1 + \min\left(1, \frac{\bar{d}_H}{\alpha}\right)\min\left(n, \sqrt{n\alpha/\bar{d}_H}\right)\right)$, which is at most $O(\sqrt{n})$. We first show several lemmata:

Lemma 5.7 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a joint strategy and $G = \mathcal{G}(\vec{S})$. Let $v \in V$ be a player who is playing a best response in $\vec{S}$. Let T be a shortest path* *tree of v in G. For any vertex $u \in V(T)$ (note that it is possible that $V(T) \neq V$), denote by k_u the number of friends of v that are descendants of u in T (including, possibly, u itself). Then:*

$$\forall u \in V(T), \alpha \geq (\delta_G(v,u) - 1)k_u \ .$$

The proof of Lemma 5.7 is done by an analysis of a specific deviation by v that builds an additional edge to u. Using Lemma 5.7 we bound the distance cost for a single player as follows.

Lemma 5.8. *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG and let $\vec{S}$ be a joint strategy. Let $v \in V$ be a player who is playing a best response in $\vec{S}$. Then:*

$$Dist^{\mathcal{N}}(\vec{S}, v) \leq 2\sqrt{(n-2)\alpha d_H(v)} + d_H(v) \ .$$

Proof: Let $G = \mathcal{G}(\vec{S})$. All $w \in N_H(v)$ must be in the same connected component as v (otherwise v could have strictly improved its cost by building edges to all disconnected friends). For any integer $h \geq 2$, we define $F_h = \{w \in N_H(v) : \delta_G(v,w) \geq h\}$, i.e., F_h is the set of friends of v in depth at least h in T, where T is the shortest path tree of v in G. Then:

$$\begin{aligned} Dist^{\mathcal{N}}(\vec{S}, v) &= \sum_{w \in N_H(v)/F_h} \delta_G(v,w) + \sum_{w \in F_h} \delta_G(v,w) \\ &\leq (h-1)(d_H(v) - |F_h|) + \sum_{w \in F_h} \delta_G(v,w) \\ &= (h-1)d_H(v) + \sum_{w \in F_h} (\delta_G(v,w) - h + 1) \ . \end{aligned} \tag{5.2}$$

For any $w \in F_h$, the term $\delta_G(v,w) - h + 1$ is exactly the number of vertices on the path between depth h in T and w. Therefore, summing this value over all $w \in F_h$ is equivalent to counting each vertex $u \in V(T)$ in depth h or greater exactly k_u times. Hence, using Lemma 5.7 and the fact that $h \geq 2$, we have:

$$\sum_{w \in F_h} \delta_G(v,w) - h + 1 = \sum_{u \in V(T):\delta_G(v,u) \geq h} k_u \leq \sum_{u \in V(T):\delta_G(v,u) \geq h} \frac{\alpha}{\delta_G(v,u) - 1} \leq \frac{(n-2)\alpha}{h-1} \ ,$$

and using this in inequality (5.2) yields:

$$Dist^{\mathcal{N}}(\vec{S}, v) \leq (h-1)d_H(v) + \frac{(n-2)\alpha}{h-1} \ ,$$

for any $h \geq 2$. Selecting $h = 1 + \lceil \sqrt{(n-2)\alpha/d_H(v)} \rceil$ completes the proof. □

Lemma 5.8, summed over all players, yields the following theorem:

Theorem 5.9 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG and $\vec{S}$ a Nash equilibrium for $\mathcal{N}$. Then:*

$$\begin{aligned} \frac{Dist^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} &\leq \frac{\bar{d}_H}{\alpha/2 + \bar{d}_H} \left(1 + \min(n-2, 2\sqrt{(n-2)\alpha/\bar{d}_H})\right) \\ &= O\left(1 + \min\left(1, \frac{\bar{d}_H}{\alpha}\right) \min\left(n, \sqrt{\frac{n\alpha}{\bar{d}_H}}\right)\right) = O(\sqrt{n}) \ . \end{aligned}$$

5.3 Upper Bound: Price Of Anarchy

We now present our main upper bound, which follows by combining Theorem 5.6 and Theorem 5.9.

Theorem 5.10 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, and assume that a Nash equilibrium exists. Then:*

$$\begin{aligned} PoA(\mathcal{N}) &\leq \begin{cases} 1 & \text{if} \quad 0 < \alpha < 1 \\ 2 - \frac{2}{\alpha+2} & \text{if} \quad 1 \leq \alpha \leq 2 \\ 3 + \frac{\bar{d}_H}{\alpha/2+\bar{d}_H} \min\left(\alpha,\ 2\log_2 n - 2 + \min\left(n-2,\ 2\sqrt{(n-2)\alpha/\bar{d}_H}\right)\right) & \text{if} \quad 2 < \alpha \end{cases} \\ &= O\left(1 + \min\left(\alpha,\ \bar{d}_H,\ \log n + \sqrt{\frac{n\alpha}{\bar{d}_H}},\ \sqrt{\frac{n\bar{d}_H}{\alpha}}\right)\right) = O(\sqrt{n}) \ . \end{aligned}$$

As stated in the theorem, in all cases we have an upper bound of $O(\sqrt{n})$ on the Price of Anarchy, but for many combinations of n, α and $\bar{d}_H$ we get a better bound. For example, we get a constant upper bound in the case that either α or $\bar{d}_H$ are constant, and similarly in the case that $\alpha = \Omega(n\bar{d}_H) = \Omega(|E_H|)$, i.e., α is very large. When α is moderately small, specifically $\alpha = O(\max(\log n, n/\bar{d}_H))$ the bound is simply $O(\min(\alpha, \bar{d}_H))$. When the friendship graph is dense enough, namely $\bar{d}_H = \Omega(n/\log n)$, and α is between $\Omega(\log n)$ and $O(\bar{d}_H \log^2 n/n) = O(\log^2 n)$ we get an upper bound of $O(\log n)$. In the remaining cases we have an upper bound of $O\left(\sqrt{n \min\left(\frac{\bar{d}_H}{\alpha}, \frac{\alpha}{\bar{d}_H}\right)}\right)$.

5.4 Additional Edge Building Cost Bounds

We show two additional improvements of our upper bound on the contribution of the edge building cost to the Price of Anarchy. The first results from a different choice of β for Lemma 5.5.

Theorem 5.11 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Denote $\gamma = 4\frac{\bar{d}_H}{\alpha}\log_2 n$. If $\gamma > 1$ then:*

$$\begin{aligned}\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} &\leq \frac{\alpha}{\alpha + 2\bar{d}_H}\left(1 + \eta\frac{\gamma}{\log_2 \gamma}\right)\\ &= O\left(\min\left(1, \frac{\bar{d}_H}{\alpha}\right)\frac{\log n}{\log\left(4\frac{\bar{d}_H}{\alpha}\log n\right)}\right) = O\left(\frac{\log n}{\log\left(4\frac{\bar{d}_H}{\alpha}\log n\right)}\right) \quad ,\end{aligned}$$

where $\eta = 2 + \frac{1}{e\ln(2)-1} \approx 3.131$.

Later we show that this result is tight (see Chapter 6). We also have the following upper bound:

Lemma 5.12 (Proof in Appendix B). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Denote by V_R the set of players that have friends that are not adjacent to them in G and by V_r the set of players that build at least one edge to a non-friend. Let $R_H(\vec{S}) = |V_R|$ and $r_H(\vec{S}) = |V_r|$. Then:*

$$|E(G)| \leq |E_H| - \frac{R_H(\vec{S}) - r_H(\vec{S})}{2} \leq |E_H| \quad ,$$

and:

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\bar{d}_H\alpha}{\alpha + 2\bar{d}_H} = O(\min(\alpha, \bar{d}_H)) \quad .$$

Note that for small values of α, specifically $\alpha = o(\log n)$, Lemma 5.12 provides a better upper bound on the contribution of the edge building cost to the Price of Anarchy than Lemma 5.6. This does not improve our asymptotic upper bound on the Price of Anarchy from Theorem 5.10, however, since that bound is already $O(\min(\alpha, \bar{d}_H))$.

Chapter 6

Lower Bound: NI-NCG with $PoA = \Omega(\log n / \log\log n)$

We now proceed to show that for large enough n there is an NI-NCG with n players for which the Price of Anarchy is bounded from below by $\Omega(\log n/\log\log n)$ (we remark that a similar lower bound result for the complete interest game is not known). Our construction is based on the observation that if the friendship graph of the game has large enough girth, then it is a Nash equilibrium network itself. The existence of a graph with large enough girth and large enough average degree follows using well known constructions of cage graphs. Our first step is to prove the aforementioned sufficient condition for Nash equilibria. Before that, we need the following lemmata:

Lemma 6.1 (Proof in Appendix C). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a joint strategy of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. Let $v \in V$ be some player and $u \in N_H(v)$ a friend of v. If v's strategy, S_v, is a best response in $\vec{S}$, then $\delta_G(v,u) \leq \min(\alpha+1, n-1)$.*

Lemma 6.2 (Proof in Appendix C). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a joint strategy and $G = \mathcal{G}(\vec{S})$. Let v be a player that is playing a best response in $\vec{S}$. If v builds an edge to u, and u is* *not a friend of v, then* $|\Gamma_G^H(v,u)| \geq 2$.

The key condition for our construction is the following:

Lemma 6.3. *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, such that $0 < \alpha < \frac{g(H)}{2} - 1$. Then H is a Nash equilibrium network, and any $\vec{S} \in \overline{\mathbf{S}}(H)$ is a Nash equilibrium.*

Proof: Let $\vec{S}$ be some sensible strategy that creates H, i.e., $\vec{S} \in \overline{\mathbf{S}}(H)$. We will show that all players are using a best response strategy in $\vec{S}$. Let $v \in V$ be a player, and assume, for contradiction, that v's strategy in $\vec{S}$ is not a best response. Then there is an alternative best response strategy $S'_v \neq S_v$ for v. Let $\vec{S'} = (\vec{S}_{-v}, S'_v)$ and $G' = \mathcal{G}(\vec{S'})$. We analyze two cases:

- S'_v is not a friendship strategy, i.e., $S'_v - N_H(v) \neq \emptyset$

  In this case there is a player $u \in S'_v$ that is not a friend of v. By Lemma 6.2 we have $|\Gamma_{G'}^H(v,u)| \geq 2$. We arbitrarily select $w, x \in \Gamma_{G'}^H(v,u)$ such that $w \neq x$. Vertices u, w, x are all in the same connected component of G' (by definition of $\Gamma_{G'}^H(v,u)$), so let P_{uw}, P_{ux} and P_{wx} be shortest paths in G' between u and w, u and x, and w and x respectively. By the triangle inequality we have $|P_{wx}| \leq |P_{uw}| + |P_{ux}|$. Also, if v were included in P_{wx}, it would split it into two shortest paths from v to w and from v to x respectively. By definition of $\Gamma_{G'}^H(v,u)$, this would imply that both these segments of P_{wx} contain u, which is a contradiction to it being a shortest path. Therefore, P_{wx} does not include v, hence it is also a path in H. We now construct a simple cycle in H by starting with the edge (v,w), following path P_{wx} to x and concluding with the edge (x,v). By definition, the length of this cycle must be at least $g(H)$, and the lemma assumes that $g(H) > 2\alpha + 2$, so:

  $$\delta_{G'}(v,w) + \delta_{G'}(v,x) = 2 + |P_{uw}| + |P_{ux}| \geq 2 + |P_{wx}| \geq g(H) > 2\alpha + 2 \ ,$$

  which is a contradiction to Lemma 6.1 for either (v,w) or (v,x).

- S'_v is a friendship strategy, i.e., $S'_v \subseteq N_H(v)$

  We know that $S'_v \neq S_v$, and S_v is a complete friendship response for v in $\vec{S}$, so v has a friend $u \in N_H(v)$ that is not adjacent to v in G'. Let P_{uv} be some shortest path from v to u in G' (such a path must exist since S'_v is a best response). Since in this case $E(G') \subseteq E_H$, P_{uv} is also fully contained in H. Adding the edge (u, v) to P_{uv} we obtain a simple cycle in H, whose size must be at least $g(H)$. Therefore:

$$\delta_{G'}(v, u) \geq g(H) - 1 > 2\alpha + 1 > \alpha + 1 \ ,$$

  which is a contradiction to Lemma 6.1.

Therefore no deviation to $S'_v \neq S_v$ can produce a best response, so S_v itself must be a best response. Since this is true for every $v \in V$, $\vec{S}$ is a Nash equilibrium. Note that the proof also holds when $g(H) = \infty$. □

We now construct games satisfying this condition, with large enough average degree (of the friendship graph). To do this we use well known constructions of cage graphs - regular graphs of a given girth and degree with a minimal number of vertices (see page 107 in [Bol78] for a history of results in this area). We shall use the following proposition (see Appendix F):

Proposition 6.4. *For any integer $g \geq 3$ and $n \geq 2^g$, there is a graph with n vertices and girth g of average degree at least $\frac{1}{2}(n/4)^{1/(g-2)}$.*

We are now ready to state our lower bound.

Theorem 6.5 (Proof in Appendix C). *For any integer $n \geq 64$ there is an NI-NCG $\mathcal{N} = \langle V, H, \alpha \rangle$ with n players, $\alpha = \Theta(\log n/ \log\log n)$, and $\bar{d}_H = \Omega(\log n)$, whose Price of Anarchy is $\Omega(\log n/ \log\log n)$.*

Our proof also implicitly shows that the contribution of the edge building cost to the Price of

Anarchy for the construction is $\Omega(\log n/\log\log n)$. Since $\alpha = o(\bar{d}_H)$, we know that for large enough n we must have $\alpha < 4\bar{d}_H < 4\bar{d}_H \log_2 n$, so Theorem 5.11 applies to our construction. An upper bound on this contribution is therefore:

$$O\left(\log n/\log\left(\frac{4\bar{d}_H}{\alpha}\log n\right)\right) = O(\log n/\log\log n) \ ,$$

since $4\bar{d}_H/\alpha > 1$, hence our upper bound from Theorem 5.11 is tight.

Note that Lemma 6.3 guarantees the existence of a Nash equilibrium for $\alpha \in (1,2)$ if $g(H) \geq 6$, and for $\alpha \in (1,1.5)$ if $g(H) = 5$. For the class of games with $g(H) \in \{3,4\}$ and $\alpha \in (1,2)$ and games with $g(H) = 5$ and $\alpha \in [1.5,2)$, the existence of Nash equilibria remains unresolved.

Chapter 7

Bounding The Price Of Stability

For the complete interest game, it was shown (in [FLM$^+$03]) that for $\alpha \leq 1$ and $\alpha \geq 2$ there is a social optimum that is also a Nash equilibrium, hence the Price of Stability is 1. In a general NI-NCG, we have the following result:

Theorem 7.1. *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG. Then:*

$$1 \leq PoS(\mathcal{N}) \leq \begin{cases} 1 & \textit{if} \quad \alpha \leq 1 \\ 2 - \frac{2}{\alpha+2} & \textit{if} \quad 1 < \alpha < 2 \\ 2 & \textit{if} \quad \alpha \geq 2 \end{cases} .$$

Proof: The case $\alpha \leq 1$ is trivial given that in this case H is both a social optimum network (Theorem 4.2) and a Nash equilibrium network (Theorem 4.4). For $1 < \alpha < 2$, the claim follows from Theorem 5.2, because the Price of Stability is not greater than the Price of Anarchy.

If $\alpha \geq 2$ then we have shown in Theorem 4.5 that there is a Nash equilibrium whose cost is at most $n(\alpha + 2\bar{d}_H)$, so using the lower bound on the social optimum cost given in Theorem 4.1, we get:

$$PoS(\mathcal{N}) \leq \frac{C^{\mathcal{N}}(\vec{S})}{n(\alpha/2 + \bar{d}_H)} \leq \frac{\alpha + 2\bar{d}_H}{\alpha/2 + \bar{d}_H} = 2 \ . \qquad \square$$

Note, however, that unlike the complete interest game, there are games for which the Price of Stability is not 1, even for non-constant α. We provide such an example for the cycle NI-NCG (see Chapter 8.2).

Chapter 8

Specific Friendship Graphs

8.1 The Forest NI-NCG

We now consider the NI-NCG where H is a forest, i.e., a collection of trees. Our main result is that for such games, H is both the unique social optimum network and the unique Nash equilibrium network, and consequently both the Price of Anarchy and the Price of Stability are 1.

Theorem 8.1 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, where H is a forest. Then H is the unique social optimum network and the unique Nash equilibrium network. Consequently $PoS(\mathcal{N}) = PoA(\mathcal{N}) = 1$.*

8.2 The Cycle NI-NCG

We now consider the special case of an NI-NCG when H is a cycle on its n vertices. It should be clear from our discussion of general NI-NCGs that for such games, both the Price of Anarchy and the Price of Stability are at most a certain constant (because the average degree of the friendship graph is 2, a constant). This section focuses on obtaining constant bounds that are as tight as possible.

8.2.1 Notation and Terminology

An NI-NCG where H is a cycle can be completely defined (up to graph isomorphism) by the number of players and α. We therefore denote such a game by $\mathcal{N} = \langle V, H = C, \alpha \rangle$, where $V = \{0, \ldots, n-1\}$ is the set of players and α is an edge building price. The friendship graph H is implied from V, and it is a cycle with edges $(v, v+1)$, where "player arithmetic" is performed modulo n. We also require that $n \geq 3$.

Given a Cycle NI-NCG we use the term *cycle edge* occasionally as a synonym for friendship edge. The term *the cycle* refers to H itself. The term *segment* refers to any simple path in the cycle, including the empty path (consisting of no vertices), any path with a single vertex, and the path that is the entire cycle. The term *line* refers to a segment with $n-1$ edges. Note that these terms refer specifically to subgraphs of H, they are not applicable to any cycle, line or path. (respectively) on the n vertices.

8.2.2 Basic Results

We begin our discussion with the social optimum. We provide a complete characterization of the social optimum based on the value of α. We first show the following two lemmata:

Lemma 8.2. *In a Cycle NI-NCG, a social optimum network is either the cycle or a tree.*

Proof: This follows trivially from Lemma 4.3. As shown, if G is a social optimum network, then $|E(G)| < |E_H|$ or $G = H$. When H is the cycle, this implies that either $|E(G)| < n$ or $G = H$. The former implies that G is a tree (since it must be connected), the latter that G is the cycle. $\square$

Lemma 8.3 (Proof in Appendix D). *In a Cycle NI-NCG, the social cost of any tree network that has a DFS traversal which discovers players in the order in which they appear in the cycle, is minimal among all networks that are trees. Any line or star network is such a minimal cost tree.*

Our social optimum characterization follows from these lemmata:

Theorem 8.4 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. The cycle is a social optimum network if and only if $\alpha \leq 2n - 4$. Any line and any star are social optimum networks if and only if $\alpha \geq 2n - 4$. Other social optimum networks may exist only if $\alpha \geq 2n - 4$, and they must be trees.*

Next we show that for a Cycle NI-NCG a Nash equilibrium always exists. We prove that for various values of α, either the cycle, a line or a star (or several of them) are Nash equilibrium networks.

Theorem 8.5 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. The cycle is a Nash equilibrium network if and only if $0 < \alpha \leq n - 2$. Any line is a Nash equilibrium network if and only if $\alpha \geq n - 2$. Any star is a Nash equilibrium network if $\alpha \geq 2$. Consequently, a Nash equilibrium always exists for a Cycle NI-NCG.*

8.2.3 Upper Bound On The Price Of Anarchy

We proceed to show an upper bound on the Price of Anarchy for a Cycle NI-NCG. As a function of α, our upper bound is 1 for $\alpha < 1$, increases for $1 \leq \alpha < n - 2$, reaches its maximum value of $2 - \frac{2}{n}$ when $\alpha = n - 2$, and decreases for $\alpha > n - 2$, going asymptotically to 1. The proof relies on the following lemma, which provides a characterization of all possible Nash equilibria for Cycle NI-NCGs:

Lemma 8.6 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Then there is a segment P of the cycle (which may contain anywhere between 0 and n vertices) such that:*

1. *G contains P, and players in P only build edges to other players in P.*

2. *For any $v \in V - V(P)$, v's strategy in $\vec{S}$ builds a single edge to player u that is not a friend of v, so that u is contained in some shortest path (in G) between v's friends, and v is not contained in any such shortest path.*

Note that P may also be empty (i.e., containing no vertices at all), and that this is not the same as P that consists of a single vertex. The Nash equilibria permitted by Lemma 8.6 include the cycle (if P is the entire cycle) and the line (if P is a line). Note also that apart from the case $P = H$ or $V(P) = \emptyset$ all other Nash equilibrium networks are trees (because there are $n-1$ edges). The following lemma further characterizes Nash equilibria based on the value of α.

Lemma 8.7 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, $G = \mathcal{G}(\vec{S})$, and P as in Lemma 8.6. If $\alpha > n-2$ then $V(P) \neq \emptyset$, $P \neq H$, and G is a tree.*

This characterization of Nash equilibria in a Cycle NI-NCG allows us to bound the Price of Anarchy:

Theorem 8.8 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. Then:*

$$PoA(\mathcal{N}) \leq \begin{cases} 1 & \textit{if} \quad \alpha < 1 \\ 2 - \frac{2}{\alpha+2} & \textit{if} \quad 1 \leq \alpha \leq n-2 \\ 1 + \frac{n-1}{\alpha+2} - \frac{1}{n} & \textit{if} \quad n-2 < \alpha \leq 2n-4 \\ 1 + \frac{n-2}{\alpha+4} & \textit{if} \quad 2n-4 < \alpha \end{cases},$$

which is less than 2 for $1 \leq \alpha < 2n-4$, less than $3/2$ for $\alpha \geq 2n-4$, and goes to 1 as $\alpha \to \infty$.

8.2.4 Lower Bound On The Price Of Anarchy

To show a lower bound on the Price of Anarchy for a Cycle NI-NCG, we present an explicit construction of a Nash equilibrium.

Definition 8.9 (Stretched Star $X^{k,n}$)**.** Let $\mathcal{N} = \langle V, H = C, \alpha\rangle$ be a Cycle NI-NCG, and recall that the players are $V = \{0, 1, ..., n-1\}$. Let $2 \leq k < n$ be an integer, such that $n - k - 1 \equiv 0(\text{mod } 4)$. The joint strategy $X^{k,n}$ is defined as follows (Figure 8.1): players $0, ..., k-1$ each build a single edge to their successor, i.e., player $i \in [0, k)$ builds an edge to player $i+1$. Player k does not build any edges. Starting from players $k+1$ and $k+2$, each pair of players build a single edge to either player 1 or player $k-1$ in alternating order: players $k+1$ and $k+2$ both build a single edge to player 1, players $k+3$ and $k+4$ both build a single edge to player $k-1$, players $k+5$ and $k+6$ both build a single edge to player 1, etc..., until all players are exhausted.

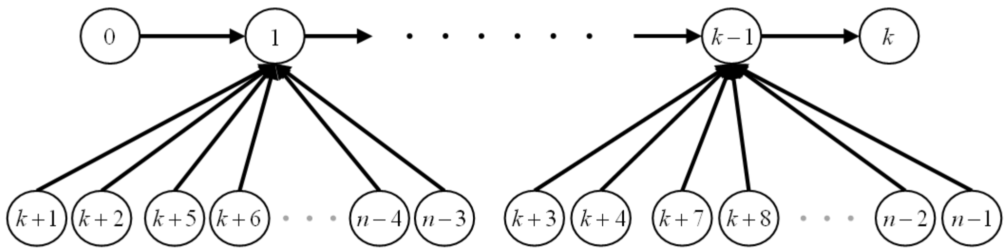

Figure 8.1: $X^{k,n}$

Note that for $k = 2$, $X^{k,n}$ is a star, and for $k = n-1$, $X^{k,n}$ is a line. Next, we show that that $X^{k,n}$ is a Nash equilibrium for certain values of α.

Lemma 8.10 (Proof in Appendix D). *Let $\mathcal{N} = \langle V, H = C, \alpha\rangle$ be a Cycle NI-NCG and k an integer such that $2 \leq k < n$ and $n - k - 1 \equiv 0$(mod 4). If $\alpha \geq k$ then $X^{k,n}$ is a Nash equilibrium for $\mathcal{N}$.*

We now use our construction to prove lower bounds on the Price of Anarchy.

Theorem 8.11 (Proof in Appendix D). *For infinitely many n, for any $\alpha \geq (n+1)/2$, and any*

$\mathcal{N} = \langle V, H = C, \alpha \rangle$ *(of n players and edge building price α), we have:*

$$PoA(\mathcal{N}) \geq \begin{cases} 1 - \frac{1}{n} + \frac{1}{4}\frac{n+1}{\alpha+2} & \mathit{for} \quad \alpha < 2n-4 \\ 1 - \frac{2}{\alpha+2} + \frac{n+1}{4(\alpha+4)} & \mathit{for} \quad \alpha \geq 2n-4 \end{cases} .$$

Additionally, for any n and any $\mathcal{N} = \langle V, H = C, \alpha \rangle$ with $2 \leq \alpha \leq (n+1)/2$, if α is an integer for which $\alpha \equiv n-1$ (mod 4), then:

$$PoA(\mathcal{N}) \geq 2 - \frac{2}{\alpha+2} - \frac{\alpha-2}{n} - \frac{6}{(\alpha+2)n} .$$

Note that there is a gap between the lower and upper bounds. For large n the gap is roughly 0 for $\alpha = 2$, increases to $\frac{3}{4}$ as α goes from 2 to $n-2$, and then decreases, going to 0 as $\alpha/n \to \infty$ and $n \to \infty$. Figure 8.2 displays a chart of our bounds and the gap between them.

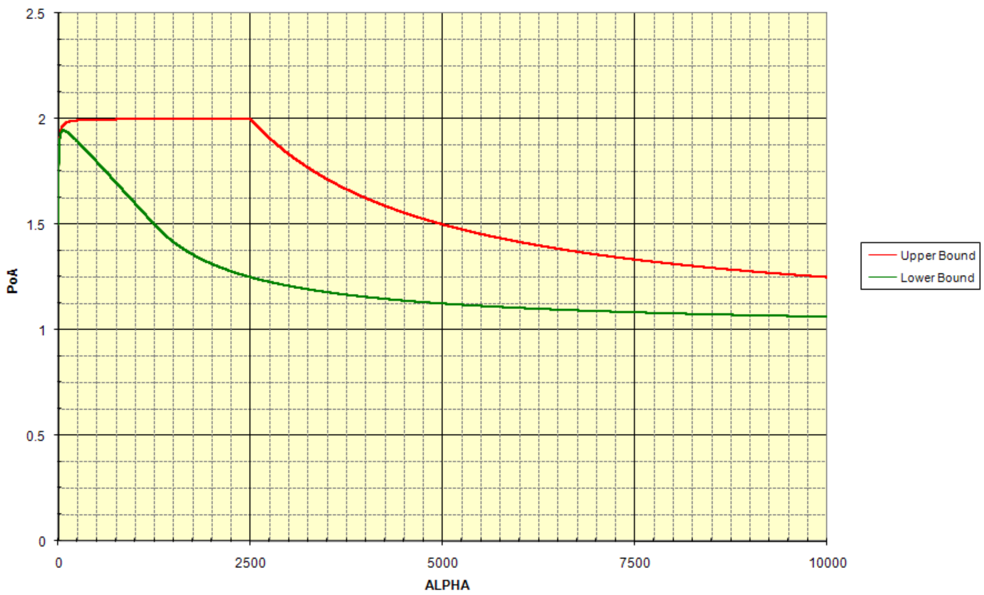

Figure 8.2: Cycle NI-NCG PoA Bounds ($n = 2500$)

8.2.5 The Price Of Stability

The Price of Stability for a Cycle NI-NCG is simple to analyze:

Theorem 8.12. *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. Then:*

$$PoS(\mathcal{N}) = \begin{cases} 1 & \textit{if} \quad \alpha \leq n-2 \\ \left(1 + \frac{2}{\alpha+2}\right)\left(1 - \frac{1}{n}\right) & \textit{if} \quad n-2 < \alpha < 2n-4 \\ 1 & \textit{if} \quad \alpha \geq 2n-4 \end{cases} .$$

Proof: The claim is trivial for $\alpha \leq n-2$ and $\alpha \geq 2n-4$, since in these cases, by Theorem 8.4 and Theorem 8.5 the social optimum is also a Nash equilibrium. For $n-2 < \alpha < 2n-4$, all Nash equilibria must be trees by Lemma 8.7. By Theorem 8.5, any line network is such a Nash equilibrium, and by Lemma 8.3 it has lower social cost than any other tree network, hence it is the Nash equilibrium of minimal social cost. The social cost of a line network is $(\alpha+4)(n-1)$, so the Price of Stability is exactly:

$$PoS(\mathcal{N}) = \frac{(\alpha+4)(n-1)}{(\alpha+2)n} = \left(1 + \frac{2}{\alpha+2}\right)\left(1 - \frac{1}{n}\right) . \qquad \square$$

Note that it is possible to achieve a Price of Stability as close as we wish to $1 + \frac{1}{n} - \frac{2}{n^2}$ by selecting $\alpha \approx n-2$. This proves the existence of NI-NCGs with non-constant α for which the Price of Stability is not 1.

Chapter 9

A Weighted Interests Network Creation Game

Our model for an NI-NCG contains several restrictions that we would like to eliminate. The fact that for a given player $v \in V$, the distance cost is the straightforward sum of the distances to all its friends, implies that the economic significance of efficient communications between any two friends is identical. For example, it means that any player v is equally interested in efficient communications with all its friends. It also means that if u_1 is v_1's friend and u_2 is v_2's friend, then it is equally important for v_1 to communicate efficiently with u_1 as it is for v_2 to communicate efficiently with u_2. Finally, our model is symmetric: if v is interested in communicating with u, then u is equally interested in communicating with v. In real world network creation scenarios we can expect several or all of these assumptions to be wrong.

We therefore study an extension of an NI-NCG where each player associates a non-negative *friendship weight* with every other player. In our new game, the distance cost of each player is the weighted sum of distances to all other players, instead of the straightforward sum of distances. We develop upper bounds on the Price of Anarchy for this variant. Through our analysis of this new game, we also obtain results for additional variants, including an NI-NCG where friendship may be asymmetric, and a game where the distance cost of a player is the expected distance to its friends under some given distribution.

A weighted network creation game was also studied by Albers et. al. [AEED+06], but their model did not permit zero friendship weights. Moreover, unless α is very large, their upper bound approaches $O(n)$ (a trivial bound) as w_{min}, the minimum friendship weight, approaches 0. Our results allow 0 weights, are always $O(\sqrt{n})$, and are not as sensitive to outlier weights.

9.1 The Model

Formally, a *network creation game with weighted interests* (WI-NCG), is a tuple $\mathcal{N} = \langle V, f, \alpha \rangle$ where $V = \{1 \dots n\}$ is the set of players, $f : V \times V \to \mathbb{R}^+ \cup \{0\}$ is a *friendship weight function* and $\alpha \in \mathbb{R}^+$ is an *edge building price*. Given a WI-NCG, we say that player u is player v's friend iff $f(v, u) > 0$. Note that unlike the NI-NCG, we do not require friendship to be symmetric. The *friendship graph* of $\mathcal{N}$, denoted $H = \mathcal{H}(\mathcal{N})$, is the directed graph $H = (V, E_H)$ whose vertices are the players V and edges are $E_H = \{(v, u) : f(v, u) > 0\}$. Player v's out-neighborhood in H, denoted $N_H^+(v) = \{u : (v, u) \in E_H\}$, is also called the *friend set* of v. We say that v and u are *semi-friends* if v is a friend of u or u is a friend of v. As before, we limit our discussion to NI-NCGs for which H has no isolated vertices.

The definitions of strategies and costs of a WI-NCG are identical to the NI-NCG, except that the *distance cost* for player $v \in V$ is defined as $Dist^{\mathcal{N}}(\vec{S}, v) = \sum_{u \in V} f(v, u)\delta_G(v, u)$. All definitions of social optima, Nash equilibria, Price of Anarchy, Price of Stability, sensible joint strategies and complete friendship responses remain as before. We use the same graph-theoretical notation as before, but add the following notation. For any directed graph $H = (V_H, E_H)$, we denote by $d_H^+(v) = |N_H^+(v)|$ the out-degree of vertex v in H, and by $\bar{d}_H^+ = \sum_{v \in V_H} d_H^+(v)/|V_H|$ the average vertex out-degree of H. Additionally, we denote by $F = \sum_{v,u \in V, v \neq u} f(v,u)$ the sum of all friendship weights, by $\bar{f}_V = F/n$ the average friendship weight per player, and by $\bar{f}_H = F/|E_H| = \bar{f}_V/\bar{d}_H^+$ the average friendship weight over all edges in H. Note that since H has no isolated vertices, we must have $F > 0$, $\bar{f}_V > 0$, and $\bar{f}_H > 0$.

Clearly, an NI-NCG is a special case of a WI-NCG where all friendship weights are 0 or 1 and friendship weights are symmetric. We also study certain additional special cases of a WI-NCG: (1) an *Asymmetric NI-NCG*, where friendship weights are 0 or 1 but may be asymmetric, and (2) a *Distributional WI-NCG* where the sum of friendship weights per player must be 1, i.e., $\forall v \in V, \sum_{u \in V - \{v\}} f(v,u) = 1$. A distributional WI-NCG models a scenario where each player has a probability distribution over all other players, and its distance cost is the expected distance to the other players under this distribution. In a distributional WI-NCG the game is unbiased towards any specific player, while in a general WI-NCG, the distances of players with higher outgoing friendship weights have a larger impact on cost (both individual and social).

9.2 Upper Bound On The Price of Anarchy

We now present several upper bounds for the Price of Anarchy of a WI-NCG. Note that unlike the NI-NCG, we do not show existence of Nash equilibria for a general WI-NCG. Hence, our upper bound is contingent on the assumption that Nash equilibria exist. We first show that our previous upper bound on the Price of Anarchy for the NI-NCG from Theorem 5.10 still holds, if we replace α by $\alpha/\bar{f}_H$ and $\bar{d}_H$ by $\bar{d}_H^+$, and disregard the special cases for $\alpha \leq 2$. The proof has the same structure as our proof of Theorem 5.10.

Theorem 9.1 (Proof in Appendix E). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, and assume that a*

Nash equilibrium exists. Let $\hat{\alpha} = \alpha/\bar{f}_H$. Then:

$$PoA(\mathcal{N}) \leq 3 + \frac{\bar{d}_H^+}{\hat{\alpha}/2 + \bar{d}_H^+} \min\left(\hat{\alpha},\ 2\log_2 n - 2 + \min\left(n-2,\ 2\sqrt{(n-2)\hat{\alpha}/\bar{d}_H^+}\right)\right)$$
$$= O\left(1 + \min\left(\hat{\alpha},\ \bar{d}_H^+,\ \log n + \sqrt{\frac{n\hat{\alpha}}{\bar{d}_H^+}},\ \sqrt{\frac{n\bar{d}_H^+}{\hat{\alpha}}}\right)\right) = O(\sqrt{n})\ .$$

Note that given a WI-NCG, multiplying both α and f by the same positive constant c yields a game which is equivalent in the sense that its player and social costs are c times the costs of the original game. The use of $\hat{\alpha}$ instead of α in our bound therefore makes sense, since it is independent of this "scaling factor".

The bound in Theorem 9.1 is, however, somewhat awkward, because it produces a bound that is far too large when the game has many small positive friendship weights. Specifically, the use of the parameters $\bar{d}_H^+$ and $\hat{\alpha}$ means that by modifying certain friendship weights that are zero to be positive and very small, we can increase the bound substantially. Note that we sometimes indeed expect a "jump" in the social cost of the social optimum and worst Nash equilibrium when we apply such a modification, since there is a fundamental difference between zero weights and nonzero weights. While a zero friendship weight between players allows them to be in different connected components in the network, a nonzero weight (as small as it may be) requires them to at least be in the same connected component. Therefore, in some cases, changing zero weights to small positive weights may require the construction of additional edges in the network, regardless of the actual value of the new weights, producing a "jump" in the social costs of the social optimum and Nash equilibria. However, our claim is that Theorem 9.1 produces a "jump" that is much too large. We explain this behavior through the following example.

Example 9.2. Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG with $2k$ players ($k \geq 2$) denoted $\{1, 2, \ldots, 2k-1, 2k\}$, such that $\alpha = 1/2$ and the players are divided into disjoint pairs of friends, i.e., players $2i-1$ and $2i$ are friends, for any $i \in \{1, k\}$. By Theorem 4.2 and Theorem 4.4 the unique social optimum network and unique Nash equilibrium network are both H itself, hence $PoA(\mathcal{N}) = 1$. Clearly, $\mathcal{N}$ can also be viewed as a WI-NCG $\mathcal{N} = \langle V, f, \alpha \rangle$, in which case every edge in H is replaced by two directed edges in opposite directions, and f assigns weight 1 to every such edge and weight 0 to any other pair of players. Viewed as a WI-NCG, we have $\bar{f}_H = 1$, $\bar{d}_H^+ = 1$ and $\hat{\alpha} = 1/2$. Theorem 9.1 therefore produces an upper bound $PoA(\mathcal{N}) = O(1)$. We now construct a new WI-NCG, $\mathcal{N}' = \langle V, f', \alpha \rangle$, where:

$$f'(v,u) = \begin{cases} 1 & \text{if } (v,u) \in E_H \\ \\ \epsilon & \text{if } v \equiv u \equiv 0 \pmod 2 \end{cases} ,$$

for some $\epsilon > 0$. In other words, $\mathcal{N}'$ is $\mathcal{N}$ after modifying the 0 weights between all pairs of even players to ϵ. The cost of the social optimum for $\mathcal{N}'$ is at least $\alpha(2k-1)+2k$ since all players must be in the same connected component. For small enough ϵ, it is easy to see that any Nash equilibrium network of $\mathcal{N}'$ must still be a tree containing H. Therefore the cost of any Nash equilibrium for $\mathcal{N}'$ is at most $\alpha(2k-1) + 2k + \epsilon(2k-1)k(k-1)$. Therefore we have $PoA(\mathcal{N}) \leq 1 + \epsilon k(k-1)$, and selecting $\epsilon < \frac{1}{k(k-1)}$ ensures $PoA(\mathcal{N}) < 2$. Note that while there is a jump in the costs, the effect on the Price of Anarchy is small.

On the other hand, let $H' = \mathcal{H}(\mathcal{N}')$ be the friendship graph of $\mathcal{N}'$. Clearly, $\bar{d}_{H'}^+ = \frac{k+1}{2}$, and:

$$\bar{f'}_{H'} = \frac{F'}{k(k+1)} = \frac{2k + \epsilon k(k-1)}{k(k+1)} = \frac{2}{k+1} + \epsilon\left(1 - \frac{2}{k+1}\right) ,$$

so by selecting ϵ to be small enough we can have $\bar{f'}_{H'}$ as close as we wish to $\frac{2}{k+1}$, and $\hat{\alpha}$ as close as we wish to $\frac{k+1}{4}$. Specifically for $\epsilon < \frac{1}{k(k-1)}$ we have $\frac{k+1}{5} < \hat{\alpha} < \frac{k+1}{4}$. By Theorem 9.1 in such a case we will have:

$$PoA(\mathcal{N}') = O(\sqrt{k}) .$$

Therefore, for small ϵ, we have a very large gap between the Price of Anarchy for $\mathcal{N}'$ and our bound from Theorem 9.1. Moreover, the bound itself exhibits a very large difference between $\mathcal{N}$ and $\mathcal{N}'$, even though the game parameters are as similar as we wish.

It is also possible to construct a similar example where both games have a single connected component, i.e., the difference between zero weights and nonzero weights is not as fundamental, but the bound "jumps" around zero weights nonetheless. We can trace this behavior of Theorem 9.1 to our proof of the weighted equivalent of Lemma 5.1. The proof of this lemma relies on an analysis of the complete friendship response for player $v \in V$. However, when v has many friends with very small friendship weights, the complete friendship response is obviously a bad choice of alternative strategy for v, since building edges to all these friends costs much but has a small impact on distance cost. What we would like to do is devise some kind of partial friendship response that takes this into account. This is achieved by the following lemma.

Lemma 9.3. *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG and $\vec{S}$ a Nash equilibrium for $\mathcal{N}$. Let $H_h = (V, E_h)$ be a directed graph where $E_h = \{(v,u) \in E_H : f(v,u) > \frac{\alpha}{n-2}\}$, i.e., H_h is the subgraph of $H = \mathcal{H}(\mathcal{N})$ with all friendship edges of "heavy" weight (greater than $\frac{\alpha}{n-2}$). Denote $\bar{l}_V = \frac{1}{n}\sum_{(v,u)\in E_H - E_h} f(v,u)$, i.e., $\bar{l}_V$ is the average cumulative "light" friendship weight per player. Then:*

$$C^{\mathcal{N}}(\vec{S}) < n(\alpha(\bar{d}^{+}_{H_h} + 1) + \bar{f}_V + (n-2)\bar{l}_V) \ .$$

Proof: Let $G = \mathcal{G}(\vec{S})$, and T a spanning forest for G. Let $v \in V$ be any player, and denote $T_v = S_v \cap N_T(v)$, i.e., T_v is a strategy of v that builds only the edges in S_v that are also in T. We define an alternative strategy S'_v as follows:

$$S'_v = T_v \cup N^{+}_{H_h}(v) \ ,$$

i.e., S'_v builds edges to all friends of v with friendship weight higher than $\frac{\alpha}{n-2}$, and also all edges built in S_v that belong to T. Clearly $|S'_v| \le |T_v| + d^{+}_{H_h}(v)$. Let $\vec{S}' = (\vec{S}_{-v}, S'_v)$ be the deviation of v to S'_v in $\vec{S}$, and $G' = \mathcal{G}(\vec{S}')$. Clearly G' contains T. Since v must be in the same connected component as all its friends in G, we conclude that it must be in the same connected component as all its friends in G', hence the distance in G' between v and any friend is at most $n-1$. The distance in G' between v and any friend $u \in N^+_{H_h}(v)$ is exactly 1. Denoting $f_v = \sum_{u \in N^+_H(v)} f(v,u)$ and $l_v = \sum_{u \in N^+_H(v) - N^+_{H_h}(v)} f(v,u)$, we conclude that v's cost in $\vec{S}'$ is at most:

$$C^{\mathcal{N}}(\vec{S}', v) \le \alpha|T_v| + \alpha d^+_{H_h}(v) + \sum_{u \in N^+_{H_h}(v)} f(v,u) + \sum_{u \in N^+_H(v) - N^+_{H_h}(v)} f(v,u)(n-1)$$

$$= \alpha|T_v| + \alpha d^+_{H_h}(v) + f_v + l_v(n-2) \ .$$

Since v is playing a best response in $\vec{S}$ this must be at least v's cost in $\vec{S}$. Summing over all players $v \in V$ we get:

$$C^{\mathcal{N}}(\vec{S}) \le \alpha \sum_{v \in V} |T_v| + \alpha n \bar{d}^+_{H_h} + n\bar{f}_V + n(n-2)\bar{l}_V \ ,$$

and since for every edge $(v,u) \in E(T)$ we have either $u \in T_v$ or $v \in T_u$, but not both (since $\vec{S}$ must be sensible), we conclude that $\sum_{v \in V} |T_v| \le n-1 < n$, hence:

$$C^{\mathcal{N}}(\vec{S}) < n(\alpha(\bar{d}^+_{H_h} + 1) + \bar{f}_V + (n-2)\bar{l}_V) \ . \qquad \square$$

Note that the choice of our weight threshold, $\frac{\alpha}{n-2}$, has no impact on the proof, but it achieves the minimal bound, and also guarantees that this bound is continuous. Using Lemma 9.3 we can now revise our upper bound.

Theorem 9.4 (Proof in Appendix E). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, and assume that a Nash equilibrium exists. Then:*

$$PoA(\mathcal{N}) \le 3 + \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V} \min\left(\alpha \frac{\bar{d}^+_{H_h}}{\bar{f}_V} + (n-3)\frac{\bar{l}_V}{\bar{f}_V},\ 2\log_2 n - 2 + 2\sqrt{(n-2)\alpha/\bar{f}_V}\right)$$

$$= O(\sqrt{n}) \ .$$

Applied to Example 9.2, for $k \ge 4$ and $\epsilon < 1/k(k-1)$, we have $\bar{d}^+_{H_h} = 1$, $\bar{f}_V = 1 + \epsilon(k-1)/2 > 1$, and $\bar{l}_V = \epsilon(k-1)/2 < \frac{1}{2k}$. We therefore have $\alpha\frac{\bar{d}^+_{H_h}}{\bar{f}_V} + (n-3)\frac{\bar{l}_V}{\bar{f}_V} \le \frac{1}{2} + 2k\frac{1}{2k} = \frac{3}{2}$, and therefore Theorem 9.4 produces a bound of $O(1)$, as we wanted.

For an asymmetric NI-NCG we have $\bar{f}_V = \bar{d}^+_H$. If $\alpha < n-2$ then all friendship edges are "heavy", hence $\bar{d}^+_{H_h} = \bar{d}^+_H$ and $\bar{l}_V = 0$. If $\alpha \geq n-2$ then all friendship edges are "light", hence $\bar{d}^+_{H_h} = 0$ and $\bar{l}_V = \bar{f}_V$. We therefore have the following corollary:

Corollary 9.5 (Asymmetric NI-NCG PoA Upper Bound). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be an asymmetric NI-NCG. Then:*

$$\begin{aligned} PoA(\mathcal{N}) &\leq 3 + \frac{\bar{d}^+_H}{\alpha/2 + \bar{d}^+_H} \min\left(\alpha,\ n-2,\ 2\log_2 n - 2 + 2\sqrt{(n-2)\alpha/\bar{d}^+_H}\right) \\ &= O(\sqrt{n}) \ . \end{aligned}$$

Note that for an NI-NCG (symmetric), Corollary 9.5 provides a better numeric upper bound than Theorem 5.10 when $\alpha \geq (n-2)\max(1, \bar{d}_H/4)$. For a distributional WI-NCG, since $\bar{f}_V = 1$, we have:

Corollary 9.6 (Distributional WI-NCG PoA Upper Bound). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a distributional WI-NCG. Then:*

$$\begin{aligned} PoA(\mathcal{N}) &\leq 3 + \frac{1}{\alpha/2 + 1} \min\left(\alpha\bar{d}^+_{H_h} + (n-3)\bar{l}_V,\ 2\log_2 n - 2 + 2\sqrt{(n-2)\alpha}\right) \\ &= O(\sqrt{n}) \ . \end{aligned}$$

Chapter 10

Topics for Further Research

Any improvement of our upper and lower bounds on the Price of Anarchy would be good news. The techniques used in [AEED$^+$06] to achieve an $O(n^{1/3})$ upper bound on the Price of Anarchy of the complete interest game do not seem directly applicable to an NI-NCG since they rely heavily on large intersections of the friend sets of the players. We would also like to see a resolution of the question of Nash equilibrium existence for $\alpha \in (1, 2)$ and $g(H) \in \{3, 4\}$, or $\alpha \in [1.5, 2)$ and $g(H) = 5$, as well as the WI-NCG. There are several additional open questions and topics that may merit further research:

- We have studied the loss of efficiency in a fully uncoordinated and selfish setting. Studying the loss of efficiency in other settings seems desirable. Especially interesting is the solution concept of a *Strong Equilibrium* and the related *Strong Price of Anarchy*, which models the loss of efficiency when player coalitions are allowed. These concepts were studied in [AFM07] for the complete interest game.

- We have studied the specific classes of a cycle and a forest. We also have partial results for friendship graphs with large girth and degree in Chapter 6. Additional interesting classes of friendship graphs, that are relevant for various real world scenarios, include grids, degenerate graphs, hypercubes, random graphs and expanders.

- Our model, while nonuniform in the players' friend sets and (in a WI-NCG) the friendship weights, still employs a uniform edge building cost function and edge efficiency function. Various extensions of the model in this respect are interesting, including nonuniform distances for the network edges, nonuniform edge building costs, and more elaborate payment functions where edge efficiency is determined by the amount paid for it.

- Most research for this class of network creation games has been focused, so far, on the theoretical bounds for loss of efficiency. There are various algorithmic, computational and mechanism design topics related to our model of network creation by distributed agents, that have received little attention. A progression to research of these issues would be both natural and useful.

Bibliography

[AART03] B. Awerbuch, Y. Azar, Y. Richter, and D. Tsur. Tradeoffs in worst-case equilibria. *Proceedings of 1st WAOA*, pages 41–52, 2003.

[ADK+04] E. Anshelevich, A. Dasgupta, J. Kleinberg, E. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. *Foundations of Computer Science, 2004. Proceedings. 45th Annual IEEE Symposium on*, pages 295–304, 2004.

[ADTW03] E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Proceedings of the thirty-fifth ACM symposium on Theory of computing*, pages 511–520, 2003.

[AEED+06] S. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On nash equilibria for a network creation game. *Proceedings of the seventeenth annual ACM-SIAM symposium on Discrete Algorithms*, pages 89–98, 2006.

[AFM07] N. Andelman, M. Feldman, and Y. Mansour. Strong Price of Anarchy. *Proceedings of the eithteenth annual ACM-SIAM symposium on Discrete Algorithms*, 2007.

[AHL02] N. Alon, S. Hoory, and N. Linial. The Moore Bound for Irregular Graphs. *Graphs and Combinatorics*, 18(1):53–57, 2002.

[Aum59] R.J. Aumann. Acceptable points in general cooperative n-person games. *Contributions to the Theory of Games*, 4:287–324, 1959.

[Aum74] R.J. Aumann. Subjectivity and correlation in randomized strategies. *Journal of Mathematical Economics*, 1(1):67–96, 1974.

[BB94] N. Biggs and N.L. Biggs. *Algebraic Graph Theory.* Cambridge University Press, Cambridge, second edition, 1994.

[BG00] V. Bala and S. Goyal. A Noncooperative Model of Network Formation. *Econometrica*, 68(5):1181–1229, 2000.

[BH83] N. Biggs and ML Hoare. The sextet construction for cubic graphs. *Combinatorica*, 3(2):153–165, 1983.

[Bol78] B. Bollobas. *Extremal graph theory.* Academic Press, London, 1978.

[CKV02] A. Czumaj, P. Krysta, and B. Vocking. Selfish traffic allocation for server farms. *Conference Proceedings of the Annual ACM Symposium on Theory of Computing*, pages 287–296, 2002.

[CP05] J. Corbo and D. Parkes. The price of selfish behavior in bilateral network formation. *Proceedings of the twenty-fourth annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 99–107, 2005.

[CSSM04] J.R. Correa, A.S. Schulz, and N.E. Stier-Moses. Selfish Routing in Capacitated Networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

[CV02] A. Czumaj and B. Vöcking. Tight bounds for worst-case equilibria. *Proceedings of the thirteenth annual ACM-SIAM symposium on Discrete algorithms*, pages 413–420, 2002.

[ES63] P. Erdos and H. Sachs. Regulare Graphe gegebener Taillenweite mit minimaler Knotenzahl. *Wiss. Z. Univ. Halle-Wittenberg, Math. Nat*, 12(3):251–258, 1963.

[FKK+02] D. Fotakis, S. Kontogiannis, E. Koutsoupias, M. Mavronicolas, and P. Spirakis. The Structure and Complexity of Nash Equilibria for a Selfish Routing Game. *Proceedings of the 29th International Colloquium on Automata, Languages and Programming*, pages 123–134, 2002.

[FLM+03] A. Fabrikant, A. Luthra, E. Maneva, C.H. Papadimitriou, and S. Shenker. On a network creation game. *Proceedings of the twenty-second annual symposium on Principles of Distributed Computing*, pages 347–351, 2003.

[GMV05] M. Goemans, V. Mirrokni, and A. Vetta. Sink equilibria and convergence. *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 142–151, 2005.

[GST04] A. Gupta, A. Srinivasan, and E. Tardos. Cost-sharing mechanisms for network design. *Proc. 7th APPROX*, pages 139–150, 2004.

[HS03] H. Haller and S. Sarangi. Nash Networks with Heterogeneous Agents. *German Institute for Economic Research: Discussion Paper, March*, 2003.

[Imr84] W. Imrich. Explicit construction of regular graphs without small cycles. *Combinatorica*, 2:53–59, 1984.

[JMT06] R. Johari, S. Mannor, and J.N. Tsitsiklis. A Contract-Based Model for Directed Network Formation. *Games and Economic Behaviour*, 2006.

[JV01] K. Jain and V. Vazirani. Applications of approximation algorithms to cooperative games. *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 364–372, 2001.

[KP99] E. Koutsoupias and C. Papadimitriou. Worst-case equilibria. *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.

[LPS88] A. Lubotzky, R. Phillips, and P. Sarnak. Ramanujan graphs. *Combinatorica*, 8:261–277, 1988.

[LU95] F. Lazebnik and V.A. Ustimenko. Explicit Construction of Graphs with an Arbitrary Large Girth and of Large Size. *Discrete Applied Mathematics*, 60(1-3):275–284, 1995.

[LUW95] F. Lazebnik, V.A. Ustimenko, and A.J. Woldar. A new series of dense graphs of large girth. *Bull. Amer. Math. Soc*, 32(1):73–79, 1995.

[Mar82] G. A. Margulis. Explicit constructions of graphs without short cycles and low density codes. *Combinatorica*, 2(1):71–78, 1982.

[Mar88] G. A. Margulis. Explicit group theoretical constructions of combinatorial schemes and their application to the design of expanders and concentrators. *Problems of Information Transmission*, 24(1):39–46, 1988.

[MS01] M. Mavronicolas and P. Spirakis. The price of selfish routing. *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 510–519, 2001.

[Nas51] J. Nash. Non-Cooperative Games. *The Annals of Mathematics*, 54(2):286–295, 1951.

[Rou02a] T. Roughgarden. *Selfish Routing*. PhD thesis, Cornell University, 2002.

[Rou02b] T. Roughgarden. The price of anarchy is independent of the network topology. *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 428–437, 2002.

[RT02] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM (JACM)*, 49(2):236–259, 2002.

[Sau67] N. Sauer. Extremaleigenschaften regularer Graphen gegebener Taillenweite I and II. *Sitzungberichte Osterreich Akad. Wiss. Math. Vatur. Kl. SB II*, 176:9–25, 1967.

[Vet02] A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. *Foundations of Computer Science, 2002. Proceedings. The 43rd Annual IEEE Symposium on*, pages 416–425, 2002.

[Wal65a] H. Walther. Eigenschaften von regularen Graphen gegebener Taillenweite und minimaler Knotenzahl. *Wiss. Z. Techn. Hochsch. Ilmenau*, 11:167–168, 1965.

[Wal65b] H. Walther. Uber regulare Graphen gegebener Taillenweite und minimaler Knotenzahl. *Wiss. Z. Techn. Hochsch. Ilmenau*, 11:93–96, 1965.

[Wei84] A. Weiss. Girths of bipartite sextet graphs. *Combinatorica*, 4:241–245, 1984.

Appendix A

Proofs From Chapter 4

Theorem 4.2 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG with $\alpha \leq 2$. Then H is a social optimum network, and any sensible joint strategy $\vec{S}$ that creates it (i.e., $\vec{S} \in \overline{\mathbf{S}}(H)$) is a social optimum. Furthermore, for $\alpha < 2$, H is the unique social optimum network.*

Proof: Let $\vec{S}$ be a social optimum. We first prove for $\alpha < 2$. If there is a pair of friends $v, u \in V, (v, u) \in E_H$ that are not adjacent in $\mathcal{G}(\vec{S})$, then modifying $\vec{S}$ by building the edge between v and u (in any direction) increases the social edge building cost by α, and reduces the social distance cost by at least 2 (1 for every direction of the friendship relation), which yields a social cost reduction, contradicting the optimality of $\vec{S}$. Therefore $H \subseteq \mathcal{G}(\vec{S})$, which implies $H = \mathcal{G}(\vec{S})$ (since removing any excess edges decreases the cost), proving that H is a unique social optimum network (there must be at least one social optimum, of course).

Similarly, for $\alpha = 2$, if there is a pair of friends $v, u \in V, (v, u) \in E_H$ that are not adjacent in $\mathcal{G}(\vec{S})$, then modifying $\vec{S}$ by building the edge between v and u cannot increase the social cost, so the resulting joint strategy must also be a social optimum. Therefore by repeatedly adding missing edges from H to the joint strategy, starting from $\vec{S}$, we always obtain a social optimum. We can add edges this way until we get some social optimum $\vec{S'}$, for which $H \subseteq G(\vec{S'})$. This will necessarily happen after a finite number of steps, and then we must also have $H = G(\vec{S'})$, proving that H is a social optimum network (albeit, not necessarily unique). $\square$

Lemma 4.3 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a social optimum for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Then $|E(G)| \leq |E_H|$ with equality possible if and only if $G = H$.*

Proof: Let $\vec{S}_H \in \overline{\mathbf{S}}(H)$ be any sensible joint strategy that creates the network H. The joint strategy $\vec{S}_H$ achieves the absolute minimum social distance cost, hence $Dist^{\mathcal{N}}(\vec{S}) \geq Dist^{\mathcal{N}}(\vec{S}_H)$. Since $\vec{S}$ is a social optimum, we must have $C^{\mathcal{N}}(\vec{S}) \leq C^{\mathcal{N}}(\vec{S}_H)$. These two inequalities imply that $B^{\mathcal{N}}(\vec{S}) \leq B^{\mathcal{N}}(\vec{S}_H)$, and since $\vec{S}_H$ builds exactly $|E_H|$ edges, we have $|E(G)| \leq |E_H|$.

If equality holds, then $B^{\mathcal{N}}(\vec{S}) = B^{\mathcal{N}}(\vec{S}_H)$, which implies that $Dist^{\mathcal{N}}(\vec{S}) = Dist^{\mathcal{N}}(\vec{S}_H)$ (otherwise $\vec{S}$ would not be a social optimum). This is only possible if $\vec{S}$ achieves distance 1 between any pair of friends, i.e., $E_H \subseteq E(G)$. Since $|E(G)| = |E_H|$ we must have $G = H$. $\square$

Theorem 4.4 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG with $\alpha \leq 1$. Then H is a Nash equilibrium network, and any joint strategy $\vec{S} \in \overline{\mathbf{S}}(H)$ is a Nash equilibrium. Furthermore, for $\alpha < 1$, H is the unique Nash equilibrium network.*

Proof: We start with the case $\alpha \leq 1$. Let $\vec{S}$ be any sensible joint strategy that creates H. For any $v \in V$, v's strategy in $\vec{S}$ is a complete friendship response, which achieves the smallest possible distance cost. Therefore v can only hope to reduce its cost by a unilateral deviation if it deviates to a strategy that builds less edges. But if v builds k fewer edges, it will have at least k friends that are non-adjacent in the new network. The distance to each one of these will be at least 2, which is an increase of at least 1, for an overall increase of distance cost of at least k. The decrease in edge building cost is $\alpha k \leq k$, so no cost reduction is possible. This completes the proof for $\alpha \leq 1$.

If $\alpha < 1$, let $\vec{S}$ be any Nash equilibrium for the game and let $G = \mathcal{G}(\vec{S})$ be the network it creates. If (v, u) is a friendship edge not included in G, then v would prefer to deviate by building this additional edge to u, since the reduction in its distance cost would be at least 1, and $\alpha - 1 < 0$. This contradicts the fact that v is playing a best response in $\vec{S}$, so we conclude that $H \subseteq G$. If $G \neq H$ then there is an edge $(v, u) \in E(G) - E_H$. Assume w.l.o.g that this edge is built by v in $\vec{S}$. Then v would benefit from removing this edge, since it has distance 1 to all its friends even without it (because $H \subseteq G$ and $(v, u) \notin E_H$). This is a contradiction, hence $G = H$, which completes the proof for $\alpha < 1$. □

Theorem 4.5 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG with $\alpha \geq 2$. There exists a Nash equilibrium for $\mathcal{N}$ whose cost is at most $n(\alpha + 2\bar{d}_H)$.*

Proof: We show an explicit construction of a Nash equilibrium. We denote by m the number of connected components in H, and by V_i the set of vertices in component i. Let $\vec{S}^*$ be the joint strategy defined by the following process, performed independently for each connected component of H: for connected component $i \in \{1 \ldots m\}$, we initialize $U_i = V_i$, where U_i denotes the set of players we need to assign a strategy in $\vec{S}^*$. At each iteration, we search for two players $v, u \in U_i$ such that u is v's only friend in U_i, i.e., $|N_H(v) \cap U_i| = \{u\}$ (we choose an arbitrary pair if more than one exists). We remove v from U_i, and set v's strategy in $\vec{S}^*$ to be $S^*_v = \{u\}$ (i.e., v builds a single edge to u). We repeat this process until we cannot find such a pair. Note that before the first step $U_i \neq \emptyset$, and after every step we have $u \in U_i$, hence necessarily $U_i \neq \emptyset$ when the process ends. We now select an arbitrary $r_i \in U_i$, and let $A_i = U_i - \{r_i\}$. For any player $v \in A_i$, we set the strategy of v in $\vec{S}^*$ to be $S^*_v = \{r_i\}$ (i.e., v builds a single edge to r_i). Finally, we set the strategy of r_i to be $S^*_{r_i} = \emptyset$. We denote $U = \bigcup_{i=1}^{m} U_i$ (at the end of the construction), and by $G^* = \mathcal{G}(\vec{S}^*)$ the network created by $\vec{S}^*$.

Our construction process uses arbitrary choices, so it actually generates a family of joint strategies and not a single joint strategy. Our claims and proofs hold for any $\vec{S}^*$ in this family. The following characteristics of $\vec{S}^*$ and G^* are easily verifiable:

- Each connected component of G^* is a rooted tree with a root r_i, and hence G^* is a forest.
- The players r_i do not build edges. All other players build a single edge to their parent in G^*. Players in $V - U$ build this single edge to a friend.
- All players in $V - U$ are using complete friendship responses in $\vec{S}^*$.
- The distance in G^* between any two friends is 1 or 2.
- For any component i and player $v \in A_i$, v has at least two friends in $U_i = A_i \cup \{r_i\}$.

We claim that $\vec{S}^*$ is a Nash equilibrium. Let $v \in V$ be a player, and let S_v^* be v's strategy in $\vec{S}^*$ and S_v' be some alternative strategy for v. Let i indicate the connected component of H that v belongs to. We split our analysis into several cases:

1. $v = r_i$

   In this case, v has distance 1 from all its friends in the network G^* without building any edges, and so no unilateral deviation can reduce its cost.

2. $v \in V - U$

   In this case S_v^* is a complete friendship response for v in $\vec{S}^*$. Therefore v can benefit by unilateral deviation to S_v' only if S_v' builds fewer edges, i.e., $S_v' = \emptyset$. This would cause v to become disconnected from a friend of his (specifically, his parent in G^*), yielding infinite cost. Therefore, no cost reduction is possible.

3. $v \in A_i$

   We denote by x the number of new edges built in S_v' (i.e., $x = |S_v' - S_v^*|$). Since the distance to a friend is at most 2, each such edge can reduce the network distance to at most one friend of v, by at most 1, for a total distance reduction of at most x. If S_v' does not remove any edges built in S_v^*, then the edge building cost for v increases by $x\alpha \geq x$, so no cost reduction is possible. If S_v' only removes the edge in S_v^* (without building any new edges), then since v must have friends in A_i, S_v' disconnects it from at least one friend and yields infinite cost, so no cost reduction is possible. Therefore, the only possibility for cost reduction is if S_v' removes the edge in S_v^* but builds at least one new edge (i.e., $x \geq 1$). In this case, the edge building cost increase is exactly $(x-1)\alpha$. We have two sub-cases:

- v and r_i are friends

  In this case, the distance cost reduction is at most $x - 1$ (at most x for the new edges, minus at least 1 for the distance increase to r_i). Therefore the total cost increase is at least $(x-1)(\alpha-1)$. Since $\alpha \geq 1$ and $x \geq 1$ this yields no total cost reduction.

- v and r_i are not friends

  The distance cost reduction is at most x, so the total cost increase is at least $(x-1)\alpha - x \geq x-2$, since $\alpha \geq 2$. Therefore, if $x \geq 2$ there is no cost reduction. If $x = 1$ then v improves the distance cost to at most one friend by at most 1. But by doing so while removing the edge to r_i, it increases the distance to any other friend in A_i by at least 1. Since v must have at least one such additional friend in A_i (it has at least 2 friends in U_i, none of which can be r_i in this case), we conclude that the distance cost cannot be reduced. Since the edge building cost is unchanged, no total cost reduction is possible for $x = 1$ either.

Therefore no alternative strategy for any player v yields a total cost improvement, so $\vec{S}^*$ is a Nash equilibrium. Since every player builds at most one edge, and the distance between any two friends is at most 2, the cost of $\vec{S}^*$ is at most $n\alpha + 4|E_H| = n(\alpha + 2\bar{d}_H)$. □

Appendix B

Proofs From Chapter 5

Theorem 5.2 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, and assume that a Nash equilibrium exists. Then:*

$$PoA(\mathcal{N}) \leq \begin{cases} 1 & \textit{if} \quad 0 < \alpha < 1 \\ 2 - \frac{2}{\alpha+2} & \textit{if} \quad 1 \leq \alpha \leq 2 \\ 1 + \frac{\bar{d}_H \alpha}{\alpha/2 + \bar{d}_H} & \textit{if} \quad 2 < \alpha \end{cases}$$
$$= O(1 + \min(\alpha, \bar{d}_H)) \ .$$

Proof: For $\alpha < 1$ the Price of Anarchy must obviously be 1 since H is the unique Nash equilibrium network (Theorem 4.4) and also a social optimum network (Theorem 4.2). For $1 \leq \alpha \leq 2$, H is a social optimum network (Theorem 4.2), so using Lemma 5.1 we have:

$$PoA(\mathcal{N}) \leq \frac{(\alpha+1)n\bar{d}_H}{\alpha n \bar{d}_H/2 + n\bar{d}_H} = \frac{\alpha+1}{\alpha/2+1} = 2 - \frac{2}{\alpha+2} \ .$$

For $\alpha > 2$, we use our lower bound on social optimum cost from Theorem 4.1 to obtain:

$$PoA(\mathcal{N}) \leq \frac{(\alpha+1)n\bar{d}_H}{n(\alpha/2+\bar{d}_H)} \leq 1 + \frac{\bar{d}_H\alpha}{\alpha/2+\bar{d}_H} \leq 1 + \min(\alpha, 2\bar{d}_H) = O(1+\min(\alpha,\bar{d}_H)) \ . \qquad \square$$

Lemma 5.4 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. If v builds an edge to u in $\vec{S}$ (i.e., $u \in S_v$), then $\alpha \leq w(v,u)(\delta_{G'}(v,u) - 1)$, where G' is the graph G without the edge (v,u).*

Proof: If v and u become disconnected in G' then the claim is obvious (since then $\delta_{G'}(v,u) = \infty$). Let us therefore assume that they remain connected in G'. Let $\vec{S'} = (\vec{S}_{-v}, S_v - \{u\})$ be the unilateral deviation of v obtained by removing the edge to u. Clearly $\vec{S'}$ creates the network G'. To bound the change in distance cost for v, we notice that the set of friends whose distance from v changes as a result of the deviation is exactly $\Gamma_G^H(v,u)$. For any such friend $x \in \Gamma_G^H(v,u)$, we must have $\delta_G(u,x) = \delta_{G'}(u,x)$. Using this and the triangle inequality we have:

$$\delta_{G'}(v,x) \leq \delta_{G'}(v,u) + \delta_{G'}(u,x) = \delta_{G'}(v,u) + \delta_G(u,x) = \delta_{G'}(v,u) + \delta_G(v,x) - 1,$$

and therefore $\delta_{G'}(v,x) - \delta_G(v,x) \leq \delta_{G'}(v,u) - 1$. Summing this over all $x \in \Gamma_G^H(v,u)$ we get that the change in distance cost for v is at most $w(v,u)(\delta_{G'}(v,u) - 1)$. Since v is playing a best response in $\vec{S}$, the increase in distance cost must be at least as much as the reduction in edge building cost for v, and the latter is precisely α, which completes the proof. □

Theorem 5.6 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, and $\vec{S}$ a Nash equilibrium of $\mathcal{N}$. Then:*

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq 3 + \frac{\bar{d}_H}{\alpha/2 + \bar{d}_H}(2\log_2 n - 3) = O\left(1 + \min\left(1, \frac{\bar{d}_H}{\alpha}\right)\log n\right) = O(\log n) \ .$$

Proof: Let $G = \mathcal{G}(\vec{S})$. We select $\beta = \frac{\alpha}{2\log_2 n}$ and using Lemma 5.5 we get:

$$\bar{d}_G \leq 1 + \frac{4\bar{d}_H \log_2 n}{\alpha} + n^{1/\log_2 n} = 3 + \frac{4\bar{d}_H \log_2 n}{\alpha} \ .$$

The number of edges built by $\vec{S}$ is $\frac{1}{2}n\bar{d}_G$, so using the lower bound on the social optimum cost from

Theorem 4.1 we obtain:

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\alpha n \bar{d}_G/2}{n(\alpha/2+\bar{d}_H)} \leq \frac{\frac{3\alpha}{2}+2\bar{d}_H\log_2 n}{\alpha/2+\bar{d}_H} = 3 + \frac{\bar{d}_H}{\alpha/2+\bar{d}_H}(2\log_2 n - 3)$$
$$\leq 3 + 2\min\left(1, \frac{2\bar{d}_H}{\alpha}\right)\log_2 n = O\left(1+\min\left(1,\frac{\bar{d}_H}{\alpha}\right)\log n\right) = O(\log n) \ . \qquad \square$$

Lemma 5.7 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a joint strategy and $G = \mathcal{G}(\vec{S})$. Let $v \in V$ be a player who is playing a best response in $\vec{S}$. Let T be a shortest path tree of v in G. For any vertex $u \in V(T)$ (note that it is possible that $V(T) \neq V$), denote by k_u the number of friends of v that are descendants of u in T (including, possibly, u itself). Then:*

$$\forall u \in V(T), \alpha \geq (\delta_G(v,u)-1)k_u \ .$$

Proof: If $\delta_G(v,u) \leq 1$ then the claim is trivial. We therefore assume that u is not v and not adjacent to v in G. Let us consider a deviation for v in $\vec{S}$, in which it builds an additional edge to u. Any vertex w that is a descendant of u in T has a shortest path from v that goes through u, so in the alternative strategy it will have a new, shorter path, that uses the new edge instead of the original segment from v to u. For each such w the reduction in distance is at least $\delta_G(v,u)-1$, and for any other vertex the distance cannot increase, so the total decrease in distance cost for v is at least $(\delta_G(v,u)-1)k_u$. Since v is playing a best response this must be at most the increase in v's edge building cost, which is precisely α. $\square$

Theorem 5.9 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG and $\vec{S}$ a Nash equilibrium for $\mathcal{N}$. Then:*

$$\frac{Dist^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\bar{d}_H}{\alpha/2+\bar{d}_H}\left(1+\min(n-2, 2\sqrt{(n-2)\alpha/\bar{d}_H})\right)$$
$$= O\left(1+\min\left(1,\frac{\bar{d}_H}{\alpha}\right)\min\left(n, \sqrt{\frac{n\alpha}{\bar{d}_H}}\right)\right) = O(\sqrt{n})$$

Proof: Summing the bound in Lemma 5.8 over all $v \in V$ yields:

$$Dist^{\mathcal{N}}(\vec{S}) \leq n\bar{d}_H + 2\sqrt{(n-2)\alpha}\sum_{v\in V}\sqrt{d_H(v)} \ .$$

Using the Cauchy-Schwartz inequality we obtain:

$$\sum_{v\in V}\sqrt{d_H(v)} \le \sqrt{n\sum_{v\in V} d_H(v)} = n\sqrt{\bar{d}_H} \ ,$$

hence:

$$Dist^{\mathcal{N}}(\vec{S}) \le n\bar{d}_H + 2n\sqrt{(n-2)\alpha\bar{d}_H} \ .$$

Since all friend pairs must be connected in $\vec{S}$, the distance between any two friends is at most $n-1$, hence $Dist^{\mathcal{N}}(\vec{S}) \le n\bar{d}_H(n-1)$. Combining the two bounds we get:

$$Dist^{\mathcal{N}}(\vec{S}) \le n\bar{d}_H\left(1+\min\left(n-2, 2\sqrt{(n-2)\alpha/\bar{d}_H}\right)\right) \ ,$$

and dividing by the lower bound on social optimum (from Theorem 4.1) we get:

$$\begin{aligned}
\frac{Dist^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} &\le \frac{\bar{d}_H}{\alpha/2+\bar{d}_H}\left(1+\min(n-2, 2\sqrt{(n-2)\alpha/\bar{d}_H})\right)\\
&\le 1+\min\left(1,\frac{2\bar{d}_H}{\alpha}\right)\min\left(n-2, 2\sqrt{\frac{(n-2)\alpha}{\bar{d}_H}}\right)\\
&= O\left(1+\min\left(1,\frac{\bar{d}_H}{\alpha}\right)\min\left(n,\sqrt{\frac{n\alpha}{\bar{d}_H}}\right)\right) \ .
\end{aligned}$$

□

Theorem 5.10 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, and assume that a Nash equilibrium exists. Then:*

$$PoA(\mathcal{N}) \le \begin{cases} 1 & \textit{if} \quad 0<\alpha<1\\ 2-\frac{2}{\alpha+2} & \textit{if} \quad 1\le\alpha\le 2\\ 3+\frac{\bar{d}_H}{\alpha/2+\bar{d}_H}\min\left(\alpha,\ 2\log_2 n-2+\min\left(n-2,\ 2\sqrt{(n-2)\alpha/\bar{d}_H}\right)\right) & \textit{if} \quad 2<\alpha \end{cases}$$

$$= O\left(1+\min\left(\alpha,\ \bar{d}_H,\ \log n+\sqrt{\frac{n\alpha}{\bar{d}_H}},\ \sqrt{\frac{n\bar{d}_H}{\alpha}}\right)\right) = O(\sqrt{n}) \ .$$

Proof: For $\alpha \le 2$ the claim follows trivially from Theorem 5.2. For $\alpha > 2$, combining our bounds from Theorem 5.9 and Theorem 5.6 we get:

$$PoA(\mathcal{N}) \le 3+\frac{\bar{d}_H}{\alpha/2+\bar{d}_H}\left(2\log_2 n-2+\min\left(n-2,\ 2\sqrt{(n-2)\alpha/\bar{d}_H}\right)\right) \ .$$

Combining the bound from Theorem 5.2 we have:

$$
\begin{aligned}
PoA(\mathcal{N}) &\leq 3 + \frac{\bar{d}_H}{\alpha/2 + \bar{d}_H} \min\left(\alpha,\ 2\log_2 n - 2 + \min\left(n-2,\ 2\sqrt{(n-2)\alpha/\bar{d}_H}\right)\right) \\
&\leq 3 + \min(1,\ 2\bar{d}_H/\alpha) \min\left(\alpha,\ 2\log_2 n + n - 4,\ 2\log_2 n - 2 + 2\sqrt{(n-2)\alpha/\bar{d}_H}\right) \\
&= O\left(1 + \min(1,\ \bar{d}_H/\alpha) \min\left(\alpha,\ n,\ \log n + \sqrt{n\alpha/\bar{d}_H}\right)\right) \\
&= O\left(1 + \min\left(\alpha,\ \bar{d}_H,\ n,\ \frac{n\bar{d}_H}{\alpha},\ \log n + \sqrt{n\alpha/\bar{d}_H},\ \frac{\bar{d}_H}{\alpha}\log n + \sqrt{n\bar{d}_H/\alpha}\right)\right) \quad . \qquad \text{(B.1)}
\end{aligned}
$$

The third case in the minimum, n, is never minimal since $\bar{d}_H < n$. We now want to show that the term $\frac{\bar{d}_H}{\alpha}\log n$ can be removed from the last case in the minimum. For $\alpha < \bar{d}_H$ we have that the fifth case in the minimum is smaller than the last, i.e., $\log n + \sqrt{n\alpha/\bar{d}_H} < \frac{\bar{d}_H}{\alpha}\log n + \sqrt{n\bar{d}_H/\alpha}$. For $\alpha \geq \bar{d}_H$ and $n \geq 16$ we have that:

$$
\left(\frac{\bar{d}_H}{\alpha}\log n\right)^2 \leq \frac{\bar{d}_H}{\alpha}\log^2 n \leq \frac{n\bar{d}_H}{\alpha} \quad ,
$$

hence $\frac{\bar{d}_H}{\alpha}\log n \leq \sqrt{n\bar{d}_H/\alpha}$ so the term $\frac{\bar{d}_H}{\alpha}\log n$ is insignificant in (B.1) in any case and can be removed.

For the fourth case in the minimum in (B.1), $n\bar{d}_H/\alpha$, we notice that for any $x \geq 0$ we have $1 + x \geq \frac{1}{2}(1+\sqrt{x})$, and setting $x = n\bar{d}_H/\alpha$ gives that $1 + n\bar{d}_H/\alpha = \Omega(1 + \sqrt{n\bar{d}_H/\alpha})$, so removing this case from the minimum does not change the bound. These simplifications yield the following upper bound:

$$
PoA(\mathcal{N}) = O\left(1 + \min\left(\alpha,\ \bar{d}_H,\ \log n + \sqrt{\frac{n\alpha}{\bar{d}_H}},\ \sqrt{\frac{n\bar{d}_H}{\alpha}}\right)\right) \quad . \qquad \square
$$

Theorem 5.11 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha\rangle$ be an NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Denote $\gamma = 4\frac{\bar{d}_H}{\alpha}\log_2 n$. If $\gamma > 1$ then:*

$$
\begin{aligned}
\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} &\leq \frac{\alpha}{\alpha + 2\bar{d}_H}\left(1 + \eta\frac{\gamma}{\log_2 \gamma}\right) \\
&= O\left(\min\left(1, \frac{\bar{d}_H}{\alpha}\right)\frac{\log n}{\log\left(4\frac{\bar{d}_H}{\alpha}\log n\right)}\right) = O\left(\frac{\log n}{\log\left(4\frac{\bar{d}_H}{\alpha}\log n\right)}\right),
\end{aligned}
$$

where $\eta = 2 + \frac{1}{e\ln(2)-1} \approx 3.131$.

Proof: First, notice that for any $x > 0$ we have $x/e \geq \ln x$. Hence, for $\gamma \geq 1$ we obtain using $x = \log_2 \gamma$ that:

$$\log_2 \gamma - \log_2 \log_2 \gamma \geq x - \ln x/\ln(2) \geq x\left(1 - \frac{1}{e\ln(2)}\right) = \frac{1}{\eta - 1}\log_2 \gamma > 0 \ . \tag{B.2}$$

Therefore, selecting:

$$\beta = 2\bar{d}_H \frac{\log_2 \gamma - \log_2 \log_2 \gamma}{\gamma}$$

guarantees $\beta > 0$. Using Lemma 5.5 with this β (and using inequality (B.2) again) we obtain:

$$\begin{aligned}\bar{d}_G &\leq 1 + \frac{\gamma}{\log_2 \gamma - \log_2 \log_2 \gamma} + n^{\frac{\log_2 \gamma - \log_2 \log_2 \gamma}{\log_2 n}} = 1 + \frac{\gamma}{\log_2 \gamma - \log_2 \log_2 \gamma} + \frac{\gamma}{\log_2 \gamma} \\ &\leq 1 + \frac{\gamma}{\log_2 \gamma}\left(\frac{1}{1/(\eta - 1)} + 1\right) = 1 + \eta\frac{\gamma}{\log_2 \gamma} \ .\end{aligned}$$

We now use the lower bound on the social optimum from Theorem 4.1 to obtain:

$$\begin{aligned}\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} &\leq \frac{\alpha n \bar{d}_G/2}{n(\alpha/2 + \bar{d}_H)} \leq \frac{\alpha}{\alpha + 2\bar{d}_H}\left(1 + \eta\frac{\gamma}{\log_2 \gamma}\right) \\ &= O\left(\min\left(1, \frac{\bar{d}_H}{\alpha}\right)\frac{\log n}{\log\left(\frac{4\bar{d}_H}{\alpha}\log n\right)}\right) = O\left(\frac{\log n}{\log\left(\frac{4\bar{d}_H}{\alpha}\log n\right)}\right) \ .\end{aligned}$$

Note that we can't simply eliminate the 4 constant, since we need a guarantee that the log in the denominator is positive. Requiring $\gamma > 4$ (instead of $\gamma > 1$) would allow us to eliminate this constant. $\square$

Lemma 5.12 (Restated). *Let* $\mathcal{N} = \langle V, H, \alpha \rangle$ *be an NI-NCG,* $\vec{S}$ *a Nash equilibrium for* $\mathcal{N}$*, and* $G = \mathcal{G}(\vec{S})$*. Denote by* V_R *the set of players that have friends that are not adjacent to them in* G *and by* V_r *the set of players that build at least one edge to a non-friend. Let* $R_H(\vec{S}) = |V_R|$ *and* $r_H(\vec{S}) = |V_r|$*. Then:*

$$|E(G)| \leq |E_H| - \frac{R_H(\vec{S}) - r_H(\vec{S})}{2} \leq |E_H| \ ,$$

and:

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\bar{d}_H \alpha}{\alpha + 2\bar{d}_H} = O(\min(\alpha, \bar{d}_H)) \ .$$

Proof: Obviously if a player v builds an edge to a non-friend u, then it cannot be adjacent in G to all its friends, because in that case it would benefit from not building the edge (v, u). This implies that $V_r \subseteq V_R$.

We also denote by $k_v = |S_v - N_H(v)|$ the number of edges v builds to non-friends, and by $m_v = |N_H(v) - N_G(v)|$ the number of friends of v that are not adjacent to it in G. Clearly $k_v = 0$ iff $v \notin V_r$ and $m_v = 0$ iff $v \notin V_R$. Also, $\sum_{v \in V} k_v = |E(G) - E_H|$ (because $\vec{S}$ must be sensible) and $\sum_{v \in V} m_v = 2|E_H - E(G)|$.

For any $v \in V_r$, we consider the set:

$$W_v = \bigcup_{u \in S_v - N_H(v)} \Gamma_G^H(v, u) \quad ,$$

i.e., the set of all w that are friends of v for which there is a single edge built by v to a non-friend u, that is included in all shortest paths from v to w. The sets $\Gamma_G^H(v, u)$ in the union are pairwise disjoint (by definition of $\Gamma_G^H(v, u)$), so the size of W_v is the sum of their sizes. There are k_v sets in the union, and for each set we must have $|\Gamma_G(v, u)| \geq 2$, by Lemma 6.2 (since v builds an edge to u and u is not a friend of v). Therefore $|W_v| \geq 2k_v$. On the other hand, all members of W_v are friends of v that are not adjacent to v in G, so $|W_v| \leq m_v$. Combining these two bounds we get $2k_v \leq m_v$. Summing over all $v \in V_r$ we get:

$$2|E(G) - E_H| \leq \sum_{v \in V_r} m_v = \sum_{v \in V_R} m_v - \sum_{v \in V_R - V_r} m_v = \sum_{v \in V} m_v - \sum_{v \in V_R - V_r} m_v$$

For any $v \in V_R - V_r$ we know that $m_v \geq 1$. So we get:

$$2|E(G) - E_H| \leq 2|E_H - E(G)| - (R_H(\vec{S}) - r_H(\vec{S}))$$

or

$$|E(G) - E_H| \leq |E_H - E(G)| - \frac{R_H(\vec{S}) - r_H(\vec{S})}{2} \leq |E_H - E(G)|$$

where the right inequality is due to $V_r \subseteq V_R$. Adding $|E(G) \cap E_H|$ to all sides proves the first part of the lemma. The second part follows by $B^{\mathcal{N}} = |E(G)|\alpha \leq n\bar{d}_H\alpha/2$ and Theorem 4.1. $\square$

Appendix C

Proofs From Chapter 6

Lemma 6.1 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a joint strategy of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. Let $v \in V$ be some player and $u \in N_H(v)$ a friend of v. If v's strategy, S_v, is a best response in $\vec{S}$, then $\delta_G(v,u) \leq \min(\alpha + 1, n-1)$.*

Proof: If v and u are adjacent in G the claim is trivial. Let us therefore assume they are not adjacent in G. Clearly v and all its friends must be in the same connected component in G (otherwise v could strictly improve its cost by adding edges to all its disconnected friends). Therefore $\delta_G(v,u) \leq n-1$. Moreover, since v's cost is finite, if v were to build an additional edge to u it would exhibit an additional edge building cost of α, but reduce its distance cost by at least $\delta_G(v,u) - 1$. Since v is playing a best response in $\vec{S}$ we conclude that necessarily $\alpha - (\delta_G(v,u) - 1) \geq 0$, which completes the proof. □

Lemma 6.2 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, $\vec{S}$ a joint strategy and $G = \mathcal{G}(\vec{S})$. Let v be a player that is playing a best response in $\vec{S}$. If v builds an edge to u, and u is not a friend of v, then $|\Gamma_G^H(v,u)| \geq 2$.*

Proof: If $|\Gamma_G^H(v,u)| = 0$ then v can deviate by removing the edge to u, thus reducing its edge building cost without increasing its distance cost, yielding a total cost reduction, contradicting the fact that v is playing a best response in $\vec{S}$. If $|\Gamma_G^H(v,u)| = 1$ then let w be the single element of $\Gamma_G^H(v,u)$. It is impossible that $u = w$ because u is not a friend of v, so $\delta_G(v,w) > 1$. Therefore, adding the edge (v,w) strictly reduces the distance cost, and subsequently removing the edge (v,u) does not increase the distance to any friend, and maintains that the edge building cost did not change. Therefore, if v deviates by replacing the edge (v,u) with (v,w), it achieves cost reduction, hence v cannot be playing a best response in $\vec{S}$. Therefore, we must have $|\Gamma_G^H(v,u)| \geq 2$. □

Theorem 6.5 (Restated). *For any integer $n \geq 64$ there is an NI-NCG $\mathcal{N} = \langle V, H, \alpha \rangle$ with n players, $\alpha = \Theta(\log n / \log\log n)$, and $\bar{d}_H = \Omega(\log n)$, whose Price of Anarchy is $\Omega(\log n / \log\log n)$.*

Proof: Our goal is to construct a graph of large enough girth and average degree, which meets the conditions of Lemma 6.3. We select:

$$\alpha = \frac{\log(n/4)}{2\log\left(\frac{1}{2}\log(n/4)\right)} - 1 \quad \text{and} \quad g = \lceil 2\alpha + 3 \rceil.$$

For $n \geq 64$, clearly $g \geq 3$, and $2^g \leq 2^{2\alpha+4} \leq 4 \cdot 2^{\log(n/4)} = n$. Therefore, by Proposition 6.4, there is a graph H with n vertices and girth g, such that:

$$\bar{d}_H \geq \frac{1}{2}\left(\frac{n}{4}\right)^{1/\lceil 2\alpha+1 \rceil} \geq \frac{1}{2}\left(\frac{n}{4}\right)^{\frac{1}{2\alpha+2}} = \frac{1}{2}\left(\frac{n}{4}\right)^{\frac{\log\left(\frac{1}{2}\log(n/4)\right)}{\log(n/4)}} = \frac{1}{4}\log(n/4) = \Omega(\log n).$$

Clearly, by our selection of g, we have $\alpha < \frac{g}{2} - 1$, so this graph meets the conditions of Lemma 6.3, hence H is a Nash equilibrium network for the NI-NCG $\mathcal{N} = \langle V, H, \alpha \rangle$. Let $\vec{S} \in \overline{\mathbf{S}}(H)$ be a Nash equilibrium that creates the network H. Then the social cost of $\vec{S}$ is $C^{\mathcal{N}}(\vec{S}) = \alpha|E_H| + 2|E_H| = \frac{1}{2}(\alpha+2)n\bar{d}_H$, and using the upper bound on social cost from Theorem 4.1, we have:

$$\begin{aligned} PoA(\mathcal{N}) &\geq \frac{(\alpha+2)n\bar{d}_H}{4n(\alpha/2 + \bar{d}_H)} \geq \frac{1}{\frac{2}{\bar{d}_H} + \frac{4}{\alpha+1}} \geq \frac{1}{\frac{8}{\log(n/4)} + \frac{8\log\left(\frac{1}{2}\log(n/4)\right)}{\log(n/4)}} \\ &= \frac{\log(n/4)}{8\log\log(n/4)} = \Omega(\log n / \log\log n) \ . \end{aligned}$$

□

Appendix D

Proofs From Chapter 8

Theorem 8.1 (Restated). *Let $\mathcal{N} = \langle V, H, \alpha \rangle$ be an NI-NCG, where H is a forest. Then H is the unique social optimum network and the unique Nash equilibrium network. Consequently $PoS(\mathcal{N}) = PoA(\mathcal{N}) = 1$.*

Proof: We begin by proving that H is the unique social optimum network. Let $\vec{S}$ be a social optimum of $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. If $|E(G)| < |E_H|$ then G must have a larger number of connected components than H (since H is a forest), which implies that there are two friends that are disconnected in G, hence $C^{\mathcal{N}}(\vec{S}) = \infty$, contradicting the optimality of $\vec{S}$. Therefore $|E(G)| \geq |E_H|$, and by Lemma 4.3 this implies that $G = H$.

To prove that H is the unique Nash equilibrium network, let $\vec{S}$ be a Nash equilibrium of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. We must have $|E(G)| \geq |E_H|$, otherwise, G has too many connected components and we have a player with an infinite cost. Therefore we can use Lemma 5.12 to obtain:

$$|E_H| \leq |E(G)| \leq |E_H| - \frac{R_H(\vec{S}) - r_H(\vec{S})}{2} \quad ,$$

which implies $R_H(\vec{S}) \leq r_H(\vec{S})$. Recall that $R_H(\vec{S}) = |V_R|$ is the number of players with non-adjacent friends in G, and $r_H(\vec{S}) = |V_r|$ is the number of players building an edge to a non-friend.

By definition, $r_H(\vec{S}) \leq |E(G) - E_H|$, because each player counted in $r_H(\vec{S})$ builds at least one directed edge to a non-friend. Moreover, subtracting $|E(G) \cap E_H|$ from both sides of the inequality in Lemma 5.12 yields $|E(G) - E_H| \leq |E_H - E(G)|$. Combining these inequalities we obtain:

$$R_H(\vec{S}) \leq |E_H - E(G)| \ . \tag{D.1}$$

Let us define $\overline{H-G}$ as the subgraph of H with vertices V_R and edges $E_H - E(G)$. This is a valid graph since by definition of V_R, for every edge $(v,u) \in E_H - E(G)$ both $v \in V_R$ and $u \in V_R$. By inequality (D.1), $\overline{H-G}$ has at least as many edges as it has vertices. This implies that if $\overline{H-G}$ is not empty then it must contain a cycle, which is a contradiction since $\overline{H-G}$ is a subgraph of H, and H does not contain cycles. Therefore we must have $R_H(\vec{S}) = 0$, which implies (by definition of $R_H(\vec{S})$) that $E_H \subseteq E(G)$. Since by by Lemma 5.12 $|E(G)| \leq |E_H|$, we conclude that $G = H$.

This shows that H is the only possible Nash equilibrium network. Existence of a Nash equilibrium is guaranteed by Lemma 6.3, since $g(H) = \infty$. □

Lemma 8.3 (Restated). *In a Cycle NI-NCG, the social cost of any tree network that has a DFS traversal which discovers players in the order in which they appear in the cycle, is minimal among all networks that are trees. Any line or star network is such a minimal cost tree.*

Proof: Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG, and $\vec{S}$ a joint strategy for $\mathcal{N}$ that creates a tree network $T = \mathcal{G}(\vec{S})$. Let P be a tour of the players over the tree, starting from some player i, following the unique path in T to $i+1$, then to $i+2$, and so on until we reach player $i-1$ (cyclically), and finally back to player i. Clearly, the number of edges in this tour is half the social distance cost of $\vec{S}$. Let us view T as rooted on player i. Let $(v,u) \in E(T)$ be some edge in T and assume w.l.o.g. that v is u's parent in T. Clearly (v,u) must be traversed at least twice, since the tour must reach u, and once it reaches u for the first time, it must leave u's subtree in T at some point (to return eventually to player i). Therefore, the length of the tour is at least $2(n-1)$ and the social distance cost at least $4(n-1)$. If T has a DFS traversal as required, then this traversal defines such a tour P for which every edge is traversed exactly twice, hence the social distance cost is exactly $4(n-1)$. Since the social edge building cost of any sensible joint strategy creating a tree network is identical, we conclude that the social cost of trees with DFS traversals as required have minimal social cost among all sensible joint strategies creating tree networks. Since any line and any star has a DFS traversal that discovers players in the order in which they appear in the cycle, the proof is complete. □

Theorem 8.4 (Restated). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. The cycle is a social optimum network if and only if $\alpha \leq 2n-4$. Any line and any star are social optimum networks if and only if $\alpha \geq 2n-4$. Other social optimum networks may exist only if $\alpha \geq 2n-4$, and they must be trees.*

Proof: Any sensible joint strategy that creates the cycle network has social cost $(\alpha+2)n$. Any sensible joint strategy that creates the line network has social cost $(\alpha+4)(n-1)$. For $\alpha < 2n-4$, therefore, the cycle is strictly better than any line and by Lemma 8.3, strictly better than any tree. By Lemma 8.2, it is therefore the only social optimum network. For $\alpha > 2n-4$, any line and any star are strictly better than the cycle, and therefore the cycle is not a social optimum network. By Lemma 8.2, only tree social optimum networks exist, and by Lemma 8.3, the line and star are two of those. For $\alpha = 2n-4$, the cycle, line and star are all social optimum networks. □

Theorem 8.5 (Restated). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. The cycle is a Nash equilibrium network if and only if $0 \leq \alpha \leq n-2$. Any line is a Nash equilibrium network if and only if $\alpha \geq n-2$. Any star is a Nash equilibrium network if $\alpha \geq 2$. Consequently, a Nash equilibrium always exists for a Cycle NI-NCG.*

Proof: Let $\vec{S}$ be any sensible joint strategy that creates the cycle. All players in $\vec{S}$ are using a complete friendship response. Therefore a player can only reduce its cost by building less edges. This is impossible for players that build no edges. For a player that builds a single edge, building no edges yields a cost increase of $n - 2 - \alpha$. For a player that builds two edges, building no edges yields infinite cost, while building one edge yields a cost increase of $n - 2 - \alpha$, regardless of which edge is built. Therefore, if $n - 2 - \alpha \geq 0$, $\vec{S}$ is a Nash equilibrium. Conversely, if there is some $\vec{S}$ that creates the cycle and is a Nash equilibrium, then it must have at least one vertex that builds one or two edges, so necessarily $n - 2 - \alpha \geq 0$.

Let $\vec{S}$ now be a sensible joint strategy that creates some line. It is obvious that all the internal players (in the line) are using best response strategies - they are using complete friendship responses, but building fewer edges would necessarily disconnect them from at least one friend. A deviation for any of the players at any end of the line cannot reduce its cost if it doesn't build additional edges. The best deviation for such a player, among those that do build additional edges, is the one that builds a single edge to the other end of the line. Such a deviation yields a cost reduction of $n - 2 - \alpha$. Therefore, for $\alpha < n - 2$, $\vec{S}$ cannot be a Nash equilibrium, and for $\alpha \geq n - 2$, $\vec{S}$ is a Nash equilibrium.

Finally we notice that for H that is a cycle, the graph constructed in the proof of Theorem 4.5 is a star graph, and moreover, any star network may be constructed using the process described in the proof. Therefore, any star is a Nash equilibrium network for $\alpha \geq 2$. $\square$

Lemma 8.6 (Restated). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, and $G = \mathcal{G}(\vec{S})$. Then there is a segment P of the cycle (which may contain anywhere between 0 and n vertices) so that:*

1. *G contains P, and players in P only build edges to other players in P.*
2. *For any $v \in V - V(P)$, v's strategy in $\vec{S}$ builds a single edge to player u that is not a friend*

of v, so that u is contained in some shortest path (in G) between v's friends, and v is not contained in any such shortest path.

Proof: We first analyze the possible best responses for a player $v \in V$. If we denote by k_v the number of friends that build edges to v in $\vec{S}$, then v's best response cannot build more than $2 - k_v$ edges (because a complete friendship response would be better). Moreover, if it builds exactly $2 - k_v$ edges then it must use a complete friendship response (building these edges to the friends that have not built edges to v). This implies that all the non-friendship edges built in $\vec{S}$ are built by players $v \in V$ for which $k_v = 0$. Moreover, these players build exactly a single such edge.

Let us denote by $Q \subseteq V$ the set of these players (that build a single edge to a non-friend). We analyze 3 cases to prove the first part of the lemma:

1. $Q = \emptyset$

   In this case only friendship edges are built, and we must have that G is either a line or the entire cycle, so defining $P = G$ satisfies the first part of the lemma.

2. $Q = V$

   In this case all players build a single non-friendship edge, so defining P to be the empty graph satisfies the first part of the lemma.

3. $Q \neq V$ and $Q \neq \emptyset$

   In this case we define $U = V - Q$ and P as the induced subgraph of H over U. P has $n - |Q|$ vertices and it is acyclic, so it has at most $n - 1 - |Q|$ edges. Let us denote by W the set of friendship edges built in $\vec{S}$. Clearly if a friendship edge is built in $\vec{S}$, both its ends must be in U, because a player cannot belong to Q if a friendship edge is built by it or to it. Therefore, $W \subseteq E(P)$, so $|W| \leq |E(P)| \leq n - 1 - |Q|$. However, we must also have

$|W| + |Q| = |E(G)| \geq n - 1$ (otherwise G is disconnected). Combining these inequalities we conclude that $|W| = |E(P)| = n - 1 - |Q|$, so P must be connected, and since it is a subgraph of H, it must be a segment. We also conclude that $W = E(P)$ (because $W \subseteq E(P)$), so P is contained in G. Since all friendship edges built in $\vec{S}$ are in P, and all non-friendship edges built in $\vec{S}$ are not built by players in P, we conclude that players in P only build edges to other players in P.

For the second part of the lemma, for every player $v \in V - V(P)$ (i.e., $v \in Q$), we already know that v builds a single edge to a non-friend u. We have $|\Gamma_G^H(v, u)| \geq 2$ by Lemma 6.2, and since v has exactly two friends, we conclude that any shortest path in G between v and its friends must begin with the edge (v, u). Denoting the friends of v by x and w, this implies that no shortest path between x and w in G may contain v, since splitting such a path into two parts by v produces two shortest paths, from v to each friend, and it is not possible that both parts contain (v, u). Moreover, since any shortest path between v and x or v and w begins with (v, u), we conclude that v's distance cost is $2 + \delta_G(u, x) + \delta_G(u, w)$, so v must select u to minimize this sum, which happens only when u is contained in a shortest path between x and w. □

Lemma 8.7 (Restated). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$, $G = \mathcal{G}(\vec{S})$, and P as in Lemma 8.6. If $\alpha > n - 2$ then $V(P) \neq \emptyset$, $P \neq H$, and G is a tree.*

Proof: From Theorem 8.5 $P \neq H$. To see why $V(P) = \emptyset$ is impossible, assume for contradiction $V(P) = \emptyset$. We notice that the joint strategy graph $\mathcal{G}^*(\vec{S})$ must contain a single cycle C (since it has exactly n edges and n vertices). Moreover, C must consist of edges built in a single cycle orientation, otherwise there is a player that builds at least 2 edges, contradicting Lemma 8.6. For any player $v \in V$ and friend $u \in N_H(v)$, any shortest path from v to u must begin with the edge v built, otherwise v would have preferred to build an edge to his other friend (recall that since

$V(P) = \emptyset$, all players build a single edge to a non-friend). Since this also holds for u, we conclude that any shortest path between friends has two edges directed at each other, hence, no two friends may be members of C (since it has a single edge orientation). Therefore the length of C is at most $\lfloor n/2 \rfloor$. For any player $v \in V(C)$, using Lemma 5.4 on the single edge it builds we now get $\alpha \leq 2(|C| - 2) \leq n - 4 < \alpha$, which is a contradiction.

The only cases from Lemma 8.6 where G may have more than $n - 1$ edges are $V(P) = \emptyset$ or $P = H$. We have shown that both are impossible for $\alpha > n - 2$, and G must be connected, so it must be a tree. $\square$

Theorem 8.8 (Restated). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG. Then:*

$$PoA(\mathcal{N}) \leq \begin{cases} 1 & \text{if} \quad \alpha < 1 \\ 2 - \frac{2}{\alpha+2} & \text{if} \quad 1 \leq \alpha \leq n-2 \\ 1 + \frac{n-1}{\alpha+2} - \frac{1}{n} & \text{if} \quad n-2 < \alpha \leq 2n-4 \\ 1 + \frac{n-2}{\alpha+4} & \text{if} \quad 2n-4 < \alpha \end{cases},$$

which is less than 2 for $1 \leq \alpha < 2n - 4$, less than $3/2$ for $\alpha \geq 2n - 4$, and goes to 1 as $\alpha \to \infty$.

Proof: For $\alpha < 1$ the Price of Anarchy is 1 (by Theorem 5.2). For $1 \leq \alpha \leq n - 2$, the cycle is a social optimum network (Theorem 8.4), and has cost $(\alpha + 2)n$. Lemma 5.1 gives an upper bound of $2(\alpha + 1)n$ on the social cost of any Nash equilibrium. Hence in this case:

$$PoA(\mathcal{N}) \leq \frac{2(\alpha+1)n}{(\alpha+2)n} = 2 - \frac{2}{\alpha+2} < 2 \, .$$

For $\alpha > n - 2$, let $\vec{S}$ be any Nash equilibrium, $G = \mathcal{G}(\vec{S})$, and P as defined in Lemma 8.6. From Lemma 8.7 we know that $V(P) \neq \emptyset$, $P \neq H$, and G is a tree, hence it has $n - 1$ edges. We therefore analyze the distance cost as follows:

1. For any player $v \in V - V(P)$, let x and w be v's two friends. We know that v builds a single edge to some non-friend u. Let D be v's distance cost in $\vec{S}$. By Lemma 8.6, u lies on a shortest path between x and w, and any shortest path from v to x or w begins with (v, u), so we must have $D = 2 + \delta_G(x, w)$. Since x and w must be connected and v is not on the shortest path between them, we have $\delta_G(x, w) \leq n - 2$, so we obtain $D \leq n$. Therefore the total distance cost contribution of these players is at most of $(n - |V(P)|)n$.

2. For every friendship edge in P, the distance is 1. We have a total distance cost contribution of $2(|V(P)| - 1)$ for these friendship edges (recall that $|V(P| \geq 1)$.

3. The only friendship edges we have not accounted for are those not in P, that are incident on one of the players at any end of P. Note that if P is a line, there is one such friendship edge, and we have not counted its distance in any direction so far. If P is not a line then there are 2 such edges, and we have counted each in one direction, but not in the other. The distance between the friends on these edges is at most $n - 1$ for a total distance cost contribution of at most $2(n - 1)$ (recall that P is not the entire cycle).

Therefore we get an upper bound of:

$$\begin{aligned} Dist^{\mathcal{N}}(\vec{S}) &\leq 2(|V(P)| - 1) + 2(n - 1) + (n - |V(P)|)n \\ &= n^2 + 2n - 4 - |V(P)|(n - 2) \leq n^2 + n - 2 \ , \end{aligned}$$

since $|V(P)| \geq 1$. Adding the edge building cost of $(n - 1)\alpha$ we get:

$$C^{\mathcal{N}}(\vec{S}) \leq \alpha n + n^2 + n - 2 - \alpha \ . \tag{D.2}$$

For $n - 2 < \alpha \leq 2n - 4$, the cycle is the social optimum. Therefore:

$$PoA(\mathcal{N}) \leq \frac{\alpha n + n^2 + n - 2 - \alpha}{(\alpha + 2)n} = 1 + \frac{n - 1}{\alpha + 2} - \frac{1}{n} < 2 - \frac{2}{n} < 2 \ .$$

For $\alpha > 2n-4$, the line is a social optimum network with cost $(\alpha+4)(n-1)$, so we get:

$$PoA(\mathcal{N}) \le \frac{\alpha n + n^2 + n - 2 - \alpha}{(\alpha+4)(n-1)} = 1 + \frac{n-2}{\alpha+4} < \frac{3}{2} - \frac{1}{n} < \frac{3}{2} \ . \qquad \square$$

Lemma 8.10 (Restated). *Let $\mathcal{N} = \langle V, H = C, \alpha \rangle$ be a Cycle NI-NCG and k an integer such that $2 \le k < n$ and $n-k-1 \equiv 0 (mod\ 4)$. If $\alpha \ge k$ then $X^{k,n}$ is a Nash equilibrium for $\mathcal{N}$.*

Proof: Let $G = \mathcal{G}(X^{k,n})$. It is obvious that no player may benefit from building fewer edges, since this would disconnect it from at least one friend. We now consider the players one by one. Players $1, ..., k-1$ are all using best response strategies, since they are using complete friendship responses while building a single edge (that cannot be removed).

For player 0, its distance to its friend player 1 is exactly 1, and its distance to its other friend, player $n-1$, is exactly $1+k-1$ because player $n-1$ must be connected to player $k-1$ (this is true even if $k = n-1$). Player 0 therefore has distance cost of $1+k$. Deviating by building a single edge to any other node cannot reduce the distance cost since player 1 and $n-1$ are at distance $k-1$. The best deviation that builds more than one edge is the complete friendship response, building a single edge to each friend, and this deviation increases the cost by at least $(2\alpha+2) - (\alpha+1+k)$ which is positive for $\alpha \ge k$. Hence, player 0 is using a best response. A similar consideration applies to player k. Its only viable alternative to not building any edges (its strategy in $X^{k,n}$) is a complete friendship response, building an edge directly to player $k+1$, which yields a cost increase of $(\alpha+2) - (k+1)$, which is, once again, positive for $\alpha \ge k$. Hence, player k is also using a best response.

Each player i, for $k < i < n$, builds a single edge to some player $j \in \{1, k-1\}$. It is easy to see that in any case, one of i's friends is connected to the same player j, while the other is connected to the other player $k-j$ (it is possible that both j and $k-j$ are the same player if $k=2$). Note that this also holds for players $i = k+1$ and $i = n-1$. The distance cost for player i is therefore

$2 + k$. The possible deviations for player i are:

1. Not building any edges, which yields infinite cost.

2. Building a single edge to a different player, which cannot decrease the distance cost, since in $X^{k,n}$ player i builds an edge to a point on the shortest path between its friends, $i - 1$ and $i + 1$.

3. Of the deviations that build more than one edge, the best is necessarily the complete friendship response, which yields a cost increase of $(2\alpha+2)-(\alpha+k+2)$ which is non-negative for $\alpha \geq k$.

Therefore, for any $k < i < n$ player i is also playing a best response, which completes the proof. □

Theorem 8.11 (Restated). *For infinitely many n, for any $\alpha \geq (n+1)/2$, and any $\mathcal{N} = \langle V, H = C, \alpha\rangle$ (of n players and edge building price α), we have:*

$$PoA(\mathcal{N}) \geq \begin{cases} 1 - \frac{1}{n} + \frac{1}{4}\frac{n+1}{\alpha+2} & \textit{for} \quad \alpha < 2n-4 \\ 1 - \frac{2}{\alpha+2} + \frac{n+1}{4(\alpha+4)} & \textit{for} \quad \alpha \geq 2n-4 \end{cases} .$$

Additionally, for any n and any $\mathcal{N} = \langle V, H = C, \alpha\rangle$ with $2 \leq \alpha \leq (n+1)/2$, if α is an integer for which $\alpha \equiv n-1$ (mod 4), then:

$$PoA(\mathcal{N}) \geq 2 - \frac{2}{\alpha+2} - \frac{\alpha-2}{n} - \frac{6}{(\alpha+2)n} .$$

Proof: Lemma 8.10 shows that $X^{k,n}$ is a Nash equilibrium for $\alpha \geq k$. Its cost is given by:

$$C^{\mathcal{N}}(X^{k,n}) = \alpha(n-1) + 2k + 2k + (n-k-1)(2+k) = (\alpha+2)(n-1) + k(n+1) - k^2 ,$$

which, as a function of k (over the real numbers), is monotonously increasing for $k <= (n+1)/2$, and monotonously decreasing for $k > (n+1)/2$. Our goal is to select an integer k so that $\alpha \geq k$

and $n-k-1 \equiv 0 \pmod 4$, that maximizes the cost $C^{\mathcal{N}}(X^{k,n})$. For the first case in the theorem, if $2 \le \alpha \le (n+1)/2$ and α is an integer such that $\alpha \equiv n-1 (\mathrm{mod}\ 4)$ then we select $k=\alpha$, and get:

$$C^{\mathcal{N}}(X^{k,n}) = (\alpha+2)(n-1) + \alpha(n+1) - \alpha^2 = (2\alpha+4)n - 2n - (\alpha+2)(\alpha-2) - 6 \ ,$$

and therefore the Price of Anarchy is at least:

$$PoA(\mathcal{N}) \ge 2 - \frac{2}{\alpha+2} - \frac{\alpha-2}{n} - \frac{6}{(\alpha+2)n} \ .$$

For the second case, let n be any integer so that $n \equiv 3 \pmod 8$. Selecting $k = (n+1)/2$ guarantees that $n-k-1 \equiv 0 \pmod 4$ and $\alpha \ge k$, so $X^{k,n}$ is a Nash equilibrium. The cost is:

$$C^{\mathcal{N}}(X^{k,n}) = (\alpha+2)(n-1) + (n+1)^2/2 - (n+1)^2/4 = (\alpha+2)(n-1) + (n+1)^2/4 \ .$$

The Price of Anarchy in this case for $\alpha < 2n-4$ is therefore at least:

$$PoA(\mathcal{N}) \ge 1 - \frac{1}{n} + \frac{1}{4}\frac{(n+1)^2}{(\alpha+2)n} \ge 1 - \frac{1}{n} + \frac{n+1}{4(\alpha+2)} \ ,$$

and for $\alpha \ge 2n-4$ it is at least:

$$PoA(\mathcal{N}) \ge 1 - \frac{2}{\alpha+2} + \frac{1}{4}\frac{(n+1)^2}{(\alpha+4)(n-1)} \ge 1 - \frac{2}{\alpha+2} + \frac{n+1}{4(\alpha+4)} \ . \qquad \square$$

Appendix E

Proofs From Chapter 9

Theorem 9.1 (Restated). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, and assume that a Nash equilibrium exists. Let $\hat{\alpha} = \alpha / \bar{f}_H$. Then:*

$$\begin{aligned} PoA(\mathcal{N}) &\leq 3 + \frac{\bar{d}_H^+}{\hat{\alpha}/2 + \bar{d}_H^+} \min\left(\hat{\alpha},\ 2\log_2 n - 2 + \min\left(n-2,\ 2\sqrt{(n-2)\hat{\alpha}/\bar{d}_H^+}\right)\right) \\ &= O\left(1 + \min\left(\hat{\alpha},\ \bar{d}_H^+,\ \log n + \sqrt{\frac{n\hat{\alpha}}{\bar{d}_H^+}},\ \sqrt{\frac{n\bar{d}_H^+}{\hat{\alpha}}}\right)\right) = O(\sqrt{n}) \ . \end{aligned}$$

To prove this theorem, we first prove the following lemmas, the equivalents of our results from Chapter 5.

Lemma E.1 (Weighted Equivalent of Theorem 4.1). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG. Then:*

$$n(\alpha/2 + \bar{f}_V) \leq OPT(\mathcal{N}) < 2n(\alpha/2 + \bar{f}_V) \ .$$

Proof: Let G be some social optimum network. As for an NI-NCG, G must have at least $n/2$ edges, otherwise it has infinite cost. The distance between any pair of semi-friends is at least 1, hence the social distance cost is at least F. This proves the lower bound. The upper bound is obtained by analyzing a sensible joint strategy that creates the star network, which builds $n-1$ edges and has distance at most 2 between any two semi-friends. $\square$

Lemma E.2 (Weighted Equivalent of Lemma 5.1). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG and $\vec{S}$ a Nash equilibrium of $\mathcal{N}$. Then:*

$$C^{\mathcal{N}}(\vec{S}) \leq n\bar{d}_H^+ \bar{f}_H(\hat{\alpha} + 1) \ .$$

Proof: Every player $v \in V$ is playing a best response in $\vec{S}$, hence its strategy S_v is no worse than the complete friendship response for v. The cost of the complete friendship response for v is at most $\alpha d_H^+(v) + \sum_{u \neq v} f(v, u)$, and summing over all players completes the proof. □

Lemma E.3 (Weighted Equivalent of Theorem 5.2). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, and assume that a Nash equilibrium exists. Then:*

$$PoA(\mathcal{N}) \leq 1 + \frac{\bar{d}_H^+ \hat{\alpha}}{\hat{\alpha}/2 + \bar{d}_H^+} = O(1 + \min(\hat{\alpha}, \bar{d}_H^+)) \ .$$

Proof: Using Lemma E.2 and our lower bound on social optimum cost from Lemma E.1 we get:

$$PoA(\mathcal{N}) \leq \frac{n\bar{d}_H^+ \bar{f}_H(\hat{\alpha} + 1)}{n(\alpha/2 + \bar{f}_V)} \leq 1 + \frac{\bar{d}_H^+ \hat{\alpha}}{\hat{\alpha}/2 + \bar{d}_H^+} \leq 1 + \min(\hat{\alpha}, 2\bar{d}_H^+) = O(1 + \min(\hat{\alpha}, \bar{d}_H^+)) \ . \qquad \square$$

Definition E.4 (Weighted Equivalent of Definition 5.3)**.** Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, $\vec{S}$ a joint strategy of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. We define $\Gamma_G^H(v, u)$ exactly as before, as the set of friends of v whose distance from v would increase should the edge (v, u) be removed from the network:

$$\Gamma_G^H(v, u) = \{x \in N_H^+(v) : \quad \Psi_G(v, x) \neq \emptyset \quad \textbf{and} \quad \forall p \in \Psi_G(v, x), u \in p\} \ .$$

For any $v \in V$ and $u \in S_v$ we define $w_f(v, u) = \sum_{x \in \Gamma_G^H(v,u)} f(v, x)$, and $W_f = \sum_{v \in V, u \in S_v} w_f(v, u)$.

Lemma E.5 (Weighted Equivalent of Lemma 5.4). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, $\vec{S}$ a Nash equilibrium for $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. If v builds an edge to u in $\vec{S}$ (i.e., $u \in S_v$), then $\alpha \leq w_f(v, u)(\delta_{G'}(v, u) - 1)$, where G' is the graph G without the edge (v, u).*

Proof: As in the proof of Lemma 5.4, if v and u become disconnected in G' then the claim is obvious. Otherwise, let $\vec{S'} = (\vec{S}_{-v}, S_v - \{u\})$ be the unilateral deviation of v obtained by removing the edge to u. The set of friends whose distance from v changes as a result of the deviation is exactly $\Gamma_G^H(v,u)$. For any such friend $x \in \Gamma_G^H(v,u)$, we have $\delta_{G'}(v,x) - \delta_G(v,x) \leq \delta_{G'}(v,u) - 1$. Hence:

$$\begin{aligned} Dist^{\mathcal{N}}(\vec{S}', v) - Dist^{\mathcal{N}}(\vec{S}, v) &= \sum_{x \in \Gamma_G^H(v,u)} f(v,x)(\delta_{G'}(v,x) - \delta_G(v,x)) \\ &\leq \sum_{x \in \Gamma_G^H(v,u)} f(v,x)(\delta_{G'}(v,u) - 1) = w_f(v,u)(\delta_{G'}(v,u) - 1) \ . \end{aligned}$$

Since v is playing a best response in $\vec{S}$, this must be at least α (the reduction in v's edge building cost), completing the proof. □

Lemma E.6 (Weighted Equivalent of Lemma 5.5). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, $\vec{S}$ a Nash equilibrium of $\mathcal{N}$ and $G = \mathcal{G}(\vec{S})$. Then for any $\beta > 0$*

$$\bar{d}_G \leq 1 + \frac{2\bar{f}_V}{\beta} + n^{2\beta/\alpha} \ .$$

Proof: Our first step is to transform G by removing all edges $(v,u) \in E(G)$ for which $w_f(v,u) \geq \beta$ (where v builds the edge in $\vec{S}$). We denote the resulting graph by G_β, and by m the number of removed edges. The total weight of edges removed is at least βm, and at most the total weight of edges in the original graph, W_f. Therefore, $\beta m \leq W + f$. On the other hand, for any $v \in V$ and $x \in N_H^+(v)$ (a friend of v), there is at most one $u \in S_v$ such that $x \in \Gamma_G^H(v,u)$ (these sets, for a fixed player v, must be pairwise disjoint). Therefore $f(v,x)$ is added at most once in the sum W_f, hence $W_f \leq F$. Combining these inequalities we get $|E(G_\beta)| \geq |E(G)| - F/\beta$. After dividing by $n/2$, we have

$$\bar{d}_{G_\beta} \geq \bar{d}_G - \frac{2\bar{f}_V}{\beta} \ . \tag{E.1}$$

For $\bar{d}_{G_\beta} \leq 2$ this yields $\bar{d}_G \leq \frac{2\bar{f}_V}{\beta} + 2$ and the claim holds trivially since $n^{2\beta/\alpha} > 1$. We therefore assume that $\bar{d}_{G_\beta} > 2$, in which case G_β has a cycle. Using Lemma E.5 (as in the proof of Lemma

5.5) and the bound of Alon et. al. [AHL02] on the number of vertices in a graph of given average degree and girth (see Appendix F), we get:

$$\bar{d}_{G_\beta} \leq 1 + n^{\left(\frac{2}{g(G_\beta)-2}\right)} \leq 1 + n^{2\beta/\alpha} \quad .$$

Using this with inequality (E.1) completes the proof. □

Lemma E.7 (Weighted Equivalent of Theorem 5.6). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, and $\vec{S}$ a Nash equilibrium of $\mathcal{N}$. Then:*

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq 3 + \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V}(2\log_2 n - 3) = O\left(1 + \min\left(1, \frac{\bar{f}_V}{\alpha}\right)\log n\right) = O(\log n) \quad .$$

Proof: Let $G = \mathcal{G}(\vec{S})$. We select $\beta = \frac{\alpha}{2\log_2 n}$ and using Lemma E.6 we get:

$$\bar{d}_G \leq 1 + \frac{4\bar{f}_V \log_2 n}{\alpha} + n^{1/\log_2 n} = 3 + \frac{4\bar{f}_V \log_2 n}{\alpha} \quad .$$

The number of edges built by $\vec{S}$ is $\frac{1}{2}n\bar{d}_G$, so using the lower bound on the social optimum cost from Lemma E.1 we obtain:

$$\begin{aligned}\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} &\leq \frac{n\alpha\bar{d}_G/2}{n(\alpha/2 + \bar{f}_V)} \leq \frac{\frac{3\alpha}{2} + 2\bar{f}_V \log_2 n}{\alpha/2 + \bar{f}_V} = 3 + \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V}(2\log_2 n - 3) \\ &\leq 3 + 2\min\left(1, \frac{2\bar{f}_V}{\alpha}\right)\log_2 n = O\left(1 + \min\left(1, \frac{\bar{f}_V}{\alpha}\right)\log n\right) = O(\log n) \quad .\end{aligned}$$ □

Lemma E.8 (Weighted Equivalent of Lemma 5.7). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, $\vec{S}$ a joint strategy and $G = \mathcal{G}(\vec{S})$. Let $v \in V$ be a player who is playing a best response in $\vec{S}$. Let T be a shortest path tree of v in G. For any vertex $u \in V(T)$ (note that it is possible that $V(T) \neq V$), denote by Λ_u the set of friends of v that are descendants of u in T (including, possibly, u itself). Let $k_u = \sum_{w \in \Lambda_u} f(v, w)$ be the total friendship weight for v in u's subtree. Then:*

$$\forall u \in V(T), \alpha \geq (\delta_G(v, u) - 1)k_u \quad .$$

Proof: If $\delta_G(v,u) \leq 1$ then the claim is trivial. Otherwise we consider a deviation for v that builds an additional edge to u. For any vertex w that is a descendant of u in T the reduction in distance is at least $\delta_G(v,u)-1$, and for any other vertex the distance cannot increase, so the total decrease in distance cost for v is at least $\sum_{w\in\Lambda_u} f(v,w)(\delta_G(v,u)-1) = (\delta_G(v,u)-1)k_u$. Since v is playing a best response this must be at most the increase in v's edge building cost, which is precisely α. $\square$

Lemma E.9 (Weighted Equivalent of Lemma 5.8). *Let $\mathcal{N} = \langle V, f, \alpha\rangle$ be a WI-NCG and let $\vec{S}$ be a joint strategy. Let $v \in V$ be a player who is playing a best response in $\vec{S}$, and denote $f_v = \sum_{w\neq v} f(v,w)$. Then:*

$$Dist^{\mathcal{N}}(\vec{S}, v) \leq 2\sqrt{(n-2)\alpha f_v} + f_v \ .$$

Proof: Let $G = \mathcal{G}(\vec{S})$. All $w \in N^+_H(v)$ must be in the same connected component as v. For any integer $h \geq 2$, we define $F_h = \{w \in N^+_H(v) : \delta_G(v,w) \geq h\}$, i.e., F_h is the set of friends of v in depth at least h in T, where T is the shortest path tree of v in G. Then:

$$\begin{aligned} Dist^{\mathcal{N}}(\vec{S}, v) &= \sum_{w\in N^+_H(v)/F_h} f(v,w)\delta_G(v,w) + \sum_{w\in F_h} f(v,w)\delta_G(v,w) \\ &\leq (h-1) \sum_{w\in N^+_H(v)/F_h} f(v,w) + \sum_{w\in F_h} f(v,w)\delta_G(v,w) \\ &= (h-1)f_v + \sum_{w\in F_h} f(v,w)(\delta_G(v,w)-h+1) \ . \end{aligned} \tag{E.2}$$

For any $w \in F_h$, the term $\delta_G(v,w)-h+1$ is exactly the number of vertices on the path between depth h in T and w. Therefore in (E.2), for every $w \in F_h$ and every vertex on the path from depth h in T to w, we add $f(v,w)$. Changing the summation order, this is equivalent to adding k_u for every $u \in V(T)$ in depth h or greater. Hence, using Lemma E.8 and the fact that $h \geq 2$, we have:

$$\sum_{w\in F_h} f(v,w)(\delta_G(v,w)-h+1) = \sum_{u\in V(T):\delta_G(v,u)\geq h} k_u \leq \sum_{u\in V(T):\delta_G(v,u)\geq h} \frac{\alpha}{\delta_G(v,u)-1} \leq \frac{(n-2)\alpha}{h-1} \ ,$$

and using this in inequality (E.2) yields:

$$Dist^{\mathcal{N}}(\vec{S}, v) \leq (h-1)f_v + \frac{(n-2)\alpha}{h-1} \quad ,$$

for any $h \geq 2$. Selecting $h = 1 + \lceil \sqrt{(n-2)\alpha/f_v} \rceil$ completes the proof. □

Lemma E.10 (Weighted Equivalent of Theorem 5.9). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG and $\vec{S}$ a Nash equilibrium for $\mathcal{N}$. Then:*

$$\frac{Dist^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V} \left(1 + \min(n-2, 2\sqrt{(n-2)\alpha/\bar{f}_V})\right)$$
$$= O\left(1 + \min\left(1, \frac{\bar{f}_V}{\alpha}\right) \min\left(n, \sqrt{\frac{n\alpha}{\bar{f}_V}}\right)\right) = O(\sqrt{n})$$

Proof: Summing the bound in Lemma E.9 over all $v \in V$ yields:

$$Dist^{\mathcal{N}}(\vec{S}) \leq F + 2\sqrt{(n-2)\alpha} \sum_{v \in V} \sqrt{f_v} \quad .$$

Using the Cauchy-Schwartz inequality we obtain:

$$\sum_{v \in V} \sqrt{f_v} \leq \sqrt{n \sum_{v \in V} f_v} = n\sqrt{\bar{f}_V} \quad ,$$

hence:

$$Dist^{\mathcal{N}}(\vec{S}) \leq F + 2n\sqrt{(n-2)\alpha \bar{f}_V} = n\bar{f}_V \left(1 + 2\sqrt{(n-2)\alpha/\bar{f}_V}\right) \quad .$$

Since all semi-friend pairs must be connected in $\vec{S}$, the distance between any two semi-friends is at most $n-1$, hence $Dist^{\mathcal{N}}(\vec{S}) \leq n\bar{f}_V(n-1)$. Combining the two bounds we get:

$$Dist^{\mathcal{N}}(\vec{S}) \leq n\bar{f}_V \left(1 + \min\left(n-2, 2\sqrt{(n-2)\alpha/\bar{f}_V}\right)\right) \quad ,$$

and dividing by the lower bound on social optimum (from Lemma E.1) we get:

$$\frac{Dist^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V} \left(1 + \min(n-2, 2\sqrt{(n-2)\alpha/\bar{f}_V})\right)$$
$$\leq 1 + \min\left(1, \frac{2\bar{f}_V}{\alpha}\right) \min\left(n-2, 2\sqrt{\frac{n-2\alpha}{\bar{f}_V}}\right)$$
$$= O\left(1 + \min\left(1, \frac{\bar{f}_V}{\alpha}\right) \min\left(n, \sqrt{\frac{n\alpha}{\bar{f}_V}}\right)\right) \quad .$$

□

We are now ready to prove Theorem 9.1.

Proof of Theorem 9.1: For any $\vec{S}$ that is a Nash equilibrium, we can rewrite Lemma E.7 using $\hat{\alpha}$ as:

$$\frac{B^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq 3 + \frac{\bar{d}_H^+}{\hat{\alpha}/2 + \bar{d}_H^+}(2\log_2 n - 3) \ ,$$

and Lemma E.10 as:

$$\frac{Dist^{\mathcal{N}}(\vec{S})}{OPT(\mathcal{N})} \leq \frac{\bar{d}_H^+}{\hat{\alpha}/2 + \bar{d}_H^+}\left(1 + \min(n-2, 2\sqrt{(n-2)\hat{\alpha}/\bar{d}_H^+})\right) \ .$$

We now notice that these rewritten lemmata, along with Lemma E.3, are exactly the same as we had for the NI-NCG in Theorem 5.6, Theorem 5.9, and Theorem 5.2 (for the case $\alpha > 2$), respectively, after replacing α by $\hat{\alpha}$ and $\bar{d}_H$ by $\bar{d}_H^+$. The rest of the proof is therefore identical to the proof of Theorem 5.10. □

Theorem 9.4 (Restated). *Let $\mathcal{N} = \langle V, f, \alpha \rangle$ be a WI-NCG, and assume that a Nash equilibrium exists. Then:*

$$\begin{aligned} PoA(\mathcal{N}) &\leq 3 + \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V} \min\left(\alpha\frac{\bar{d}_{H_h}^+}{\bar{f}_V} + (n-3)\frac{\bar{l}_V}{\bar{f}_V},\ 2\log_2 n - 2 + 2\sqrt{(n-2)\alpha/\bar{f}_V}\right) \\ &= O\left(1 + \min\left(\alpha\frac{\bar{d}_{H_h}^+}{\bar{f}_V} + n\frac{\bar{l}_V}{\bar{f}_V},\ \bar{d}_{H_h}^+ + n\frac{\bar{l}_V}{\alpha},\ \log n + \sqrt{n\alpha/\bar{f}_V},\ \sqrt{n\bar{f}_V/\alpha}\right)\right) = O(\sqrt{n}) \ . \end{aligned}$$

Proof: By Lemma 9.3 and Lemma E.1 we have:

$$\begin{aligned} PoA(\mathcal{N}) &< \frac{n(\alpha(\bar{d}_{H_h}^+ + 1) + \bar{f}_V + (n-2)\bar{l}_V)}{n(\alpha/2 + \bar{f}_V)} \leq 2 + \frac{\alpha\bar{d}_{H_h}^+ + (n-2)\bar{l}_V - \bar{f}_V}{\alpha/2 + \bar{f}_V} \\ &\leq 2 + \frac{\alpha\bar{d}_{H_h}^+ + (n-3)\bar{l}_V}{\alpha/2 + \bar{f}_V} \ , \end{aligned} \tag{E.3}$$

and by Theorem 9.1 we have:

$$\begin{aligned} PoA(\mathcal{N}) &\leq 3 + \frac{\bar{d}_H^+}{\hat{\alpha}/2 + \bar{d}_H^+} \min\left(\hat{\alpha},\ 2\log_2 n - 2 + \min\left(n-2,\ 2\sqrt{(n-2)\hat{\alpha}/\bar{d}_H^+}\right)\right) \\ &= 3 + \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V} \min\left(\alpha\frac{\bar{d}_H^+}{\bar{f}_V},\ 2\log_2 n - 2 + \min\left(n-2,\ 2\sqrt{(n-2)\alpha/\bar{f}_V}\right)\right) \ . \end{aligned} \tag{E.4}$$

To combine the two bounds we first notice that $\bar{d}^+_{H_h} \cdot \frac{\alpha}{n-2} < \bar{f}_V - \bar{l}_V$, since all edges in H_h have friendship weight greater than $\frac{\alpha}{n-2}$. Therefore, $\alpha \bar{d}^+_{H_h} + (n-3)\bar{l}_V < (n-2)\bar{f}_V$, so the $n-2$ case in (E.4) may be removed. We also have $\bar{l}_V \leq (\bar{d}^+_H - \bar{d}^+_{H_h})\frac{\alpha}{n-2}$, since all edges in $E_H - E(H_h)$ have friendship weight at most $\frac{\alpha}{n-2}$. This yields $\alpha \bar{d}^+_{H_h} + (n-3)\bar{l}_V < \alpha \bar{d}^+_H$, so the $\alpha \bar{d}^+_H / \bar{f}_V$ case in (E.4) may also be removed. The combined bound can therefore be written as:

$$
\begin{aligned}
PoA(\mathcal{N}) &\leq 3 + \frac{\bar{f}_V}{\alpha/2 + \bar{f}_V} \min\left(\alpha \frac{\bar{d}^+_{H_h}}{\bar{f}_V} + (n-3)\frac{\bar{l}_V}{\bar{f}_V},\ 2\log_2 n - 2 + 2\sqrt{(n-2)\alpha/\bar{f}_V} \right) \\
&= O(\sqrt{n}) \ . \qquad \square
\end{aligned}
$$

Appendix F

Cage Graphs

This paper uses several results regarding the connection between the girth and average degree of a graph. Specifically, we use adaptations of well-known upper and lower bounds for the number of vertices in cage graphs. A (d, g)-cage graph is a regular graph of degree d and girth g, with a minimal number of vertices. A well known lower bound on the number of vertices of cage graphs is the *Moore bound* (see [BB94], p. 180). The Moore bound states that the number of vertices in a regular graph of degree d and girth g is at least $n_0(d, g)$, where:

$$n_0(d, 2r) = 2\sum_{i=0}^{r-1}(d-1)^i \tag{F.1}$$

$$n_0(d, 2r+1) = 1 + d\sum_{i=0}^{r-1}(d-1)^i = 1 + \frac{d}{2}n_0(d, 2r) \tag{F.2}$$

A similar result was proven by Alon et. al. [AHL02] for irregular graphs as well, substituting d by the average degree $\bar{d}$, and assuming $\bar{d} \geq 2$.

For the upper bound the basic result was provided by Erdös and Sachs [ES63], and it was initially improved by Sauer [Sau67] and Walther [Wal65a, Wal65b]. Margulis [Mar82] provided the first explicit construction of a high-girth graph of high degree. Additional constructions and bound improvements are provided in [Imr84, BH83, Wei84, Mar88, LPS88, LU95, LUW95] (see page 107 in [Bol78] for a history). We use the result from Sauer [Sau67] in this paper. Denoting by $v(d,g)$ the number of vertices of a (d,g)-cage, the bound states that:

$$v(d,g) \leq \begin{cases} \frac{4}{3} + \frac{29}{12}2^{g-2} & \text{for } d = 3 \text{ and } g \text{ odd} \\ \frac{2}{3} + \frac{29}{12}2^{g-2} & \text{for } d = 3 \text{ and } g \text{ even} \\ 2(d-1)^{g-2} & \text{for } d \geq 4 \text{ and } g \text{ odd} \\ 4(d-1)^{g-3} & \text{for } d \geq 4 \text{ and } g \text{ even} \end{cases} . \tag{F.3}$$

We use a simplified (and weaker) version of this bound:

Proposition 6.4 (Restated). *For any integer $g \geq 3$ and $n \geq 2^g$, there is a graph with n vertices and girth g of average degree at least $\frac{1}{2}(n/4)^{1/(g-2)}$.*

Proof: We simplify the bound in inequality F.3 to:

$$v(d,g) \leq 4(d-1)^{g-2} \tag{F.4}$$

which holds for any $g \geq 3$ and $d \geq 3$. Now, given $g \geq 3$ and $n \geq 2^g$ we select d to be:

$$d = \left\lfloor \left(\frac{n}{4}\right)^{\frac{1}{g-2}} \right\rfloor + 1,$$

which ensures that $d \geq 3$ because $n \geq 2^g$. Using the bound (F.4) we know that there is a d-regular graph G_1 of girth g with n_1 vertices where $n_1 \leq 4(d-1)^{g-2} \leq n$. Let $m = \lfloor n/n_1 \rfloor$. Obviously $m \geq 1$. Also, since $m \geq n/n_1 - 1$, we have:

$$mn_1 \geq \frac{1}{2}mn_1 + \frac{1}{2}\left(\frac{n}{n_1} - 1\right)n_1 = \frac{n}{2} + \frac{m-1}{2}n_1 \geq \frac{n}{2} \ .$$

We construct the graph G by taking m copies of G_1 along with $n - mn_1$ single vertices and connecting them all in a line (a single edge connects each component, a copy of G_1 or a single vertex, to the next). G obviously has n vertices, and girth g. It's average degree is:

$$\bar{d}_G \geq \frac{mdn_1}{n} \geq \frac{d}{2} \geq \frac{1}{2}\left(\frac{n}{4}\right)^{\frac{1}{g-2}},$$

so G satisfies the claim. $\square$