CHAPTER 19

Network Formation Games and the Potential Function Method

Éva Tardos and Tom Wexler

Abstract

Large computer networks such as the Internet are built, operated, and used by a large number of diverse and competitive entities. In light of these competing forces, it is surprising how efficient these networks are. An exciting challenge in the area of algorithmic game theory is to understand the success of these networks in game theoretic terms: what principles of interaction lead selfish participants to form such efficient networks?

In this chapter we present a number of network formation games. We focus on simple games that have been analyzed in terms of the efficiency loss that results from selfishness. We also highlight a fundamental technique used in analyzing inefficiency in many games: the potential function method.

19.1 Introduction

The design and operation of many large computer networks, such as the Internet, are carried out by a large number of independent service providers (Autonomous Systems), all of whom seek to selfishly optimize the quality and cost of their own operation. Game theory provides a natural framework for modeling such selfish interests and the networks they generate. These models in turn facilitate a quantitative study of the trade-off between efficiency and stability in network formation. In this chapter, we consider a range of simple network formation games that model distinct ways in which selfish agents might create and evaluate networks. All of the models we present aim to capture two competing issues: players want to minimize the expenses they incur in building a network, but at the same time seek to ensure that this network provides them with a high quality of service.

There are many measures by which players might evaluate the quality of a network. In this chapter, we focus primarily on measures of distance (Section 19.2) and connectivity (Section 19.3), rather than measures based on congestion effects (as is done in Chapter 18). We also assume that players have financial considerations. In Sections 19.2 and 19.3, players seek to minimize the construction costs of the networks they create. In Section 19.4, we look at a game with a more sophisticated financial aspect: players represent service providers who set prices for users and seek to maximize their profit, namely their income from users minus the cost of providing the service.

For all of the games we consider, we use Nash equilibrium as the solution concept, and refer to networks corresponding to these equilibria as being *stable*. The models we focus on involve players who can unilaterally build edges, and thus the Nash equilibrium solution concept is appropriate.

To evaluate the overall quality of a network, we consider the *social cost*, or the sum of all players' costs. We refer to the networks that optimize social cost as *optimal* or *socially efficient*. The main goal of this chapter is to better understand the quantitative trade-off between networks that are stable and those that are socially efficient. More precisely, we are interested in bounding the price of anarchy and the price of stability (as defined in Chapter 17). The models we consider in this chapter are network formation games in which these measures are provably small.

In Section 19.2 we consider a local connection game where the nodes of the graph are players who pay for the edges that connect them directly to other nodes (incident edges). In selecting a strategy, players face two conflicting desires: to pay as little as possible, and to have short paths to all other nodes. Our goal here is to bound the efficiency loss resulting from stability. Such connection games have been extensively studied in the economics literature (see Jackson (2006) for a survey) to model social network formation, using edges to represent social relations. The local connection game can also be thought of as a simple model for the way subnetworks connect in computer networks (by establishing peering points), or as modeling the formation of subnetworks in overlay systems such as P2P (peer-to-peer) networks connecting users to each other for downloading files.

We will use a model in which players can form edges to a neighbor unilaterally, and will use Nash equilibrium as our solution concept. This differs from much of the literature in economics, where it is typically assumed that an edge between two players needs the consent or contribution from both players, and where the notion of pairwise stability is used instead of Nash equilibria. We will discuss how the results in Section 19.2 extend to models using pairwise stable equilibria in the notes in Section 19.5.1.

The model we examine was introduced by Fabrikant et al. (2003) and represents the first quantitative effort to understand the efficiency loss of stable networks. In this game, a single parameter α represents the cost of building any one edge. Each player (represented by a node) perceives the quality of a network as the sum of distances to all other nodes. Players aim to minimize a cost function that combines both network quality and building costs: they attempt to minimize the sum the building costs they incur and the distances to all other players. Thus, players use α as a trade-off parameter between their two objectives. This is perhaps the simplest way to model this type of trade-off. While the simplicity of this game makes it easy to evaluate, such a stylized model ignores a number of issues, such as varying costs and possible congestion effects. In Section 19.5.1, we discuss related models that address some of these issues.

In Section 19.3 we study a very different (and also quite simple) model of network design, introduced by Anshelevich et al. (2004), called the global connection game. Whereas players in the game of Section 19.2 only make local choices (which other nodes to link to), players in this game make global decisions, in that they may build edges throughout the network. Unlike the local connection game, this global game attempts to model players who actually build and maintain large-scale shared networks. This model also allows for greater heterogeneity in the underlying graph.

In the global connection game, a player is not associated with an individual node of the networks, but instead has certain global connectivity goals. To achieve these goals, a player may contribute money to any set of edges in the network. As before, we view connectivity as the primary measure of quality. However, players do not desire uniform connectivity; instead, each player has a subset of nodes that it needs to connect, and aims to do so as cheaply as possible. Furthermore, unlike in the local game, players are not concerned with distance, and simply want to connect their terminals.

As in the previous model, players are sensitive to costs. Edge e has a cost $c_e \geq 0$, and players who use e share this cost. In particular, we focus on a *fair sharing rule*; all players using an edge must share its cost evenly. This natural cost-sharing scheme can be derived from the Shapley value, and has many nice properties. We also examine other cost-sharing games, and discuss the role of fair sharing in the price of stability results.

A key technique used in this section is the potential function method. This method has emerged as a general technique in understanding the quality of equilibria. We review this technique in detail in Section 19.3.2. While this technique provides results only regarding the price of stability, it is interesting to note that many of the currently known price of anarchy results (e.g., most of the results in Part III of this book) are for potential games.

In Section 19.4, we consider another potential game; a facility location game with a more sophisticated cost model. In the previous two sections, players simply minimized their costs. Here, edges still have costs, but players also select prices for users so as to maximize net income: price charged minus the cost paid. We again consider a very simplified model in which players place facilities to serve clients, thereby forming a network between the providers and the clients. We show that a socially efficient network is stable (i.e., the price of stability is 1), and bound the price of anarchy.

In the context of facility location games, we also bound the quality of solutions obtained after sufficiently long selfish play, without assuming that players have yet reached an equilibrium. As we have seen in part I of this book, equilibrium solutions may be hard to find (Chapter 2), and natural game play may not converge to an equilibrium (Chapter 4). Thus it is often useful to evaluate the quality of the transient solutions that arise during competitive play. The facility location game considered in this section is one of the few classes of games for which this strong type of bound is known.

19.2 The Local Connection Game

In this section we consider the simple network formation game of Fabrikant et al. (2003), where players can form links to other players. We consider a game with n players, where each player is identified with a node. Node u may choose to build edges from u to any subset of nodes, thereby creating a network. Players have two competing goals; players want to build (and thus pay) for as few edges as possible, yet they also want to form a network that minimizes the distance from their own node to all others. Our main focus in this section is to quantitatively understand the inefficiency that results from the selfish behavior of these network builders.

19.2.1 Model

Players in the local connection game are identified with nodes in a graph G on which the network is to be built. A strategy for player u is a set of undirected edges that u will build, all of which have u as one endpoint. Given a strategy vector S, the set of edges in the union of all players' strategies forms a network $G(S)$ on the player nodes. Let $dist_S(u, v)$ be the shortest path (in terms of number of edges) between u and v in $G(S)$. We use $dist(u, v)$ when S is clear from context. The cost of building an edge is specified by a single parameter, α. Each player seeks to make the distances to all other nodes small, and to pay as little as possible. More precisely, player u's objective is to minimize the sum of costs and distances $\alpha n_u + \sum_v dist(u, v)$, where n_u is the number of edges bought by player u.

Observe that since edges are undirected, when a node u buys an edge (u, v), that edge is also available for use from v to u, and in particular, is available for node v. Thus, at Nash equilibrium at most one of the nodes u and v pay for the connecting edge (u, v). Also, since the distance $dist(u, v)$ is infinite whenever u and v are not connected, at equilibrium we must have a connected graph. We say that a network $G = (V, E)$ is *stable* for a value α if there is a stable strategy vector S that forms G.

The social cost of a network G is $SC(G) = \sum_{u \neq v} dist(u, v) + \alpha|E|$, the sum of players' costs. Note that the distance $dist(u, v)$ contributes to the overall quality twice (once for u and once for v). We will be comparing solutions that are stable to those that are optimal under this measure.

19.2.2 Characterization of Solutions and the Price of Stability

We now characterize the structure of an optimal solution as a function of α. A network is *optimal* or *efficient* if it minimizes the social cost $SC(G)$.

Lemma 19.1 *If $\alpha \geq 2$ then any star is an optimal solution, and if $\alpha \leq 2$ then the complete graph is an optimal solution.*

PROOF Consider an optimal solution G with m edges. We know $m \geq n - 1$; otherwise, the graph would be disconnected, and thus have an infinite cost. All ordered pairs of nodes not directly connected by an edge must have a distance of at least 2 from each other, and there are $n(n - 1) - 2m$ such pairs. Adding the remaining $2m$ pairs with distance 1 yields $\alpha m + 2n(n - 1) - 4m + 2m = (\alpha - 2)m + 2n(n - 1)$ as a lower bound on the social cost of G. Both a star and the complete graph match this bound. Social cost is minimized by making m as small as possible when $\alpha > 2$ (a star) and as large as possible when $\alpha < 2$ (a complete graph). □

Both the star and the complete graph can also be obtained as a Nash equilibrium for certain values of α, as shown in the following lemma.

Lemma 19.2 *If $\alpha \geq 1$ then any star is a Nash equilibrium, and if $\alpha \leq 1$ then the complete graph is a Nash equilibrium.*

PROOF First suppose $\alpha \geq 1$, and consider a star. It turns out that any assignment of edges to incident players corresponds to a Nash equilibrium, but for this result, we need only demonstrate a single solution. In particular, consider the strategy in which player 1 (the center of the star) buys all edges to the other players, while the remaining $n - 1$ leaf players buy nothing. Player 1 has no incentive to deviate, as doing so disconnects the graph and thus incurs an infinite penalty. Any leaf player can deviate only by adding edges. For any leaf player, adding k edges saves k in distance but costs αk, and thus is not a profitable deviation. Thus the star is a Nash equilibrium.

Now suppose $\alpha \leq 1$. Consider a complete graph, with each edge assigned to an incident player. A player who stops paying for a set of k edges saves αk in cost, but increases total distances by k, so this outcome is stable. □

There are other equilibria as well, some of which are less efficient (see Exercise 19.6). However, these particular Nash equilibria, in conjunction with the above optimal solutions, suffice to upper bound the price of stability.

Theorem 19.3 *If $\alpha \geq 2$ or $\alpha \leq 1$, the price of stability is 1. For $1 < \alpha < 2$, the price of stability is at most 4/3.*

PROOF The statements about $\alpha \leq 1$ and $\alpha \geq 2$ are immediate from Lemmas 19.1 and 19.2. When $1 < \alpha < 2$, the star is a Nash equilibrium, while the optimum structure is a complete graph. To establish the price of stability, we need to compute the ratio of costs of these two solutions. The worst case for this ratio occurs when α approaches 1, where it attains a value of

$$\frac{2n(n-1) - 2(n-1)}{2n(n-1) - n(n-1)/2} = \frac{4n^2 - 6n + 2}{3n^2 - 3n} < 4/3.$$

□

Exercise 19.3 shows that the complete graph is the unique equilibrium for $\alpha < 1$, so we also have that the price of anarchy is 1 in this range. We now address the price of anarchy for larger values of α.

19.2.3 The Price of Anarchy

The first bound on the price of anarchy for this game was given by Fabrikant et al. (2003), and involves two steps: bounding the diameter of the resulting graph, and using the diameter to bound the cost. We begin with the second step.

Lemma 19.4 *If a graph G at Nash equilibrium has diameter d, then its social cost is at most $O(d)$ times the minimum possible cost.*

PROOF The cost of the optimal solution is at least $\Omega(\alpha n + n^2)$, as we need to buy a connected graph, which costs at least $(n-1)\alpha$, and there are $\Omega(n^2)$ distances, each of which is at least 1. To bound the quality of the solution, consider the distance costs and edge costs separately. The distance cost is at most $n^2 d$, and thus is at most d times the minimum possible.

We now examine edge costs. First we consider *cut edges*, those edges whose removal disconnects G. There are at most $n-1$ cut edges, so the total cost of all cut edges is at most $\alpha(n-1)$, which in turn is at most the optimal solution cost. Now consider the set of all noncut edges paid for by a vertex v. We will argue that there are $O(nd/\alpha)$ such edges, with cost $O(dn)$ for node v, and thus the total cost of all noncut edges is $O(dn^2)$. This will establish that the cost of G is $O(\alpha n + dn^2)$, completing the proof.

Pick a node u, and for each edge $e = (u, v)$ paid for by node u, let V_e be the set of nodes w, where the shortest path from u to w goes through edge e. We will argue that the distance between nodes u and v with edge e deleted is at most $2d$. Thus deleting e increases the total distance from u to all other nodes by at most $2d|V_e|$. Since deleting the edge would save α in edge costs and G is stable, we must have that $\alpha \leq 2d|V_e|$, and hence $|V_e| \geq \alpha/2d$. If there are at least $\alpha/2d$ nodes in each V_e, then the number of such edges adjacent to a node v must be at most $2dn/\alpha$, as claimed.

We now bound the distance between nodes u and v with edge e deleted. Consider Figure 19.1, depicting a shortest path avoiding edge e. Let $e' = (u', v')$ be the edge on this path entering the set V_e. The segment P_u of this path from u to node u' is the shortest path from u to u' as $u' \notin V_e$, and hence deleting e does not affect the shortest path. So P_u is at most d long. The segment P_v from v' to v is at most $d-1$ long, as $P_v \cup e$ forms the shortest path between u and v'. Thus the total length is at most $2d$. □

Using this lemma, we can bound the price of anarchy by $O(\sqrt{\alpha})$.

Theorem 19.5 *The diameter of a Nash equilibrium is at most $2\sqrt{\alpha}$, and hence the price of anarchy is at most $O(\sqrt{\alpha})$.*

PROOF From Lemma 19.4, we need only prove that for any nodes u and v, $dist(u, v) < 2\sqrt{\alpha}$. Suppose for nodes u and v, $dist(u, v) \geq 2k$, for some k. The

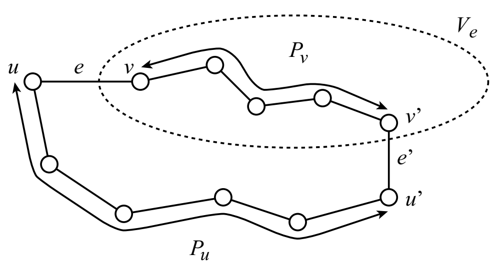

Figure 19.1. Path $P_u, (u', v')P_v$ is the u–v shortest path after edge $e = (u, v)$ is deleted.

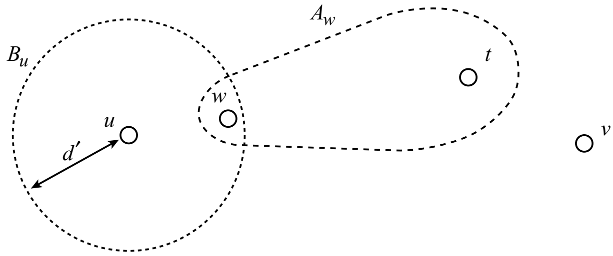

Figure 19.2. Nodes u and v that are at maximum distance d apart. B is the set of nodes at most $d' = (d-1)/4$ away from node u, and A_w is the set of nodes whose shortest path leaves B at w.

main observation is that by adding the edge (u, v), the node u would pay α and improve her distance to the nodes on the second half of the $u - v$ shortest path by $(2k-1) + (2k-3) + \cdots + 1 = k^2$. So if $dist(u, v) > 2\sqrt{\alpha}$, node u would benefit from adding the edge (u, v), a contradiction. □

We now show an $O(1)$ bound on the price of anarchy that was given by Lin (2003) (and independently also by Albers et al., 2006) for $\alpha = O(\sqrt{n})$.

Theorem 19.6 *The price of anarchy is $O(1)$ whenever α is $O(\sqrt{n})$. More generally, price of anarchy is $O(1 + \alpha/\sqrt{n})$.*

PROOF We again use Lemma 19.4, so all we have to do is improve our bound on the diameter d. Consider nodes u and v with $dist(u, v) = d$. Let $d' = \lfloor(d - 1)/4\rfloor$ and let B be the set of nodes at most d' away from u, as shown on Figure 19.2. Consider how the distance $d(v, w)$ changes for nodes $w \in B$ by adding edge (v, u). Before adding the edge $dist(v, w) \geq d - d'$. After adding (v, u), the distance decreases to at most $d' + 1$. Thus v saves at least $(d - 2d' - 1)$ in distance to all nodes in B, and hence would save at least $(d - 2d' - 1)|B| \geq (d-1)|B|/2$ in total distance costs by buying edge (v, u). If G is stable, we must have $(d-1)|B|/2 \leq \alpha$.

For a node $w \in B$ let A_w contain all nodes t for which the u–t shortest path leaves the set B after the node w. Note that if A_w is nonempty, then w must be exactly at distance d' from u. Therefore, node u would save $|A_w|(d'-1)$ in distance cost by buying edge (u, w). If the network is at equilibrium, then we must have that $|A_w|(d'-1) \leq \alpha$. There must be a node $w \in B$ that has $|A_w| \geq (n - |B|)/|B|$. Combining these, we get that

$$(d'-1)(n - |B|)/|B| \leq \alpha.$$

This implies that $|B|(1 + \alpha/(d'-1)) \geq n$, and since $\alpha > d > d'$,

$$|B| \geq n(d'-1)/2\alpha.$$

Combining this with the previous bound of $\alpha \geq (d-1)|B|/2$ yields

$$\alpha \geq (d-1)|B|/2 \geq (d-1)n(d'-1)/4\alpha \geq n(d'-1)^2/\alpha.$$

Thus $\alpha^2 \geq n(d' - 1)^2$ and hence $d \leq 4(d' + 1) + 1 \leq 4\alpha/\sqrt{n} + 9$, which implies the claimed bound by Lemma 19.4. □

19.3 Potential Games and a Global Connection Game

In this section we introduce a broad class of games known as *potential games*. This class encompasses a number of natural and well-studied network-based games. As we will see, potential games possess many nice properties; pure equilibria always exist, best response dynamics are guaranteed to converge, and the price of stability can be bounded using a technique called the *potential function method*. Our motivating example for this class of games is a network formation game called the *global connection game*, which was discussed in Chapter 17. We begin by defining this game, and present some theorems about pure equilibria and the price of stability. We then introduce potential games, and provide generalized results for this broader framework.

The network formation game discussed in Section 19.2 is local in the sense that a player can build links to other nodes, but has no direct means for affecting distant network structure. Such might be the case with social networks or peering relationships in a digital network. The global connection game, in contrast, models players who make global structural decisions; players may build edges throughout the network, and thus consider relatively complex strategies. This game might be more appropriate for modeling the actual construction and maintenance of large-scale physical networks.

Beyond the varying scope of players' strategies, there are two additional features that differentiate these network formation games. First, in exchange for the global connection game's broader strategy space, we consider a relatively simplified player objective function. In particular, we assume that players are unconcerned with their distance to other nodes in the network, and instead want only to build a network that connects their terminals as cheaply as possible. The second notable distinction is that the global connection game supports cooperation, in that multiple players may share the cost of building mutually beneficial links. In the local connection game, an edge might benefit multiple players, and yet the edge's cost is always covered fully by one of the two incident players. We now give a formal description of the global connection game.

19.3.1 A Global Connection Game

We are given a directed graph $G = (V, E)$ with nonnegative edge costs c_e for all edges $e \in E$. There are k players, and each player i has a specified source node s_i and sink node t_i (the same node may be a source or a sink for multiple players). Player i's goal is to build a network in which t_i is reachable from s_i, while paying as little as possible to do so. A strategy for player i is a path P_i from s_i to t_i in G. By choosing P_i, player i is committing to help build all edges along P_i in the final network. Given a strategy for each player, we define the constructed network to be $\cup_i P_i$.

It remains to allocate the cost of each edge in this network to the players using it, as this will allow players to evaluate the utility of each strategy. In principle, there are a vast number of possible cost-sharing mechanisms, each of which induces a distinct network formation game. We will briefly touch on this large space of games at the end of the section, but for now, our primary focus will be on a single cost-sharing mechanism with a number of nice properties, that is both simple and easy to motivate.

In particular, we consider the mechanism that splits the cost of an edge evenly among all players whose path contains it. More concretely, if k_e denotes the number of players whose path contains edge e, then e assigns a cost share of c_e/k_e to each player using e. Thus the total cost incurred by player i under a strategy vector S is given by

$$\text{cost}_i(S) = \sum_{e \in P_i} c_e/k_e.$$

Note that the total cost assigned to all players is exactly the cost of the constructed network. This equal-division mechanism was suggested by Herzog et al. (1997), and has a number of basic economic motivations. Moulin and Shenker prove that this mechanism can be derived from the Shapley (2001) value, and it can be shown to be the unique cost-sharing scheme satisfying a number of natural sets of axioms (see Feigenbaum et al., 2001; Moulin and Shenker, 2001). We refer to it as the *fair* or *Shapley cost-sharing mechanism*. The social objective for this game is simply the cost of the constructed network.

One may view this game as a competitive version of the generalized Steiner tree problem; given a graph and pairs of terminals, find the cheapest possible network connecting all terminal pairs. Indeed, an optimal generalized Steiner tree is precisely the outcome against which we will compare stable solutions in evaluating the efficiency of equilibria. This connection highlights an important difference between this game and routing games; in routing games such as those discussed in Chapter 18, players are sensitive to congestion effects, and thus seek sparsely used paths. But in the global connection game, as with the Steiner forest problem, the objective is simply to minimize costs, and thus sharing edges is in fact encouraged.

The two examples in Chapter 17 provide a few useful observations about this game. Example 17.2 (see Figure 19.3(a)) shows that even on very simple networks, this game has multiple equilibria, and that these equilibria may differ dramatically in quality. There are two equilibria with costs k and 1 respectively. Since the latter is also optimal

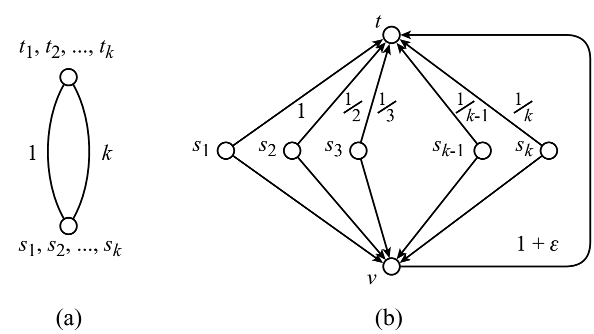

Figure 19.3. An instance of the global connection game with price of anarchy k (a) and an instance with price of stability $\mathcal{H}_k$ (b).

solution, the price of anarchy is k, while the price of stability is 1. It is not hard to show that the price of anarchy can never exceed k on any network (see Exercise 19.9), and thus this simple example captures the worst-case price of anarchy. Our primary goal will be to bound the price of stability in general.

Example 17.3 (see Figure 19.3(b)) shows that the price of stability can indeed exceed 1; this network has a unique Nash equilibrium with cost $\mathcal{H}_k$, the kth harmonic number, while the optimal solution has a cost of $1 + \epsilon$. Thus, the price of stability on this network is roughly $\mathcal{H}_k$. Our aim is to prove that pure equilibria always exist and provide an upper bound the price of stability. Both of these results make use of a *potential function*, which we will formally introduce in Section 19.3.2.

Consider an instance of the global connection game, and a strategy vector $S = (P_1, P_2, \ldots, P_k)$ containing an $s_i - t_i$ path for each player i. For each edge e, define a function $\Psi_e(S)$ mapping strategy vectors to real values as

$$\Psi_e(S) = c_e \cdot \mathcal{H}_{k_e},$$

where k_e is the number of players using edge e in S, and $\mathcal{H}_k = \sum_{j=1}^{k} 1/j$ is the kth harmonic number. Let $\Psi(S) = \sum_e \Psi_e(S)$. While this function does not obviously capture any important feature of our game, it has the following nice property.

Lemma 19.7 *Let $S = (P_1, P_2, \ldots, P_k)$, let $P_i' \neq P_i$ be an alternate path for some player i, and define a new strategy vector $S' = (S_{-i}, P_i')$. Then*

$$\Psi(S) - \Psi(S') = u_i(S') - u_i(S).$$

PROOF This lemma states that when a player i changes strategies, the corresponding change in $\Psi(\cdot)$ exactly mirrors the change in i's utility. Let k_e be the number of players using e under S. For any edge e that appears in both or neither of P_i and P_i', the cost paid by i toward e is the same under S and S'. Likewise, $\Psi_e(\cdot)$ has the same value under S and S'. For an edge e in P_i but not in P_i', by moving from S to S', i saves (and thus increases her utility by) c_e/k_e, which is precisely the decrease in $\Psi_e(\cdot)$. Similarly, for an edge e in P_i' but not in P_i, player i incurs a cost of $c_e/(k_e + 1)$ in switching from S to S', which matches the increase in $\Psi_e(\cdot)$. Since $\Psi(\cdot)$ is simply the sum of $\Psi_e(\cdot)$ over all edges, the collective change in player i's utility is exactly the negation of the change in $\Psi(\cdot)$. □

We also note that $\Psi(S)$ is closely related to $cost(S)$, the cost of the network generated by S. More precisely, consider any edge e used by S. The function $\Psi_e(S)$ is at least c_e (any used edge is selected by at least 1 player), and no more than $\mathcal{H}_k c_e$ (there are only k players). Thus we have

Lemma 19.8 $\text{cost}(S) \leq \Psi(S) \leq \mathcal{H}_k \text{cost}(S)$.

These two lemmas are used to prove the following two theorems, which will follow from Theorems 19.11, 19.12, and 19.13.

Theorem 19.9 *Any instance of the global connection game has a pure Nash equilibrium, and best response dynamics always converges.*

Theorem 19.10 *The price of stability in the global connection game with k players is at most $\mathcal{H}_k$, the* kth *harmonic number.*

Since the proofs of these two results actually apply to a much broader class of games (i.e., potential games), we now introduce these games and prove the corresponding results in this more general context.

19.3.2 Potential Games and Congestion Games

For any finite game, an *exact potential function* Φ is a function that maps every strategy vector S to some real value and satisfies the following condition: If $S = (S_1, S_2, \ldots, S_k)$, $S'_i \neq S_i$ is an alternate strategy for some player i, and $S' = (S_{-i}, S'_i)$, then $\Phi(S) - \Phi(S') = u_i(S') - u_i(S)$. In other words, if the current game state is S, and player i switches from strategy S_i to strategy S'_i, then the resulting savings i incurs exactly matches the decrease in the value of the potential function. Thus Lemma 19.7 simply states that Ψ is an exact potential function for the global connection game.

It is not hard to see that a game has at most one potential function, modulo addition by a constant. A game that does possess an exact potential function is called an *exact potential game*. For the remainder of this chapter, we will drop the word "exact" from these terms (see Exercise 19.13 for an inexact notion of a potential function). A surprising number of interesting games turn out to be potential games, and this structure has a number of strong implications for the existence of and convergence to equilibria.

Theorem 19.11 *Every potential game has at least one pure Nash equilibrium, namely the strategy S that minimizes $\Phi(S)$.*

PROOF Let Φ be a potential function for this game, and let S be a pure strategy vector minimizing $\Phi(S)$. Consider any move by a player i that results in a new strategy vector S'. By assumption, $\Phi(S') \geq \Phi(S)$, and by the definition of a potential function, $u_i(S') - u_i(S) = \Phi(S) - \Phi(S')$. Thus i's utility can not increase from this move, and hence S is stable. □

Going one step further, note that any state S with the property that Φ cannot be decreased by altering any one strategy in S is a Nash equilibrium by the same argument. Furthermore, best response dynamics simulate local search on Φ; improving moves for players decrease the value of the potential function. Together, these observations imply the following result.

Theorem 19.12 *In any finite potential game, best response dynamics always converge to a Nash equilibrium.*

Note that these two results imply Theorem 19.9.

A less abstract characterization of potential games can be found in a class of games called *congestion games* (Rosenthal, 1973). A congestion game has k players and n resources. Player i has a set $\mathcal{S}_i$ of allowable strategies, each of which specifies a subset of resources. Each resource j has a load-dependent cost function $c_j(x)$, indicating the cost incurred by any player i whose chosen strategy includes resource j if there are x such players in total. The total cost charged to player i who chooses a strategy S_i is simply the sum of the costs incurred from each resource in S_i. Thus if the total load on link j is x_j, then i pays $\sum_{j \in S_i} c_j(x_j)$. The Global Connection game is clearly a congestion game; edges are resources, $s_i - t_i$ paths are allowable strategies for player i, and the cost functions are $c_e(x) = c_e/x$.

Rosenthal (1973) proved that any congestion game is a potential game (see Exercise 19.15). Monderer and Shapley (1996) proved the converse; for any potential game, there is a congestion game with the same potential function.

We now present a generic upper bound on the price of stability for an arbitrary potential game.

19.3.3 The Potential Function Method and the Price of Stability

Suppose that we have a potential game G with a potential function $\Phi(S)$ and social cost function $c(S)$. If $\Phi(S)$ and $c(S)$ are similar, then the price of stability must be small. We make this precise in the following theorem.

Theorem 19.13 *Suppose that we have a potential game with potential function Φ, and assume further that for any outcome S, we have*

$$\frac{\text{cost}(S)}{A} \leq \Phi(S) \leq B \cdot \text{cost}(S)$$

for some constants $A, B > 0$. Then the price of stability is at most AB.

PROOF Let S^N be a strategy vector that minimizes $\Phi(S)$. From Theorem 19.11, S^N is a Nash equilibrium. It suffices to show that the actual cost of this solution is not much larger than that of a solution S^* of minimal cost. By assumption, we have that $\frac{\text{cost}(S^N)}{A} \leq \Phi(S^N)$. By the definition of S^N, we have that $\Phi(S^N) \leq \Phi(S^*)$. Finally, the second inequality of our assumption implies that $\Phi(S^*) \leq B \cdot \text{cost}(S^*)$. Stringing these inequalities together yields $\text{cost}(S^N) \leq AB \cdot \text{cost}(S^*)$, as desired. □

Note that this result, taken together with Lemma 19.8, directly implies Theorem 19.10. This technique for bounding the price of stability using a potential function is known as the *potential function method*.

In general, outcomes that minimize the potential function may not be the best Nash equilibrium, and thus this bound is not always tight (see Exercise 19.14). However, in the case of the global connection game, we have seen that the price of stability is at least $\mathcal{H}_k$. Thus, for this class of games, the bound given by the potential function method is the best possible.

Notice that we have essentially already seen the potential function method used in the nonatomic selfish routing game of Chapter 18. For this routing game, all equilibria have the same social value, and hence the price of anarchy and the price of stability are the same. Because of this, Theorem 18.16 is phrased as a statement about the price of anarchy, but we can still view this result as an application the potential function method. In the last section of this chapter, we will see yet another application of this technique for a potential game that models competitive facility location.

We have seen that potential games have pure equilibria, and that the price of stability can be bounded via the potential function method. We now consider the complexity of finding these equilibria in general potential games.

19.3.4 Finding Nash Equilibria in Potential Games

Theorem 19.12 provides an algorithmic means of reaching pure equilibria in potential games. Unfortunately, this theorem makes no claim regarding the rate of this convergence. In some games, best response dynamics always converges quickly, but in many games it does not. In some games, the potential function Φ can be minimized in polynomial time, but in others the minimization problem is NP-hard. To get a better handle on the complexity of finding pure equilibria in potential games, we consider the closely related problem of finding local optima in optimization problems.

The class of *Polynomial Local Search* problems (PLS) was defined by Johnson et al. (1988) as an abstract class of local optimization problems. First, let us define a general *optimization problem* (say a minimization problem) as follows. We have a set of instances I, and for each instance $x \in I$ a set of feasible solutions $F(x)$ and a cost function $c_x(s)$ defined on all $s \in F(x)$. We also have an oracle (or a polynomial-time algorithm) that takes an instance x and a candidate solution s, and checks whether s is a feasible solution ($s \in F(x)$). If it is, the oracle computes the cost of that solution, $c_x(s)$. The optimization problem is to find a solution $s \in F(x)$ with minimum cost $c_x(s)$ for a given instance $x \in I$.

To define a *local optimization problem*, we must also specify a neighborhood $N_x(s) \subset F(x)$ for each instance $x \in I$ and each solution $s \in F(x)$. A solution $s \in F(x)$ is *locally optimal* if $c_x(s) \leq c_x(s')$ for all $s' \in N_x(s)$. The local optimization problem is to find a local optimum $s \in F(x)$ for a given instance $x \in I$. A *local optimization problem is in PLS* if we have an oracle that, for any instance $x \in I$ and solution $s \in F(x)$, decides whether s is locally optimal, and if not, returns $s' \in N_x(s)$ with $c_x(s') < c_x(s)$.

Fabrikant et al. (2004) show that finding a Nash equilibrium in potential games is PLS-complete, assuming that the best response of each player can be found in polynomial time. To see that the problem belongs to PLS, we will say that the neighbors $N_x(s)$ of a strategy vector s are all the strategy vectors s' that can be obtained from s by a single player changing his or her strategy. By definition, a potential function Φ is locally optimal for cost function $c_x(s) = \Phi(s)$ if and only if it is a pure Nash equilibrium, so finding a pure Nash equilibrium is in PLS.

A problem is PLS-complete if it is in PLS and there is a polynomial time reduction from all other problems in PLS such that local optima of the target problem correspond to local optima of the original one. Since the introduction of this class in Johnson et al. (1988), many local search problems have been shown to be PLS-complete, including the weighted versions of satisfiability (Krentel, 1989). The *weighted satisfiability problem* is defined by a formula in conjunctive normal form $C_1 \wedge \ldots \wedge C_n$, with a nonnegative weight w_j for each clause C_j. Solutions s are truth assignments of variables, and the associated cost $c(s)$ is the sum of the weights of the unsatisfied clauses. The neighbors of a truth assignment s are the assignments obtained by flipping a single variable in s.

Here we show via a reduction from this weighted satisfiability problem that finding a pure Nash equilibrium in potential games is PLS complete.

Theorem 19.14 *Finding a pure Nash equilibrium in potential games, where best response can be computed in polynomial time, is PLS complete.*

PROOF We have argued that finding a pure Nash equilibrium in such games is in PLS. To see that the problem is PLS complete, we use a reduction from the weighted satisfiability problem. Consider a weighted satisfiability instance with k variables $x_1, \ldots, x_k$, and n clauses $C_1, \ldots, C_n$ with weight w_j for clause C_j. Our congestion game will have one player for each variable, and one resource for each clause. Player i, associated with variable x_i, has two possible strategies: it can either select the set of resources S_i consisting of all clauses that contain the term x_i, or $\bar{S}_i$, which includes all clauses containing the term $\bar{x}_i$. Selecting S_i corresponds to setting x_i to *false*, while selecting $\bar{S}_i$ corresponds to setting x_i to *true*.

The main observation is that a clause C_j with k_j literals is false if and only if the corresponding element has congestion k_j. Let C_j be a clause with k_j literals and weight w_j. We define the congestion cost of the element j corresponding to the clause C_j as $c_j(\xi) = 0$ if $\xi < k_j$ and $c_j(k_j) = w_j$. For the strategy vector corresponding to the truth assignment s, the potential function has value $\Phi(s) = \sum_j c_j(\xi_j)$, where ξ_j is the number of false literals in C_j. The weight of assignment s is exactly $\Phi(s)$, and thus the equilibria of this game are precisely the local optima of the satisfiability problem. □

19.3.5 Variations on Sharing in the Global Connection Game

We now return to our motivating example, the global connection game. By definition, this game requires that the cost of any built edge be shared equally among all players using that edge. This sharing rule is natural, arguably fair, and as we have seen, implies a number of nice properties. But is this really the best possible sharing rule? Could perhaps another sharing rule induce even better outcomes? We can view this question as a problem of mechanism design, although here we use the term more broadly than in Chapter 9; instead of seeking to elicit "truthful" behavior, we simply want to guarantee that stable outcomes exist and are reasonably efficient.

If we want to design games to induce better outcomes, we must first decide to what extent we will allow ourselves, as mechanism designers, to alter the game. After all, suppose that we define a game in which players receive a large penalty for taking any path that does not conform with a particular optimal solution. Such a game has pure equilibria, and the price of anarchy is trivially 1. But intuitively, this is not a satisfying solution; this game is too restrictive and fails to capture the decentralized spirit of our earlier network formation games. Therefore, our first hurdle is to specify the class of "reasonable" games that are open for consideration.

To this end, Chen et al. (2006) introduce the class of *cost-sharing games*. This class includes the global connection game, as well as similar games with other cost-sharing rules. A cost-sharing game is played on a graph with edge costs and terminals s_i, t_i for each player i. A strategy for player i is an $s_i - t_i$ path. Given a strategy vector S, cost shares are assigned to players on an edge-by-edge basis as specified by a *cost-sharing method* ξ_e for each edge e. In particular, if S_e is the set of players whose path includes e under S, then $\xi_e(i, S_e) \geq 0$ is cost share assigned to i for e. The total cost incurred by player i is the sum of i's cost shares. We require that any cost-sharing method satisfy two basic properties:

- Fairness: For all i, e we have $\xi_e(i, S_e) = 0$ if $i \notin S_e$.
- Budget-balance: For all e we have $\sum_i \xi_e(i, S_e) = c_e$.

A *cost-sharing scheme* specifies a cost-sharing method per edge given a network, a set of players, and a strategy vector. This definition allows cost-sharing schemes to make use of global information, and thus we also consider the special case of *oblivious* cost-sharing schemes, in which cost-sharing methods depend only on c_e and S_e. Note that the Shapley network formation game is an oblivious cost-sharing game, with the cost-sharing method $\xi_e(i, S_e) = c_e/|S_e|$ for $i \in S_e$.

We now return to our question regarding the relative efficiency of the Shapley scheme. In particular, we will show that nonoblivious cost-sharing schemes can provide far better guarantees than the Shapley scheme.

Theorem 19.15 *For any undirected network in which all players seek to reach a common sink, there is a nonoblivious cost-sharing scheme for which the price of anarchy is at most 2.*

PROOF We define a nonoblivious cost-sharing scheme for which players at equilibrium may be viewed as having simulated Prim's MST heuristic for approximating a min cost Steiner tree. Since this heuristic is 2-approximation algorithm, such a scheme suffices. More concretely, if t is the common sink, we order players as follows. Let player 1 be a player whose source s_1 is closest to t, let player 2 be a player whose source s_2 is closest to $\{t, s_1\}$, and so on. Define a cost-sharing method that assigns the full cost of e to the player in S_e with the smallest index. Since player 1 pays fully for her path regardless of the other players' choices, at equilibrium player 1 must choose a shortest path from s_1 to t, and inductively, the remaining players effectively simulate Prim's algorithm as well. □

On the other hand, if we restrict our attention to oblivious schemes, Chen, Roughgarden, and Valiant prove that for general networks, we cannot do better than the Shapley cost-sharing scheme in the worst case. More precisely, they argue that any oblivious cost-sharing scheme either fails to guarantee the existence of pure equilibria or has a price of stability that is at least $\mathcal{H}_k$ for some game. Thus we have an answer to our original question; while there may be nonoblivious schemes that perform better than Shapley cost-sharing, no oblivious scheme offers a smaller price of stability in the worst case. See the notes on this chapter (Section 19.5.2) for a brief discussion of research concerning other cost-sharing approaches.

19.4 Facility Location

In the models we have considered so far, players construct networks so as to achieve certain connectivity-based goals. Intuitively, these goals are meant to capture players' desires to provide service for some implicit population of network users. Given this perspective, we might then ask what happens when we instead view players as financially motivated agents; after all, service providers are primarily concerned with maximizing profits, and only maintain networks for this purpose. This suggests a model in which players not only build networks but also charge for usage, while network users spur competition by seeking the cheapest service available.

We will consider here a pricing game introduced by Vetta (2002) that is based on the facility location problem. In the facility location problem, we want to locate k facilities, such as Web servers or warehouses, so as to serve a set of clients profitably. Our focus here will be to understand the effect of selfish pricing on the overall efficiency of the networks that players form.

We first present Vetta's competitive facility location problem, in which players place facilities so as to maximize their own profit. We then show that this facility location game is a potential game, and prove that the price of anarchy for an even broader class of games is small.

19.4.1 The Model

Suppose that we have a set of users that need a service, and k service providers. We assume that each service provider i has a set of possible locations A_i where he can locate his facility.

Define $A = \cup_i A_i$ to be the set of all possible facility locations. For each location $s_i \in A_i$ there is an associated cost c_{js_i} for serving customer j from location s_i. We can think of these costs as associated with edges of a bipartite graph that has all users on one side and all of A on the other, as shown on Figure 19.4. A strategy vector

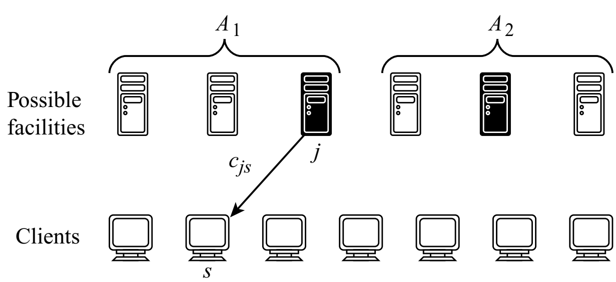

Figure 19.4. The bipartite graph of possible locations and clients. Selected facilities are marked in black.

$s = \{s_1, \ldots, s_k\}$ can be thought of as inducing a subgraph of this graph consisting of the customers and the selected location nodes (marked as black on Figure 19.4).

Our goal is to maximize social welfare, rather than simply minimizing the cost of the constructed network. We assume that customer j has a value π_j for service, and gathers $\pi_j - p$ benefit by receiving service at a price $p < \pi_j$. Locating a facility s_i is free, but that service provider i must pay c_{js_i} to serve client j from location s_i. Doing so generates a profit of $p - c_{js_i}$. If provider i services customer j from location s_i, then this arrangement creates a social value (or surplus) of $\pi_j - c_{js_i}$, the value π_j of service minus the cost c_{js_i} at which the service is provided. Note that this social surplus is independent of the price $p = p_{ij}$ charged; varying p_{ij} simply redistributes welfare between the customer and the provider. We define the *social welfare* $V(S)$ to be the total social value over all providers and customers.

To simplify notation, we assume that $\pi_j \geq c_{js_i}$ for all j, i, and $s_i \in A_i$. To see that this requires no loss of generality, note that decreasing c_{js_i} to be at most π_j does not change the value of any assignment: when $\pi_j < c_{js_i}$ customer j cannot be served from location s_i, while $\pi_j = c_{js_i}$ allows us to serve customer j from location s_i at cost. In either case, the assignment of serving client j from facility s_i results in 0 social value.

To complete the game, we must specify how prices are set and assignments are determined. Given a strategy vector s, we assume that each customer is assigned to a facility that can serve for the lowest cost. The price p_{ij} charged to a customer j using player i's facility s_i is the cost of the second cheapest connection available to j, i.e., $\min_{i' \neq i} c_{js'_i}$. Intuitively, this is the highest price i could expect to get away with charging j; charging any more would give some player i' an incentive to undercut i.

Indeed, we can construct an equivalent interpretation of this game in which prices are selected strategically. Consider a three-stage game where both providers and customers are strategic agents. In the first stage, providers select facility locations. In the second stage, providers set prices for users. And, in the last stage, users select a provider for service, and pay the specified price.

As we saw in Chapter 1, subgame perfect equilibrium is a natural solution concept for multistage games. We will use here a further refinement of this concept, the *trembling hand perfect equilibrium* for extensive form games (see Mas-Colell et al., 1995). Assume that with probability $\epsilon > 0$, each player picks a strategy chosen uniformly at random, and chooses a best strategy with the remaining $(1 - \epsilon)$ probability. We use the notion of subgame perfect equilibrium for this ϵ-perturbed game. A *trembling hand perfect equilibrium* is an equilibrium that can be reached as the limit of equilibria in the ϵ-perturbed game as ϵ approaches 0. This stronger notion of stability is required to prevent providers from offering unprofitably low prices and thereby forcing other providers to artificially lower their own prices.

19.4.2 Facility Location as a Potential Game

We start by proving that the facility location game is a potential game.

Theorem 19.16 *The facility location game is a potential game with social value $V(s)$ as the potential function.*

PROOF We need to argue that if a provider i changes her selected location, then the change in social welfare $V(s)$ is exactly the change in the provider's welfare. To show this, we imagine provider i choosing to "drop out of the game" and show that the change in social welfare $V(s)$ is exactly i's profit.

If provider i "drops out," each client j that was served by provider i switches over to his second best choice. Recall that p_{ij} is exactly the cost of this choice. Thus the client will be served at cost p_{ij} rather than c_{js_i}, so the increase in cost is $p_{ij} - c_{js_i}$, exactly the profit provider i gathers from j.

To prove the statement about provider i changing his strategy, we can think of the change in two steps: first the provider leaves the game, and then reenters with a different strategy. The change in social welfare is the difference between the profit of provider i in the two strategies. □

Corollary 19.17 *There exists a pure strategy equilibrium, and furthermore, all efficient outcomes of the facility location game are stable. Thus, the price of stability is 1. Finally, best response dynamics converge to an equilibrium, but this equilibrium may not be socially optimal.*

Our next goal is to prove that the price of anarchy for this facility location game is small. However, it turns out that the proof applies to a much broader class of games, which we present now.

19.4.3 Utility Games

Vetta (2002) introduced the facility location game as one example of a large class of games called *utility games*. In a utility game, each player i has a set of available strategies A_i, which we will think of as locations, and we define $A = \cup_i A_i$. A social welfare function $V(S)$ is defined for all $S \subseteq A$. Observe that welfare is purely a function of the selected locations, as is the case with the facility location game. In defining the socially optimum set, we will consider only sets that contain one location from each strategy set A_i. However, various structural properties of the function $V(S)$ will be assumed for all $S \subseteq A$. For a strategy vector s, we continue to use $V(s)$ as before, and let $\alpha_i(s)$ denote the welfare of player i. A game defined in this manner is said to be a utility game if it satisfies the following three properties.

- **(i)** $V(S)$ is *submodular*: for any sets $S \subset S' \subset A$ and any element $s \in A$, we have $V(S + s) - V(S) \geq V(S' + s) - V(S')$. In the context of the facility location game, this states that the marginal benefit to social welfare of adding a new facility diminishes as more facilities are added.
- **(ii)** The total value for the players is less than or equal to the total social value: $\sum \alpha_i(s) \leq V(s)$.
- **(iii)** The value for a player is at least his added value for the society: $\alpha_i(s) \geq V(s) - V(s - s_i)$.

A utility game is *basic* if property (iii) is satisfied with equality, and *monotone* if for all $S \subseteq S' \subseteq A$, $V(S) \leq V(S')$.

To view the facility location game as a utility game, we consider only the providers as players. We note that the social welfare $V(S) = \sum_j (\pi_j - \min_{a \in S} c_{ja})$ is indeed purely a function of the selected locations.

Theorem 19.18 *The facility location problem is a monotone basic utility game.*

PROOF Property (ii) is satisfied essentially by definition, and we used the equality of property (iii) property in proving Theorem 19.16. To show property (i), notice that adding a new facility decreases the cost of serving some of the clients. The magnitude of this decrease can only become smaller if the clients are already choosing from a richer set of facilities. Finally, adding a facility cannot cause the cost of serving a client to increase, and thus the facility location game is monotone. □

19.4.4 The Price of Anarchy for Utility Games

Since the facility location game is a potential game with the social welfare as the potential function, the price of stability is 1. In fact, this applies for any basic utility game (any utility game with $\alpha_i(s) = V(s) - V(s - s_i)$ for all strategy vectors s and players i). Unfortunately, the increased generality of utility games comes at a cost; these games are not necessarily potential games, and indeed, pure equilibria do not always exist. However, we now show that for monotone utility games that do possess pure equilibria (such as the facility location game), the price of anarchy is at most 2.

Theorem 19.19 *For all monotone utility games the social welfare of any pure Nash equilibrium is at least half the maximum possible social welfare.*

PROOF Let S be the set of facilities selected at an equilibrium, and O be the set of facilities in a socially optimal outcome. We first note that $V(O) \le V(S \cup O)$ by monotonicity. Let O^i denote the strategies selected by the first i players in the socially optimal solution. That is, $O^0 = \emptyset$, $O^1 = \{o_1\}, \ldots, O^k = O$. Now

$$V(O) - V(S) \le V(S \cup O) - V(S) = \sum_{i=0}^{n} [V(S \cup O^i) - V(S \cup O^{i-1})].$$

By submodularity (property (i))

$$V(S \cup O^i) - V(S \cup O^{i-1}) \le V(S + o_i - s_i) - V(S - s_i)$$

for all i. Using property (iii), we can further bound this by $\alpha_i(S + o_i - s_i)$. Since S is an equilibrium, $\alpha_i(S + o_i - s_i) \le \alpha_i(S)$. Together these yield

$$V(O) - V(S) \le V(O \cup S) - V(S) \le \sum_i \alpha_i(S).$$

Finally, property (ii) implies that $\sum_i \alpha_i(S) \le V(S)$, so $V(O) \le 2V(S)$, and hence the price of anarchy is at most 2. □

19.4.5 Bounding Solution Quality without Reaching an Equilibrium

For any monotone basic utility game, one can also bound the quality of the solution without assuming that players reach an equilibrium, as was shown in a sequence of two papers by Mirrokni and Vetta (2004) and Goemans et al. (2005).

Theorem 19.20 *Consider an arbitrary solution in a monotone basic utility game. Suppose that at each time step, we select a player at random and make a best response move for that player. For any constant $\epsilon > 0$ the expected social value of the solution after $O(n)$ such moves is at least $1/2 - \epsilon$ times the maximum possible social value.*[1]

PROOF Let S be a state, and O be an socially optimal strategy vector. We will prove that the expected increase in social welfare in one step is at least $\frac{1}{n}(V(O) - 2V(S))$, which implies the claimed bound after $O(n)$ steps.

Let β_i be the maximum possible increase in the value for player i. Thus the expected increase in value is $\frac{1}{n}\sum_i \beta_i$. Selecting strategy o_i is an available move, so $\beta_i \geq \alpha_i(S - s_i + o_i) - \alpha_i(S)$, and by basicness, $\beta_i \geq V(S - s_i + o_i) - V(S - s_i) - \alpha_i(S)$.

The rest of the proof mirrors the price of anarchy proof above. We have

$$V(O) - V(S) \leq \sum_{i=0}^{n} [V(S - s_i + o_i) - V(S - s_i)]$$

as before. We bound $V(S + o_i - s_i) - V(S - s_i) \leq \alpha_i(S) + \beta_i$. Using this with property (ii) yields

$$V(O) - V(S) \leq \sum_i (\alpha_i(S) + \beta_i) \leq V(S) + \sum_i \beta_i.$$

Thus $\sum_i \beta_i \geq V(O) - 2V(S)$, and the expected increase in $V(S)$ is $\frac{1}{n}(V(O) - 2V(S))$. The difference $V(O) - 2V(S)$ is expected to decrease by a factor of $(1 - \frac{2}{n})$ each step. After $n/2$ steps, the difference is expected to decrease by a factor of e, and after $\log(\epsilon^{-1})n$ steps shrinks to an ϵ factor. □

19.5 Notes

19.5.1 Local Connection Game

Network formation games have a long history in the social sciences, starting with the work of Myerson (1977, 1991). A standard example of such games can be found in Jackson and Wolinsky (1996) (see Jackson (2006) for a more comprehensive survey). These network formation games are often used to model the creation of social networks, and aim to capture pairwise relations between individuals who may locally form direct links to one another. In other contexts, these games might model peering relations

[1] The constant in the $O(.)$ notation depends on $\log \epsilon^{-1}$.

between pairs of Autonomous Systems (Johari et al., 2006; Fabrikant et al., 2003), or bilateral contracts between systems as with P2P overlay networks (Chun et al., 2004; Christin et al., 2004). Most network formation games in the economics literature use a bilateral process, in which edges only form between two agents with the consent of both parties, unlike the unilateral process of Section 19.2.

Jackson and Wolinsky (1996) examine the trade-off between efficient and stable networks by studying various network formation games and specifying the conditions under which some (or all) stable outcomes are socially efficient, as done in Section 19.2.2. Section 19.2.3 explores *how* efficient nonoptimal stable outcomes may be.

Corbo and Parkes (2005) study a bilateral variant of local connection game. In the *bilateral network formation game*, two nodes must pay the α cost to form a connecting edge. Thus edges represent bilateral agreements in which players agree to evenly share the edge cost (which is effectively 2α). This contrasts with the unilateral edge formation used in the local connection game. Otherwise, the games are the same; players have the same strategy sets, and evaluate the resulting network in the same manner.

Nash equilibria do not appear to be well-suited for modeling bilateral agreements; for a graph to be stable, we need only ensure that no node wants to drop edges, since a player cannot singlehandedly add an edge. For example, the empty graph is always a Nash equilibrium in the bilateral game, and hence the price of anarchy is very high.

Jackson and Wolinsky (1996) suggest using the notion of *pairwise stable equilibrium*; no user u wants to drop any adjacent edge $e = (u, v)$, and no pair of users u and v wants to add the connecting edge (u, v). This stability concept is closely related to a variant of Nash equilibrium in which we allow coalitions of two players to deviate together (u and v may drop any subset of edges adjacent to them, and possibly add the edge (u, v) connecting them, if this is beneficial to both players). This is the solution concept used in the stable matching problem (see Chapter 10), where the natural deviation for a matching that is not stable is by a "blocking pair": a man and a woman who prefer each other to their current partners.

The optimal network structure is the same as in the unilateral game with edge cost 2α. The proof of Theorem 19.1 can be modified to show that when $\alpha \geq 1$, the star is pairwise stable, and when $\alpha \leq 1$ the complete graph is pairwise stable. Note that in both cases, these networks are also efficient, so the price of stability is 1. One can also extend the bounds of Lemma 19.4 and Theorems 19.5 and 19.6 to bound the quality of a worst pairwise stable equilibrium (see Exercise 19.8).

Andelman et al. (2007) consider the effect of coalitions in the unilateral game. Recall from Chapter 1 that a strong Nash equilibrium is one where no coalition has a joint deviation that is profitable for all members. Andelman et al. show that when $\alpha \in (1, 2)$, there is no stable network resisting deviations by coalitions of size 3, and also that when $\alpha \geq 2$, all strong Nash equilibria have cost at most twice the optimum, i.e., the strong price of anarchy is at most 2.

There are many other natural and relevant variations to the discussed network formation games. One important aspect of the model suggested by Jackson (2006) and Jackson and Wolinsky (1996), is that nodes are not required to reach all other nodes in the network. Instead, node u has a value w_{uv} for connecting to another node v, and this benefit decays exponentially with distance. In this game, pairwise stable equilibria may not be connected.

Chun et al. (2004) introduce a variant of the unilateral network formation game to model overlay networks. They allow the cost incurred by node u for adding a directed edge (u, v) to depend upon v and the degree of v, thereby modeling some congestion effects. The authors also extended the notion of distance beyond hop-count, and consider restricting the set of possible connections available to each player. Using Nash equilibria as their solution concept, they study the quantitative trade-offs between cost, node degree, and path length in an experimental setting. Christin et al. (2004) also use these models, and argue that using approximate (rather than exact) equilibria can improve the predictive power of the model and accommodate small errors in modeling and decision making.

Johari et al. (2006) introduced a related game for modeling bilateral contracts between systems. In this game players form directed edges to carry traffic, and the payments along the links are negotiated, in that players can make offers and demands. Anshelevich et al. (2006) propose a variant of this model with fixed routing that includes both directed links and symmetric peering links, and show that in this model, there exists an efficient solution that is approximately stable in some sense.

Open Problems

We have given a bound of $O(\sqrt{\alpha})$ for the price of anarchy of the local connection game, and improved this bound to $O(1)$ for small α. Also, Albers et al. (2006) proved an $O(1)$ bound for the case $\alpha > 12n \log n$ (see also Exercise 19.7 for the case $\alpha > n^2$). It is an open problem whether a constant bound holds for all values of α.

The local connection game is an extremely simple model of network formation. It would be valuable to understand to what extent the price of anarchy bounds apply to broader classes of games. Albers et al. (2006) extend the price of anarchy bound to games where traffic (which affects distance costs multiplicatively) is not uniform, but edge costs remain uniformly α as before. Unfortunately, these bounds depend on the traffic weights, and are only $O(1)$ when these weights are relatively small ($n^2 w_{\max} \leq \alpha$). Is there a natural characterization of all traffic patterns that support a constant price of anarchy? And, can the price of anarchy results extend to models where edge costs vary over the network?

As mentioned, Christin et al. (2004) argue that approximate equilibria are better models of natural network formation. Can we extend our price of anarchy bounds to approximate equilibria?

So far we have been concerned with the quality of equilibria, and did not consider the network formation process. Does "natural" game play of these local connection games converge to an equilibrium efficiently? Bala and Goyal (2000) show that in their model, game play does converge to an equilibrium in some cases. Is this also true in broader class of games? In cases when natural game play does not converge, or only converges slowly, can one bound the quality of the solution after a "long enough" game play, as we have seen in Section 19.4.5?

The network formation process of Bala and Goyal (2000) is very uniform, and leads to networks with extremely simple structure (such as a cycle, star or wheel).

Newman (2003) and Jackson and Rogers (2006) introduce more complex network-formation process based on a random graph generation process that results in graphs that have a number of real-world network properties, such as the power-law degree distribution, clustering, etc. Unfortunately, this process is exogenous, and not really based on personal incentives. One exciting open challenge is to develop an incentive-based and endogenous model of network formation that generates more heterogenous and realistic networks.

19.5.2 Potential Games and a Global Connection Game

The global connection game is related to a large body of work on cost-sharing (see Feigenbaum et al., 2001; Herzog et al., 1997; and the references therein). Much of this work is not game-theoretic; the network is typically assumed to be fixed, and the goal is to compute cost shares with certain properties. Chapter 15 considers cost sharing in a game-theoretic context by assuming the existence of a central authority who must compute cost shares for nodes, each of which has a private utility for inclusion in the network. Thus, the focus is on developing a cost sharing mechanism that induces nodes to reveal their true valuations.

Our general results for potential games suggest some natural extentions to the global connection game. For example, if we consider the global connection game played on undirected networks, then $\Psi(S)$ is still a potential function. Thus we again have that pure equilibria exist and the price of stability is at most $\mathcal{H}_k$. We can also generalize the global connection game by allowing players to have more than two terminals they wish to connect. In such a game, players would select trees spanning their terminals rather than paths. Again, it is easily verified that $\Psi(S)$ is a potential function, so the same results apply. Furthermore, we assumed that the cost of each edge c_e is independent of the number of users. Consider the case when the cost $c_e(k_e)$ of the edge e depends on the number of players (k_e) that use the edge e. The same analysis also extends to this version, assuming the function $c_e(k_e)$ is concave, that is, the cost exhibits an "economy of scale" property; adding a new user is cheaper when a larger population is using the edge.

Anshelevich et al. (2003) consider an unrestricted variant of the global connection game. In this game, players select not only a path but also cost shares for each edge on that path. If the combined shares for an edge cover its cost, that edge is built. Players are assumed to be unhappy if their path is not fully built, and otherwise aim to minimize their cost shares. This game does not necessarily have pure equilibria, and even when it does, even the price of stability may be $O(k)$. However, in the special case of single source games (all players seek connection to a common terminal), the price of stability is shown to be 1.

Chen and Roughgarden (2006) study a weighted generalization of the global connection game, in which each player has a weight, and costs shares are assigned in proportion to these weights. This turns out not to be a potential game, and further, the authors provide an instance in which no pure equilibrium exists. This paper focuses on finding outcomes that are both approximate equilibria and close to optimal. Another similar weighted game is presented by Libman and Orda (2001), with a different mechanism for distributing costs among users. They do not consider the quality of equilibria, but instead study convergence in parallel networks.

Milchtaich (1996) considers a generalization of congestion games in which each player has her own payoff function. Equilibria are shown to exist in some class of these games even though a potential function may not.

Open Problems

Recall the network shown in Figure 19.3(b), which shows that the price of stability may $\mathcal{H}_k$ for the global connection game. Note, however, that if the edges are undirected, then the price of stability falls to 1. The actual worst-case price of anarchy for undirected graphs remains an open question.

There are a wide variety of cost-sharing schemes, as defined by Chen and Roughgarden (2006), that might be relevant either for practical reasons (such as being more fair), or because they induce better outcomes for certain specific classes of networks. Many such schemes, including weighted fair sharing, do not yield exact potential games. For a large number of these cost-sharing games, the price of anarchy, the price of stability, and even the existence of pure equilibria remain unresolved.

More generally, the class of games we consider aims to model situations where users are building a global shared network and care about global properties of the network they build. Our focus was on requiring connectivity (of a terminal set) and aiming to minimize cost. More generally, it would be valuable to understand which type of utility measures yield games with good price of stability properties. For example, we might consider users who are allowed to leave some terminals unconnected, or who care about other properties of the resulting network, such as distances, congestion, etc.

Potential functions are an important tool in understanding the price of anarchy and stability in games. A recent survey of Roughgarden (2006) shows that one can also understand the price of anarchy analysis of resource allocation problems (see Chapter 21) via the potential function method. Surprisingly, many of the price of anarchy and stability results known to date are for potential games (and their weighted variants); the routing games of Chapter 18, the facility location game of Section 19.4, and the load balancing problems of Chapter 20. In a number of these cases, the analysis of the price of anarchy or stability uses alternative techniques to derive stronger bounds than could have been obtained using the potential function method (e.g., bounding the price of anarchy with multiple equilibria, or analyzing weighted variants of these games). However, one wonders if potential functions still play a role here that we do not fully understand.

19.5.3 Facility Location Game

There is a large body of literature dedicated to understanding the effects of pricing in games. Much of this work focuses on establishing the existence of equilibria, and considering qualitative properties of equilibria (such as whether improved service leads to improved profit, or if selfish pricing leads to socially efficient outcomes).

Our focus with the facility location game is to understand the effect of selfish pricing on the overall efficiency of a network. In many settings, selfish pricing leads to a significant reduction in social welfare, and may also yield models with no pure equilibria. An example of this issue is the pricing game of Example 8 in Chapter 1.

See also Chapter 22 for a discussion of these issues in the context of communication networks.

Our price of anarchy bound requires that social welfare be monotone in the set of facilities selected. It is natural to try to extend this game to a scenario in which facilities cost money: in addition to paying the service cost c_{js_i} for servicing a client j from a facility s_i, the provider also pays an installation cost $f(s_i)$ for building at s_i. Unfortunately, there is no constant bound for the price of anarchy for this case. See Exercise 19.17, which observes that when investment costs are large, noncooperative players do not always make the right investments, and thus equilibria may be far from optimal.

Utility games defined in Section 19.4.3 have a wide range of applications, including routing (Vetta, 2002) (see Exercise 19.18), and a market sharing game introduced by Goemans et al. (2006) in the context of content distribution in ad-hoc networks (see Exercises 19.16 for a special case).

In Section 19.4 we bounded the price of anarchy only for pure equilibria. Recall, however, that general utility games may not have pure equilibria. Theorem 19.19 bounding the quality of equilibria also holds for mixed equilibria (Vetta, 2002) and thus is applicable in a much broader context.

Section 19.4.5 showed that in basic utility games, we can bound the quality of solutions without reaching an equilibrium. Such bounds would be even more valuable for general utility games, as these might not have any pure equilibria. Goemans et al. (2005) provide such bounds for a few other games, including some routing games. Unfortunately, the quality of a solution in a general utility game can be very low even after infinitely long game play, as shown by Mirrokni and Vetta (2004).

Open Problems

Many pricing games fail to have Nash equilibria (Example 8 from Chapter 1) and others have equilibria with very low social value (high price of anarchy and stability). The facility location games give a class of examples where pure Nash equilibria exist, and the price of anarchy is small. It would be great to understand which other classes of pricing games share these features.

It will also be extremely important to understand which other classes of games admit good-quality bounds after limited game play, as shown in Section 19.4.5 for facility location games.

Acknowledgments

We thank Ramesh Johari, Tim Roughgarden, Elliot Anshelevich, and Vahab Mirrokni for their valuable comments and suggestions on an early draft of this chapter.

Bibliography

A. Albers, S. Eilts, E. Even-Dar, Y. Mansour, and L. Roditty. On Nash equilibria for a network creation game. In *Proc. ACM-SIAM Symp. Discrete Algorithms*, pp. 89–98, 2006.

N. Andelman, M. Feldman, and Y. Mansour. Strong Price of Anarchy. In *Proc. ACM-SIAM Symp. on Discrete Algorithms*, 2007.

E. Anshelevich, A. Dasgupta, J. Kleinberg, Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocatio. In *Proc. IEEE Symp. on Fdns. of Computer Science*, pp. 295–304, 2004.

E. Anshelevich, A. Dasgupta, E. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proc. ACM Symp. Theory of Computing*, pp. 511–520, 2003.

E. Anshelevich, B. Shepherd, and G. Wilfong. Strategic network formation through peering and service agreements. In *Proc. IEEE Symp. on Fdns. of Computer Science*, 2006.

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

H.-L. Chen and T. Roughgarden. Network design with weighted players. In *Proc. ACM Symp. on Parallel Algorithms and Architecture*, pp. 29–38, 2006.

H.-L. Chen, T. Roughgarden, and G. Valiant. *Designing Networks with Good Equilibria*, unpublished manuscript, 2006.

N. Christin, J. Grossklags, and J. Chuang. Near rationality and competitive equilibria in networked systems. In *ACM SIGCOMM'04 Workshop on Practice and Theory of Incentives in Networked Systems (PINS)*, August 2004.

B.-G. Chun, R. Fonseca, I. Stoica, and J. Kubiatowicz. Characterizing selfishly constructed overlay networks. In *Proc. IEEE INFOCOM'04*, Hong Kong, March 2004.

J. Corbo and D. Parkes. The price of selfish behavior in bilateral network formation. In *Proc. ACM Symp. Princ. of Distributed Systems*, pp. 99–107, 2005.

A. Fabrikant, A. Luthra, E. Maneva, C. Papadimitriou, and C. Shenker, On a network creation game. In *Proc. ACM Symp. Princ. of Distributed Systems*, pp. 247–351, 2003.

A. Fabrikant, C. Papadimitriou, and K. Talwar. The complexity of pure nash equilibria. In *Proc. ACM Symp. of Theory of Computing*, pp. 604–612, 2004.

J. Feigenbaum, C. Papadimitriou, and S. Shenker. Sharing the cost of multicast transmissions. *J. Comp. and Syst. Sci.*, 63:21–41, 2001.

M. Goemans, L. Li, V. Mirrokni, and M. Thottan. Market sharing games applied to content distribution in ad-hoc networks. In *Proc. ACM Intl. Symp. on Mobile Ad Hoc Networking and Computing*, pp. 1020–1033, 2004.

M. Goemans, V. S. Mirrokni, and S. Vetta. Sink Equilibria and Convergence. In *Proc. IEEE Symp. Fdns. of Comp. Sci.*, pp. 142–154, 2005.

S. Herzog, S. Shenker, and D. Estrin. Sharing the "Cost" of multicast trees: An axiomatic analysis. *IEEE/ACM Trans. on Networking*, Dec. 1997.

M. Jackson. A survey of models of network formation: Stability and efficiency. In G. Demange and M. Wooders (eds), *Group Formation Economics: Networks, Clubs and Coalitions*. Cambridge University Press, Cambridge, UK, in press.

M. Jackson and B. Rogers. Meeting strangers and friends of friends: How random are social networks. In press. *Amer. Econ. Rev.*, 2006.

M. Jackson and A. Wolinsky. A strategic model of social and economic networks. *J. Econ. Theory*, 71(1):44–74, 1996.

R. Johari, S. Mannor, and J. N. Tsitsiklis. A contract-based directed network formation. *Games Econ. Behav.*, 56:201–224, 2006.

D.S. Johnson, C.H. Papadimitriou, and M. Yannakakis. How easy is local search? *J. Computer System Sci.*, 37:79–100, 1988.

M.W. Krentel. Structure in locally optimal solutions. In *Proc. IEEE Fdns. of Comp. Sci.*, pp. 216–221, 1989.

L. Libman and A. Orda. Atomic resource sharing in noncooperative networks. *Telecommunication Systems*, 17:4, 385–409, 2001.

H. Lin. *On the Price of Anarchy of a Network Creation Game*. Unpublished Manuscript December 2003.

A. Mas-Colell, M.D. Whinston and J. R.Green *Microeconomic Theory*. Oxford University Press, 1995.

I. Milchtaich. Congestion games with player-specific payoff functions. *Games Econ. Behav.*, 111–124, 1996.

V.S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proc. APPROX 2004*.

V.S. Mirrokni and S. Vetta. Convergence issues in competitive games. In *International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, Springer, 2004.

D. Monderer and L. Shapley. Potential games. *Games Econ. Behav.*, 14:124–143, 1996.

H. Moulin and S. Shenker. Strategyproof sharing of submodular costs: Budget balance versus efficiency. *Econ. Theory*, 18:511–533, 2001.

R.B. Myerson. Graphs and cooperative games. *Math. Operat. Res.*, 2(3), 1977.

R.B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

M. Newman. The structure and function of complex networks. *SIAM Review*, 45:167–256, 2003.

R.W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *Intl. J. Game Theory*, 2:65–67, 1973.

T. Roughgarden. Potential functions and the inefficiency of equilibria. In *Proc. Intl. Congress of Mathematicians (ICM)*, 2006.

A. Vetta. Nash equilibria in competitive societies, with applications to facility location, traffic routing and auctions. In *Proc. IEEE Symp. on Fdns. of Computer Science*, pp. 416–425, 2002.

Exercises

19.1 Consider the local connection game from Section 19.2. In Lemma 19.1, we saw that the star is an optimal solution for $\alpha \geq 2$, and the complete graph is an optimal solution for $\alpha \leq 2$. Prove that if $\alpha \neq 2$, then these are in fact the only optimal solutions.

19.2 Give a complete characterization of all optimal networks for $\alpha = 2$.

19.3 Show that when $\alpha < 1$ the complete graph is the only equilibrium.

19.4 Show that a sufficiently long path cannot be a Nash equilibrium of the local connection game from Section 19.2.

19.5 Show that any path can be a pairwise stable network for a large enough value of α in the bilateral network formation game introduced in Section 19.5.1.

19.6 Construct a Nash equilibrium that is not a star for $\alpha > 2$.

19.7 Show that when $\alpha > n^2$ all Nash equilibria of the local connection game are trees and the price of anarchy is bounded by a constant.

19.8 Prove that the bounds of Lemma 19.4 and Theorems 19.5 and 19.6 are also valid for the worst possible quality of a pairwise stable equilibria of the bilateral version of the game (where an edge needs to be selected, and paid for by both endpoints to be included in G).

19.9 Prove that in the global connection game, the price of anarchy can never exceed k, the number of players.

19.10 Consider the following weighted generalization of the global connection game. For each player i, we have a weight $w_i > 0$. As before, each player selects a single path connecting her source and sink. But instead of sharing edge costs equally, players are now assigned cost shares in proportion to their weight. In particular, for a strategy vector S and edge e, let S_e denote those players whose path contains e, and let $W_e = \sum_{i \in S_e} w_i$ be the total weight of these players. Then player i pays $c_e w_i / W_e$ for each edge $e \in P_i$. Note that if all players have the same weight, this is the original game. Show that, in general, this game does not have an exact potential function.

19.11 In the global network formation game, edge costs reflect fixed building expenses, and thus each player's share for an edge e decreases as more players use e. We might also consider a model with the opposite behavior, i.e., a model in which the cost of using e increases with the number of players. This would be more appropriate for modeling latency or similar effects that make congestion undesirable.

Consider a game played on a network G with k players. Player i has a source s_i and a sink t_i. Each edge $e \in G$ also has a nondecreasing latency function $\ell_e(x)$, indicating the cost incurred by *each* player on e if there are x of these players. A strategy for i is a path from s_i to t_i, and choosing a path P_i incurs a total cost of

$$\text{cost}(P_i) = \sum_{e \in P_i} \ell_e(k_e),$$

where k_e is the number of players using e.

(a) Prove that this game has an exact potential function.

(b) Suppose that we also give each player i an integral weight $w_i \geq 1$.

A strategy for i is a multiset S_i of w_i paths from s_i to t_i. Notice that we do not insist that these paths be disjoint, or even distinct. Costs are now assigned in a natural way; we first compute the cost that each individual path would be charged if each corresponded to a distinct player. Then each player i is charged the sum of the costs of all paths in S_i. Prove that if the latency functions $\ell_e(x)$ are linear for all e, then this game has an exact potential function.

(c) Show that if $\ell_e(x)$ is not linear, then there may not be an exact potential function.

19.12 One problem with using best response dynamics to find pure equilibria in potential games such as the global connection game is that the running time may be exponential. One natural way to deal with this problem is to run best response dynamics, but to consider only moves that provide a substantial decrease in the potential function. In particular, for a constant $\epsilon > 0$, we say a best response move is *substantial* if it decreases the potential function by at least an ϵ/k fraction of its current value. We consider the process of making substantial best response moves until none are available.

(a) Prove that this process terminates in time that is polynomial in n, k, and $\log(\epsilon^{-1})$.

(b) Show that the resulting outcome is not necessarily an approximate equilibrium. That is, show that there may be players who can decrease their costs by an arbitrarily large factor.

19.13 Suppose that we have a game G and a function $\Phi(S)$ mapping game states to reals with the following property: for any strategy vectors $S = (S_1, S_2, \ldots, S_k)$, and any alternate strategy $S_i' \neq S_i$ for some player i, then if $S' = (S_{-i}, S_i')$, we have that $\Phi(S) - \Phi(S')$ and $u_i(S') - u_i(S)$ share the same sign. Thus $\Phi(S)$ behaves like an exact potential function, except instead of tracking a player's improvement exactly, it simply tracks the direction of the improvement; when a player makes an improving move, the potential function decreases. We call such a function an *ordinal potential function*, and G an *ordinal potential game*.

(a) Prove that if G is an ordinal potential game, then best response dynamics always converge to a Nash equilibrium.

(b) Prove that the converse is also true; if, from any starting configuration, best response dynamics always converge to an equilibrium, then G is an ordinal potential game.

19.14 Give an example of the global connection game for which the best Nash equilibrium does not minimize the potential function Ψ.

19.15 Prove that any congestion game is an exact potential game.

19.16 Consider the following location game. We have an unweighted, undirected network G and k players. Each player selects a node in G as their location. Each node v has one unit of wealth that it uniformly distributes to all players in $N[v]$, the closed neighborhood of v. If there are no players in $N[v]$, this wealth is lost. For example, if v has neighbors u and x, 2 players locate at v, 3 players locate at u, and no one locates at x, then v awards 1/5 to each of these 5 players. The utility of a player is simply the sum of the value awarded to it by all nodes. We define the social utility of this game as the number of nodes that have at least one player located in their closed neighborhood.

(a) Prove that the price of anarchy of this game can be arbitrarily close to 2.

(b) Prove that this location game is a valid utility game.

19.17 In theorem 19.19 we showed that if the facilities cost 0, then the social welfare of any Nash equilibrium is at least 1/2 of the maximum possible social welfare of any solution. In this problem, we consider a variant where facilities cost money; each possibly facility s_i has a cost $f(s_i)$, to be paid by a player who locates a facility at s_i.

(a) Is the same bound on the quality of a Nash equilibrium also true for the variant of this game that facilities cost money? Prove or give an example where it is not true.

(b) Let F denote the total facility cost of at a Nash equilibrium S, i.e., the sum $\sum_{s_i \in S} f_{s_i}$. Show that we can bound the optimum $V(O)$ by $2V(S) + F$.

19.18 We now consider a variant of the selfish routing game of Chapter 18 with k players. We have a graph G and a delay function $\ell_e(x)$ that is monotone increasing and convex for each edge $e \in E$. Player i has a source s_i and a destination t_i, and must select an $s_i - t_i$ path P_i on which to route 1 unit of traffic. Player i will tolerate up to d_i delay. Player i picks a path from s_i to t_i with minimum delay, or no path at all if this delay exceeds d_i.

(a) Show that this game always has a pure (deterministic) Nash equilibrium.

(b) The traditional way to evaluate such routing games is with the sum of all delays as cost. However, in this version, the cost may be low simply because few players get routed. Thus we can instead consider the value gathered by each player; d_i minus the delay incurred if i does route her traffic, and 0 if she doesn't. By definition, all players routed have nonnegative value. The total value of a solution is simply the sum of player values. Show that this is a utility game.

(c) Is this game a monotone utility game?