

Routing Games

Tim Roughgarden

Abstract

This chapter studies the inefficiency of equilibria in noncooperative routing games, in which self-interested players route traffic through a congested network. Our goals are threefold: to introduce the most important models and examples of routing games; to survey optimal bounds on the price of anarchy in these models; and to develop proof techniques that are useful for bounding the inefficiency of equilibria in a range of applications.

18.1 Introduction

A majority of the current literature on the inefficiency of equilibria concerns routing games. One reason for this popularity is that routing games shed light on an important practical problem: how to route traffic in a large communication network, such as the Internet, that has no central authority. The routing games studied in this chapter are relevant for networks with “source routing,” in which each end user chooses a full route for its traffic, and also for networks in which traffic is routed in a distributed, congestion-sensitive manner. Section 18.6 contains further details on these applications.

This chapter focuses on two different models of routing games, although the inefficiency of equilibria has been successfully quantified in a range of others (see Section 18.6). The first model, *nonatomic selfish routing*, is a natural generalization of Pigou’s example (Example 17.1) to more complex networks. The modifier “nonatomic” refers to the assumption that there are a very large number of players, each controlling a negligible fraction of the overall traffic. We also study *atomic selfish routing*, where each player controls a nonnegligible amount of traffic. We single out these two models for three reasons. First, both models are conceptually simple but quite general. Second, the price of anarchy is well understood in both of these models. Third, the two models are superficially similar, but different techniques are required to analyze the inefficiency of equilibria in each of them.

The chapter proceeds as follows. Section 18.2 introduces nonatomic and atomic selfish routing games and explores several examples. Section 18.3 studies the existence

and uniqueness of equilibria in routing games. It also offers a glimpse of the potential function method, a technique that will be developed further in Chapter 19. Section 18.4 proves tight upper bounds on the price of anarchy in nonatomic and atomic selfish routing games. Section 18.5 proposes two ways to reduce the price of anarchy in nonatomic selfish routing games. Section 18.6 concludes with bibliographic notes.

18.2 Models and Examples

18.2.1 Nonatomic Selfish Routing

To introduce nonatomic selfish routing games, we recall the essential features of Pigou's example (Example 17.1 and Figure 17.1). First, we are given a network describing the routes available to the players. In Pigou's example, there are two parallel routes, each a single edge, that connect a source vertex s to a sink vertex t . Each edge has a cost that is a function of the amount of traffic that uses the edge. We assume that selfish players choose routes to minimize the cost incurred; in an equilibrium outcome, all players choose a path of minimum cost. In the equilibrium in Pigou's example, all players choose the second edge, and the cost of this edge in this outcome is 1.

More generally, a selfish routing game occurs in a *multicommodity flow network*, or simply a *network*. A network is given by a directed graph $G = (V, E)$, with vertex set V and directed edge set E , together with a set $(s_1, t_1), \dots, (s_k, t_k)$ of source–sink vertex pairs. We also call such pairs *commodities*. Each player is identified with one commodity; note that different players can originate from different source vertices and travel to different sink vertices. We use \mathcal{P}_i to denote the s_i – t_i paths of a network. We consider only networks in which $\mathcal{P}_i \neq \emptyset$ for all i , and define $\mathcal{P} = \cup_{i=1}^k \mathcal{P}_i$. We allow the graph G to contain parallel edges, and a vertex can participate in multiple source–sink pairs.

We describe the routes chosen by players using a *flow*, which is simply a nonnegative vector indexed by the set \mathcal{P} of source–sink paths. For a flow f and a path $P \in \mathcal{P}_i$, we interpret f_P as the amount of traffic of commodity i that chooses the path P to travel from s_i to t_i . Traffic is “inelastic,” in that there is a prescribed amount r_i of traffic identified with each commodity i . A flow f is *feasible* for a vector r if it routes all of the traffic: for each $i \in \{1, 2, \dots, k\}$, $\sum_{P \in \mathcal{P}_i} f_P = r_i$. In particular, we do not impose explicit edge capacities.

Finally, each edge e of a network has a *cost function* $c_e : \mathcal{R}^+ \rightarrow \mathcal{R}^+$. We always assume that cost functions are nonnegative, continuous, and nondecreasing. All of these assumptions are reasonable in applications where cost represents a quantity that only increases with the network congestion; delay is one natural example. When we study the price of anarchy in Section 18.4, we also explore more severe assumptions on the network cost functions. We define a *nonatomic selfish routing game*, or simply a *nonatomic instance*, by a triple of the form (G, r, c) .

Next we formalize the notion of equilibrium in nonatomic selfish routing games. Define the cost of a path P with respect to a flow f as the sum of the costs of the constituent edges: $c_P(f) = \sum_{e \in P} c_e(f_e)$, where $f_e = \sum_{P \in \mathcal{P} : e \in P} f_P$ denotes the amount of traffic using paths that contain the edge e . Since we expect selfish traffic to attempt to minimize its cost, we arrive at the following definition.

Definition 18.1 (Nonatomic equilibrium flow) Let f be a feasible flow for the nonatomic instance (G, r, c) . The flow f is an *equilibrium flow* if, for every commodity $i \in \{1, 2, \dots, k\}$ and every pair $P, \tilde{P} \in \mathcal{P}_i$ of s_i - t_i paths with $f_P > 0$,

$$c_P(f) \leq c_{\tilde{P}}(f).$$

In other words, all paths in use by an equilibrium flow f have minimum-possible cost (given their source, sink, and the congestion caused by f). In particular, all paths of a given commodity used by an equilibrium flow have equal cost. Section 18.3.1 proves that every nonatomic instance admits at least one equilibrium flow, and that all equilibrium flows of a nonatomic instance have equal cost.

In Pigou's example, routing all of the traffic on the second link defines an equilibrium flow; only one path carries flow, and the only alternative has equal cost. Splitting the traffic equally between the two links defines a flow that is not an equilibrium flow; the first link carries a strictly positive amount of traffic and its cost is 1, but there is a strictly cheaper alternative (the second link, with cost 1/2).

Remark 18.2 Our description of nonatomic selfish routing games and their equilibria does not parallel that of simultaneous-move games in Chapter 1. For example, we have not explicitly defined the set of players. While more general types of nonatomic games are frequently defined explicitly in terms of player sets, strategy profiles, and player payoff functions, selfish routing games possess special structure. In particular, the cost incurred by a player depends only on its path and the amount of flow on the edges of its path, rather than on the identities of any of the players. Games of this type are often called *congestion games*. Because of this structure, it is sufficient and simpler to work directly with flows in nonatomic selfish routing games.

When we quantify the inefficiency of equilibrium flows in Section 18.4, we consider only the utilitarian objective of minimizing the total cost incurred by traffic. (Other objectives have been studied; see Section 18.6.) Precisely, since the cost incurred by a player choosing the path P in the flow f is $c_P(f)$, and f_P denotes the amount of traffic choosing the path P , we define the *cost* of a flow f as

$$C(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P. \tag{18.1}$$

Expanding $c_P(f)$ as $\sum_{e \in P} c_e(f_e)$ and reversing the order of summation in (18.1) gives a useful alternative definition of the cost of a flow:

$$C(f) = \sum_{e \in E} c_e(f_e) f_e. \tag{18.2}$$

For an instance (G, r, c) , we call a feasible flow *optimal* if it minimizes the cost over all feasible flows.

As in Chapter 17, the *price of anarchy* of a nonatomic selfish routing game, with respect to this objective, is the ratio between the cost of an equilibrium flow and that of an optimal flow. We can use the cost of an arbitrary equilibrium flow in lieu of that of a worst equilibrium flow (cf. Chapter 17), since all equilibrium flows of a nonatomic




Figure 18.1. A nonlinear variant of Pigou's example (Example 18.3).

instance have equal cost (Section 18.3.1). In Pigou's example, the equilibrium flow routes all of the traffic on the second link and has cost 1. As we will see in Section 18.3.1, the optimal flow splits the traffic equally between the two links and has cost $3/4$. The price of anarchy in Pigou's example is therefore $4/3$.

We conclude this section with two more important examples of nonatomic selfish routing networks.

Example 18.3 (Nonlinear Pigou's example) The inefficiency of the equilibrium flow in Pigou's example can be amplified with a seemingly minor modification to the network. Suppose that we replace the previously linear cost function $c(x) = x$ on the lower edge with the highly nonlinear one $c(x) = x^p$ for p large (Figure 18.1). As in Pigou's example, the cost of the unique equilibrium flow is 1. The optimal flow routes a small ϵ fraction of the traffic on the upper edge and has cost $\epsilon + (1 - \epsilon)^{p+1}$, where ϵ tends to 0 as p tends to infinity. Precisely, Section 18.3.1 shows that $\epsilon = 1 - (p + 1)^{-1/p}$. As p tends to infinity, the cost of the optimal flow approaches 0 and the price of anarchy grows without bound. Exercise 18.1 shows that this rate of growth is roughly $p / \ln p$ as $p \rightarrow \infty$.

While the price of anarchy in our final example is no larger than in Pigou's example, it is arguably a more shocking display of the inefficiency of equilibria in selfish routing networks.

Example 18.4 (Braess's Paradox) Consider the four-node network shown in Figure 18.2(a). There are two disjoint routes from s to t , each with combined cost $1 + x$, where x is the amount of traffic that uses the route. Assume that there is one unit of traffic. In the equilibrium flow, the traffic is split evenly between the two routes, and all of the traffic experiences $3/2$ units of cost.

Now suppose that, in an effort to decrease the cost encountered by the traffic, we build a zero-cost edge connecting the midpoints of the two existing routes. The new network is shown in Figure 18.2(b). What is the new equilibrium flow?

The previous equilibrium flow does not persist in the new network: the cost of the new route $s \rightarrow v \rightarrow w \rightarrow t$ is never worse than that along the two original paths, and it is strictly less whenever some traffic fails to use it. As a consequence, the unique equilibrium flow routes all of the traffic on the new route. Because of the ensuing heavy congestion on the edges (s, v) and (w, t) , all of the traffic now




Figure 18.2. Braess's Paradox. The addition of an intuitively helpful edge can adversely affect all of the traffic.

experiences two units of cost. Braess's Paradox thus shows that the intuitively helpful action of adding a new zero-cost edge can increase the cost experienced by all of the traffic!

Braess's Paradox also has remarkable analogues in several physical systems; see Section 18.6 for details.

The optimal flow in the second network of Example 18.4 is the same as the equilibrium flow in the first network. The price of anarchy in the second network is therefore $4/3$, the same as that in Pigou's example. This is not entirely a coincidence; in Section 18.4.1 we prove that no nonatomic instance with cost functions of the form $ax + b$ has a price of anarchy larger than $4/3$.

While this chapter does not explicitly study Braess's Paradox, we obtain bounds on the worst-case severity of the paradox as a consequence of our results on the price of anarchy (Remark 18.22).

18.2.2 Atomic Selfish Routing

An *atomic selfish routing game* or *atomic instance* is defined by the same ingredients as a nonatomic one: a directed graph $G = (V, E)$, k source–sink pairs $(s_1, t_1), \dots, (s_k, t_k)$, a positive amount r_i of traffic for each pair (s_i, t_i) , and a nonnegative, continuous, nondecreasing cost function $c_e : \mathcal{R}^+ \rightarrow \mathcal{R}^+$ for each edge e . We also denote an atomic instance by a triple (G, r, c) . The intuitive difference between a nonatomic and an atomic instance is that in the former, each commodity represents a large population of individuals, each of whom controls a negligible amount of traffic; in the latter, each commodity represents a single player who must route a significant amount of traffic on a single path.

More formally, atomic instances are finite simultaneous-move games in the sense of Chapter 1. There are k players, one for each source–sink pair. Different players can have identical source–sink pairs. The strategy set of player i is the set \mathcal{P}_i of s_i – t_i paths, and if player i chooses the path P , then it routes its r_i units of traffic on P . A flow is now a nonnegative vector indexed by players and paths, with $f_P^{(i)}$ denoting the amount of traffic that player i routes on the s_i – t_i path P . A flow f is *feasible* for an

atomic instance if it corresponds to a strategy profile: for each player i , $f_P^{(i)}$ equals r_i for exactly one s_i - t_i path and equals 0 for all other paths. The cost $c_P(f)$ of a path P with respect to a flow f and the cost $C(f)$ of a flow f are defined as in Section 18.2.1.

An equilibrium flow of an atomic selfish routing game is a feasible flow such that no player can strictly decrease its cost by choosing a different path for its traffic.

Definition 18.5 (Atomic equilibrium flow) Let f be a feasible flow for the atomic instance (G, r, c) . The flow f is an *equilibrium flow* if, for every player $i \in \{1, 2, \dots, k\}$ and every pair $P, \tilde{P} \in \mathcal{P}_i$ of s_i - t_i paths with $f_P^{(i)} > 0$,

$$c_P(f) \leq c_{\tilde{P}}(\tilde{f}),$$

where \tilde{f} is the flow identical to f except that $\tilde{f}_P^{(i)} = 0$ and $\tilde{f}_{\tilde{P}}^{(i)} = r_i$.

We have defined equilibrium flows to correspond to pure-strategy Nash equilibria (see Chapter 1). Flows corresponding to mixed-strategy Nash equilibria have also been studied (see Section 18.6), but we will not consider them in this chapter.

While the definitions of nonatomic and atomic instances are very similar, the two models are technically quite different. The next example illustrates two of these differences. First, different equilibrium flows of an atomic instance can have different costs; as claimed in Section 18.2.1 and proved in Section 18.3.1, all equilibrium flows of a nonatomic instance have equal cost. Second, the price of anarchy in atomic instances can be larger than in their nonatomic counterparts. The following atomic instance has affine cost functions – of the form $ax + b$ – and its price of anarchy is $5/2$; in every nonatomic instance with affine cost functions, the price of anarchy is at most $4/3$ (Section 18.4.1). We call this the AAE example, after the initials of its discoverers (see Section 18.6).

Example 18.6 (AAE example) Consider the bidirected triangle network shown in Figure 18.3. We assume that there are four players, each of whom needs to route one unit of traffic. The first two have source u and sinks v and w , respectively; the third has source v and sink w ; and the fourth has source w and sink v . Each player has two strategies, a one-hop path and a two-hop path. In the optimal flow, all players route on their one-hop paths, and the cost of this flow is 4. This flow is also an equilibrium flow. On the other hand, if all players route on their two-hop paths, then we obtain a second equilibrium flow. Since the first two players each incur three units of cost and the last two players each incur two units of cost, this equilibrium flow has a cost of 10. The price of anarchy of this instance is therefore $10/4 = 2.5$.

Exercise 18.2 explores variants of the AAE example.

Next we study the even more basic issue of the existence of equilibrium flows. Recall that equilibrium flows for atomic instances correspond to pure-strategy Nash equilibria, which do not always exist in arbitrary finite games (see Chapter 1). Do they always exist in atomic selfish routing games? Our second example answers this question in the negative.




Figure 18.3. The AAE example (Example 18.6). In atomic instances with affine cost functions, different equilibrium flows can have different costs, and the price of anarchy can be as large as $5/2$.

Example 18.7 (Nonexistence in weighted atomic instances) Consider the network shown in Figure 18.4. Extend this network to an atomic selfish routing game by adding two players, both with source s and sink t , with traffic amounts $r_1 = 1$ and $r_2 = 2$.

We claim that there is no equilibrium flow in this atomic instance. To prove this, let P_1, P_2, P_3 , and P_4 denote the paths $s \rightarrow t$, $s \rightarrow v \rightarrow t$, $s \rightarrow w \rightarrow t$, and $s \rightarrow v \rightarrow w \rightarrow t$, respectively. The following four statements then imply the claim.

- (1) If player 2 takes path P_1 or P_2 , then the unique response by player 1 that minimizes its cost is the path P_4 .
- (2) If player 2 takes path P_3 or P_4 , then the unique best response by player 1 is the path P_1 .
- (3) If player 1 takes the path P_4 , then the unique best response by player 2 is the path P_3 .
- (4) If player 1 takes the path P_1 , then the unique best response by player 2 is the path P_2 .

We leave verification of (1)–(4) to the reader.

On the other hand, Section 18.3.2 proves that every atomic instance in which all players route the same amount of traffic admits at least one equilibrium flow. We call




Figure 18.4. An atomic instance with no equilibrium flow (Example 18.7).

instances of this type *unweighted*. Example 18.6 is an unweighted instance, while Example 18.7 is not.

18.3 Existence, Uniqueness, and Potential Functions

This section collects existence and uniqueness results about equilibrium flows in nonatomic and atomic selfish routing games. We also introduce the potential function method, a fundamental proof technique.

18.3.1 Nonatomic Selfish Routing: Existence and Uniqueness

Our next goal is to show that in nonatomic selfish routing games, equilibrium flows always exist and are essentially unique. By “essentially unique,” we mean that all equilibrium flows of a nonatomic instance have the same cost. In particular, the price of stability (Section 17.1) and the price of anarchy coincide in every nonatomic instance. Formally, our aim is to prove the following theorem.

Theorem 18.8 (Existence and uniqueness of equilibrium flows) *Let (G, r, c) be a nonatomic instance.*

- (a) *The instance (G, r, c) admits at least one equilibrium flow.*
- (b) *If f and \tilde{f} are equilibrium flows for (G, r, c) , then $c_e(f_e) = c_e(\tilde{f}_e)$ for every edge e .*

Part (b) of the theorem and Definition 18.1 easily imply that two equilibrium flows of a nonatomic instance have equal cost.

We prove Theorem 18.8 with the *potential function method*. The idea of this method is to exhibit a real-valued “potential function,” defined on the outcomes of a game, such that the equilibria of the game are precisely the outcomes that optimize the potential function. Potential functions are useful because they enable the application of optimization techniques to the study of equilibria. When a game admits a potential function, there are typically consequences for the existence, uniqueness, and inefficiency of equilibria.

To motivate the potential functions corresponding to nonatomic selfish routing games, we present a characterization of optimal flows in such games. To state this characterization cleanly, we assume that for every edge e of the given nonatomic instance, the function $x \cdot c_e(x)$ is continuously differentiable and convex. Note that $x \cdot c_e(x)$ is the contribution to the social cost function (18.2) by traffic on the edge e . Let $c_e^*(x) = (x \cdot c_e(x))' = c_e(x) + x \cdot c_e'(x)$ denote the *marginal cost function* for the edge e . For example, if $c(x)$ denotes the cost function $c(x) = ax^p$ for some $a, p \geq 0$, then the corresponding marginal cost function is $c^*(x) = (p + 1)ax^p$. Let $c_P^*(f) = \sum_{e \in P} c_e^*(f)$ denote the sum of the marginal costs of the edges in the path P with respect to the flow f . The characterization follows.

Proposition 18.9 (Characterization of optimal flows) *Let (G, r, c) be a nonatomic instance such that, for every edge e , the function $x \cdot c_e(x)$ is convex and continuously differentiable. Let c_e^* denote the marginal cost function of the*

edge e . Then f^* is an optimal flow for (G, r, c) if and only if, for every commodity $i \in \{1, 2, \dots, k\}$ and every pair $P, \tilde{P} \in \mathcal{P}_i$ of s_i - t_i paths with $f_P^* > 0$,

$$c_P^*(f^*) \leq c_{\tilde{P}}^*(f^*).$$

Proposition 18.9 follows immediately from the first-order conditions of a convex optimization problem with nonnegativity constraints. We omit the details and focus instead on how the proposition leads to a potential function for equilibrium flows in nonatomic instances, and on the implications of this potential function for the existence and uniqueness of equilibrium flows.

Definition 18.1 and Proposition 18.9 immediately imply that equilibrium flows and optimal flows are the same thing, just with respect to different sets of cost functions.

Corollary 18.10 (Equivalence of equilibrium and optimal flows) *Let (G, r, c) be a nonatomic instance such that, for every edge e , the function $x \cdot c_e(x)$ is convex and continuously differentiable. Let c_e^* denote the marginal cost function of the edge e . Then f^* is an optimal flow for (G, r, c) if and only if it is an equilibrium flow for (G, r, c^*) .*

For instance, in Pigou’s example (Example 17.1), the marginal cost functions of the two edges are $c^*(x) = 1$ and $c^*(x) = 2x$. The equilibrium flow with respect to the marginal cost functions splits the traffic equally between the two links, equalizing their marginal costs at 1; by Corollary 18.10, this flow is optimal in the original network. In the nonlinear variant of Pigou’s example (Example 18.3), the marginal cost functions are $c^*(x) = 1$ and $c^*(x) = (p + 1)x^p$; the optimal flow therefore routes $(p + 1)^{-1/p}$ units of traffic on the second link and the rest on the first. In Braess’s Paradox with the zero-cost edge added (Example 18.4 and Figure 18.2(b)), routing half of the traffic on each of the paths $s \rightarrow v \rightarrow t$ and $s \rightarrow w \rightarrow t$ equalizes the marginal costs of all three paths at 2, and therefore provides an optimal flow.

To construct a potential function for equilibrium flows, we need to “invert” Corollary 18.10: of what function do equilibrium flows arise as the global minima? The answer is simple: to recover Definition 18.1 as an optimality condition, we seek a function $h_e(x)$ for each edge e – playing the previous role of $x \cdot c_e(x)$ – such that $h'_e(x) = c_e(x)$. Setting $h_e(x) = \int_0^x c_e(y) dy$ for each edge e thus yields the desired potential function. Moreover, since c_e is continuous and nondecreasing for every edge e , every function h_e is both continuously differentiable and convex.

Precisely, call

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} c_e(x) dx \tag{18.3}$$

the *potential function* of a nonatomic instance (G, r, c) . Invoking Proposition 18.9, with each function $x \cdot c_e(x)$ replaced by $h_e(x) = \int_0^x c_e(y) dy$, yields the same condition as in Definition 18.1; we have therefore characterized equilibrium flows as the global minimizers of the potential function Φ .

Proposition 18.11 (Potential function for equilibrium flows) *Let (G, r, c) be a nonatomic instance. A flow feasible for (G, r, c) is an equilibrium flow if and only if it is a global minimum of the corresponding potential function Φ given in (18.3).*

Theorem 18.8 now follows from Proposition 18.11 and routine calculus.

PROOF OF THEOREM 18.8 We first note that, by definition, the set of feasible flows of (G, r, c) can be identified with a compact (i.e., closed and bounded) subset of $|\mathcal{P}|$ -dimensional Euclidean space. Since edge cost functions are continuous, the potential function is a continuous function on this set. By Weierstrass's Theorem from elementary mathematical analysis, the potential function Φ achieves a minimum value on this set. By Proposition 18.11, every point at which Φ attains its minimum corresponds to an equilibrium flow of (G, r, c) .

For part (b), recall that each cost function is nondecreasing, and hence each summand on the right-hand side of (18.3) is convex. Hence, the potential function Φ is a convex function.

Now suppose that f and \tilde{f} are equilibrium flows for (G, r, c) . By Proposition 18.11, both f and \tilde{f} minimize the potential function Φ . We consider all convex combinations of f and \tilde{f} – that is, all vectors of the form $\lambda f + (1 - \lambda)\tilde{f}$ for $\lambda \in [0, 1]$. All of these vectors are feasible flows. Since Φ is a convex function, a chord between two points on its graph cannot pass below its graph. In algebraic terms, we have

$$\Phi(\lambda f + (1 - \lambda)\tilde{f}) \leq \lambda\Phi(f) + (1 - \lambda)\Phi(\tilde{f}) \quad (18.4)$$

for every $\lambda \in [0, 1]$. Since both f and \tilde{f} are global minima of Φ , the inequality (18.4) must hold with equality for all of their convex combinations. Since every summand of Φ is convex, this can occur only if every summand $\int_0^x c_e(y) dy$ is linear between the values f_e and \tilde{f}_e . In turn, this implies that every cost function c_e is constant between f_e and \tilde{f}_e . \square

18.3.2 Atomic Selfish Routing: Existence

We now consider equilibrium flows in atomic instances. The AAE example (Example 18.6) suggests that no interesting uniqueness results are possible in such instances, so we focus instead on the existence of equilibrium flows. Similarly, Example 18.7 demonstrates that a general atomic instance need not admit an equilibrium flow. There are two approaches to circumventing this counterexample. The first, taken in this section, is to place additional restrictions on atomic instances so that equilibrium flows are guaranteed to exist. The second approach, discussed in Remark 18.26, is to relax the equilibrium concept so that an equilibrium exists in every atomic instance.

The key result in this section is the following theorem, which establishes the existence of equilibrium flows in atomic instances in which all players control the same amount of traffic.

Theorem 18.12 (Equilibrium flows in unweighted atomic instances) *Let (G, r, c) be an atomic instance in which every traffic amount r_i is equal to a common positive value R . Then (G, r, c) admits at least one equilibrium flow.*

PROOF We obtain Theorem 18.12 by discretizing the potential function (18.3) for nonatomic instances and the proof of Theorem 18.8(a). Assume for simplicity that $R = 1$. Set

$$\Phi_a(f) = \sum_{e \in E} \sum_{i=1}^{f_e} c_e(i) \quad (18.5)$$

for every feasible flow f . Note that Φ_a is the same as the previous potential function Φ for nonatomic instances, except that the integral $\int_0^{f_e} c(x) dx$ has been replaced by the sum $\sum_{i=1}^{f_e} c_e(i)$.

Since the atomic instance (G, r, c) has a finite number of players, and each of these has a finite number of strategies, there are only a finite number of possible flows. One of these, call it f , is a global minimum of the potential function Φ_a . We claim that f is an equilibrium flow for (G, r, c) . To prove it, assume for contradiction that in f , the player i could strictly decrease its cost by deviating from the path P to the path \tilde{P} , yielding the new flow \tilde{f} . In other words, we assume that

$$0 > c_{\tilde{P}}(\tilde{f}) - c_P(f) = \sum_{e \in \tilde{P} \setminus P} c_e(f_e + 1) - \sum_{e \in P \setminus \tilde{P}} c_e(f_e). \quad (18.6)$$

On the other hand, consider the impact of player i 's deviation on the potential function Φ_a : for edges in $\tilde{P} \setminus P$, the corresponding sum in (18.5) acquires the extra term $c_e(f_e + 1)$; for edges in $P \setminus \tilde{P}$, the corresponding sum sheds the term $c_e(f_e)$; and for edges of $P \cap \tilde{P}$, the corresponding sum remains the same. Thus, $\Phi_a(\tilde{f}) - \Phi_a(f)$ is precisely the third expression of (18.6). Since this expression is negative, the potential function value of \tilde{f} is strictly less than that of f , which contradicts our choice of f . \square

Remark 18.13 The proof of Theorem 18.12 establishes a remarkable property of the potential function Φ_a : it “tracks” the change in cost experienced by a deviating player. More formally, for every flow, every player, and every deviation by a player, the change in the player’s cost is identical to the change in the potential function. This property has consequences beyond the existence result of Theorem 18.12. For example, it implies that “best-response dynamics” are guaranteed to converge to an equilibrium flow. See Chapter 19 for further details.

Remark 18.14 The proof of Theorem 18.12 did not use any assumptions about the edge cost functions. In particular, it is also valid when cost functions are not nondecreasing. This property will be crucial for some of the network design games studied in Chapter 19, which can be viewed as atomic selfish routing games with *decreasing* cost functions.

The next theorem guarantees the existence of equilibrium flows under a different restriction – affine cost functions. (Recall that a cost function $c_e(x)$ is *affine* if it has the form $a_e x + b_e$; we always assume that $a_e, b_e \geq 0$.)

Theorem 18.15 (Equilibrium flows with affine cost functions) *Let (G, r, c) be an atomic instance with affine cost functions. Then (G, r, c) admits at least one equilibrium flow.*

The proof of Theorem 18.15 follows the same outline as that of Theorem 18.12, and uses a variant of the potential function method. See Exercise 18.4 for further details.

18.4 The Price of Anarchy of Selfish Routing

18.4.1 Nonatomic Selfish Routing: The Price of Anarchy

This section gives an essentially complete analysis of the price of anarchy in nonatomic selfish routing games. As we know from the nonlinear variant of Pigou’s example (Example 18.3), the price of anarchy depends on “nonlinearity” of the network cost functions. Our goal is to show that it depends on nothing else – not the network size, the network structure, nor the number of commodities. More precisely, we show that for every conceivable restriction on the cost functions of a network, the price of anarchy is maximized (over all multicommodity networks) by the network that best “simulates” Pigou’s example and its nonlinear variants.

As an aside, we note that the potential function characterization of nonatomic equilibrium flows (Proposition 18.11) already gives a good, but not optimal, upper bound on the price of anarchy. The intuitive explanation is simple: if equilibrium flows exactly optimize a potential function (18.3) that is a good approximation of the objective function (18.2), then they should also be approximately optimal.

Theorem 18.16 (Potential function upper bound) *Let (G, r, c) be a nonatomic instance, and suppose that $x \cdot c_e(x) \leq \gamma \cdot \int_0^x c_e(y) dy$ for all $e \in E$ and $x \geq 0$. Then the price of anarchy of (G, r, c) is at most γ .*

PROOF Let f and f^* be equilibrium and optimal flows for (G, r, c) , respectively. Since cost functions are nondecreasing, the cost of a flow (18.2) is always at least its potential function value (18.3). The hypothesis ensures that the cost of a flow is at most γ times its potential function value. The theorem follows by writing

$$C(f) \leq \gamma \cdot \Phi(f) \leq \gamma \cdot \Phi(f^*) \leq \gamma \cdot C(f^*),$$

with the second inequality following from Proposition 18.11. \square

Theorem 18.16 implies that the price of anarchy of selfish routing is large only in networks with “highly nonlinear” cost functions. For example, if c_e is a polynomial function with degree at most p and nonnegative coefficients, then $x \cdot c_e(x) \leq (p + 1) \int_0^x c_e(y) dy$ for all $x \geq 0$. Theorem 18.16 then shows that the price of anarchy in nonatomic instances with such cost functions is at most linear in p .

Corollary 18.17 (Potential function bound for polynomials) *If (G, r, c) is a nonatomic instance with cost functions that are polynomials with nonnegative coefficients and degree at most p , then the price of anarchy of (G, r, c) is at most $p + 1$.*

This upper bound is nearly matched by Example 18.3, although the upper and lower bounds differ by roughly a $\ln p$ multiplicative factor (Exercise 18.1). We close this gap using a different and important proof technique, which is driven by *variational inequalities*.

We first formalize a natural lower bound on the price of anarchy based on “Pigou-like examples.”

Definition 18.18 (Pigou bound) Let \mathcal{C} be a nonempty set of cost functions. The *Pigou bound* $\alpha(\mathcal{C})$ for \mathcal{C} is

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{x, r \geq 0} \frac{r \cdot c(r)}{x \cdot c(x) + (r - x)c(r)}, \quad (18.7)$$

with the understanding that $0/0 = 1$.

The point of the Pigou bound is that it lower bounds the price of anarchy in instances with cost functions in \mathcal{C} .

Proposition 18.19 *Let \mathcal{C} be a set of cost functions that contains all of the constant cost functions. Then the price of anarchy in nonatomic instances with cost functions in \mathcal{C} can be arbitrarily close to $\alpha(\mathcal{C})$.*

PROOF Fix a choice of $c \in \mathcal{C}$ and $x, r \geq 0$. We can complete the proof by exhibiting a selfish routing network with cost functions in \mathcal{C} and price of anarchy at least $c(r)r/[c(x)x + (r - x)c(r)]$. Since c is nondecreasing, this expression is at most 1 if $x \geq r$; we can therefore assume that $x < r$.

Let G be a two-vertex, two-edge network as in Figure 18.1. Give the lower edge the cost function $c_1(y) = c(y)$ and the upper edge the constant cost function $c_2(y) = c(r)$. By assumption, both of these cost functions lie in \mathcal{C} . Set the traffic rate to r . Routing all of the traffic on the lower edge yields an equilibrium flow with cost $c(r)r$. Routing x units of traffic on the lower edge and $r - x$ units of traffic on the upper edge gives a feasible flow with cost $[c(x)x + (r - x)c(r)]$. The price of anarchy in this instance is thus at least $c(r)r/[c(x)x + (r - x)c(r)]$, as desired. \square

While Proposition 18.19 assumes that the set \mathcal{C} includes all of the constant cost functions, its conclusion holds whenever \mathcal{C} is *inhomogeneous* in the sense that $c(0) > 0$ for some $c \in \mathcal{C}$ (Exercise 18.5).

We next show that, even though the Pigou bound is based only on Pigou-like examples, it is also an upper bound on the price of anarchy in general multicommodity flow networks. The proof requires the following variational inequality characterization of equilibrium flows.

Proposition 18.20 (Variational inequality characterization) *Let f be a feasible flow for the nonatomic instance (G, r, c) . The flow f is an equilibrium flow if and only if*

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

for every flow f^* feasible for (G, r, c) .

PROOF Fix f and define the function H_f on the set of feasible flows by

$$H_f(f^*) = \sum_{i=1}^k \sum_{P \in \mathcal{P}_i} c_P(f) f_P^* = \sum_{e \in E} c_e(f_e) f_e^*;$$

the same reversal of sums used to prove the equivalence of (18.1) and (18.2) shows that these two definitions of $H_f(f^*)$ agree. The value $H_f(f^*)$ denotes the cost of a flow f^* after the cost function of each edge e has been changed to the constant function everywhere equal to $c_e(f_e)$. By the second definition of H_f , the proposition is equivalent to the assertion that a flow f is an equilibrium flow if and only if it minimizes $H_f(\cdot)$ over all feasible flows.

Examining the first definition of H_f shows that a flow f^* minimizes H_f if and only if, for every commodity i , $f_p^* > 0$ only for paths P that minimize $c_P(f)$ over all s_i - t_i paths. Since the flow f satisfies this condition if and only if it is an equilibrium flow, the proof is complete. \square

We now show that the Pigou bound is tight.

Theorem 18.21 (Tightness of the Pigou bound) *Let \mathcal{C} be a set of cost functions and $\alpha(\mathcal{C})$ the Pigou bound for \mathcal{C} . If (G, r, c) is a nonatomic instance with cost functions in \mathcal{C} , then the price of anarchy of (G, r, c) is at most $\alpha(\mathcal{C})$.*

PROOF Let f^* and f be optimal and equilibrium flows, respectively, for a nonatomic instance (G, r, c) with cost functions in the set \mathcal{C} . The theorem follows by writing

$$\begin{aligned} C(f^*) &= \sum_{e \in E} c_e(f_e^*) f_e^* \\ &\geq \frac{1}{\alpha(\mathcal{C})} \sum_{e \in E} c_e(f_e) f_e + \sum_{e \in E} (f_e^* - f_e) c_e(f_e) \\ &\geq \frac{C(f)}{\alpha(\mathcal{C})}, \end{aligned}$$

where the first inequality follows from Definition 18.18, applied to each edge e with $x = f_e^*$ and $r = f_e$, and the second inequality follows from Proposition 18.20. \square

Proposition 18.19 and Theorem 18.21 show that, for essentially every fixed restriction on the allowable cost functions, the price of anarchy is maximized by Pigou-like

examples. Determining the largest-possible price of anarchy in Pigou-like examples (i.e., the Pigou bound) is a tractable problem in many cases. For example, it is precisely $4/3$ when \mathcal{C} is the set of affine cost functions (Exercise 18.6), and more generally is $[1 - p \cdot (p + 1)^{-(p+1)/p}]^{-1} \approx p / \ln p$ when \mathcal{C} is the set of polynomials with degree at most p and nonnegative coefficients. In these cases, the maximum price of anarchy (among all multicommodity instances) is achieved by the instances in Examples 17.1 and 18.3. The Pigou bound is also known for several other classes of cost functions; see Section 18.6 for references.

Remark 18.22 (Bounds on Braess's Paradox) Braess's Paradox (Example 18.4) shows that adding edges to a network can increase the cost of its equilibrium flow. Since the equilibrium flow in the original network is a candidate for the optimal flow in the second network, the ratio between the costs of the new and original equilibrium flows is a lower bound on the price of anarchy in the latter network.

On the other hand, Theorem 18.21 and Exercise 18.6 show that the price of anarchy is at most $4/3$ in every network with affine cost functions. Thus, adding edges to a network with affine cost functions cannot increase the cost of its equilibrium flow by more than a $4/3$ factor. Example 18.4 is therefore a worst-case manifestation of Braess's Paradox in networks with affine cost functions. Similar bounds also apply to the physical analogues of Braess's Paradox that are described in Section 18.6.

18.4.2 Atomic Selfish Routing: The Price of Anarchy

We now consider atomic selfish routing games. We again obtain tight bounds on the price of anarchy, at least for polynomial cost functions, but the discrete nature of atomic instances complicates the analysis.

We first note that the potential function method, which gave nontrivial bounds on the price of anarchy for nonatomic instances (Theorem 18.16), cannot be used for atomic instances. The difficulty stems from the non-uniqueness of equilibrium flows in atomic instances (Example 18.6). Recall that a bound on the price of anarchy is a guarantee that *all* equilibrium flows of an instance are nearly optimal. Reviewing the proof of Theorem 18.16, we observe that the potential function method argues about only *one* equilibrium flow – the one with minimum potential function value. As a result, the potential function method is directly useful only for bounding the price of *stability* rather than the price of anarchy. While these two quantities coincide in nonatomic selfish routing games, they are generally different in atomic ones. (See Section 18.6 for results on the price of stability in atomic selfish routing games.)

We instead rely on proof techniques that are partially inspired by the variational inequality of Proposition 18.20. This inequality expresses the fact that equilibrium flows route all traffic on shortest paths, with respect to the induced edge costs. We derive a similar, if more complicated, condition for atomic instances. To keep the proofs as transparent as possible, we focus on atomic instances with affine cost functions. Recall from Theorem 18.15 that every such instance admits at least one equilibrium flow. The

analysis can also be extended to other cost functions and other equilibrium concepts; see Remark 18.26 and Section 18.6 for more details.

Our goal is the following theorem.

Theorem 18.23 (The price of anarchy in affine atomic instances) *If (G, r, c) is an atomic instance with affine cost functions, then the price of anarchy of (G, r, c) is at most $(3 + \sqrt{5})/2 \approx 2.618$.*

A variant of the AAE example (Example 18.6) shows that the upper bound in Theorem 18.23 is the best possible if different players can control different amounts of flow (Exercise 18.2(a)). If all of the players control the same amount of flow, then a variant of the following proof gives an improved upper bound of $5/2$, which matches the lower bound furnished by the AAE example (Exercise 18.7).

We build up to Theorem 18.23 in a sequence of steps. We begin with a lemma that follows immediately from the definition of an equilibrium flow.

Lemma 18.24 (Equilibrium condition) *Let (G, r, c) be an atomic instance in which each edge e has an affine cost function $c_e(x) = a_e x + b_e$ with $a_e, b_e \geq 0$. Let f and f^* be equilibrium and optimal flows, respectively, for (G, r, c) . Let player i use the path P_i in f and the path P_i^* in f^* . Then*

$$\sum_{e \in P_i} [a_e f_e + b_e] \leq \sum_{e \in P_i^*} [a_e (f_e + r_i) + b_e]. \quad (18.8)$$

Our second step is to combine the inequalities of Lemma 18.24 – one per player – to relate the cost of an arbitrary equilibrium flow to that of an optimal flow.

Lemma 18.25 (Equilibrium inequality) *With the same assumptions and notation as in Lemma 18.24,*

$$C(f) \leq C(f^*) + \sum_{e \in E} a_e f_e f_e^*. \quad (18.9)$$

PROOF For each player i , multiply the inequality (18.8) by r_i . Summing up the resulting k inequalities, we obtain

$$\begin{aligned} C(f) &\leq \sum_{i=1}^k r_i \left(\sum_{e \in P_i^*} a_e (f_e + r_i) + b_e \right) \\ &\leq \sum_{i=1}^k r_i \left(\sum_{e \in P_i^*} a_e (f_e + f_e^*) + b_e \right) \\ &= \sum_{e \in E} [a_e (f_e + f_e^*) + b_e] f_e^*, \end{aligned}$$

where the equality follows by reversing the order of summation. Since the final expression equals the right-hand side of (18.9), the proof is complete. \square

To complete the proof of Theorem 18.23, we upper bound the magnitude of the “error term” in (18.9) relative to the costs of the equilibrium and optimal flows.

PROOF OF THEOREM 18.23 Let f and f^* denote equilibrium and optimal flows, respectively, for the atomic instance (G, r, c) . Assume that edge e has the cost function $c_e(x) = a_e x + b_e$ for $a_e, b_e \geq 0$. Apply the Cauchy–Schwarz Inequality to the vectors $\{\sqrt{a_e} f_e\}_{e \in E}$ and $\{\sqrt{a_e} f_e^*\}_{e \in E}$ to obtain

$$\sum_{e \in E} a_e f_e f_e^* \leq \sqrt{\sum_{e \in E} a_e f_e^2} \cdot \sqrt{\sum_{e \in E} a_e (f_e^*)^2} \leq \sqrt{C(f)} \cdot \sqrt{C(f^*)}.$$

Combining this with the Equilibrium Inequality (18.9), dividing through by $C(f^*)$, and rearranging gives

$$\frac{C(f)}{C(f^*)} - 1 \leq \sqrt{\frac{C(f)}{C(f^*)}}.$$

Squaring both sides and solving the corresponding quadratic inequality $x^2 - 3x + 1 \leq 0$, we find that

$$\frac{C(f)}{C(f^*)} \leq \frac{3 + \sqrt{5}}{2} \approx 2.618,$$

as claimed. \square

Theorem 18.23 can be extended to atomic instances with cost functions that are polynomials with nonnegative coefficients and degree at most a parameter p . However, the upper bound on the price of anarchy increases with p roughly in proportion to the exponential function p^p – much faster than in nonatomic instances. This exponential dependence is not an artifact of the above proof approach, as nearly matching lower bounds on the price of anarchy are known (Section 18.6).

Remark 18.26 Strictly speaking, the price of anarchy is not always defined in general atomic instances, where equilibrium flows need not exist (Example 18.7). Nevertheless, Theorem 18.23 has been extended to atomic instances with polynomial cost functions in three different ways. First, when such an instance *does* admit at least one equilibrium flow, then all such flows have cost at most $p^{O(p)}$ times that of an optimal flow. Second, by Nash’s Theorem (Chapters 1 and 2), every such instance admits a *mixed*-strategy Nash equilibrium, and the expected cost of every such equilibrium is at most $p^{O(p)}$ times that of an optimal flow. Finally, similar upper bounds have been proved for “sink equilibria,” an equilibrium concept that always exists in finite games and is motivated by convergence issues.

18.5 Reducing the Price of Anarchy

As we have seen, the price of anarchy can be large in both nonatomic and atomic selfish routing games when cost functions are highly nonlinear. This motivates a question first posed in Section 17.3: how can we design or modify a selfish routing network, without explicitly imposing an optimal solution, to minimize the inefficiency of its equilibria? Can modest intervention significantly reduce the price of anarchy? We briefly discuss two techniques for mitigating the inefficiency of selfish routing in nonatomic instances: influencing traffic with edge taxes (Subsection 18.5.1) and increasing the capacity of the network (Subsection 18.5.2).

18.5.1 Marginal Cost Pricing

Our first approach to reducing the price of anarchy in nonatomic selfish routing games is to use *marginal cost taxes* on the edges of the network. The idea of marginal cost pricing is to charge each network user on each edge for the additional cost its presence causes for the other users of the edge. To discuss this idea formally, we allow each edge e of a nonatomic selfish routing network to possess a nonnegative *tax* τ_e . We denote a nonatomic instance (G, r, c) with edge taxes τ by $(G, r, c + \tau)$. An equilibrium flow for such an instance $(G, r, c + \tau)$ is defined as in Definition 18.1, with all traffic traveling on routes that minimize the sum of the edge costs and edge taxes. Equivalently, it is an equilibrium flow for the nonatomic instance (G, r, c^τ) , where the cost function c_e^τ is a shifted version of the original cost function c_e : $c_e^\tau(x) = c_e(x) + \tau_e$ for all $x \geq 0$.

The principle of marginal cost pricing asserts that for a flow f feasible for a nonatomic instance (G, r, c) , the tax τ_e assigned to the edge e should be $\tau_e = f_e \cdot c'_e(f_e)$, where c'_e denotes the derivative of c_e . (Assume for simplicity that the cost functions are differentiable.) The term $c'_e(f_e)$ corresponds to the marginal increase in cost caused by one user of the edge, and the term f_e is the amount of traffic that suffers from this increase. We can also interpret the marginal cost tax τ_e using Corollary 18.10: τ_e is precisely the “extra term” in the marginal cost function that is absent from the original cost function. These taxes correct for the failure of selfish users to account for the second, “altruistic” term of the marginal cost function. Formally, Corollary 18.10 easily implies the following guarantee.

Theorem 18.27 *Let (G, r, c) be a nonatomic instance such that, for every edge e , the function $x \cdot c_e(x)$ is convex and continuously differentiable. Let f^* be an optimal flow for (G, r, c) and let $\tau_e = f_e^* \cdot c'_e(f_e^*)$ denote the marginal cost tax for edge e with respect f^* . Then f^* is an equilibrium flow for $(G, r, c + \tau)$.*

Marginal cost taxes thus induce an optimal flow as an equilibrium flow; in this sense, such taxes reduce the price of anarchy to 1. Theorem 18.27 also holds with weaker assumptions on the cost functions; in particular, the convexity hypothesis is not needed. For further discussion of pricing problems in routing games, see Chapter 22.

18.5.2 Capacity Augmentation

Our final result is a novel type of bound on the inefficiency of equilibrium flows in nonatomic selfish routing games with arbitrary cost functions. This bound does not involve the price of anarchy, which is unbounded in such networks (Example 18.3), and instead shows that the cost of an equilibrium flow is at most that of an optimal flow that is forced to route twice as much traffic between each source–sink pair. As we will see, this result implies that in lieu of centralized control, the inefficiency of selfish routing can be offset by a moderate increase in link speed.

Example 18.28 Consider the nonlinear variant of Pigou’s example (Example 18.3). When there is one unit of traffic, the equilibrium flow routes all of the flow on the lower edge, while the optimal flow routes ϵ units of flow on the upper edge and the rest on the lower edge (where $\epsilon \rightarrow 0$ as $p \rightarrow \infty$). When the amount r of traffic to be routed exceeds one, an optimal flow assigns the additional $r - 1$ units of traffic to the upper link, incurring a cost that tends to $r - 1$ as $p \rightarrow \infty$. In particular, for every p an optimal flow feasible for twice the original traffic amount ($r = 2$) has cost at least 1, the cost of the equilibrium flow in the original instance.

We now show that the upper bound stated in Example 18.28 for the nonlinear variant of Pigou’s example holds in every nonatomic instance.

Theorem 18.29 *If f is an equilibrium flow for (G, r, c) and f^* is feasible for $(G, 2r, c)$, then $C(f) \leq C(f^*)$.*

PROOF Let f and f^* denote an equilibrium flow for (G, r, c) and a feasible flow for $(G, 2r, c)$, respectively. For each commodity i , let d_i denote the minimum cost of an s_i – t_i path with respect to the flow f . Definition 18.1 and the definition of cost (18.1) imply that $C(f) = \sum_i r_i d_i$.

The key idea is to define a set of cost functions \bar{c} that satisfies two properties: lower bounding the cost of f^* relative to that of f is easy with respect to \bar{c} ; and the new cost functions \bar{c} approximate the original ones c . Specifically, we set $\bar{c}_e(x) = \max\{c_e(f_e), c_e(x)\}$ for each edge e . Let $\bar{C}(\cdot)$ denote the cost of a flow in the instance (G, r, \bar{c}) . Note that $\bar{C}(f^*) \geq C(f^*)$ while $\bar{C}(f) = C(f)$.

We first upper bound the amount by which the new cost $\bar{C}(f^*)$ of f^* can exceed its original cost $C(f^*)$. For every edge e , $\bar{c}_e(x) - c_e(x)$ is zero for $x \geq f_e$ and bounded above by $c_e(f_e)$ for $x < f_e$, so $x(\bar{c}_e(x) - c_e(x)) \leq c_e(f_e)f_e$ for all $x \geq 0$. Thus

$$\bar{C}(f^*) - C(f^*) = \sum_{e \in E} f_e^*(\bar{c}_e(f_e^*) - c_e(f_e^*)) \leq \sum_{e \in E} c_e(f_e)f_e = C(f). \quad (18.10)$$

In other words, evaluating f^* with cost functions \bar{c} , rather than c , increases its cost by at most an additive $C(f)$ factor.

Now we lower bound $\bar{C}(f^*)$. By construction, the modified cost $\bar{c}_e(\cdot)$ of an edge e is always at least $c_e(f_e)$, so the modified cost $\bar{c}_P(\cdot)$ of a path $P \in \mathcal{P}_i$ is always at

least $c_P(f)$, which in turn is at least d_i . The modified cost $\bar{C}(f^*)$ therefore equals

$$\sum_{P \in \mathcal{P}} \bar{c}_P(f^*) f_P^* \geq \sum_{i=1}^k \sum_{P \in \mathcal{P}_i} d_i f_P^* = \sum_{i=1}^k 2r_i d_i = 2C(f). \quad (18.11)$$

The theorem now follows immediately from inequalities (18.10) and (18.11). \square

Another interpretation of Theorem 18.29 is that the benefit of centralized control is equaled or exceeded by the benefit of a sufficient improvement in link technology.

Corollary 18.30 *Let (G, r, c) be a nonatomic instance and define the modified cost function \tilde{c}_e by $\tilde{c}_e(x) = c_e(x/2)/2$ for each edge e . Let \tilde{f} be an equilibrium flow for (G, r, \tilde{c}) with cost $\tilde{C}(\tilde{f})$, and f^* a feasible flow for (G, r, c) with cost $C(f^*)$. Then $\tilde{C}(\tilde{f}) \leq C(f^*)$.*

Simple calculations show that Theorem 18.29 and Corollary 18.30 are equivalent; see Exercise 18.8(a).

Corollary 18.30 takes on a particularly nice form in instances in which all cost functions are *M/M/1 delay functions*. Such a cost function has the form $c_e(x) = (u_e - x)^{-1}$, where u_e can be interpreted as an edge capacity or a queue service rate; the function is defined to be $+\infty$ when $x \geq u_e$. (Rigorously allowing infinite costs in this selfish routing model requires some care; we ignore these issues in this chapter.) In this case, the modified function \tilde{c}_e of Corollary 18.30 is $\tilde{c}_e(x) = 1/2(u_e - x/2) = 1/(2u_e - x)$. Corollary 18.30 thus suggests the following design principle for selfish routing networks with M/M/1 delay functions: *to outperform optimal routing, just double the capacity of every edge.*

18.6 Notes

18.6.1 Nonatomic Selfish Routing

Nonatomic selfish routing was first studied in the context of transportation networks. Pigou (1920) informally discussed Pigou's example in his 1920 book, *The Economics of Welfare*, in order to illustrate the inefficiency of equilibria. He also anticipated the principle of marginal cost pricing discussed in Theorem 18.27; indeed, marginal cost taxes are sometimes called *Pigouvian taxes*. The model was first formally defined by Wardrop (1952). For this reason, equilibrium flows in nonatomic selfish routing games are often called *Wardrop equilibria*. We use the term "equilibrium flow" so that the terminology for nonatomic and atomic selfish routing games is the same.

Beckmann et al. (1956) proved a number of fundamental results for the nonatomic model. Theorem 18.8, Proposition 18.9, Corollary 18.10, Proposition 18.11, and Theorem 18.27 were first proved in Beckmann et al. (1956), via proofs essentially identical to the ones given here. Details on first-order conditions for convex programming problems can be found in Bertsekas (1999, Chapter 2). Schmeidler (1973) founded the theory of general noncooperative nonatomic games.

Two decades after nonatomic selfish routing games were first defined, researchers began to use them to model the routing of data through communication networks. Nonatomic selfish routing is immediately relevant for networks that employ so-called *source routing*, meaning that each sender is responsible for selecting a full path of links to the receiver. Assuming that senders seek paths of minimum cost, senders of data in such networks correspond to the users of a selfish routing network.

In large networks such as the Internet, *distributed shortest-path routing* is typically used instead of source routing. In distributed shortest-path routing, each link is given a positive *length*, and data are forwarded along a path of minimum total length to its destination. Shortest-path routing leaves a key parameter unspecified: the length of each edge. A direct correspondence between selfish routing and shortest-path routing exists if and only if the edge cost functions coincide with the lengths used to define shortest paths. In other words, when an x fraction of the overall network traffic is using an edge with cost function $c(\cdot)$, then the corresponding shortest-path routing algorithm should define the length of the edge as the number $c(x)$. If the cost function c is nonconstant, then this is a congestion-dependent definition of the edge length. In this case, shortest-path routing will route traffic exactly *as if* it is a network with selfish routing (or source routing). For details on this equivalence, see the textbook by Bertsekas and Tsitsiklis (1989). See Qiu et al. (2003), for example, for a more recent paper that studies selfish routing from a computer networking perspective.

Braess's Paradox was discovered by Braess (1968). The variant in Example 18.4 was noted by L. Schulman (personal communication, October 1999). For surveys on the large literature inspired by Braess's Paradox, see Roughgarden (2006) and D. Braess's home page (Braess, 2007).

Cohen and Horowitz (1991) noted that Braess's Paradox has startling analogues in physical systems. For instance, Example 18.4 can be simulated in the following system of strings and springs. One end of a spring is attached to a fixed support, and the other end to a very short string. A second identical spring is hung from the free end of the string and carries a heavy weight. Finally, strings are connected, with very little slack, from the support to the upper end of the second spring and from the lower end of the first spring to the weight. Assuming that the springs are ideally elastic, the stretched length of a spring is a linear function of the force applied to it. We can therefore view the network of strings and springs as a selfish routing game, where force corresponds to traffic and physical distance corresponds to cost. Remarkably, severing the very short taut string causes the weight to levitate away from the ground! The rise in the weight is the same as the improvement in the equilibrium flow obtained by deleting the zero-cost edge of Figure 18.2(b) to recover the network of Figure 18.2(a).

The price of anarchy in nonatomic selfish routing games was first studied by Roughgarden and Tardos (2002). The nonlinear variant of Pigou's example (Example 18.3) is from Roughgarden and Tardos (2002), as is Theorem 18.16. Roughgarden and Tardos (2002) also proved the special case of Theorem 18.21 for networks with affine cost functions (where the price of anarchy is at most $4/3$). Roughgarden (2003) introduced the Pigou bound and proved Theorem 18.21 under the same convexity hypothesis used in Theorem 18.9. The solution to Exercise 18.5 can also be found in Roughgarden (2003). A. Ronen (personal communication, March 2002) suggested using the variational inequality in Proposition 18.20, which was first proved by Smith

(1979). Correa et al. (2004) proved Theorem 18.21 without any convexity assumptions. This theorem has been generalized to wider classes of nonatomic games; see Roughgarden (2005a) for a survey, as well as a discussion of the price of anarchy of nonatomic selfish routing games with nonutilitarian objectives.

Finally, Theorem 18.29 is due to Roughgarden and Tardos (2002). A proof of Corollary 18.30 and a counterexample to Theorem 18.29 in atomic instances can be found in Roughgarden (2005a). For extensions of Theorem 18.29 to networks with restricted cost functions, including a solution to Exercise 18.8(e), see Chakrabarty (2004) and Correa et al. (2005).

18.6.2 Atomic Selfish Routing

Atomic selfish routing games were first considered by Rosenthal (1973), who proved Theorem 18.12 with the potential function method. Rosenthal also introduced the concept of “congestion games” (Remark 18.2). Monderer and Shapley (1996) undertook a more general study of “potential games” – games that admit a potential function, which in turn can be used to prove that best-response dynamics converge to an equilibrium (Remark 18.13). Potential games are now studied in their own right; see Voorneveld et al. (1999) and Roughgarden (2005a, Section 4.8) for surveys of this literature.

Rosenthal (1973) showed that equilibrium flows need not exist in weighted multi-commodity atomic instances. Example 18.7 is due to Goemans et al. (2005). Fotakis et al. (2005) proved Theorem 18.15 for weighted instances with affine cost functions.

The price of anarchy of atomic instances was first studied by Suri et al. (2007) in the context of the asymmetric scheduling games described in Exercise 18.3 below. Among other results, they proved an upper bound of $5/2$ on the price of anarchy in such games when each player controls one unit of traffic and when all cost functions are affine. This paper also introduced the proof structure used to prove Theorem 18.23 in this chapter.

Awerbuch et al. (2005) significantly generalized the results in Suri et al. (2007). They proved Theorem 18.23, as well as the refinement discussed in Exercise 18.7. The AAE example and the variant in Exercise 18.2(a) are from Awerbuch et al. (2005), as are the exponential (in the degree bound p) upper and lower bounds on the price of anarchy for polynomial cost functions with nonnegative coefficients. For refined versions of these upper and lower bounds, see Olver (2006). Awerbuch et al. (2005) extended all of their upper bounds to mixed-strategy Nash equilibria. Goemans et al. (2005) extended the upper bounds to “sink equilibria,” a notion of equilibrium that is motivated by best-response dynamics and that always exists in finite noncooperative games.

For unweighted instances and pure-strategy equilibrium flows, the results in Awerbuch et al. (2005) were obtained independently by Christodoulou and Koutsoupias (2005b). The proofs in Christodoulou and Koutsoupias (2005b) extend without much difficulty to weighted instances and mixed-strategy Nash equilibria. Christodoulou and Koutsoupias (2005b) also studied the price of anarchy with respect to the egalitarian objective (see Section 17.1) and provide solutions to parts (b) and (c) of Exercise 18.2.

Caragiannis et al. (2006) provide a solution to Exercise 18.3(b), as well as numerous other results about the price of anarchy and stability in different classes of asymmetric scheduling instances. For results on the price of stability in atomic selfish routing

games, see Anshelevich et al. (2004), Christodoulou and Koutsoupias (2005a), and Caragiannis et al. (2006).

Finally, several researchers have studied selfish routing in the *atomic splittable* model. This model is similar to the atomic selfish routing games studied in this chapter; the key difference is that a player i is permitted to route its r_i units of traffic *fractionally* over the s_i-t_i paths of the network. This model is also different from nonatomic selfish routing games; for example, if there is only one player controlling all of the traffic in the network, then the player will minimize its cost by routing this traffic optimally. More generally, a player takes into account the congestion it causes for its own traffic, while ignoring the congestion it creates for other players.

Equilibrium flows in the atomic splittable model can behave in counterintuitive ways (see Exercise 18.9, taken from Catoni and Pallottino, 1991), and the price of anarchy in this model is not well understood. It was initially claimed that the upper bounds on the price of anarchy for nonatomic instances carry over to atomic splittable ones (Roughgarden, 2005b; Correa et al., 2005), but Cominetti et al. (2006) recently gave counterexamples to these claims in multicommodity networks. Obtaining tight bounds on the price of anarchy in this model remains an important open question.

Bibliography

- E. Anshelevich, A. Dasgupta, J. Kleinberg, É. Tardos, T. Wexler, and T. Roughgarden. The price of stability for network design with fair cost allocation. In *Proc. 45th Symp. Fdns. of Computer Science*, pp. 295–304, 2004.
- B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proc. 37th Symp. Theory of Computing*, pp. 57–66, 2005.
- M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- D. P. Bertsekas. *Nonlinear Programming*, 2nd ed. Athena Scientific, 1999.
- D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*, 2nd ed. Prentice-Hall, 1989. Athena Scientific, 1997.
- D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968. English translation in Braess (2005).
- D. Braess. On a paradox of traffic planning. *Transport. Sci.*, 39(4):446–450, 2005.
- D. Braess. <http://homepage.ruhr-uni-bochum.de/Dietrich.Braess/>, Homepage, 2007.
- I. Caragiannis, M. Flammini, C. Kaklamanis, P. Kanellopoulos, and L. Moscardelli. Tight bounds for selfish and greedy load balancing. In *Proc. 33rd Annual Intl. Colloq. in Automata, Languages, and Programming*, LNCS 4051:311–322, 2006.
- S. Catoni and S. Pallottino. Traffic equilibrium paradoxes. *Transport. Sci.*, 25(3):240–244, 1991.
- D. Chakrabarty. Improved bicriteria results for the selfish routing problem. Unpublished manuscript, 2004.
- G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. In *Proc. 13th Euro. Symp. on Algorithms (ESA)*, pp. 59–70, 2005a.
- G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proc. 37th Symp. on Theory of Computing*, pp. 67–73, 2005b.
- J. E. Cohen and P. Horowitz. Paradoxical behavior of mechanical and electrical networks. *Nature*, 352(8):699–701, 1991.
- R. Cominetti, J. R. Correa, and N.E.S. Moses. Network games with atomic players. In *Proc. 33rd Intl. Colloq. in Automata, Languages, and Programming*, LNCS 4051:525–536, 2006.

- J. R. Correa, A. S. Schulz, and N.E.S. Moses. Selfish routing in capacitated networks. *Math. Operat. Res.*, 29(4):961–976, 2004.
- J. R. Correa, A. S. Schulz, and N.E.S. Moses. On the inefficiency of equilibria in congestion games. In *Proc. 11th Conf. on Integer Programming and Combinatorial Optimization*, pp. 167–181, 2005.
- D. Fotakis, S. C. Kontogiannis, and P. G. Spirakis. Selfish unsplittable flows. *Theor. Comput. Sci.*, 348(2–3):226–239, 2005.
- M. X. Goemans, V. S. Mirrokni, and A. Vetta. Sink equilibria and convergence. In *Proc. 46th Symp. on Foundations of Computer Science*, pp. 142–151, 2005.
- D. Monderer and L. S. Shapley. Potential games. *Games Econ. Behav.*, 14(1):124–143, 1996.
- N. Olver. *The Price of Anarchy and a Priority-Based Model of Routing*. M.S. thesis, McGill University, 2006.
- A. C. Pigou. *The Economics of Welfare*. Macmillan, 1920.
- L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker. On selfish routing in Internet-like environments. In *Proc. SIGCOMM*, pp. 151–162, 2003.
- R. W. Rosenthal. The network equilibrium problem in integers. *Networks*, 3(1):53–59, 1973.
- T. Roughgarden. The price of anarchy is independent of the network topology. *J. Comput. System Sci.*, 67(2):341–364, 2003.
- T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005a.
- T. Roughgarden. Selfish routing with atomic players. In *Proc. 16th Symp. Discrete Algorithms*, pp. 1184–1185, 2005b.
- T. Roughgarden. On the severity of Braess’s Paradox: Designing networks for selfish users is hard. *J. Computer System Sci.*, 72(5):922–953, 2006.
- T. Roughgarden and É. Tardos. How bad is selfish routing? *J. ACM*, 49(2):236–259, 2002.
- D. Schmeidler. Equilibrium points of nonatomic games. *J. Statist. Phys.*, 7(4):295–300, 1973.
- M. J. Smith. The existence, uniqueness and stability of traffic equilibria. *Transport. Res., Part B*, 13(4):295–304, 1979.
- S. Suri, C. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 47(1): 79–96, 2007.
- M. Voorneveld, P. Borm, F. van Meegen, S. Tijs, and G. Facchini. Congestion games and potentials reconsidered. *Intl. Game Theory Rev.*, 1(3–4):283–299, 1999.
- J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proc. Institute of Civil Engineers, Pt. II*, volume 1, pp. 325–378, 1952.

Exercises

- 18.1** Recall the nonlinear variant of Pigou’s example (Example 18.3). Prove that as the degree p of the cost function of the second link tends to infinity, the price of anarchy tends to infinity as $p/\ln p$.
- 18.2** This exercise explores lower bounds on the price of anarchy in atomic selfish routing games with affine cost functions.
- (a) Modify the players’ weights in the AAE example (Example 18.6) so that the price of anarchy in the resulting weighted atomic instance is precisely $(3 + \sqrt{5})/2 \approx 2.618$.
- (b) Can you devise an unweighted atomic instance with 3 players, affine cost functions, and price of anarchy equal to $5/2$? Can you achieve a price of anarchy of $(3 + \sqrt{5})/2$ using 3 players and variable weights?