Selfish Routing and the Price of Anarchy

Tim Roughgarden*

January 7, 2006

Abstract

Selfish routing is a classical mathematical model of how self-interested users might route traffic through a congested network. The outcome of selfish routing is generally inefficient, in that it fails to optimize natural objective functions. The price of anarchy is a quantitative measure of this inefficiency.

We survey recent work that analyzes the price of anarchy of selfish routing. We also describe related results on bounding the worst-possible severity of a phenomenon called Braess's Paradox, and on three techniques for reducing the price of anarchy of selfish routing. This survey concentrates on the contributions of the author's PhD thesis, but also discusses several more recent results in the area.

1 Introduction

Over the past several years, there has been a tremendous surge of activity at the interface of computer science and economics. This survey is a brief introduction to two intertwined facets of this emerging research area, *the price of anarchy* and *selfish routing*. The price of anarchy, first defined by Koutsoupias and Papadimitriou [59, 69], measures the extent to which competition approximates cooperation. It is motivated by the well-known fact that noncooperative equilibria can be *inefficient*, in that they need not optimize natural objective functions [33, 75]. Selfish routing refers to a mathematical model of traffic in a congested network. This model has a long history in the transportation science literature [8, 13, 72, 100] and has also been widely studied by the computer networking community (see e.g. [11, 16, 42, 43, 67, 73]). The price of anarchy has recently been extensively studied in this model.

This survey concentrates on the contributions of the author's PhD thesis [82], but also discusses several more recent results on the price of anarchy of selfish routing. In most cases, we provide self-contained proofs. Many more details, results, and references can be found in the recent book [86], which is an expanded and revised version of [82].

*Department of Computer Science, Stanford University, 462 Gates Building, 353 Serra Mall, Stanford, CA 94305. Supported in part by ONR grant N00014-04-1-0725, DARPA grant W911NF-04-9-0001, and an NSF CAREER Award. Email: tim@cs.stanford.edu.

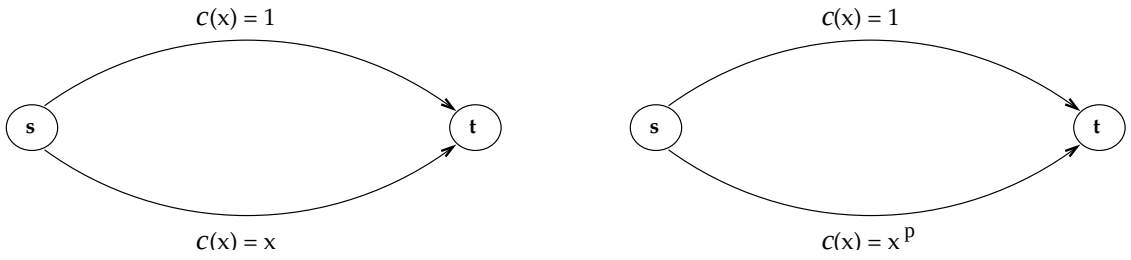

(a) Pigou's example

(b) A nonlinear variant

Figure 1: Pigou's example and a nonlinear variant. The cost function $c(x)$ describes the cost incurred by users of an edge, as a function of the amount of traffic routed on the edge.

1.1 Two Motivating Examples

We now introduce selfish routing and motivate the results described in this survey by informally exploring two important examples. Pigou discovered the first example in 1920 [72]; Braess found the second in 1968 [14].

Example 1.1 (Pigou's example [72]) Consider the simple network shown in Figure 1(a). Two disjoint edges connect a source vertex s to a sink vertex t. Each edge is labeled with a *cost function* $c(\cdot)$, which describes the cost (e.g., travel time) incurred by users of the edge, as a function of the amount of traffic routed on the edge. The upper edge has the constant cost function $c(x) = 1$, and thus represents a route that is relatively long but immune to congestion. The cost of the lower edge, by contrast, is governed by the function $c(x) = x$ and thus increases as the edge gets more congested. In particular, the lower edge is cheaper than the upper edge if and only if less than one unit of traffic uses it.

Suppose there is one unit of traffic, representing a very large population of network users, and that each user chooses independently between the two routes from s to t. Assuming that each network user aims to minimize its cost, we should expect all traffic to follow the lower edge. Indeed, each network user should reason as follows: the lower route is never worse than the upper one, even when it is fully congested, and it is superior whenever some of the other users are foolish enough to take the upper route. In the "selfish routing outcome", we therefore expect all networks users to incur one unit of cost.

Now suppose that, by whatever means, we can choose how the traffic is routed. Can we leverage this power to improve over the selfish routing outcome? To see that we can, consider assigning half of the traffic to each of the two routes. The network users forced onto the upper edge experience one unit of cost, and are thus no worse off than in the previous outcome. On the other hand, users permitted on the lower edge now enjoy lighter traffic conditions, and incur a mere $1/2$ unit of cost. We have therefore lowered the cost of half of the users while making no one worse off. Moreover, the average cost incurred by traffic has decreased from 1 to $3/4$.

Pigou's example demonstrates that selfish routing need not produce an optimal outcome. This phenomenon can be amplified with a seemingly minor modification to Example 1.1. Suppose we replace the previously linear cost function $c(x) = x$ with the highly nonlinear one $c(x) = x^p$ for p large (Figure 1(b)). As in Example 1.1, selfish users will all travel on the lower route, incurring a cost of 1. On the other hand, if we could force a small ϵ fraction of the traffic to travel along the upper route, then the average cost would drop to $\epsilon + (1-\epsilon)^{p+1}$, which approaches 0 as ϵ tends to 0 and p tends to infinity.

In Section 2, we will define the price of anarchy of selfish routing as the average cost of traffic in a selfish outcome divided by the minimum-possible average cost. If the price of anarchy of a network is close to 1, then we conclude that the negative impact of selfish routing is relatively small. The price of anarchy in Example 1.1 is at least 4/3, and it tends to infinity with p in the nonlinear variant of Pigou's example.

The price of anarchy of selfish routing can therefore be large if the network cost functions are "sufficiently nonlinear". Pigou's example and its nonlinear variant motivate the following questions, which are central to Section 2 of this survey. Can the price of anarchy be large even when cost functions are "not too nonlinear"? Is the price of anarchy larger in bigger, more complicated networks? Is it larger in multicommodity networks, where traffic emanates from and terminates at multiple locations? In Section 2 we will prove that the answer to all of these questions is "no"—in fact, Pigou's example and simple variants are in some sense universal bad examples for the price of anarchy of selfish routing.

While the price of anarchy in our next example is no larger than in Pigou's example, it is arguably a more startling and unintuitive display of the suboptimality of selfish routing.

Example 1.2 (Braess's Paradox [14]) Consider the four-node network shown in Figure 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. The routes are therefore identical, and selfish traffic should split evenly between them. Assuming that there is one unit of traffic, all network users experience 3/2 units of cost in the selfish routing outcome.

Now suppose that, in an effort to decrease the cost encountered by the traffic, we build a short, high-capacity edge connecting the midpoints of the two existing routes. The new network is shown in Figure 2(b), with the new edge (v, w) possessing the constant cost function $c(x) = 0$. How will selfish traffic react?

We cannot expect the previous traffic pattern to persist in the new network. As in Pigou's example, the cost of the new route $s \to v \to w \to t$ is never worse than that along the two original paths, and it is strictly less whenever some traffic fails to use it. We therefore expect all network users to deviate to the new route. Because of the ensuing heavy congestion on the edges (s, v) and (w, t), all of the traffic now 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!

Example 1.2 shows that adding a new edge to a network can increase the cost incurred by selfish traffic. Equivalently, *removing* one edge from a network with linear cost functions can decrease this cost by a factor of at least 4/3. Can removing edges from a network decrease the cost incurred by selfish traffic by a larger factor in larger networks, or with nonlinear

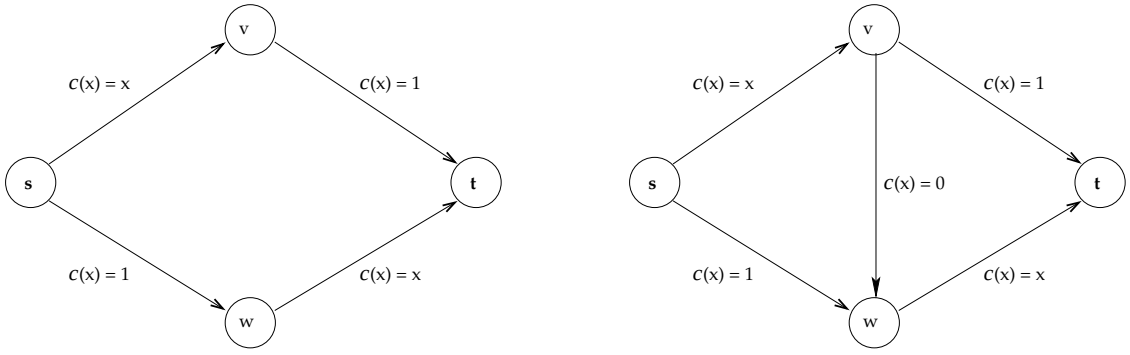

(a) Initial network

(b) Augmented network

Figure 2: Braess's Paradox. The addition of an intuitively helpful edge can adversely affect all of the traffic.

cost functions, or with multiple commodities, or with multiple edge removals allowed? If so, by how much? In Section 3 we give precise answers to all of these questions.

1.2 Overview

We begin in Section 2 by proving matching upper and lower bounds on the price of anarchy of selfish routing. After defining the classical model of selfish routing that we study, we formalize the lower bound on the price of anarchy provided by simple variants of Pigou's example. As suggested by Example 1.1 and the subsequent nonlinear variants, this lower bound will depend on the set of allowable edge cost functions. We then show a matching upper bound for essentially every set of allowable cost functions. For example, the price of anarchy in every multicommodity network with linear cost functions—functions of the form $ax + b$ with $a, b \geq 0$—is at most $4/3$. Thus the price of anarchy in such networks is maximized by Pigou's example (Example 1.1). Similarly, the price of anarchy of multicommodity networks with cost functions that are polynomials with nonnegative coefficients and degree at most p is maximized by the nonlinear variant of Pigou's example shown in Figure 1(b). We also explicitly compute the largest-possible price of anarchy with respect to several different types of cost functions.

Section 3 studies the worst-possible severity of Braess's Paradox. We show that Braess's Paradox can be arbitrarily severe, even in single-commodity networks, provided nonlinear cost functions, large networks, and multiple edge removals are permitted. Precisely, for every $n \geq 2$, there is a single-commodity, n-vertex network such that removing $\lfloor n/2 \rfloor - 1$ edges decreases the cost incurred by selfish traffic by an $\lfloor n/2 \rfloor$ factor. We also show that this construction is optimal in several senses, discuss extensions to multicommodity networks, and show that Braess's Paradox is impossible to detect efficiently (assuming $P \neq NP$).

Section 4 tackles the problem of reducing the price of anarchy in networks where it is unacceptably high. It surveys positive results for three distinct approaches: increasing the

network capacity, routing a small portion of the traffic centrally, and influencing network users by taxing network edges.

Finally, Section 5 describes the broader research context for the results of this survey, and discusses other recent work that quantifies the inefficiency of noncooperative equilibria.

2 Bounding the Price of Anarchy

This section formally defines selfish routing networks, equilibria, and the price of anarchy (Subsection 2.1); introduces a simple lower bound on the price of anarchy that is based on Pigou's example (Subsection 2.2); and proves a matching upper bound on the price of anarchy (Subsection 2.3).

2.1 Preliminaries

Selfish Routing Networks

We begin by reviewing the terminology of classical multicommodity flow networks. See [2], for example, for more details and for historical notes on network flows. A multicommodity flow network is described by a directed graph $G = (V, E)$, with vertex set V and edge set E, and a set $(s_1, t_1), \ldots, (s_k, t_k)$ of source-sink vertex pairs, also called *commodities*. Parallel edges are allowed, and a vertex can participate in multiple source-sink pairs.

For a multicommodity network G, let $\mathcal{P}_i$ denote the set of simple s_i-t_i paths and $\mathcal{P}$ the union $\cup_{i=1}^{k} \mathcal{P}_i$. We always assume that $\mathcal{P}_i \neq \emptyset$ for every i. A *flow* in G is a nonnegative vector, indexed by $\mathcal{P}$. 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 navigate from s_i to t_i. A flow f induces a *flow on edges* $\{f_e\}_{e \in E}$, where $f_e = \sum_{P \in \mathcal{P}: e \in P} f_P$ denotes the total amount of flow using the edge e. Finally, we use r to denote a nonnegative vector of *traffic rates*, indexed by the commodities of G. A flow f in G is *feasible* for r if it routes all of the prescribed traffic: for each $i \in \{1, 2, \ldots, k\}$, $\sum_{P \in \mathcal{P}_i} f_P = r_i$.

To model the negative consequences of increasing congestion, we give each edge e of a network G a nonnegative, continuous, and nondecreasing *cost function* c_e. A cost function $c_e(\cdot)$ denotes the cost (e.g. travel time) incurred by traffic that traverses edge e, as a function of the edge congestion f_e. A selfish routing network is then given by a triple of the form (G, r, c), where G is a multicommodity flow network, r is a vector of traffic rates, and c is a vector of cost functions, indexed by the edges of G. We often call such a triple an *instance*.

Equilibria

We next discuss equilibria in selfish routing networks. Let f be a flow feasible for the instance (G, r, c). The overall cost $c_P(f)$ incurred by traffic on the path P in the flow f is defined as the sum of the costs of the constituent edges: $c_P(f) = \sum_{e \in P} c_e(f_e)$. Naturally, we expect selfish traffic to attempt to minimize its cost. This leads to the following definition, which was first formulated by Wardrop [100].

Definition 2.1 ([100]) Let f be a feasible flow for the instance (G, r, c). The flow f is a *Wardrop equilibrium* if, for every commodity $i \in \{1, 2, \ldots, k\}$ 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 a Wardrop equilibrium 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 a Wardrop equilibrium have equal cost. In the theoretical computer science literature, Wardrop equilibria are also called *Nash flows*. Haurie and Marcotte [47] formalized the precise correspondence between Wardrop equilibria and Nash equilibria of finite normal-form games [66].

Remark 2.2 In Definition 2.1, we are implicitly assuming that every network user controls a negligible portion of the overall traffic, so that the actions of an individual user have essentially no effect on the network congestion. In the game theory literature, games with this property are called *nonatomic* [91]. Several recent papers have analyzed the price of anarchy in atomic variants of the selfish routing model studied in this survey; see Section 5 for references.

Beckmann, McGuire, and Winsten [8] resolved the important issues of existence and uniqueness of Wardrop equilibria.

Proposition 2.3 ([8]) *Let (G, r, c) be an instance.*

(a) The instance (G, r, c) admits at least one Wardrop equilibrium.

(b) If f and $\tilde{f}$ are Wardrop equilibria for (G, r, c), then $c_e(f_e) = c_e(\tilde{f}_e)$ for every edge e.

The first part of Proposition 2.3 guarantees that a Wardrop equilibrium exists in every instance. The second part states that every two Wardrop equilibria induce identical edge costs. While two Wardrop equilibria need not induce identical flows on edges, Proposition 2.3(b) is strong enough for our purposes.

The proof of Proposition 2.3 in [8] is remarkable. Beckmann, McGuire, and Winsten [8] showed, by invoking the Karush-Kuhn-Tucker conditions (see e.g. [71]), that the Wardrop equilibria of an instance (G, r, c) are precisely the flows that minimize the *potential function*

$$\Phi(f) = \sum_{e \in E} \int_0^{f_e} c_e(x)dx \tag{1}$$

over all feasible flows for (G, r, c). Since cost functions are continuous and the space of all flows is compact, Weierstrass's Theorem then implies Proposition 2.3(a). Since cost functions are nondecreasing, the function Φ is convex, and Proposition 2.3(b) then follows without much difficulty. This use of a potential function has been influential in both game theory and theoretical computer science. Led by the work of Rosenthal [77] and Monderer and Shapley [65], potential functions have become a standard tool in noncooperative game theory for proving the existence of pure-strategy Nash equilibria in certain classes of games. In theoretical computer science, potential functions have been used to bound the price of

anarchy in several applications [3, 52, 88, 89]. Intuitively, if equilibria optimize a potential function that is "close to" the objective function, then equilibria cannot be too inefficient. Indeed, the proximity between the potential function Φ in (1) and our objective function (3) below implies near-optimal upper bounds on the price of anarchy of selfish routing [89]. In this survey we focus only on optimal bounds, however, which follow from a different proof approach.

Speaking of which, the following *variational inequality* characterization of Wardrop equilibria, due to Smith [93], will play a crucial role in our upper bound on the price of anarchy.

Proposition 2.4 ([93]) *A flow f feasible for (G, r, c) is a Wardrop equilibrium 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).

Proposition 2.4 can easily be derived as an optimality condition for minimizers of the potential function (1). For simplicity, we instead give a short direct proof.

Proof: Definition 2.1 easily implies that a flow f is a Wardrop equilibrium if and only if

$$\sum_{P \in \mathcal{P}} c_P(f) f_P \leq \sum_{P \in \mathcal{P}} c_P(f) f_P^* \tag{2}$$

for every flow f^* feasible for (G, r, c). Writing $c_P(f) = \sum_{e \in P} c_e(f_e)$ and reversing the order of summation on both sides of (2) then proves the proposition. ■

The Price of Anarchy

We conclude the preliminaries by defining the price of anarchy. Since this definition aims to quantify the inefficiency of an equilibrium, it requires an objective function. We adopt the usual objective function from min-cost network flow, and define the *cost* $C(f)$ of a flow f in (G, r, c) as

$$C(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P = \sum_{e \in E} c_e(f_e) f_e. \tag{3}$$

The first equality in (3) is a definition; the second follows from the same reversal of sums as in the proof of Proposition 2.4. A flow feasible for an instance (G, r, c) is *optimal* if it minimizes the cost over all feasible flows. Because cost functions are continuous and the space of flows is compact, every instance admits an optimal flow.

We now define the price of anarchy as the ratio between the cost of a Wardrop equilibrium and of an optimal flow.

Definition 2.5 ([59, 69]) The *price of anarchy* $\rho(G, r, c)$ of an instance (G, r, c) is

$$\rho(G, r, c) = \frac{C(f)}{C(f^*)},$$

where f is a Wardrop equilibrium and f^* is an optimal flow for (G, r, c). The *price of anarchy* $\rho(\mathcal{I})$ of a non-empty set $\mathcal{I}$ of instances is $\sup_{(G,r,c) \in \mathcal{I}} \rho(G, r, c)$.

Definition 2.1 and Proposition 2.3(b) easily imply that all Wardrop equilibria have equal cost, and thus the price of anarchy of an instance is well defined unless there is a flow with zero cost. In this case, all Wardrop equilibria also have zero cost, and we define the price of anarchy of the instance to be 1.

2.2 The Pigou Bound

Definition of the Pigou Bound

Pigou's example and its nonlinear variant (Subsection 1.1) show that the price of anarchy of selfish routing depends, at the very least, on the type of cost functions allowed. We will therefore aim for a bound on the price of anarchy that is parameterized by the set of allowable cost functions, and that is optimal for each such set. Common examples of sets of cost functions include linear functions, polynomials, and queueing delay functions.

For every set $\mathcal{C}$ of allowable cost functions, Pigou-like examples provide a natural lower bound on the price of anarchy of instances with cost functions in $\mathcal{C}$. Specifically, suppose $\mathcal{C}$ contains all of the constant cost functions, and choose a cost function $c_2 \in \mathcal{C}$ and a traffic rate $r \geq 0$. Let $c_1 \in \mathcal{C}$ denote the cost function everywhere equal to $c_2(r)$. Consider the usual two-node, two-link network of Pigou's example (Figure 1), give the upper and lower edges the cost functions c_1 and c_2, respectively, and set the traffic rate to be r. Routing all traffic on the lower edge yields a Wardrop equilibrium with cost $c_2(r)r$. The price of anarchy in this instance is thus

$$\max_{0 \leq x \leq r} \frac{r \cdot c_2(r)}{x \cdot c_2(x) + (r-x)c_2(r)}.$$

Definition 2.6 below uses this expression but does not constrain x from above by r; since c_2 is nondecreasing, this modification does not affect the value of the maximum.

We can now obtain a lower bound on the price of anarchy by choosing the cost function c_2 and the traffic rate r in the most pernicious way possible.

Definition 2.6 ([27, 83]) Let $\mathcal{C}$ be a nonempty set of cost functions. The *Pigou bound* $\alpha(\mathcal{C})$ for $\mathcal{C}$ is

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{x,r \geq 0} \frac{r \cdot c(r)}{x \cdot c(x) + (r-x)c(r)}, \tag{4}$$

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

Examples

While the defining equation (4) of the Pigou bound may appear fearsome to evaluate, it simplifies to a closed-form expression for many interesting sets of cost functions.

Example 2.7 ([83, 89]) If $\mathcal{C} = \{ax + b : a, b \geq 0\}$ is the set of linear cost functions, then elementary calculations show that $\alpha(\mathcal{C}) = 4/3$.

Thus Example 1.1 determines the Pigou bound for linear cost functions.

Example 2.8 ([27]) Similarly, if $\mathcal{C}$ is the set of concave cost functions, then $\alpha(\mathcal{C}) = 4/3$.

Example 2.9 ([83]) If $\mathcal{C}$ is the set of polynomials with nonnegative coefficients and degree at most p, then

$$\alpha(\mathcal{C}) = [1 - p \cdot (p+1)^{-(p+1)/p}]^{-1}. \tag{5}$$

As p grows large, the right-hand side of (5) tends to infinity as $p/\ln p$ [83, 99].

The right-hand side of (5) is simply the price of anarchy in the nonlinear variant of Pigou's example discussed in Subsection 1.1. The Pigou bound for (nondecreasing) bounded-degree polynomials with arbitrary coefficients is not well understood, though partial results have recently been obtained by So [94].

Remark 2.10 One of the most popular types of cost functions in transportation science applications is quartic functions with nonnegative coefficients (see e.g. Sheffi [92]). The Pigou bound (5) for such functions is roughly 2.15.

Our final example is for the delay functions of M/M/1 queues—queues with Poisson arrivals and exponentially distributed service times—which are common in computer network applications (see e.g. [10, 11]). These delay functions correspond to cost functions of the form $c(x) = 1/(u-x)$, where u can be interpreted as an edge capacity or a queue service rate. The value of such a function is defined to be $+\infty$ when $x \geq u$. Allowing infinite costs requires some technical modifications to the selfish routing model that we ignore in this survey; see [82] for more details.

The Pigou bound for the set of M/M/1 delay functions is $+\infty$ [42]. Intuitively, this follows from Example 2.9 because an M/M/1 delay function behaves like a polynomial with arbitrarily large degree when it is nearly saturated. In analogy to restricting the polynomial degree in Example 2.9, we impose a lower bound u_{min} on all queue service rates and an upper bound R_{max} on the value that the traffic rate r can take on in (4).

Example 2.11 ([83]) Suppose $R_{max} < u_{min}$ and let $\mathcal{C} = \{(u-x)^{-1} : u \geq u_{min}\}$ be the set of M/M/1 delay functions with service rate at least u_{min}. Let $\alpha(\mathcal{C})$ denote the largest-possible price of anarchy in Pigou-like networks with cost functions in $\mathcal{C}$ and traffic rate at most R_{max}. (Formally, $\alpha(\mathcal{C})$ is given by (4) with the additional restriction that $x, r \leq R_{max}$.) Then

$$\alpha(\mathcal{C}) = \frac{1}{2}\left(1 + \sqrt{\frac{u_{min}}{u_{min} - R_{max}}}\right). \tag{6}$$

The right-hand side of (6) tends to infinity as $R_{max} \to u_{min}$, but is bounded by a constant if R_{max} is at most a constant fraction of u_{min}.

Simple Worst-Case Networks

The Pigou bound uses only simple networks to provide a lower bound on the price of anarchy. Specifically, the next proposition follows immediately from the definition of the bound.

Proposition 2.12 ([83]) *Let $\mathcal{C}$ be a set of cost functions that includes all of the constant functions, and let $\mathcal{I}$ denote the single-commodity instances with a two-node, two-link network and cost functions in $\mathcal{C}$. Then*

$$\rho(\mathcal{I}) \geq \alpha(\mathcal{C}).$$

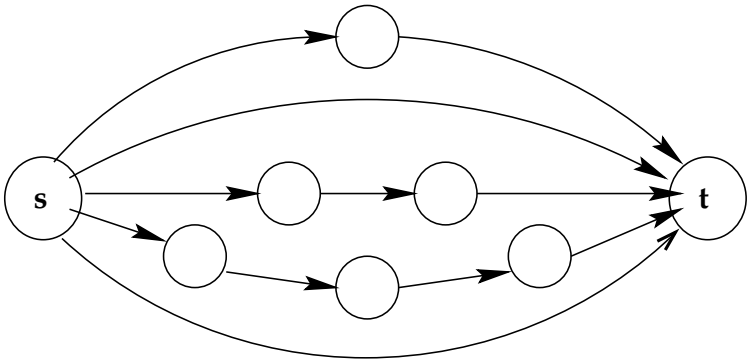

Figure 3: Worst-case networks for inhomogeneous sets of cost functions. The number of paths and the number of edges in each path can be arbitrarily large.

If the set $\mathcal{C}$ does not contain all of the constant cost functions, then we can obtain similar results using modestly more complex networks. For example, suppose the set $\mathcal{C}$ of cost functions is *diverse* in the sense that $\{c(0) : c \in \mathcal{C}\} = [0, \infty)$. Then an edge with the constant cost function $c(x) = a$ can, for all practical purposes, be "simulated" by a large number of parallel edges that each have a cost function satisfying $c(0) = a$. This observation means that a Pigou-like network with a constant cost function can be replaced by a network with two nodes, (an unrestricted number of) parallel links, and cost functions in $\mathcal{C}$ without affecting the price of anarchy.

Proposition 2.13 ([83]) *Let $\mathcal{C}$ be a diverse set of cost functions, and let $\mathcal{I}$ denote the single-commodity instances with a network of parallel links and cost functions in $\mathcal{C}$. Then*

$$\rho(\mathcal{I}) \geq \alpha(\mathcal{C}).$$

The set of cost functions in Example 2.11 is not diverse when $u_{min} > 0$, but it is *inhomogeneous* in the sense that it contains a function c with $c(0) > 0$. As in Proposition 2.13, the Pigou bound remains valid for such sets of cost functions provided we allow somewhat more complex networks. Specifically, let a *union of paths* mean a network with one source, one sink, and an arbitrarily large number of internally vertex-disjoint paths directed from the source to the sink (Figure 3).

Proposition 2.14 ([83]) *Let $\mathcal{C}$ be an inhomogeneous set of cost functions, and let $\mathcal{I}$ denote the single-commodity instances with a network that is a union of paths and with cost functions in $\mathcal{C}$. Then*

$$\rho(\mathcal{I}) \geq \alpha(\mathcal{C}).$$

The idea of the proof of Proposition 2.14 is to impose diversity by considering the closure $\overline{\mathcal{C}}$ of $\mathcal{C}$ under multiplication by positive scalars, apply Proposition 2.13, and use multiple copies of edges with cost functions in $\mathcal{C}$ to simulate edges with cost functions in $\overline{\mathcal{C}}$.

Remark 2.15 The Pigou bound does not apply to homogeneous sets $\mathcal{C}$ of cost functions, where $c(0) = 0$ for all $c \in \mathcal{C}$. The price of anarchy of selfish routing is not completely

understood for such sets; the upper bound in the next subsection holds for these sets, but it is not optimal. See [28] for refined upper bounds on the price of anarchy with respect to homogeneous sets of sufficiently low-degree polynomials.

2.3 Optimality of the Pigou Bound

With all the preliminaries in place, we can now easily prove an upper bound on the price of anarchy of selfish routing that matches the Pigou bound. For convenience, we first state a lemma that follows immediately from Definition 2.6.

Lemma 2.16 *Let $\mathcal{C}$ be a set of cost functions and $\alpha(\mathcal{C})$ the Pigou bound for $\mathcal{C}$. For $c \in \mathcal{C}$ and $x, r \geq 0$,*

$$x \cdot c(x) \geq \frac{r \cdot c(r)}{\alpha(\mathcal{C})} + (x - r)c(r).$$

We now use this lemma and the variational inequality of Proposition 2.4 to prove the optimality of the Pigou bound.

Theorem 2.17 ([27, 83]) *Let $\mathcal{C}$ be a set of cost functions and $\alpha(\mathcal{C})$ the Pigou bound for $\mathcal{C}$. If (G, r, c) is an instance with cost functions in $\mathcal{C}$, then*

$$\rho(G, r, c) \leq \alpha(\mathcal{C}).$$

Proof: Let f^* and f be an optimal flow and a Wardrop equilibrium, respectively, for an 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 for the first inequality we have applied Lemma 2.16 to each edge e with $x = f_e^*$ and $r = f_e$, and the second inequality follows from Proposition 2.4. ■

Theorem 2.17 implies that the lower bounds on the price of anarchy in Examples 2.7–2.11 are the best possible. Thus the price of anarchy of networks with linear (or concave) cost functions is precisely 4/3; the price of anarchy of networks with cost functions that are polynomials with nonnegative coefficients and degree at most p is precisely the right-hand side of (5); and the price of anarchy of instances with sum of all traffic rates at most R_{max}, cost functions that are M/M/1 delay functions, and service rates bounded below by $u_{min} > R_{max}$ is precisely the right-hand side of (6).

Moreover, since the Pigou bound is based only on the simplest of instances, the matching upper bound of Theorem 2.17 implies that simple networks always furnish worst-possible examples of the inefficiency of selfish routing. Precisely, Propositions 2.12–2.14 and Theorem 2.17 give the following corollary.

Corollary 2.18 *Let $\mathcal{C}$ be a set of cost functions.*

(a) If $\mathcal{C}$ contains the constant functions, then the price of anarchy of instances with cost functions in $\mathcal{C}$ is achieved, up to an arbitrarily small factor, by a single-commodity instance with a two-node, two-link network.

(b) If $\mathcal{C}$ is diverse, then the price of anarchy of instances with cost functions in $\mathcal{C}$ is achieved, up to an arbitrarily small factor, by a single-commodity instance with a network of parallel links.

(c) If $\mathcal{C}$ is inhomogeneous, then the price of anarchy of instances with cost functions in $\mathcal{C}$ is achieved, up to an arbitrarily small factor, by a single-commodity instance with a network that is a union of paths.

Informally, Corollary 2.18 states that the price of anarchy is controlled *only* by the set of allowable cost functions, and is essentially independent of the number of commodities and of the complexity of the allowable network topologies.

Remark 2.19 Theorem 2.17 has undergone several iterations in just a few short years. It was first proved for the special case of linear cost functions in Roughgarden and Tardos [89]. Roughgarden [79] then proved Theorem 2.17 for bounded-degree polynomials with nonnegative coefficients. The proof in [79] was fairly complex and did not explicitly take advantage of the variational inequality given in Proposition 2.4. Roughgarden [81] extended this proof and established Theorem 2.17 for all sets of cost functions that satisfy a weak technical condition (met by essentially all cost functions that arise in applications). Ronen [76] pointed out that Proposition 2.4 could be used to vastly simplify the proof of Theorem 2.17, under the same technical condition. This revised analysis appears in [83]. Correa, Schulz, and Stier Moses [27] then showed that, once the proof is based on Proposition 2.4, it can be modified so that no technical conditions whatsoever are needed. The proof of Theorem 2.17 given above is taken from [27]. More recently, two more proofs of Theorem 2.17 have been given by Tardos [96] and Correa, Schulz, and Stier Moses [28].

3 Bounding Braess's Paradox

This section studies the worst-possible severity of Braess's Paradox. Subsection 3.1 gives a construction that shows that the severity of Braess's Paradox can grow with the network size when nonlinear cost functions and multiple edge removals are permitted. Subsection 3.2 proves matching upper bounds for single-commodity networks. Subsection 3.3 gives a brief overview of Braess's Paradox in multicommodity networks. Finally, Subsection 3.4 presents negative results for the computational problem of efficiently detecting Braess's Paradox.

3.1 A Bigger Braess's Paradox

The discovery of Braess's Paradox [14] immediately intrigued researchers and catalyzed numerous research directions (see [78] for a survey). However, nearly all of this work focused

on Braess's original four-node network (Figure 2) and variants thereof. We next show that Braess's original example is merely the tip of the iceberg: Braess's Paradox can be arbitrarily severe in large single-commodity networks.

We measure the severity of Braess's Paradox with the *Braess ratio*—the factor by which the cost of a Wardrop equilibrium exceeds that of an equilibrium in a subnetwork.

Definition 3.1 The *Braess ratio* $\beta(G, r, c)$ of a single-commodity instance (G, r, c) is

$$\beta(G, r, c) = \max_{H \subseteq G} \frac{C(f)}{C(f^H)}, \tag{7}$$

where H ranges over subnetworks of G that contain an s-t path, and f and f^H denote Wardrop equilibria for (G, r, c) and (H, r, c), respectively.

As with Definition 2.5, the Braess ratio of an instance (G, r, c) is well defined unless it admits a flow with zero cost, in which case we define $\beta(G, r, c)$ to be 1.

Remark 3.2 For now, we only define the Braess ratio for single-commodity networks. There are multiple ways to extend Definition 3.1 to multicommodity networks; see Subsection 3.3 for details.

The Braess ratio in Example 1.2 is 4/3. No larger Braess ratio is possible in single-commodity networks with linear cost functions. This fact is a consequence of the following close connection between the price of anarchy and the Braess ratio.

Proposition 3.3 *If (G, r, c) is a single-commodity instance, then*

$$\beta(G, r, c) \leq \rho(G, r, c).$$

Proof: For every subgraph H of G, a Wardrop equilibrium f^H of (H, r, c) is a feasible flow for (G, r, c); by the definition of the price of anarchy, the cost of f^H is at most a $\rho(G, r, c)$ factor less than that of a Wardrop equilibrium for (G, r, c). ■

As promised, Theorem 2.17 and Proposition 3.3 imply that every single-commodity instance with linear cost functions has a Braess ratio of at most 4/3. The upper bound in Proposition 3.3 is also tight, up to constant factors, for many other types of cost functions (see Remark 3.6 below).

Exhibiting a family of instances with arbitrarily large Braess ratios requires a new, more complicated construction than those we have seen so far. Proposition 3.3 implies that such a family must make use of cost functions drawn from a sufficiently rich set (such as polynomials with unbounded degree). We encountered one such family in the nonlinear variant of Pigou's example (Subsection 1.1), but it is easy to see that all of these instances have a Braess ratio of 1. There is also an analogous nonlinear variant of Example 1.2, obtained by replacing the linear cost functions on the edges (s, v) and (w, t) with the functions $c(x) = x^p$ for p large. The Braess ratios of these instances approach 2 as $p \to \infty$. As we will see in Subsection 3.2, a Braess ratio larger than 2 cannot arise without allowing larger networks and multiple edge removals.

Our main result in this subsection is the following.

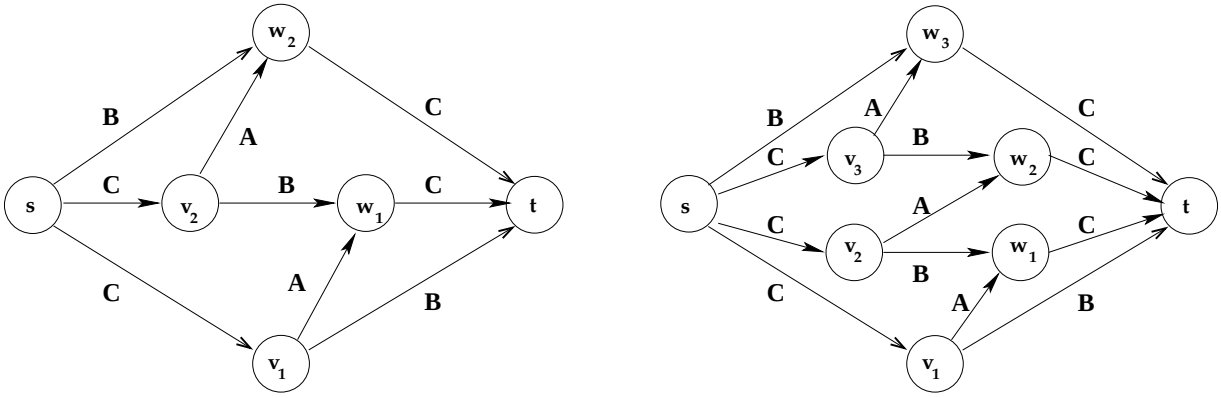

(a) B^2 (b) B^3

Figure 4: The second and third Braess graphs. Edges are labeled with their types.

Theorem 3.4 ([78]) *For every $n \geq 2$, there is a single-commodity instance (G, r, c) with n vertices and*

$$\beta(G, r, c) \geq \left\lfloor \frac{n}{2} \right\rfloor.$$

Proof: We can assume that n is even, since the odd case reduces to the even case by adding an isolated vertex. We can also assume that n is at least 4. Write $n = 2k + 2$ for a positive integer k.

We next define the kth *Braess graph* B^k. Start with a set of $2k + 2$ vertices $V^k = \{s, v_1, \ldots, v_k, w_1, \ldots, w_k, t\}$. The edge set E^k is the union of the sets $\{(s, v_i), (v_i, w_i), (w_i, t) : 1 \leq i \leq k\}$, $\{(v_i, w_{i-1}) : 2 \leq i \leq k\}$, and $\{(v_1, t)\} \cup \{(s, w_k)\}$ (see Figure 4). Call edges of the form (v_i, w_i) the *type A edges*, edges of the form (v_i, w_{i-1}), (s, w_k), and (v_1, t) the *type B edges*, and edges of the form (s, v_i) and (w_i, t) the *type C edges* (see Figure 4). Note that B^1 is the graph in the original Braess's Paradox (Figure 2(b)).

Define cost functions on the edges of B^k as follows.

(A) Type A edges are given the cost function $c_e^k(x) = 0$.

(B) Type B edges are given the cost function $c_e^k(x) = 1$.

(C) For each $i \in \{1, 2, \ldots, k\}$, the type C edges (w_i, t) and (s, v_{k-i+1}) are given a continuous, nondecreasing cost function $c_e^k(x)$ with $c_e^k(k/(k+1)) = 0$ and $c_e^k(1) = i$.

For $i = 1, \ldots, k$, let P_i denote the path $s \to v_i \to w_i \to t$. For $i = 2, \ldots, k$, let Q_i denote the path $s \to v_i \to w_{i-1} \to t$. Define Q_1 to be the path $s \to v_1 \to t$ and Q_{k+1} the path $s \to w_k \to t$. On one hand, routing one unit of flow on each of $P_1, \ldots, P_k$ yields a Wardrop equilibrium f for (B^k, k, c^k) in which all traffic incurs cost $k+1$ (Figure 5(a)). On the other hand, if H is the subgraph obtained from B^k by deleting the k type A edges, then routing $k/(k+1)$ units of flow on each of $Q_1, \ldots, Q_{k+1}$ yields a Wardrop equilibrium f^H for (H, k, c^k)

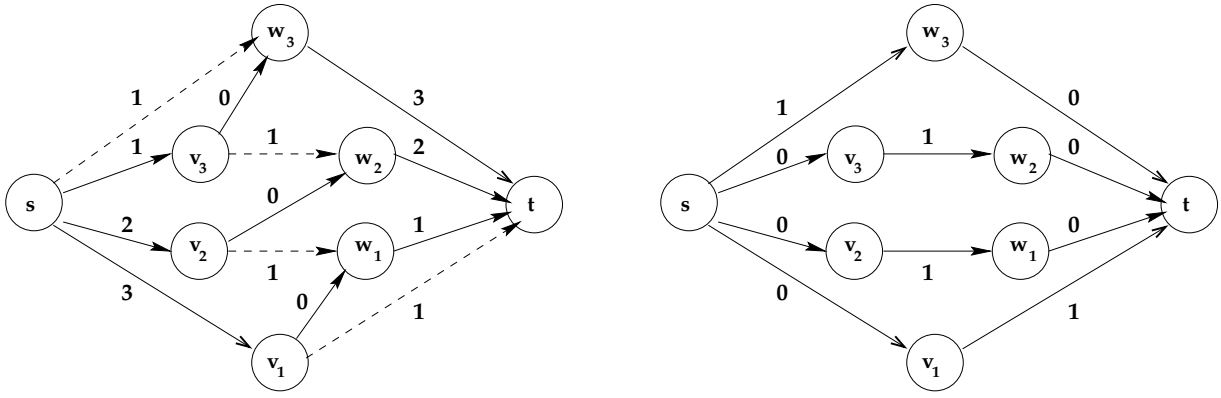

(a) Wardrop equilibrium in $(B^3, 3, c^3)$

(b) Wardrop equilibrium in the optimal subgraph

Figure 5: Proof of Theorem 3.4, when $k = 3$. Solid edges carry traffic in the Wardrop equilibrium, dashed edges do not. Edge costs are with respect to the Wardrop equilibrium.

in which all traffic incurs only one unit of cost (Figure 5(b)). Thus

$$\beta(G, r, c) \geq C(f)/C(f^H) = k + 1 = n/2,$$

completing the proof. ■

Remark 3.5 In the proof of Theorem 3.4, the subgraph H was obtained from B^k by removing k edges. Thus, for every positive integer k, there is a single-commodity instance for which removing k edges can decrease the cost of a Wardrop equilibrium by a factor of $k+1$.

Remark 3.6 The construction in the proof of Theorem 3.4 can also be adapted to scenarios where arbitrary cost functions are not allowed. For example, suppose cost functions are restricted to be polynomials with nonnegative coefficients and degree at most p. Consider the instance (B^k, k, c), where $k \approx p/\ln p$, and where the cost functions c for B^k are identical to those in the proof of Theorem 3.4, except that a type C edge of the form (w_i, t) or (s, v_{k-i+1}) receives the cost function ix^p. Arguing as in the proof of Theorem 3.4 shows that the Braess ratio of (B^k, k, c) is $\Omega(k) = \Omega(p/\ln p)$ as $p \to \infty$. This Braess ratio matches, up to a constant factor, the upper bound for this set of cost functions that follows from Theorem 2.17 and Proposition 3.3. See [78] for more details and further examples.

3.2 A Matching Upper Bound

This subsection shows that, among single-commodity networks, the Braess ratio is maximized by the networks constructed in the proof of Theorem 3.4.

Theorem 3.7 ([78]) *If (G, r, c) is a single-commodity instance with n vertices, then*

$$\beta(G, r, c) \leq \left\lfloor \frac{n}{2} \right\rfloor.$$

Following Lin, Roughgarden, and Tardos [61], we will obtain Theorem 3.7 as a consequence of a more general theorem. The statement of this more general result uses the following definition.

Definition 3.8 Let (G, r, c) be a single-commodity instance and S a subset of the edges of G. The set S is *sparse* if no two edges of S share an endpoint, and in addition no edge of S is incident to s or t.

In other words, a set of edges is sparse if and only if they form an (undirected) matching of $V \setminus \{s, t\}$.

Our general bound on Braess's Paradox states that the size of the largest sparse set removed controls how much the cost of a Wardrop equilibrium can decrease.

Theorem 3.9 ([61]) *Let (G, r, c) be a single-commodity instance, H a subgraph of G, and f and $\tilde{f}$ Wardrop equilibria for (G, r, c) and (H, r, c), respectively. Let S denote the edges in G but not H. If every sparse subset of S contains at most k edges, then*

$$C(f) \leq (k+1) \cdot C(\tilde{f}).$$

Before proving Theorem 3.9, we show that it easily implies Theorem 3.7, as well as an upper bound on the severity of Braess's Paradox that is parameterized by the number of edges removed.

Proof of Theorem 3.7: Since there are only $n-2$ vertices of G that are not s or t, every sparse set of edges has at most $\lfloor (n-2)/2 \rfloor = \lfloor n/2 \rfloor - 1$ edges. Theorem 3.9 now implies the theorem. ■

The next corollary implies that the only way to achieve arbitrarily large Braess ratios is to allow an unlimited number of edge removals, answering a question of Kameda [53].

Corollary 3.10 ([61]) *Removing k edges from a single-commodity network decreases the cost of a Wardrop equilibrium by at most a factor of $k+1$.*

Proof: Obvious from Definition 3.8 and Theorem 3.9. ■

In particular, we noted earlier that simple nonlinear variants on Braess's original example achieve a Braess ratio arbitrarily close to 2; if only a single edge removal is allowed, then no single-commodity instance has a larger Braess ratio. More generally, the construction in the proof of Theorem 3.4 matches the bound of Corollary 3.10 for every k (see Remark 3.5).

We now turn toward the proof of Theorem 3.9. This proof will be more delicate than the upper bounds on the price of anarchy given in Section 2. In particular, our proof must be sensitive to the number of network vertices, whereas two-node networks typically determine the price of anarchy (Corollary 2.18). Because of this, our techniques will have a much stronger combinatorial flavor than those in Section 2.

Because cost functions are arbitrary in Theorem 3.9, we take a modest approach to lower bounding the cost of a Wardrop equilibrium f in the original network relative to that of a Wardrop equilibrium $\tilde{f}$ in a subnetwork. We will identify edges on which $\tilde{f}$ routes at least as much traffic as f. Since cost functions are nondecreasing, the cost incurred by $\tilde{f}$ on these edges is at least that incurred by f. The next definition is largely motivated by this idea.

Definition 3.11 Let f and $\tilde{f}$ be feasible flows for the instance (G, r, c).

(a) An edge e of G is $(f, \tilde{f})$-*light* if $f_e \leq \tilde{f}_e$ and $\tilde{f}_e > 0$, $(f, \tilde{f})$-*heavy* if $f_e > \tilde{f}_e$, and $(f, \tilde{f})$-*useless* if $f_e = \tilde{f}_e = 0$.

(b) An undirected path is $(f, \tilde{f})$-*alternating* if it comprises only forward $(f, \tilde{f})$-light edges and backward $(f, \tilde{f})$-heavy edges.

When the context is clear, we drop the dependence on f and $\tilde{f}$ from the terms in Definition 3.11.

Example 3.12 Consider the Braess's Paradox network (Figure 2(b)). Let f be the Wardrop equilibrium and $\tilde{f}$ the optimal flow, which splits the traffic evenly between the paths $s \to v \to t$ and $s \to w \to t$. Then, edges (s, v), (v, w), and (w, t) are $(f, \tilde{f})$-heavy while edges (s, w) and (v, t) are $(f, \tilde{f})$-light. The unique $(f, \tilde{f})$-alternating s-t path is $s \to w \to v \to t$.

The next lemma states that, for every pair of feasible flows, an s-t alternating path exists. It is an easy consequence of flow conservation arguments.

Lemma 3.13 *Let f and $\tilde{f}$ be flows feasible for the single-commodity instance (G, r, c). Then, there is an $(f, \tilde{f})$-alternating s-t path. Moreover, if f is directed acyclic, then every such path begins and ends with an $(f, \tilde{f})$-light edge.*

Proof: Suppose for contradiction that there is no $(f, \tilde{f})$-alternating s-t path and let S denote the set of nodes reachable from s via such paths. The set S contains s and, by assumption, does not contain t. Since S is an s-t cut, the net f-flow and $\tilde{f}$-flow exiting S is precisely r.

Since vertices in S can be reached from s via $(f, \tilde{f})$-alternating paths and vertices outside S cannot, edges that exit S cannot be $(f, \tilde{f})$-light, and edges that enter S cannot be $(f, \tilde{f})$-heavy. Since the net flow across S is positive (assuming $r > 0$), some non-useless (and thus $(f, \tilde{f})$-heavy) edge exits S. Taken together, these facts imply that the net f-flow exiting S is strictly greater than the net $\tilde{f}$-flow exiting S, a contradiction.

Moreover, if f is directed acyclic, then it sends no flow into s or out of t. Thus, the first and last edges of every $(f, \tilde{f})$-alternating s-t path must be $(f, \tilde{f})$-light. ■

Our proof of Theorem 3.9 will proceed by induction along an alternating path, repeatedly using the shortest-path structure of a Wardrop equilibrium. This structure is summarized by the following characterization of such equilibria.

Lemma 3.14 *Let f be a flow feasible for the single-commodity instance (G, r, c). For a vertex v in G, let $d(v)$ denote the length, with respect to edge lengths $c_e(f_e)$, of a shortest s-v path in G. Then*

$$d(w) - d(v) \leq c_e(f_e)$$

for every edge $e = (v, w)$, and f is a Wardrop equilibrium if and only if equality holds whenever $f_e > 0$.

Lemma 3.14 follows from Definition 2.1 and basic properties of shortest paths.

Lemma 3.14, or the fact that Wardrop equilibria minimize the potential function in (1), easily implies the following strengthening of Proposition 2.3 for single-commodity instances.

Lemma 3.15 *Every single-commodity instance admits a Wardrop equilibrium that is a directed acyclic flow.*

Details of the proofs of Lemmas 3.14 and 3.15 can be found in [78, 86]. We are finally prepared to prove Theorem 3.9.

Proof of Theorem 3.9: Let f be a directed acyclic Wardrop equilibrium for (G, r, c) and $\tilde{f}$ a Wardrop equilibrium for (H, r, c). We view $\tilde{f}$ as a flow in the larger network G in the obvious way. For a vertex v, let $d(v)$ denote the shortest-path distance from s to v with respect to edge lengths $c_e(f_e)$ in G, and $\tilde{d}(v)$ the s-v distance with respect to $c_e(\tilde{f}_e)$ in H. Note that Definition 2.1 and the definition of cost (3) imply that $C(f) = r \cdot d(t)$ and $C(\tilde{f}) = r \cdot \tilde{d}(t)$, so the theorem reduces to proving that $d(t) \leq (k+1) \cdot \tilde{d}(t)$, where k is the size of some sparse set of edges present in G but not H.

Let P be an $(f, \tilde{f})$-alternating s-t path, which exists by Lemma 3.13. A *segment* of P is a maximal subpath of P that contains only $(f, \tilde{f})$-light or only $(f, \tilde{f})$-heavy edges. Edges that are in G but not H are called *absent*. Since $\tilde{f}_e > 0$ on $(f, \tilde{f})$-light edges, absent edges can only reside in $(f, \tilde{f})$-heavy segments. The key claim is that if v is a vertex at the end of a segment of P, and i (heavy) segments of P between s and v contain an absent edge, then

$$d(v) \leq \tilde{d}(v) + i \cdot \tilde{d}(t). \tag{8}$$

This claim implies the theorem. To see why, first apply (8) to t to obtain

$$d(t) \leq \tilde{d}(t) + k \cdot \tilde{d}(t) = (k+1) \cdot \tilde{d}(t), \tag{9}$$

where k is the number of segments of P that include an absent edge. Inequality (9) reduces the proof of the theorem to exhibiting a sparse set of k absent edges. Since f is a directed acyclic flow, Definition 3.11 and Lemma 3.13 imply that the $(f, \tilde{f})$-heavy segments of P are disjoint from each other and from s and t. Picking one absent edge from each of the k $(f, \tilde{f})$-heavy segments of P that contain one thus provides the desired sparse set.

We now prove (8) by induction on the segments of P. The inequality trivially holds when $v = s$, so suppose it holds for a vertex v that is last on a segment of P, or is the source s. We wish to prove (8) for w, defined as the last vertex on the next segment. Let i denote the number of earlier segments of P that contain at least one absent edge. By the inductive hypothesis, $\tilde{d}(v) \leq d(v) + i \cdot \tilde{d}(t)$.

The inductive step has two cases. For the first case, suppose that the segment between v and w contains at least one absent edge. As absent edges can only be $(f, \tilde{f})$-heavy, this segment comprises only $(f, \tilde{f})$-heavy backward edges. Since there is a path of (heavy) edges from w to v, each carrying f-flow, Lemma 3.14 implies that $d(w) \leq d(v)$. Since the path P begins with an $(f, \tilde{f})$-light edge (Lemma 3.13), $v \neq s$ and there is an $(f, \tilde{f})$-light edge entering

v. Since $\tilde{f}$ routes flow into v, it must route flow from v to t. By Lemma 3.14, $\tilde{d}(v) \leq \tilde{d}(t)$. Combining what we know with the inductive hypothesis, the proof of the inductive step is complete:

$$d(w) \leq d(v) \leq \tilde{d}(v) + i \cdot \tilde{d}(t) \leq (i+1) \cdot \tilde{d}(t) \leq \tilde{d}(w) + (i+1) \cdot \tilde{d}(t).$$

For the second case of the inductive step for (8), suppose that the current segment $Q \subseteq P$ contains no absent edges. We will prove, by induction on the vertices of Q, that

$$d(x) \leq \tilde{d}(x) + i \cdot \tilde{d}(t) \tag{10}$$

for all vertices x of Q. The base case ($x = v$) follows from the outer inductive hypothesis (8). For the (inner) inductive step, suppose $d(x) \leq \tilde{d}(x) + i \cdot \tilde{d}(t)$ for a vertex x of Q and let y denote the next vertex on the segment.

If the edge $(x, y) \in P$ is $(f, \tilde{f})$-light, then $c_e(f_e) \leq c_e(\tilde{f}_e)$ and $\tilde{f}_e > 0$. Since f and $\tilde{f}$ are Wardrop equilibria, Lemma 3.14 and the inductive hypothesis imply that

$$d(y) \leq d(x) + c_e(f_e) \leq \tilde{d}(x) + c_e(\tilde{f}_e) + i \cdot \tilde{d}(t) = \tilde{d}(y) + i \cdot \tilde{d}(t),$$

which establishes (10) for the vertex y.

If the edge $e = (y, x) \in P$ is $(f, \tilde{f})$-heavy, then

$$\begin{aligned} d(y) &= d(x) - c_e(f_e) && (11)\\ &\leq \tilde{d}(x) + i \cdot \tilde{d}(t) - c_e(\tilde{f}_e) && (12)\\ &\leq \tilde{d}(y) + i \cdot \tilde{d}(t), && (13) \end{aligned}$$

where equation (11) follows from Proposition 3.14 and the fact that $f_e > 0$, inequality (12) follows from the inductive hypothesis and the fact that $\tilde{f}_e \leq f_e$, and inequality (13) follows from Proposition 3.14.

In either case, the inner inductive step (10) holds. This completes the proof of the outer inductive step (8) and of the theorem. ■

3.3 Multicommodity Networks

So far, this section has only studied Braess's Paradox in single-commodity networks. We next briefly survey very recent results of Lin et al. [62] on Braess's Paradox in multicommodity networks. We define the Braess ratio for such networks as follows. For a multicommodity instance (G, r, c) and a commodity i, let $d_i(G, r, c)$ denote the common cost incurred by all traffic of commodity i in a Wardrop equilibrium for (G, r, c). Note $d_i(G, r, c)$ is well defined by Definition 2.1 and Proposition 2.3.

Definition 3.16 The *Braess ratio* $\beta(G, r, c)$ of a multicommodity instance (G, r, c) is

$$\beta(G, r, c) = \max_{H \subseteq G} \min_{i=1}^{k} \frac{d_i(G, r, c)}{d_i(H, r, c)},$$

where H ranges over the subnetworks of G that contain an s_i-t_i path for each i.

Thus the Braess ratio of a multicommodity instance is large only if removing some set of edges decreases the cost incurred by the traffic of every commodity by a large amount. Definitions 3.1 and 3.16 coincide in single-commodity networks.

Remark 3.17 Since removing edges from a multicommodity network affects traffic from different commodities in different ways, there are several possible measures for the severity of Braess's Paradox in such networks. For example, one natural measure is given by the same defining equation (7) as in the definition of the Braess ratio for single-commodity networks (Definition 3.1). Unfortunately, while Proposition 3.3 still holds for this measure, no interesting bounds are possible for networks with arbitrary cost functions: even in two-commodity, three-node networks, removing a single edge can decrease the cost of a Wardrop equilibrium by an arbitrarily large factor.

The upper bound on the Braess ratio in Theorem 3.7 does not carry over to multicommodity networks: Lin et al. [62] showed that the Braess ratio can grow exponentially with the network size, even in two-commodity networks.

Theorem 3.18 ([62]) *There is a family of two-commodity networks* $\{(G^n, r^n, c^n)\}_{n=1}^{\infty}$ *such that* G_n *has* $O(n)$ *vertices and edges and* $\beta(G^n, r^n, c^n) = 2^{\Omega(n)}$ *as* $n \to \infty$.

In fact, the construction in the proof of Theorem 3.18 shows the following: adding a single edge to a two-commodity network (G^n, r^n, c^n) with $O(n)$ vertices and edges, $d_1(G^n, r^n, c^n) = 0$, and $d_2(G^n, r^n, c^n) = 1$ can increase the common cost incurred by traffic of the two commodities to roughly the $(n-1)$th and nth Fibonacci numbers, respectively.

On the other hand, the Braess ratio is always at most exponential in the network size.

Theorem 3.19 ([62]) *There is a constant* $c > 0$ *such that for every* $k, n \geq 1$ *and every instance* (G, r, c) *with* k *commodities and* n *vertices,* $\beta(G, r, c) \leq 2^{ckn}$.

The proof of Theorem 3.19 actually shows the stronger statement that if f is a Wardrop equilibrium for the k-commodity, n-vertex instance (G, r, c) and $\tilde{f}$ is feasible for (G, r, c), then the maximum cost $\max_i d_i(G, r, c)$ incurred by traffic in f is $2^{O(kn)}$ times the maximum cost incurred by traffic in $\tilde{f}$. The question of whether or not the largest-possible Braess ratio of multicommodity networks depends on the number of commodities is open.

3.4 Detecting Braess's Paradox Is Hard

Previous results of this section were devoted to the analysis of the worst-case severity of Braess's Paradox. Braess's Paradox also suggests a natural algorithmic question: given a network, is it suffering from the paradox? If so, which edges should be removed to recover the best-possible Wardrop equilibrium?

This innocuous question turns out to be extremely difficult to answer, in a sense we make precise below. To keep things simple, we will initially consider only single-commodity networks with linear cost functions. Detecting Braess's Paradox can be phrased as an optimization problem as follows: given a single-commodity instance (G, r, c) with linear cost

functions, find a subnetwork that minimizes the cost of a Wardrop equilibrium for (H, r, c) over all subnetworks $H \subseteq G$. We call this optimization problem LINEAR NETWORK DESIGN.

LINEAR NETWORK DESIGN can be solved by enumerating all subgraphs H of G, computing a Wardrop equilibrium in each, and picking the best solution. (Since Wardrop equilibria are the minima of the convex function in (1), one can be computed using convex programming.) On the other hand, there may be an exponential number of candidate subnetworks H. How well can we solve this optimization problem if we use only a reasonable amount of computational resources?

We will use basic concepts of computational complexity theory as described in, for example, Garey and Johnson [44]. Recall that a *γ-approximation algorithm* for a minimization problem runs in polynomial time and returns a solution no more than γ times as costly as an optimal solution. The value γ is the *approximation ratio* or *performance guarantee* of the algorithm.

While we would obviously like to solve LINEAR NETWORK DESIGN optimally in polynomial time, a natural weaker goal is to design a γ-approximation algorithm with γ as close to 1 as possible. Of course, even the *trivial algorithm*, which always returns the entire network G, can be viewed as an approximation algorithm for LINEAR NETWORK DESIGN. Because the Braess ratio of every network with linear cost functions is at most 4/3 (Proposition 3.3), we have the following guarantee on the trivial algorithm.

Proposition 3.20 *The trivial algorithm is a $\frac{4}{3}$-approximation algorithm for* LINEAR NETWORK DESIGN.

Needless to say, we should aspire to design better, more clever approximation algorithms. Alas, none exist, assuming $P \neq NP$.

Theorem 3.21 ([78]) *For every $\epsilon > 0$, there is no $(\frac{4}{3} - \epsilon)$-approximation algorithm for* LINEAR NETWORK DESIGN *(unless $P = NP$).*

Proof: We give a polynomial-time "gap reduction" from the NP-complete problem 2 DIRECTED DISJOINT PATHS (2DDP) [40]: given a directed graph $G = (V, E)$ and distinct vertices $s_1, s_2, t_1, t_2 \in V$, are there s_i-t_i paths P_i for $i = 1, 2$, such that P_1 and P_2 are vertex-disjoint? We can prove the theorem by showing how a $(\frac{4}{3} - \epsilon)$-approximation algorithm for LINEAR NETWORK DESIGN can be used to differentiate between "yes" and "no" instances of 2DDP in polynomial time.

Consider an instance $\mathcal{I}$ of 2DDP, as above. Augment the vertex set V by an additional source s and sink t, and include the directed edges (s, s_1), (s, s_2), (t_1, t), and (t_2, t) (see Figure 6). Denote the new network by $G' = (V', E')$ and endow the edges of E' with the following linear cost functions c: edges of E are given the cost function $c(x) = 0$, edges (s, s_2) and (t_1, t) are given the cost function $c(x) = x$, and edges (s, s_1) and (t_2, t) are given the cost function $c(x) = 1$. The instance $(G', 1, c)$ can be constructed from $\mathcal{I}$ in polynomial time.

We can complete the proof by establishing two statements: if $\mathcal{I}$ is a "yes" instance of 2DDP, then G' admits a subnetwork H such that a Wardrop equilibrium for $(H, 1, c)$ has cost 3/2; and if $\mathcal{I}$ is a "no" instance, then for every subnetwork H of G', a Wardrop equilibrium for $(H, 1, c)$ has cost at least 2.

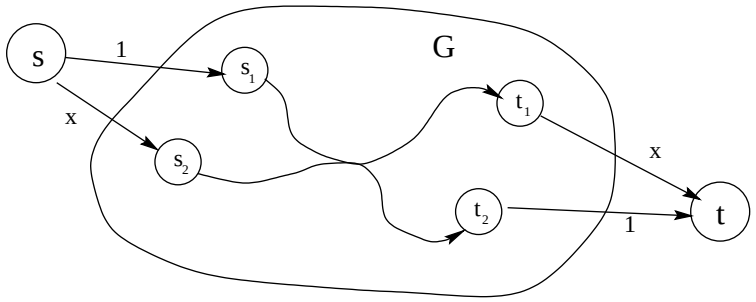

Figure 6: Proof of Theorem 3.21. In a "no" instance of 2DDP, the existence of s_1-t_1 and s_2-t_2 paths implies the existence of an s_2-t_1 path.

First suppose there are vertex-disjoint s_1-t_1 and s_2-t_2 paths P_1 and P_2 in G, respectively. Obtain H by deleting all edges of G not contained in some P_i. Then, H is a subgraph of G' with exactly two s-t paths, and routing half a unit of flow along each yields a Wardrop equilibrium with cost 3/2 (cf., Figure 2(a)).

Now suppose $\mathcal{I}$ is a "no" instance and consider a subgraph H of G'. We can assume that H contains an s-t path. If H has an s-t path P containing an s_2-t_1 path, then routing all of the flow on P yields a Wardrop equilibrium with cost 2 (cf., Figure 2(b)). Otherwise, since $\mathcal{I}$ is a "no" instance of 2DDP, two sole possibilities remain (see Figure 6): either for precisely one $i \in \{1, 2\}$, H has an s-t path P containing an s_i-t_i path, or all s-t paths P in H contain an s_1-t_2 path of G. In either case, routing one unit of flow along such a path P provides a Wardrop equilibrium with cost 2. ■

Thus, no polynomial-time algorithm for Linear Network Design has an approximation ratio superior to that of the trivial algorithm. Equivalently, it is NP-hard to distinguish between "paradox-free" instances (with Braess ratio 1) and instances suffering from the most severe manifestations of the paradox (with Braess ratio 4/3).

While we have only established the optimality of the trivial algorithm for networks with linear cost functions, similar results hold with other sets of allowable edge cost functions. For example, let General Network Design be the analogous optimization problem for single-commodity networks with arbitrary cost functions. Theorem 3.7 implies that the trivial algorithm is a $\lfloor n/2 \rfloor$-approximation algorithm for General Network Design, where n is the number of network vertices. On the other hand, the following inapproximability result holds.

Theorem 3.22 ([78]) *Assuming $P \neq NP$, for every $\epsilon > 0$ there is no $(\lfloor n/2 \rfloor - \epsilon)$-approximation algorithm for* General Network Design.

The proof of Theorem 3.22 is somewhat involved and makes use of the Braess graphs that were introduced in the proof of Theorem 3.4. For the proof, and similar results for other sets of allowable cost functions, see [78]. For analogous intractability results for multicommodity networks, which build on the two-commodity networks alluded to in Theorem 3.19, see Lin et al. [62].

4 Coping with Selfishness: How To Reduce the Price of Anarchy

We have seen that the price of anarchy of selfish routing can be large in networks with highly nonlinear cost functions, including with functions that are common in applications, such as M/M/1 delay functions. This final technical section asks: other than somehow enforcing optimal routing, what can we do about it? Can modest intervention, when feasible, significantly reduce the price of anarchy? We briefly discuss three techniques for mitigating the inefficiency of selfish routing: increasing the capacity of the network (Subsection 4.1), routing a small amount of traffic centrally (Subsection 4.2), and influencing traffic with edge taxes (Subsection 4.3).

4.1 Capacity Augmentation

For the rest of this survey we study networks with arbitrary cost functions, where the price of anarchy is unbounded. We next show that a bound on the inefficiency of selfish routing in such networks is nonetheless possible, via a so-called bicriteria approach. Specifically, our next result is that the cost of a Wardrop equilibrium is at most that of an optimal flow that is forced to route twice as much traffic between each source-sink pair. We will see that this result has the following alternative interpretation: in lieu of centralized control, the inefficiency of selfish routing can be offset by a moderate increase in link speed.

Example 4.1 Consider the nonlinear variant of Pigou's example (Figure 1(b)): a two-node, two-link network with cost functions $c(x) = 1$ and $c(x) = x^p$ for p large. Recall that with one unit of traffic, the Wardrop equilibrium routes all 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 $\epsilon \to 0$ as $p \to \infty$). When the traffic rate r 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 $p \to \infty$. In particular, for every p an optimal flow feasible for twice the original traffic rate ($r = 2$) has cost at least 1, which equals the cost of the Wardrop equilibrium in the original instance.

We now show that the bound stated in Example 4.1 holds for *all* instances.

Theorem 4.2 ([89]) *If f is a Wardrop equilibrium for (G, r, c) and f^* is feasible for $(G, 2r, c)$, then*

$$C(f) \leq C(f^*).$$

Proof: Let f and f^* denote a Wardrop equilibrium for (G, r, c) and a feasible flow for $(G, 2r, c)$, respectively. For each commodity i, let $d_i(G, r, c)$ denote the common cost incurred by the traffic of commodity i in the flow f (see Definition 3.16). Definition 2.1 and the definition of cost (3) imply that $C(f) = \sum_i r_i d_i(G, r, c)$.

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, in the sense that the cost of f^* with respect to

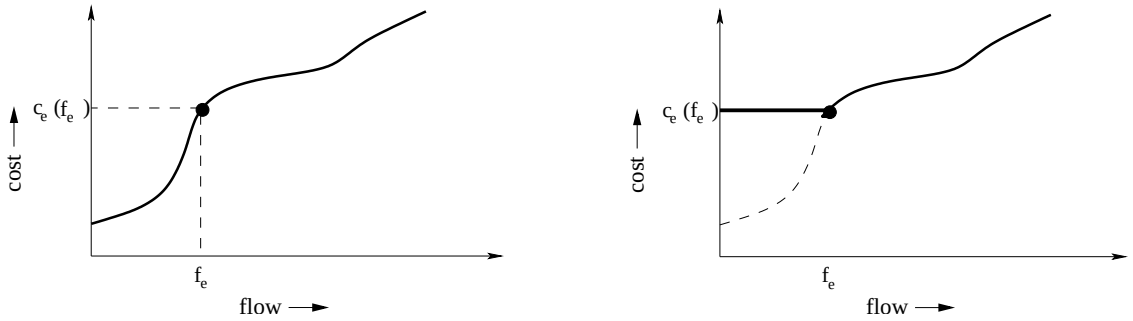

(a) Graph of the cost function c_e and its value at the flow value f_e

(b) Graph of the cost function $\bar{c}_e$

Figure 7: Construction in the proof of Theorem 4.2 of the modified cost function $\bar{c}_e$ given the original cost function c_e and the Wardrop equilibrium value f_e. Solid lines denote graphs of functions.

$\bar{c}$ is close to its original cost. Specifically, we define the cost function $\bar{c}_e$ for each edge e as follows:

$$\bar{c}_e(x) = \begin{cases} c_e(f_e) & \text{if } x \leq f_e \\ c_e(x) & \text{if } x \geq f_e. \end{cases}$$

Figure 7 illustrates this construction. 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

$$x(\bar{c}_e(x) - c_e(x)) \leq c_e(f_e)f_e \tag{14}$$

for all $x \geq 0$. The left-hand side of (14)—the discrepancy between $x\bar{c}_e(x)$ and $xc_e(x)$—is maximized when x is slightly smaller than f_e and when $c_e(x) = 0$. In this case, the value of the left-hand side of (14) is essentially the area of the rectangle enclosed by dashed lines in Figure 7(a), which in turn is the cost incurred by the Wardrop equilibrium f on the edge e. Thus

$$\bar{C}(f^*) - C(f^*) = \sum_{e \in E} f_e^*(\bar{c}_e(f_e^*) - c_e(f_e^*)) \leq \sum_{e \in E} c_e(f_e)f_e = C(f). \tag{15}$$

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 $d_i(G, r, c)$. Therefore,

$$\bar{C}(f^*) = \sum_{P \in \mathcal{P}} \bar{c}_P(f^*)f_P^* \geq \sum_{i=1}^{k} \sum_{P \in \mathcal{P}_i} d_i(G, r, c)f_P^* = \sum_{i=1}^{k} 2r_i d_i(G, r, c) = 2C(f). \tag{16}$$

The theorem now follows immediately from inequalities (15) and (16). ■

Remark 4.3 Example 4.1 shows that the bound in Theorem 4.2 is the best possible.

Another interpretation of Theorem 4.2 is that the benefit of centralized control is equaled or exceeded by the benefit of a sufficient improvement in link technology.

Corollary 4.4 ([82]) *Let (G, r, c) be an 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 a Wardrop equilibrium 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 4.2 and Corollary 4.4 are equivalent; for details, see [82] or [86].

Corollary 4.4 takes on a particularly nice form in instances in which all cost functions are M/M/1 delay functions (Example 2.11). In this case, if the cost function c_e of edge e is $c_e(x) = (u_e - x)^{-1}$, then the modified function $\tilde{c}_e$ is $\tilde{c}_e(x) = 1/2(u_e - x/2) = 1/(2u_e - x)$. Corollary 4.4 thus offers the following advice for networks where cost functions are M/M/1 delay functions and capacity is cheap: *to outperform optimal routing, just double the capacity of every edge.*

4.2 Stackelberg Routing

A second approach to reducing the price of anarchy, also explored in [82], is to allow a small portion of the network traffic to be routed centrally. We will call this *Stackelberg routing*, after a concept from noncooperative game theory called Stackelberg games [98]. In the interest of space, we only describe our model of Stackelberg routing informally, via two examples.

Example 4.5 To understand the potential power of Stackelberg routing, consider the nonlinear variant of Pigou's example (Figure 1(b)) with p large. Suppose we are granted the ability to route a $\gamma \in [0, 1]$ fraction of the traffic as we wish, knowing that the other $(1 - \gamma)$ fraction of the traffic will then choose routes selfishly, as usual. (Definition 2.1 thus governs the routes chosen by selfish traffic, but not by the centrally routed traffic.) We will call a routing of the centrally controlled traffic a *Stackelberg strategy*. Observe that for every Stackelberg strategy, the selfish traffic will use the lower edge—the upper route is never attractive to selfish users, even if the lower one is fully congested. On the other hand, if we route some traffic on the upper edge ourselves, the cost of the overall solution decreases. In particular, if γ is sufficiently large, we can mimic the optimal flow on the upper edge (routing excess traffic on the lower edge) and induce the optimal flow. Thus Stackelberg routing can decrease, or even eradicate, the inefficiency of selfish routing in this example.

Example 4.6 Stackelberg routing also has its limitations. Suppose we modify Example 4.5 by replacing the cost function $c(x) = x^p$ of the lower edge in Figure 1(b) by the cost function $c(x) = x^p/(1 - \gamma)^p$, where γ is the fraction of traffic that we are permitted to route centrally. The key observation is that no matter how the centrally controlled traffic is routed, there is enough selfish traffic to fully congest the lower edge. Therefore, Stackelberg strategies that

route at most γ units of traffic cannot produce a flow with cost less than 1. On the other hand, the optimal flow, which routes $\gamma + \epsilon$ units of flow on the upper edge and the rest on the lower edge, has cost approaching γ as $p \to \infty$ and $\epsilon \to 0$.

Stackelberg routing was first proposed by Korilis, Lazar, and Orda [58], who were motivated by so-called virtual private networks (see e.g. Birman [12] for a discussion of VPNs). The main goal in [58] was to characterize the instances in which some Stackelberg strategy induces an optimal flow (as in Example 4.5). This problem has also been studied more recently by Kaporis, Politopoulou, and Spirakis [54]. Here we follow Roughgarden [85] and seek worst-case bounds on the ratio between the cost of the best flow possible with Stackelberg routing and that of an optimal flow. Example 4.6 shows that, for each $\gamma \in (0, 1]$, this ratio can be arbitrarily close to $1/\gamma$, even in two-node, two-link networks. One of the main results of [85] is a matching upper bound for networks of parallel links.

Theorem 4.7 ([85]) *For every instance (G, r, c) with a network of parallel links and every $\gamma \in (0, 1]$, there is a Stackelberg strategy that routes γr units of traffic and yields a flow with cost at most $1/\gamma$ times the cost of an optimal flow for (G, r, c).*

Theorem 4.7 provides a smooth trade-off between optimal flows and Wardrop equilibria, as a function of the fraction of centrally controlled traffic. When $\gamma = 0$, we are stuck with a Wardrop equilibrium, which can cost arbitrarily more than an optimal flow in a network with arbitrary cost functions. When $\gamma = 1$ and we control all of the traffic, we can of course route the traffic optimally. Example 4.6 and Theorem 4.7 precisely quantify the inefficiency of selfish routing for all intermediate values of γ (in networks of parallel links).

Remark 4.8 The proof of Theorem 4.7 is constructive, and uses a simple iterative algorithm to compute a good Stackelberg strategy. This algorithm runs in polynomial time as long as the network cost functions satisfy a mild convexity condition (see [85] for details). While this algorithm is sufficient to obtain the best-possible worst-case guarantee in Theorem 4.7, it does not compute the optimal Stackelberg strategy in every instance. Indeed, the optimization problem of computing an optimal Stackelberg strategy is NP-hard [85], though it can be closely approximated in polynomial time [60].

Theorem 4.7 applies only to networks of parallel links, and the power of Stackelberg routing in more general networks is not fully understood. The $1/\gamma$ upper bound of Theorem 4.7 does not hold in general single-commodity networks, and no interesting bounds are possible in multicommodity networks [86]. Very recently, Fleischer and Swamy [39] proved an analogue of Theorem 4.7, with $1/\gamma$ replaced by a somewhat larger function of γ, for a wide class of networks, including series-parallel networks and the Braess graphs of Subsection 3.1. The question of whether or not such a result holds for general single-commodity networks is open.

4.3 Pricing Network Edges

We conclude with a third, very natural approach to reducing the price of anarchy of selfish routing: influencing selfish behavior with edge taxes. While not discussed in [82], this idea

has been extensively studied since the earliest papers on selfish routing. The literature on pricing selfish routing networks is vast, and we will confine our attention to only one classical result and two currently active research directions. See Yang and Huang [103], for example, for an introduction to this research area.

Pigou [72] suggested what are often called *marginal cost taxes* or *Pigouvian taxes*. 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 now allow each edge e of a selfish routing network to possess a nonnegative *tax* τ_e. We denote a selfish routing instance (G, r, c) with edge taxes τ by $(G, r, c + \tau)$. A Wardrop equilibrium for such an instance $(G, r, c + \tau)$ is defined as in Definition 2.1, with all traffic traveling on routes that minimize the sum of the edge costs and edge taxes. Equivalently, it is a Wardrop equilibrium for the instance (G, r, c^τ), where the cost function c^τ is a shifted version of the original cost function c_e: $c_e^\tau(x) = c_e(x) + \tau_e$ for all $x \geq 0$.

Mathematically, the principle of marginal cost pricing asserts that for a flow f feasible for an 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. Pigou [72] suggested that these taxes should eliminate all of the inefficiency of selfish routing, and Beckmann, McGuire, and Winsten [8] made this idea rigorous.

Proposition 4.9 ([8, 72]) *Let (G, r, c) be an instance with differentiable cost functions, admitting an optimal flow f^*. Let $\tau_e = f_e^* \cdot c_e'(f_e^*)$ denote the marginal cost tax for edge e with respect f^*. Then f^* is a Wardrop equilibrium for $(G, r, c + \tau)$.*

In words, marginal cost taxes induce an optimal flow as a Wardrop equilibrium.

While Proposition 4.9 may appear to be a complete solution to the problem of reducing the price of anarchy of selfish routing, it possesses several drawbacks. Two of these have recently motivated much research in the theoretical computer science and mathematical programming communities.

First, the definition of a Wardrop equilibrium in Proposition 4.9 implicitly assumes that all network users trade off cost and taxes in an identical way. For example, if edge costs represent travel time, some users might be more sensitive to time delays, while others are concerned primarily with monetary expenses. Several papers have considered the following model of *heterogeneous* traffic: each network user chooses a path that minimizes a weighted sum of the edge costs and the edge taxes. In other words, the preferences of a network user are summarized by a single scalar—the monetary value to the user of one unit of cost. This model was first proposed and studied in the transportation science literature [30, 32, 74], but analogues of Proposition 4.9 were only recently given for heterogeneous traffic. Cole, Dodis, and Roughgarden [22] proved that, in single-commodity networks with heterogeneous traffic, there is always a set of taxes that induces the optimal flow as a Wardrop equilibrium. This result was extended to multicommodity networks independently by Fleischer, Jain, and Mahdian [38], Karakostas and Kolliopoulos [56], and Yang and Huang [102].

Second, in networks where cost functions can have large derivatives, the marginal cost taxes of Proposition 4.9 can be extremely large. Several solutions have recently been proposed

for this problem, including computing a tax that induces an optimal flow but also minimizes the taxes paid [9, 48]; incorporating the taxes paid into the objective function [22, 24]; and proving worst-case bounds on the largest tax needed to induce an optimal flow [22, 37, 38].

5 Recent Related Work

This survey describes the basic results on the price of anarchy of selfish routing. However, we have only scratched the surface of a broader issue: quantifying the inefficiency of noncooperative equilibria in applications with selfish users. This fundamental problem has only recently been systematically studied, but there is already a large literature addressing many aspects of it. We conclude this survey by briefly discussing some of the recent work in this lively research area.

First, the price of anarchy has been analyzed in numerous variants and generalizations of the basic selfish routing model studied in this survey. Several recent papers have extended Theorem 2.17 to more general classes of games [17, 70, 90]. The price of anarchy of selfish routing has also been studied with objective functions other than (3) [26, 62, 80, 84, 101]; with edge capacities and other types of "side constraints" [27, 49, 55]; when the traffic rates can vary with the network congestion [17, 23]; when network users can have non-negligible size [3, 6, 7, 21, 25, 28, 41, 57, 87, 89, 95]; and with definitions of path cost $c_P(f)$ other than the sum of all edge costs [7, 23].

Second, the price of anarchy is a very general concept—applicable to every noncooperative game with a notion of equilibrium and a nonnegative objective function. In games where different equilibria can have different objective function values, the price of anarchy is usually defined as the ratio between the objective function value of the worst equilibrium and that of an optimal solution [59]. The related concept of the *price of stability* [3] instead considers the objective function value of the best equilibrium. The price of anarchy and price of stability have been successfully analyzed in a diverse array of applications with selfish users over the past few years. These include scheduling (see [29, 35] and the references therein), facility location [31, 45, 64, 97], network design [3, 4, 18, 34, 36], resource allocation [50, 52, 104], and other networking games [1, 5, 46, 51].

Third, researchers have begun to study the inefficiency of different notions of a selfish outcome. For example, Goemans, Mirrokni, and Vetta [45, 64] have extended the concept of the price of anarchy to games in which equilibria need not exist. The work in [45, 64] is also motivated by the important problem of understanding when a small price of anarchy implies that selfish users can "learn", by independent and repeated experimentation from an arbitrary initial state, an approximately optimal outcome. Another example is given by Christodoulou and Koutsoupias [20], who studied the inefficiency of correlated equilibria in scheduling games.

Finally, an emerging research direction is to use the price of anarchy as a measure for the performance of a network protocol that interacts with selfish users. This idea connects the analysis of the inefficiency of game-theoretic equilibria with *mechanism design*, a classical subfield of microeconomics that studies how to design games that possess equilibria with good properties (see e.g. [63, Chapter 23] or [68, Chapter 10]). For example, Johari [50, Chapter 5] considers a class of network resource allocation protocols, each of which can be

viewed as a game with selfish users, and proves that a natural "proportional sharing" protocol minimizes the (worst-case) inefficiency of equilibria. A second example is the recent work by Chen, Roughgarden, and Valiant [19] that analyzes how the price of stability in a class of network design games [3] depends on the choice of an underlying cost-sharing protocol.

Acknowledgements

I have been lucky to benefit from the advice and assistance of many individuals while working on the research described in this survey. The list is too long to give here, but I must mention Éva Tardos, who supervised my doctoral thesis [82] and collaborated with me on much of the work described in this survey; Leonard Schulman, who introduced me to selfish routing; and Christos Papadimitriou, who asked about the price of anarchy of selfish routing.

In addition, I thank José Correa, Henry Lin, Mukund Sundararajan, and Greg Valiant for helpful comments on an earlier draft. Finally, I thank Alberto Caprara for inviting me to write this survey, and for his patience throughout the writing process.

References

[1] D. Acemoglu and A. Ozdaglar. Flow control, routing, and performance under monopoly pricing. Technical Report WP-1696, MIT LIDS, 2003.

[2] R. K. Ahuja, T. L. Magnanti, and J. B. Orlin. *Network Flows: Theory, Algorithms, and Applications*. Prentice-Hall, 1993.

[3] 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 *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 295–304, 2004.

[4] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 511–520, 2003.

[5] J. Aspnes, K. Chang, and A. Yampolskiy. Inoculation strategies for victims of viruses and the sum-of-squares partition problem. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 43–52, 2005.

[6] B. Awerbuch, Y. Azar, and L. Epstein. The price of routing unsplittable flow. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 57–66, 2005.

[7] R. Banner and A. Orda. Bottleneck routing games in communication networks. In *Proceedings of the 25th INFOCOM Conference*, 2006. To appear.

[8] M. J. Beckmann, C. B. McGuire, and C. B. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.

[9] P. Bergendorff, D. W. Hearn, and M. V. Ramana. Congestion toll pricing of traffic networks. In P. M. Pardalos, D. W. Hearn, and W. W. Hager, editors, *Network Optimization*, pages 51–71. Springer, 1997.

[10] D. P. Bertsekas and R. G. Gallager. *Data Networks*. Prentice-Hall, 1987. Second Edition, 1991.

[11] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989. Second Edition, Athena Scientific, 1997.

[12] K. P. Birman. *Building Secure and Reliable Network Applications*. Manning, 1996.

[13] D. E. Boyce, H. S. Mahmassani, and A. Nagurney. A retrospective of Beckmann, McGuire, and Winsten's *Studies in the Economics of Transportation*. *Papers in Regional Science*, 84(1):85–103, 2005.

[14] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968. English translation in [15].

[15] D. Braess. On a paradox of traffic planning. *Transportation Science*, 39(4):446–450, 2005.

[16] D. G. Cantor and M. Gerla. Optimal routing in a packet-switched computer network. *IEEE Transactions on Computers*, 23(10):1062–1069, 1974.

[17] C. K. Chau and K. M. Sim. The price of anarchy for non-atomic congestion games with symmetric cost maps and elastic demands. *Operations Research Letters*, 31(5):327–334, 2003.

[18] H. Chen and T. Roughgarden. Network design with weighted players. Submitted, 2005.

[19] H. Chen, T. Roughgarden, and G. Valiant. Sharing costs to minimize the inefficiency of selfish network design. In preparation, 2006.

[20] G. Christodoulou and E. Koutsoupias. On the price of anarchy and stability of correlated equilibria of linear congestion games. Manuscript, 2005.

[21] G. Christodoulou and E. Koutsoupias. The price of anarchy of finite congestion games. In *Proceedings of the 37th Annual ACM Symposium on Theory of Computing (STOC)*, pages 67–73, 2005.

[22] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proceedings of the 35th Annual ACM Symposium on Theory of Computing (STOC)*, pages 521–530, 2003.

[23] R. Cole, Y. Dodis, and T. Roughgarden. Bottleneck links, variable demand, and the tragedy of the commons. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2006. To appear.

[24] R. Cole, Y. Dodis, and T. Roughgarden. How much can taxes help selfish routing? *Journal of Computer and Systems Sciences*, 2006. To appear.

[25] R. Cominetti, J. R. Correa, and N. E. Stier-Moses. Network games with atomic players. Manuscript, 2005.

[26] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Computational complexity, fairness, and the price of anarchy of the maximum latency problem. In *Proceedings of the 10th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 3064 of *Lecture Notes in Computer Science*, pages 59–73, 2004.

[27] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. Selfish routing in capacitated networks. *Mathematics of Operations Research*, 29(4):961–976, 2004.

[28] J. R. Correa, A. S. Schulz, and N. E. Stier Moses. On the inefficiency of equilibria in congestion games. In *Proceedings of the 11th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 167–181, 2005.

[29] A. Czumaj. Selfish routing on the Internet. In J. Leung, editor, *Handbook of Scheduling: Algorithms, Models, and Performance Analysis*, chapter 42. CRC Press, 2004.

[30] S. C. Dafermos. A multicriteria route-mode choice traffic equilibrium model. In *Proceedings of the Frontiers in Transportation Equilibrium and Supply Models Symposium*, 1981.

[31] N. R. Devanur, N. Garg, R. Khandekar, V. Pandit, A. Saberi, and V. V. Vazirani. Price of anarchy, locality gap, and a network service provider game. In *Proceedings of the First Annual International Workshop on Internet and Network Economics (WINE)*, volume 3828 of *Lecture Notes in Computer Science*, pages 1046–1055, 2005.

[32] R. B. Dial. A model and algorithm for multicriteria route-mode choice. *Transportation Research, Series B*, 13(4):311–316, 1979.

[33] P. Dubey. Inefficiency of Nash equilibria. *Mathematics of Operations Research*, 11(1):1–8, 1986.

[34] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. J. Shenker. On a network creation game. In *Proceedings of the 22nd ACM Symposium on Principles of Distributed Computing (PODC)*, pages 347–351, 2003.

[35] R. Feldmann, M. Gairing, T. Lücking, B. Monien, and M. Rode. Selfish routing in non-cooperative networks: A survey. In *Proceedings of the 28th International Symposium on Mathematical Foundations of Computer Science (MFCS)*, volume 2747 of *Lecture Notes in Computer Science*, pages 21–45, 2003.

[36] P. M. Ferreira. *Interconnected Communication Networks Provisioned Selfishly.* PhD thesis, Carnegie Mellon University, 2004.

[37] L. K. Fleischer. Linear tolls suffice: New bounds and algorithms for tolls in single source networks. *Theoretical Computer Science*, 348(2-3):217–225, 2005.

[38] L. K. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 277–285, 2004.

[39] L. K. Fleischer and C. Swamy. Private communication, August, 2005.

[40] S. Fortune, J. E. Hopcroft, and J. Wyllie. The directed subgraph homeomorphism problem. *Theoretical Computer Science*, 10(2):111–121, 1980.

[41] D. Fotakis, S. C. Kontogiannis, and P. G. Spirakis. Selfish unsplittable flows. *Theoretical Computer Science*, 348(2-3):226–239, 2005.

[42] E. J. Friedman. Genericity and congestion control in selfish routing. In *Proceedings of the 43rd Annual IEEE Conference on Decision and Control (CDC)*, pages 4667–4672, 2004.

[43] R. G. Gallager. A minimum delay routing algorithm using distributed computation. *IEEE Transactions on Communications*, 25(1):73–85, 1977.

[44] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. Freeman, 1979.

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

[46] M. M. Halldórsson, J. Y. Halpern, E. L. Li, and V. S. Mirrokni. On spectrum sharing games. In *Proceedings of the 23rd Annual ACM Symposium on Principles of Distributed Computing (PODC)*, pages 107–114, 2004.

[47] A. Haurie and P. Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15(3):295–308, 1985.

[48] D. W. Hearn and M. B. Yildirim. A toll pricing framework for traffic assignment problems with elastic demand. In *Transportation and Network Analysis Current Trends: Miscellanea in Honor of Michael Florian*, chapter 9. Kluwer, 2002.

[49] O. Jahn, R. Möhring, A. S. Schulz, and N. E. Stier Moses. System-optimal routing of traffic flows with user constraints in networks with congestion. *Operations Research*, 53(4):600–616, 2005.

[50] R. Johari. *Efficiency Loss in Market Mechanisms for Resource Allocation*. PhD thesis, MIT, 2004.

[51] R. Johari and J. N. Tsitsiklis. Routing and peering in a competitive Internet. Technical Report 2570, MIT LIDS, 2003.

[52] R. Johari and J. N. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 29(3):407–435, 2004.

[53] H. Kameda. Private communication, June, 2002.

[54] A. C. Kaporis, E. I. Politopoulou, and P. G. Spirakis. The price of optimum in Stackelberg games. Technical Report TR05-055, ECCC, 2005.

[55] G. Karakostas and S. G. Kolliopoulos. Selfish routing in the presence of side constraints. Technical Report CAS-03-13-GK, Department of Computing and Software, McMaster University, 2003.

[56] G. Karakostas and S. G. Kolliopoulos. Edge pricing of multicommodity networks for heterogeneous selfish users. In *Proceedings of the 45th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 268–276, 2004.

[57] S. C. Kontogiannis and P. G. Spirakis. Atomic selfish routing in networks: A survey. In *Proceedings of the First Annual International Workshop on Internet and Network Economics (WINE)*, volume 3828 of *Lecture Notes in Computer Science*, pages 989–1002, 2005.

[58] Y. A. Korilis, A. A. Lazar, and A. Orda. Achieving network optima using Stackelberg routing strategies. *IEEE/ACM Transactions on Networking*, 5(1):161–173, 1997.

[59] E. Koutsoupias and C. H. Papadimitriou. Worst-case equilibria. In *Proceedings of the 16th Annual Symposium on Theoretical Aspects of Computer Science (STACS)*, volume 1563 of *Lecture Notes in Computer Science*, pages 404–413, 1999.

[60] V. S. Anil Kumar and M. V. Marathe. Improved results for Stackelberg scheduling strategies. In *Proceedings of the 29th Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 2380 of *Lecture Notes in Computer Science*, pages 776–787, 2002.

[61] H. Lin, T. Roughgarden, and É. Tardos. A stronger bound on Braess's Paradox. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 333–334, 2004.

[62] H. Lin, T. Roughgarden, É. Tardos, and A. Walkover. Braess's Paradox, Fibonacci numbers, and exponential inapproximability. In *Proceedings of the 32nd Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3580 of *Lecture Notes in Computer Science*, pages 497–512, 2005.

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

[64] V. S. Mirrokni and A. Vetta. Convergence issues in competitive games. In *Proceedings of the 7th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 183–194, 2004.

[65] D. Monderer and L. S. Shapley. Potential games. *Games and Economic Behavior*, 14(1):124–143, 1996.

[66] J. F. Nash. Non-cooperative games. *Annals of Mathematics*, 54(2):286–295, 1951.

[67] A. Orda, R. Rom, and N. Shimkin. Competitive routing in multiuser communication networks. *IEEE/ACM Transactions on Networking*, 1(5):510–521, 1993.

[68] M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* MIT Press, 1994.

[69] C. H. Papadimitriou. Algorithms, games, and the Internet. In *Proceedings of the 33rd Annual ACM Symposium on Theory of Computing (STOC)*, pages 749–753, 2001.

[70] G. Perakis. The price of anarchy when costs are non-separable and asymmetric. In *Proceedings of the 10th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, volume 3064 of *Lecture Notes in Computer Science*, pages 46–58, 2004.

[71] A. L. Peressini, F. E. Sullivan, and J. J. Uhl, Jr. *The Mathematics of Nonlinear Programming.* Springer, 1988.

[72] A. C. Pigou. *The Economics of Welfare.* Macmillan, 1920.

[73] L. Qiu, Y. R. Yang, Y. Zhang, and S. Shenker. On selfish routing in Internet-like environments. In *Proceedings of SIGCOMM*, pages 151–162, 2003.

[74] R. E. Quandt. A probabilistic abstract mode model. In *Studies in Travel Demand, Volume II*, pages 90–113. Mathematica, 1967.

[75] A. Rapoport and A. M. Chammah. *Prisoner's Dilemma.* University of Michigan Press, 1965.

[76] A. Ronen. Private communication, March, 2002.

[77] R. W. Rosenthal. A class of games possessing pure-strategy Nash equilibria. *International Journal of Game Theory*, 2(1):65–67, 1973.

[78] T. Roughgarden. Designing networks for selfish users is hard. In *Proceedings of the 42d Annual Symposium on Foundations of Computer Science (FOCS)*, pages 472–481, 2001. Full version to appear in *Journal of Computer and Systems Sciences.*

[79] T. Roughgarden. The price of anarchy in networks with polynomial edge latency. Technical Report 2001-1847, Cornell University, 2001.

[80] T. Roughgarden. How unfair is optimal routing? In *Proceedings of the 13th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 203–204, 2002.

[81] T. Roughgarden. The price of anarchy is independent of the network topology. In *Proceedings of the 34th Annual ACM Symposium on Theory of Computing (STOC)*, pages 428–437, 2002. Preliminary version of [83].

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

[83] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.

[84] T. Roughgarden. The maximum latency of selfish routing. In *Proceedings of the 15th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 973–974, 2004.

[85] T. Roughgarden. Stackelberg scheduling strategies. *SIAM Journal on Computing*, 33(2):332–350, 2004.

[86] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.

[87] T. Roughgarden. Selfish routing with atomic players. In *Proceedings of the 16th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1184–1185, 2005.

[88] T. Roughgarden. Potential functions and the inefficiency of equilibria. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2006. To appear.

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

[90] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 49(2):389–403, 2004.

[91] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.

[92] Y. Sheffi. *Urban Transportation Networks: Equilibrium Analysis with Mathematical Programming Methods*. Prentice-Hall, 1985.

[93] M. J. Smith. The existence, uniqueness and stability of traffic equilibria. *Transportation Research, Part B*, 13(4):295–304, 1979.

[94] M. A. So. Private communication, March, 2005.

[95] S. Suri, C. Tóth, and Y. Zhou. Selfish load balancing and atomic congestion games. *Algorithmica*, 2006. To appear.

[96] É. Tardos. CS684 course notes. Cornell University, 2004.

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

[98] H. von Stackelberg. *Marktform und Gleichgewicht*. Springer, 1934.

[99] B. von Stengel. Private communication, July, 2004.

[100] J. G. Wardrop. Some theoretical aspects of road traffic research. In *Proceedings of the Institute of Civil Engineers, Pt. II*, volume 1, pages 325–378, 1952.

[101] D. Weitz. The price of anarchy. Manuscript, 2001.

[102] H. Yang and H. Huang. The multi-class, multi-criteria traffic network equilibrium and systems optimum problem. *Transportation Research, Part B*, 38(1):1–15, 2004.

[103] H. Yang and H. Huang. *Mathematical and Economic Theory of Road Pricing*. Elsevier Science, 2005.

[104] L. Zhang. The efficiency and fairness of a fixed budget resource allocation game. In *Proceedings of the 32nd Annual International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 3580 of *Lecture Notes in Computer Science*, pages 485–496, 2005.