Bottleneck Links, Variable Demand, and the Tragedy of the Commons

Richard Cole*　　Yevgeniy Dodis†　　Tim Roughgarden‡

April 8, 2005

Abstract

The price of anarchy, a measure of the inefficiency of selfish behavior, has been successfully analyzed in a diverse array of models over the past five years. The overwhelming majority of this work has studied optimization problems that sought an optimal way to allocate a fixed demand to resources whose performance degrades with increasing congestion. While fundamental, such problems overlook a crucial feature of many applications: the intrinsic coupling of the quality or cost of a resource and the demand for that resource. This coupling motivates allowing demand to vary with congestion, which in turn can lead to "the tragedy of the commons"—severe inefficiency caused by the overconsumption of a shared resource.

Allowing the demand for resources to vary with their congestion illuminates a second issue with existing studies of the price of anarchy: the standard additive method of aggregating the costs of different resources in a player's strategy is inappropriate for some important applications, including many of those with variable demand. For example, in networking applications a key performance metric is the achievable throughput along a path, which is controlled by its bottleneck (most congested) edge. This disconnect motivates consideration of nonlinear cost aggregation functions, such as the ℓ_p norms.

In this paper, we initiate the study of the price of anarchy with variable demand and with broad classes of nonlinear aggregation functions. We focus on selfish routing in single- and multicommodity networks, and on the ℓ_p norms for $1 \leq p \leq \infty$; our main results are as follows.

- For a natural "prize-collecting" objective function, the price of anarchy in multicommodity networks with variable demand is no larger than that in fixed-demand networks. Thus the inefficiency arising from the tragedy of the commons is no more severe than that from routing inefficiencies.
- Using the ℓ_p norm with $1 < p < \infty$ as a cost aggregation function can dramatically increase the price of anarchy in multicommodity networks (relative to additive aggregation), but causes no such additional inefficiency in single-commodity networks.
- Using the ℓ_∞ norm as a cost aggregation function can dramatically increase the price of anarchy, even in single-commodity networks. If attention is restricted to equilibria with additional structure, however—structure that is ensured by distributed shortest-path routing protocols—then using the ℓ_∞ norm does not increase the price of anarchy relative to additive aggregation.

*Department of Computer Science, New York University, 251 Mercer Street, New York, NY 10012. Supported in part by NSF grant CCR0105678. Email: `cole@cs.nyu.edu`.

†Department of Computer Science, New York University, 251 Mercer Street, New York, NY 10012. Supported in part by an NSF CAREER Award. Email: `dodis@cs.nyu.edu`.

‡Department of Computer Science, Stanford University, 462 Gates Building, Stanford CA 94305. Supported in part by ONR grant N00014-04-1-0725 and DARPA grant W911NF-04-9-0001. This work was done in part while visiting NYU and supported by an NSF Postdoctoral Fellowship. Email: `tim@cs.stanford.edu`.

1 Introduction

The Price of Anarchy and Variable Demand

The *price of anarchy* [25]—formally defined as the worst-case ratio between the objective function values of a Nash equilibrium and of an optimal solution—is an increasingly popular measure of the inefficiency of selfish behavior. Over the past five years, the price of anarchy has been successfully analyzed in a diverse array of applications, including scheduling (see [15] and the references therein), routing (see [33] and the references therein), facility location [39], network design [2, 3, 17], resource allocation [23], and other networking games [4, 20, 22]. Most of these previous works have identified natural classes of noncooperative games in which the price of anarchy is provably small; hence, selfish behavior results in only a modest loss of efficiency in these games.

The overwhelming majority of this work studied optimization problems of the following sort: given resources whose performance degrades with increasing congestion, allocate a fixed demand for the resources in an optimal way. While obviously fundamental, such problems overlook a crucial feature of many applications: the intrinsic coupling of the quality or cost of a resource and the demand for that resource. Put differently, we expect the demand for an uncongested resource to be relatively high, and that this demand will fall as the resource becomes more congested and expensive. Allowing the demand for a resource to vary inevitably gives rise to a tradeoff between two different quantities: the number of users that benefit from the resource, and the quality of the resource (which degrades as more and more users benefit from it). We next illustrate this tradeoff with a stark, famous example: *the tragedy of the commons* [21].

The Tragedy of the Commons

The tragedy of the commons typically refers to a strategic scenario with a shared resource that is effectively destroyed by overconsumption. In lieu of the traditional bovine example [21], we will illustrate this idea in a network routing context.

Consider a large but fixed population of agents who are each considering traversing a link from a node s to a node t. Suppose that if an x fraction of the population makes the trip, then each of the itinerant agents incurs a cost of $c(x)$ but reaps a benefit of 1. (Agents that stay home receive zero benefit and cost.) Suppose further that we instantiate the cost function as $c(x) = x^d$ for d large. Then the net benefit of making the trip is always nonnegative, even if the link is fully congested, and we expect the entire population to travel to t, resulting in zero net benefit for all.

Given dictatorial control of the population, we could implement a far superior outcome by detaining an ϵ fraction of the population; this would result in a $1 - \epsilon$ fraction of the population enjoying a net benefit of nearly 1 (for d large). In other words—and this is the tragedy of the commons—the fact that the final ϵ fraction of the population insists on making the trip congests the shared resource to the point that none of the population extracts any net benefit from it.

In all previous works on the price of anarchy of "selfish routing", agents were not permitted to refuse to travel—put differently, the amount of traffic in the network was exogenous (fixed), rather than endogenous as in the above example. In the above single-link network, there is of course no inefficiency when the amount of traffic is fixed (there is only one feasible solution). This brings us to the first goal of this paper.

(1) Quantify the inefficiency that arises from allowing the demand for a resource to vary with its congestion.

Bottleneck Links and Nonlinear Aggregation Functions

Allowing the demand for resources to vary with their congestion illuminates a second issue with most existing studies of the price of anarchy: the standard way of aggregating the costs of different resources in a player's strategy is inappropriate for some important applications, including many of those with variable demand.

To motivate this point, we recall the definition of *congestion games* [29]. In a congestion game, there is a ground set of elements (resources), and players' strategies are subsets of this ground set. Each element has a cost that is a function of its congestion—the number of players that select strategies containing it. For example, in selfish routing, the edges of the network form the ground set, and the strategies of a player are paths from a source vertex to a destination vertex. The cost of a player's strategy in a congestion game is given by an *aggregation function* of the costs of the elements in the strategy.

Almost all of the aforementioned work on the price of anarchy concerns congestion games with different objective functions and different restrictions on the game structure. All of these studies of congestion games save two [1, 6] have a key assumption in common, however: *that the aggregation function is linear*. In fact, with one exception [35], these papers assume that the aggregation function is additive—that the cost of a strategy is simply the sum of the costs of the elements it contains. (Although in several, such as the scheduling results surveyed in [15], all strategies are singletons and hence no aggregation function is needed.)

While additive aggregation functions are arguably the most natural ones, they are not well suited for all applications. For instance, when analyzing the performance of a communication network with a variable amount of traffic, a key performance metric is the achievable *throughput* along a path, which is controlled by its *bottleneck* (most congested) link; see e.g. Keshav [24]. In fact, the studies of Qiu et al. [28] and Akella, Chawla, and Seshan [1], both of which tried to adapt theoretical results for selfish routing to more faithful models of the Internet, singled out the choice of the additive aggregation function over the bottleneck link metric as a key disconnect between selfish routing and the traditional concerns of the networking community.

The bottleneck link metric corresponds to using the ℓ_∞ norm as the aggregation function. Banner and Orda [6] point out that the ℓ_∞ norm is the natural aggregation function in many additional applications. For example, in wireless networks, the transmission capability of a path is constrained by the node with the smallest lifetime, as determined by its remaining battery power and the amount of traffic that it must send [10]. The ℓ_∞ norm also arises when robustness to bursty traffic [5] or to growing demand [40] is a priority. For further discussion and examples, see [6].

Determining the price of anarchy with nonlinear aggregation functions, and in particular with the ℓ_∞ norm, is therefore essential to understanding the consequences of selfish behavior in the above applications. This is the second goal of this paper.

(2) Analyze the price of anarchy in fixed- and variable-demand congestion games with nonlinear aggregation functions, and in particular with the ℓ_∞ aggregation function.

Our Results

In this paper, we initiate the study of the price of anarchy in congestion games with variable demand and with broad classes of nonlinear aggregation functions. We focus on selfish routing in single- and multicommodity networks, and on the ℓ_p norms for $1 \leq p \leq \infty$.

For the first goal (1), we augment the basic selfish routing model in a standard way, so that each player has a fixed benefit of making the trip. If the player can travel from its source to its destination incurring cost below this benefit, it makes the trip; otherwise, it does not. In the transportation science literature, such networks are said to possess *elastic traffic* (see e.g. [18]).

As discussed at the beginning of the paper, with elastic traffic there are two quantities to optimize: we would like to maximize the benefit of the players, but also to minimize the cost they incur. To study the price of anarchy, we must aggregate these two quantities into a single objective function. The most natural ways to accomplish this—such as maximizing the *consumer surplus*, defined as the total benefit minus the total cost—result in mixed-sign objective functions. As is typical of approximation measures, non-trivial bounds on the price of anarchy for such objective functions are possible only under very strong assumptions. For example, in the single-link network above, the consumer surplus of the noncooperative equilibrium is 0, while an optimal solution has strictly positive consumer surplus (even if the link has cost $c(x) = x$).

On the other hand, mixed-sign objectives in optimization can often be transformed, in a non-approximation-preserving way, into natural same-sign objectives (see e.g. [19] and the references therein). Indeed, we show that for the usual additive aggregation function and the "prize-collecting" objective of minimizing the *lost* benefit plus the sum of the costs incurred, *there is no tragedy of the commons*. Formally, we prove that the worst-possible price of anarchy in multicommodity networks with elastic traffic—specified by an *arbitrary* continuous distribution of the benefits of traveling—and cost functions in a set $\mathcal{C}$ is no more than that of networks with a fixed amount of traffic and cost functions in $\mathcal{C}$. As a consequence, the price of anarchy in multicommodity networks with linear cost functions and arbitrary benefit distributions is at most $4/3$ [34]; if the cost functions are polynomials with nonnegative coefficients and degree at most d, the price of anarchy in such networks is $O(d/\log d)$ [31]. Thus, for this prize-collecting objective, the worst-possible inefficiency that arises from variable demand has no greater magnitude than that arising from routing inefficiencies.

For the second goal (2), we prove matching positive and negative results about the price of anarchy of selfish routing with ℓ_p aggregation functions. On the negative side, we give examples in Section 4.1 that demonstrate the following.

- For every $1 < p \leq \infty$, there is a family of two-commodity networks with linear cost functions and inelastic traffic in which the price of anarchy grows polynomially with the network size. (Cf. the $p = 1$ case, where the price of anarchy is at most $4/3$ in networks with linear cost functions, inelastic traffic, and an arbitrary number of commodities [34].) The "bicriteria bound" of [34] also fails to hold.

- For the ℓ_∞ norm, there is a family of single-commodity networks with linear cost functions and inelastic traffic in which the price of anarchy grows polynomially with the network size, and in which the bicriteria bound of [34] does not hold.

We also achieve the following matching positive results in Sections 4 and 5.

- For every $1 < p < \infty$, the price of anarchy in single-commodity networks with the ℓ_p norm is no worse than that in the $p = 1$ case. (E.g., is at most 4/3 for linear cost functions, $O(d/\log d)$ for bounded-degree polynomials, and so on.) The bicriteria bound of [34] also holds in such networks.

- For the ℓ_∞ norm and a natural *subclass* of equilibria, which we call *subpath-optimal*, the price of anarchy in single-commodity networks is no worse than that in the ℓ_1 case. The bicriteria bound of [34] also holds for this subclass of equilibria.

These positive results hold for networks with inelastic traffic (Section 4) and, under an additional technical condition, with elastic traffic as well (Section 5).

In particular, these results imply that the inefficiency of selfish routing with the ℓ_p norm with $p > 1$ is provably larger in multicommodity networks than in single-commodity ones. This separation stands in contrast to the provable equivalence of single- and multicommodity networks for the price of anarchy under the ℓ_1 norm [13, 31]. Indeed, all previously known proof techniques for bounding the price of anarchy of selfish routing (for the ℓ_1 norm) [11, 13, 14, 27, 31, 34, 38] in no way referred to the number of commodities of a network, nor to any combinatorial structure whatsoever. These proof techniques for the ℓ_1 case therefore appear to be necessarily incapable of extending to the ℓ_p case with $p > 1$—a new, fundamentally combinatorial proof technique is required. We give such a technique in Section 4.

Finally, while our positive result for a subclass of equilibria in the ℓ_∞ norm is reminiscent of analyses of the *best* equilibrium (the "price of stability") [3, 2, 13], it is much stronger in the following sense.

While studying the best equilibrium typically cannot be justified without allowing some centralized intervention [2], we show that subpath-optimal equilibria are in fact the "natural" outcome of decentralized optimization from a networking perspective. Specifically, we show that if an equilibrium is computed by a distributed shortest-path routing protocol, then it will automatically be subpath-optimal.

Further Related Work

While we are not aware of any previous work that focused on the price of anarchy with variable demand or with general nonlinear aggregation functions, there are a few papers related to the goals of the present work. First, Vetta [39] considered profit-maximization facility location games, auctions, and a variant of selfish routing. These games have both benefits and costs, and Vetta [39] considered a mixed-sign objective that is related to the consumer surplus. As noted above, approximation results are hard to come by with mixed-sign objective functions; because of this, Vetta [39] could only deduce non-trivial bounds on the price of anarchy under strong conditions. For facility location games, Vetta [39] proved that the price of anarchy is 2 only in the special case of all-zero costs. Similarly, the price of anarchy of profit-maximization selfish routing is bounded only when the benefits of routing are so large than an optimal solution routes all of the traffic; this assumption effectively rules out the tragedy of the commons *a priori*. The present paper sacrifices the mixed-sign objective for a prize-collecting one, but in exchange proves bounds on the price of anarchy without any assumptions on the relative magnitudes of the costs and benefits of a game.

Finally, as alluded to earlier, two recent papers studied aspects of selfish routing with the bottleneck link metric. First, Akella, Chawla, and Seshan [1] studied the price of anarchy under this metric for a variant of selfish routing with edge capacities and a maximization objective, but only obtained bounds on the price of anarchy that depend polynomially on the network size or on the ratio between the maximum and minimum edge capacities. Second, Banner and Orda [6] also recently studied selfish routing with the ℓ_∞ norm. However, the results of [6] primarily concern the existence and computation of equilibria, as well as the price of stability for an objective function that we do not consider in this paper. Lastly, this paper considers the "nonatomic" selfish routing model introduced by Wardrop [41], where all players are assumed to control a negligible fraction of the overall traffic, whereas [1, 6] consider "atomic" selfish routing with a finite set of players.

2 The Model

Instances. In this section we describe a model of selfish routing that includes elastic traffic and a potentially nonlinear aggregation function. By a selfish routing *instance*, we mean a triple (G, Γ, c) made up of the following ingredients. First, $G = (V, E)$ is a directed network with sources $s_1, \ldots, s_k \in V$ and sinks $t_1, \ldots, t_k \in V$. Second, Γ is a vector of nonincreasing, continuous functions indexed by source-sink pairs (or *commodities*) i; Γ_i models the distribution of the benefits of travel for the traffic of commodity i, and is assumed to be defined on the population $[0, R_i]$, where $R_i \in [0, \infty)$ is the size of the population. Assuming that Γ_i is nonincreasing amounts to ordering players according to their benefit of participating. Finally, c is a vector of nonnegative, continuous, nondecreasing *cost functions*, indexed by E.

Paths and Flows. For a network G, let $\mathcal{P}_i$ denote the s_i-t_i paths of G and let $\mathcal{P} = \cup_{i=1}^{k} \mathcal{P}_i$. A *flow* is a vector f indexed by $\mathcal{P}$. For a fixed flow f, we use r_i to denote the amount $r_i = \sum_{P \in \mathcal{P}} f_P$ of traffic of the ith commodity that is routed by f. We always assume that the flow represents those most interested in traveling, so the r_i units of traffic correspond to the subset $[0, r_i]$ of the entire population $[0, R_i]$. A flow is *feasible* for (G, Γ, c) if $r_i \leq R_i$ for all commodities i.

For a flow f, let $f_e = \sum_{P \in \mathcal{P} : e \in P} f_P$ denote the amount of traffic using the edge e. The cost of an edge e with respect to f is $c_e(f_e)$. If P is a path containing the edges $e_1, e_2, \ldots, e_m$ and f is a flow, then the *cost* $c_P(f)$ of a path P with respect to f is $\|c_{e_1}(f_{e_1}), \ldots, c_{e_m}(f_{e_m})\|$ for some aggregation function $\|\cdot\|$. In the traditional selfish routing model, $\|\cdot\|$ is the sum function. In this paper, we will allow $\|\cdot\|$ to be any ℓ_p norm $\|\cdot\|_p$ with $1 \le p \le \infty$, where, by definition, $\|v_1, \ldots, v_m\|_p = (v_1^p + \cdots v_m^p)^{1/p}$ if $p < \infty$ and $\|v_1, \ldots, v_m\|_p = \max_i v_i$ if $p = +\infty$. (Since we will only be taking the norm of nonnegative vectors, we omit the usual absolute value signs from these formulae.) We will sometimes call such an aggregation function a *path norm*.

Nash Flows. Intuitively, a flow is at *Nash equilibrium* (or is a *Nash flow*) if no player can do better by changing its mind—be it by switching paths or by switching whether or not to participate. Mathematically, we have the following definition.

Definition 2.1 A flow f that is feasible for (G, Γ, c) is at Nash equilibrium if:

(a) for every commodity i and paths $P, P' \in \mathcal{P}_i$ with $f_P > 0$, $c_P(f) \le c_{P'}(f)$;

(b) for every commodity i, the common cost $c_i(f)$ of all s_i-t_i flow paths is equal to $\Gamma(r_i)$.

Part (a) of Definition 2.1 is the usual condition that no player should be able to decrease its cost by switching paths. For part (b), first note that if f satisfies part (a), then $c_i(f)$ is well defined—if $f_P > 0$ and $f_{P'} > 0$ with $P, P' \in \mathcal{P}_i$, then $c_P(f) = c_{P'}(f)$. Part (b) then asserts that all participants enjoy benefit at least equal to their cost (since $\Gamma(a) \ge \Gamma(r_i) = c_i(f)$ for all $a \in [0, r_i]$), and similarly that all non-participants would incur at least as much cost as benefit if they did participate (since $\Gamma(a) \le \Gamma(r_i) = c_i(f)$ for all $a \notin [0, r_i]$).

Existence of Nash flows can be established in a number of ways; for example, it is a consequence of the very general results of Schmeidler [36].

Proposition 2.2 *Every instance* (G, Γ, c) *admits at least one Nash flow.*

In some parts of this paper, Nash flows will not be unique. In these cases, we will be interested in bounds on the performance of *all* of the Nash flows of an instance.

Remark 2.3 In Section 4 we will focus on instances with *inelastic traffic*. Such instances can be modeled with elastic traffic by defining the functions Γ to be sufficiently large everywhere. In Section 4 we will adopt the more direct approach of defining an instance with inelastic traffic via a triple (G, r, c), where the amount of traffic r_i routed by each commodity is now exogenous. The definition of a Nash flow is then merely part (a) of Definition 2.1.

3 No Tragedy of the Commons with Elastic Traffic

For our first main result, we will show that there is no tragedy of the commons with the usual additive aggregation function and elastic traffic, in the sense that for a natural prize-collecting objective, the price of anarchy in instances with elastic traffic is no more than that in instances with inelastic traffic. We begin with a discussion of objective functions and the price of anarchy in Section 3.1 before proceedings to the proof of this result in Section 3.2.

3.1 Preliminaries

As noted in the Introduction, in an instance (G, Γ, c) with elastic traffic there are two natural desiderata for a flow f: the cost $\sum_P c_P(f) f_P$ that f incurs should be small, while the benefit $\sum_i \int_0^{r_i} \Gamma(x)dx$ reaped should be large. As the single-link network in the Introduction demonstrates, no approximation bound is possible if one of these quantities is subtracted from the other, unless benefits are assumed to be large relative to costs as in Vetta [39]. We therefore study the *combined cost* $CC(f)$ of a flow f, defined as the cost added to the *lost* benefit:

$$CC(f) = \sum_P c_P(f) f_P + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x)dx. \tag{1}$$

This objective is inspired by the "prize-collecting" objectives that have been extensively studied in approximation algorithms; see, for example, the survey of Goemans and Williamson [19] for further background. One could also analogously consider "combined benefit"—the benefit earned plus the cost not incurred—but we believe this to be a less natural objective than the combined cost.

With the objective function now set, an *optimal flow* for an instance is simply one that minimizes the combined cost over all feasible flows. The *price of anarchy* $\rho(G, \Gamma, c)$ of an instance (G, Γ, c) is defined as the largest-possible ratio $CC(f)/CC(f^*)$, where f is a Nash flow and f^* is an optimal flow. Note that this definition makes sense even when Nash flows are not unique.

We will call instances with inelastic traffic and the additive aggregation function *basic instances*. The price of anarchy is well understood in such instances. Since we will only be studying models that generalize basic instances, our "holy grail" will be upper bounds on the price of anarchy that match those for basic instances.

We next review the known upper and lower bounds on the price of anarchy in basic instances; simple examples [34] show that such bounds must be parameterized by the allowable edge cost functions. Toward that end, define the *anarchy value* $\alpha(\mathcal{C})$ of a non-empty set of cost functions $\mathcal{C}$ to be:

$$\alpha(\mathcal{C}) = \sup_{c \in \mathcal{C}} \sup_{x, f \geq 0} \frac{f \cdot c(f)}{x \cdot c(x) + (f - x)c(f)}. \tag{2}$$

The anarchy value of $\mathcal{C}$ is essentially the worst-possible price of anarchy in a two-node, two-link network where one link has a constant cost function and the other link has a cost function in $\mathcal{C}$. Thus $\alpha(\mathcal{C})$ lower bounds the price of anarchy of basic instances with cost functions in $\mathcal{C}$, assuming only that $\mathcal{C}$ contains all of the constant cost functions. Looking ahead toward the next two sections, we note that this lower bound arises only from networks of parallel links, and that all of the ℓ_p norms coincide in such networks. Thus $\alpha(\mathcal{C})$ lower bounds the price of anarchy with respect to all ℓ_p norms, even when restricting attention to (single-commodity) networks of parallel links.

The following are known for basic instances [13, 31]. First, the price of anarchy of a (multicommodity) instance (G, r, c) with cost functions in a set $\mathcal{C}$ is at most $\alpha(\mathcal{C})$. Second, the value of $\alpha(\mathcal{C})$ is known for many natural sets $\mathcal{C}$: if $\mathcal{C}$ contains only linear or concave functions, then $\alpha(\mathcal{C}) \leq \frac{4}{3}$; if $\mathcal{C}$ contains only polynomials with nonnegative coefficients and degree at most d, then $\alpha(\mathcal{C}) = O(d/\log d)$. Qualitatively, these results imply that the price of anarchy is small in basic instances if and only if cost functions are not "extremely steep."

3.2 Bounding the Price of Anarchy with Elastic Traffic

The goal of this section is to prove that $\alpha(\mathcal{C})$, the known upper bound on the price of anarchy for basic instances with cost functions in $\mathcal{C}$, also upper bounds the price of anarchy of multicommodity instances with arbitrary elastic traffic and cost functions in $\mathcal{C}$, provided we use the combined cost objective function (1) and the ℓ_1 aggregation function.

We note that there is a naive reduction from networks with elastic traffic to networks with inelastic traffic—achieved by adding a new s_i-t_i link for each i to represent the traffic not routed—but this reduction produces a network with cost functions that depend on the distributions Γ (which in turn control the price of anarchy of the network). We will see that this dependence is unnecessary—the price of anarchy in instances with elastic traffic is determined only by the edge cost functions, and is independent of the distributions Γ.

Our proof is a generalization of the "variational inequality" proof technique of Roughgarden [31] and Correa, Schulz, and Stier Moses [13] to the elastic traffic case. Variational inequalities were first introduced as a tool to study traffic equilibria by Smith [37] and Dafermos [16]; for elastic traffic, the appropriate variational inequality is the following.

Proposition 3.1 *If f and f^* are Nash and feasible flows for (G, Γ, c), respectively, then*

$$\sum_{e \in E} c_e(f_e)[f_e^* - f_e] + \sum_{i=1}^{k} \Gamma(r_i)[r_i - r_i^*] \geq 0,$$

where $r_i = \sum_{P \in \mathcal{P}_i} f_P$ and $r_i^* = \sum_{P \in \mathcal{P}_i} f_P^*$.

Proposition 3.1 can be proved directly or via the convex programming approach pioneered by Beckmann, McGuire, and Winsten [7]. For details see, for example, Nagurney [26]. We can now prove the main result of this section. (See Appendix B for the proof.)

Theorem 3.2 *If (G, Γ, c) is an instance with elastic traffic and cost functions in the set $\mathcal{C}$, then the price of anarchy $\rho(G, \Gamma, c)$ is at most $\alpha(\mathcal{C})$.*

Remark 3.3 Theorem 3.2 also holds more generally for the nonatomic congestion games studied in [35].

Remark 3.4 The very first line of the proof of Theorem 3.2—which rewrites the first term of $CC(f^*)$ as an equivalent sum over edges—already crucially uses properties of the ℓ_1 path norm. Indeed, we will see in the next section that only restricted versions of Theorem 3.2 can hold for the ℓ_p path norms with $p > 1$, even with inelastic traffic.

Remark 3.5 The proof of Theorem 3.2 proves a stronger statement, that permits better upper bounds on the price of anarchy in instances where Nash flows do not route all of the traffic of every commodity. We defer further elaboration until the full version.

4 Nonlinear Path Norms

In this section we study the price of anarchy with nonlinear aggregation functions—in particular, with the ℓ_p path norms with $1 \leq p \leq \infty$. We will see that the price of anarchy of selfish routing behaves differently in each of the three cases of $p = 1$, $p \in (1, \infty)$, and $p = +\infty$. In Section 4.1 we exhibit two examples demonstrating the negative results promised in the Introduction; these examples will sculpt our goals for the rest of the section. In Section 4.2 we identify a natural subclass of Nash flows for the ℓ_∞ norm, justify them from a networking perspective, and prove optimal bounds on their inefficiency in single-commodity networks. In Section 4.3 we treat ℓ_p norms with $p < \infty$; while this case is technically more challenging than that of the ℓ_∞ norm, we will be rewarded with bounds on the inefficiency of arbitrary Nash flows in single-commodity networks.

Throughout this section, we consider only instances with inelastic traffic. In this context, our objective function is the *cost*, defined for a flow f feasible for an instance (G, r, c) as $C(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P$. This is the first term of the combined cost (1). The price of anarchy $\rho(G, r, c)$ is then defined in the usual way.

We will generalize the positive results of this section to networks with elastic traffic in Section 5.

4.1 Motivating Examples

We now give the two examples promised in the Introduction. The first shows that good bounds on the price of anarchy of selfish routing with the ℓ_p norm cannot be achieved in multicommodity networks, even in networks with linear cost functions and inelastic traffic, provided $p > 1$.

Example 4.1 Fix an ℓ_p norm with $1 < p \leq \infty$ and consider the two-commodity network shown in Figure 1 in Appendix A. For a parameter $k \geq 1$, there are k internally disjoint paths $s_1 \to v_i \to w_i \to t_1$ ($i \in \{1, 2, \ldots, k\}$). Edges (v_i, w_i) have the cost function $c(x) = x$; other edges in these paths cost 0. There are $k - 1$ *cross edges* (w_i, v_{i+1}) ($i \in \{1, 2, \ldots, k-1\}$), each with cost 0. The second source s_2 is connected to v_1 with a zero-cost edge, and w_k is connected to t_2 with a zero-cost edge. Finally, there is a direct s_2-t_2 edge with constant cost $c(x) = (r_2 + 1)k^{1/p}$, where r_2 is the traffic rate of the second commodity, which is a function of k and p that we will define shortly. (If $p = \infty$, we interpret $1/p$ as 0.) The traffic rate r_1 of the first commodity is k.

First, consider the flow f^* that routes the traffic of the first commodity evenly across the k s_1-t_1 paths, and routes the second commodity's traffic on the direct s_2-t_2 link. The cost of f^* is $k + r_2 \cdot (r_2 + 1)k^{1/p}$. Next, by the choice of the cost of the direct s_2-t_2 link, the following flow f is at Nash equilibrium: route the first commodity's traffic evenly across the s_1-t_1 paths and the second commodity's traffic on the s_2-t_2 path containing all of the cross edges. The cost of f is $k \cdot (r_2 + 1) + r_2 \cdot (r_2 + 1)k^{1/p}$. The price of anarchy in the network is at least $C(f)/C(f^*)$; choosing r_2 so that $r_2(r_2 + 1) = k^{1-1/p}$, this ratio is $\Omega(k^{(1-1/p)/2})$. Since $n = O(k)$, this ratio grows polynomially in the network size for every fixed $p > 1$.

Finally, note that doubling the traffic rates increases the cost of the optimal flow by only a constant factor, so the bicriteria bound of [34]—stating that a Nash flow is no more expensive than an optimal flow at double the traffic rates, even with arbitrary cost functions—does not hold in this network.

Our second example shows that even in single-commodity networks with linear cost functions and inelastic traffic, there are no good bounds on the price of anarchy for the ℓ_∞ norm.

Example 4.2 Suppose we modify the network of Example 4.1 by removing s_2, t_2, and edges incident to them. This yields the network of Figure 2. There is a unique s-t path that contains all of the cross edges; call it the *zigzag path*. With respect to the ℓ_∞ norm, the flow f that routes all traffic on the zigzag path is at Nash equilibrium—all s-t paths have cost k with respect to f and the ℓ_∞ norm—and has cost k^2. On the other hand, routing traffic evenly among the k three-hop paths provides a flow with cost k. (This flow is also at Nash equilibrium.) The price of anarchy in this network is therefore at least k.

As with Example 4.1, the bicriteria bound of [34] also fails in this example.

Examples 4.1 and 4.2 justify restricting our attention to single-commodity networks and, for the ℓ_∞ norm, to natural subclasses of equilibria.

4.2 The ℓ_∞ Norm and Subpath-Optimal Nash Flows

We next consider single-commodity networks with the ℓ_∞ norm. Example 4.2 shows that additional restrictions are needed to prove a good bound on the price of anarchy. We will require only a modest extra condition on Nash flows, stating the Nash flow condition holds not only for the destination t, but also for all intermediate nodes v.

Definition 4.3 Suppose (G, r, c) is a single-commodity instance with inelastic traffic and the ℓ_∞ path norm. Let f be a flow feasible for (G, r, c) and let $d(v)$ denote the minimum cost, with respect to f and the ℓ_∞ norm, of an s-v path. The flow f is a *subpath-optimal Nash flow* if whenever an s-t path $P \in \mathcal{P}$ with $f_P > 0$ includes a vertex v, the s-v subpath of P has ℓ_∞ norm $d(v)$.

To see that a subpath-optimal Nash flow is indeed a Nash flow, take $v = t$. Note also that the zigzag Nash flow of Example 4.2 is not subpath-optimal, while the optimal flow is. Finally, notions similar to Definition 4.3 were also proposed, for different purposes and without any networking justification, in [1, 6].

The next proposition is meant to suggest that no extra notion of coordination or centralized intervention is required to justify subpath-optimal Nash flows—indeed, processes that result in a Nash flow automatically ensure the subpath-optimality condition. In the interests of space, we will state and discuss the proposition somewhat informally.

Proposition 4.4 *A fixed point of a distributed shortest-path routing protocol is a subpath-optimal Nash flow.*

By a distributed shortest-path routing protocol, we mean a Bellman-Ford-type shortest-path algorithm—in networking jargon, a "distance vector protocol" such as OSPF (see e.g. [24]). Proposition 4.4 assumes that the cost functions $\{c_e(\cdot)\}$ are used for edge lengths, and that path lengths are evaluated using the ℓ_∞ norm. Such a shortest-path routing protocol naturally computes not just shortest s-t paths, but also shortest s-v paths for all intermediate vertices v, which is precisely the subpath-optimality condition. We defer further details until the full version. For more details on distributed shortest-path routing protocols and on results along the lines Proposition 4.4, see Bertsekas and Tsitsiklis [8].

We now prove bounds on the inefficiency of subpath-optimal Nash flows. Examples 4.1 and 4.2 show that our proof techniques must make crucial use of both the single-commodity and the subpath-optimal assumptions. Since all previous proof techniques for bounding the price of anarchy of selfish routing (with the ℓ_1 path norm) worked equally well for single-commodity networks and the much more general nonatomic congestion games [11, 13, 14, 27, 31, 34, 35, 38], we will require an intrinsically more combinatorial argument.

We first prove a lemma that identifies a type of "minimal cut" with respect to a subpath-optimal Nash flow. Loosely, we will then treat the edges crossing this cut as a network of parallel links, which will enable us to prove both bounds on the price of anarchy as well as an analogue of the bicriteria bound of [34]. We defer all of the proofs until Appendix B.2.

In the statement of the lemma, we use the notation $\delta^+(S)$ ($\delta^-(S)$), where S is a set of vertices, to denote the edges with tail (head) in S and head (tail) outside S.

Lemma 4.5 *Let (G, r, c) be a single-commodity instance with inelastic traffic and the ℓ_∞ path norm. Let f be a subpath-optimal Nash flow for (G, r, c) in which all flow paths of f have cost $c(f)$, and let S be the set of vertices reachable from the source s via edges with cost strictly less than $c(f)$. Then:*

(a) *S is an s-t cut;*

(b) $c_e(f_e) \geq c(f)$ *for all* $e \in \delta^+(S)$*;*

(c) $c_e(f_e) = c(f)$ *for all* $e \in \delta^+(S)$ *with* $f_e > 0$*;*

(d) $f_e = 0$ *for all* $e \in \delta^-(S)$.

Theorem 4.6 *Let (G, r, c) be a single-commodity instance with cost functions in $\mathcal{C}$, f a subpath-optimal Nash flow under the ℓ_∞ norm, and f^* a feasible flow. Then $C(f) \leq \alpha(\mathcal{C}) \cdot C(f^*)$.*

As discussed in Section 3.1, the upper bound in Theorem 4.6 is the best possible. Simple examples [34] also show that following bicriteria bound is also optimal.

Theorem 4.7 *Let (G, r, c) be a single-commodity instance and f a subpath-optimal Nash flow under the ℓ_∞ norm. If f^* be feasible for $(G, 2r, c)$, then $C(f) \leq C(f^*)$.*

4.3 The ℓ_p Norms

In this section we extend Theorems 4.6 and 4.7 to the ℓ_p path norms with $p < \infty$. The proofs for the $p < \infty$ case are more involved than those for the $p = \infty$ case. In particular, Lemma 4.5—which essentially reduced the problem of bounding the inefficiency of a Nash flow to that of bounding its inefficiency across a single, well-behaved cut—has only weak analogues for the ℓ_p norms with $p < \infty$. Specifically, we will need to bound the cost of a Nash flow across *many* cuts, and then aggregate the results into a bound on the overall cost of the flow. Because we work with fairly general path norms, the aggregation step is somewhat delicate. On the other hand, since Nash flows under the ℓ_p norm with $p < \infty$ are automatically subpath-optimal, we can prove bounds on their cost without any extra restrictions.

Due to space constraints, we state the main results and defer proof sketches to Appendix B.2.

Theorem 4.8 *Let (G, r, c) be a single-commodity instance with the ℓ_p norm ($p < \infty$) and cost functions in $\mathcal{C}$. If f and f^* are Nash and feasible flows for (G, r, c), respectively, then $C(f) \le \alpha(\mathcal{C}) \cdot C(f^*)$.*

Theorem 4.9 *Let (G, r, c) be a single-commodity instance with the ℓ_p norm ($p < \infty$). If f and f^* are Nash and feasible flows for (G, r, c) and $(G, 2r, c)$, respectively, then $C(f) \le C(f^*)$.*

Remark 4.10 Theorems 4.8 and 4.9 can be generalized to all path norms that satisfy certain symmetry properties, but we not aware of any compelling norms that satisfy these properties other than the ℓ_p norms.

5 No Tragedy of the Commons with Nonlinear Path Norms

In this final section we investigate whether or not the positive results of Sections 3 and 4 can be combined. Our answer will be somewhat incomplete. In particular, we will answer in the affirmative only when a Nash flow sends at least as much traffic through the network as an optimal flow. Simple examples—including networks of parallel links, and by definition all networks that illustrate the tragedy of the commons—satisfy this condition. Unfortunately, examples related to Braess's Paradox [9, 30] show that it does not hold for arbitrary single-commodity networks. We do not know how to analyze the price of anarchy for instances with elastic traffic and nonlinear path norms in which this technical condition does not hold.

Proposition 5.1 *Let (G, Γ, c) be a single-commodity instance with elastic traffic, the ℓ_∞ norm, and cost functions in the set $\mathcal{C}$. If f and f^* are subpath-optimal Nash and optimal flows for (G, Γ, c), respectively, that satisfy $\sum_P f_P \ge \sum_P f^*_P$, then $C(f) \le \alpha(\mathcal{C}) \cdot C(f^*)$.*

Proposition 5.2 *Let (G, Γ, c) be a single-commodity instance with elastic traffic, the ℓ_p norm ($p < \infty$), and cost functions in the set $\mathcal{C}$. If f and f^* are Nash and optimal flows for (G, Γ, c), respectively, that satisfy $\sum_P f_P \ge \sum_P f^*_P$, then $C(f) \le \alpha(\mathcal{C}) \cdot C(f^*)$.*

Propositions 5.1 and 5.2 follow fairly easily from Theorems 4.6 and 4.8, respectively, from a common proof.

Proof of Propositions 5.1 and 5.2: Consider an instance (G, Γ, c) and flows f, f^* of the specified form. By assumption, $\sum_P f_P = r \ge r^* = \sum_P f^*_P$. Let Γ' denote Γ restricted to $[0, r]$. Then f and f^* are Nash and feasible flows for (G, Γ', c), respectively, and the cost of each has dropped by an additive factor of $\int_r^R \Gamma(x)dx$. Hence the price of anarchy of (G, Γ', c) can only be larger than that of (G, Γ, c).

Next, replace Γ' be the function $\overline{\Gamma}$ that is everywhere equal to $\Gamma(r)$. The flow f is still at Nash equilibrium for $(G, \overline{\Gamma}, c)$ and its cost is unchanged. Since $\overline{\Gamma} \le \Gamma'$, the cost of f^* can only decrease and the price of anarchy again only increases.

Finally, reduce $(G, \overline{\Gamma}, c)$ to a network (G', r, c) with inelastic traffic as discussed at the beginning of Section 3.2. (That discussion was in the context of the ℓ_1 path norm, but the reduction works equally well for all ℓ_p norms.) Note that G' is obtained from G be adding a direct s-t link with constant cost $\Gamma(r)$. Thus (G', r, c) is a network with inelastic traffic and cost functions that are either constant or in $\mathcal{C}$. It is easy to verify that adding constant functions to a set does not affect its α-value [31], and thus Theorem 4.6 or 4.8 implies that the cost of f in (G', r, c) is at most $\alpha(\mathcal{C})$ times that of f^* in (G', r, c). Since the above reductions only increased the ratio of costs of these two flows, the cost of f in (G, Γ, c) is at most $\alpha(\mathcal{C})$ times that of f^* in (G, Γ, c). ■

References

[1] A. Akella, S. Chawla, and S. Seshan. Realistic models for selfish routing in the Internet. Manuscript, 2003.

[2] 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.

[3] 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.

[4] 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)*, 2005. To appear.

[5] A. Banerjee and J. Yoo. Minimizing maximum logical link congestion in packet switched optimal networks. In *Proceedings of the 1997 IEEE International Conference on Communcations (ICC)*, volume 3, pages 1298–1302, 1997.

[6] R. Banner and A. Orda. Bottleneck games in noncooperative networks. Technical Report CCIT-510, Technion University, 2004.

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

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

[9] D. Braess. Über ein Paradoxon aus der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968.

[10] J. H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proceedings of INFOCOM*, pages 22–31, 2000.

[11] 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.

[12] 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.

[13] 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.

[14] 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)*, 2005. To appear.

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

[16] S. C. Dafermos. Traffic equilibrium and variational inequalities. *Transportation Science*, 14(1):42–54, 1980.

[17] 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.

[18] N. H. Gartner. Optimal traffic assignment with elastic demands: A review; Part I. Analysis framework. *Transportation Science*, 14(2):174–191, 1980.

[19] M. X. Goemans and D. P. Williamson. The primal-dual method for approximation algorithms and its application to network design problems. In D. S. Hochbaum, editor, *Approximation Algorithms for NP-Hard Problems*, chapter 4, pages 144–191. PWS Publishing Company, 1997.

[20] 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.

[21] G. Hardin. The tragedy of the commons. *Science*, 162:1243–1248, December 13 1968.

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

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

[24] S. Keshav. *An Engineering Approach to Computer Networking*. Addison-Wesley, 1997.

[25] 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.

[26] A. Nagurney. *Network Economics: A Variational Inequality Approach*. Kluwer, 1993. Second Edition, 1999.

[27] 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.

[28] 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.

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

[30] 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.

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

[32] 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.

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

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

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

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

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

[38] É. Tardos. Course notes. Cornell CS684, 2004.

[39] 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.

[40] Y. Wang and Z. Wang. Explicit routing algorithms for Internet traffic engineering. In *Proceedings of the 8th International Conference on Computer Communications and Networks (ICCN)*, pages 582–588, 1999.

[41] 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.

A Figures

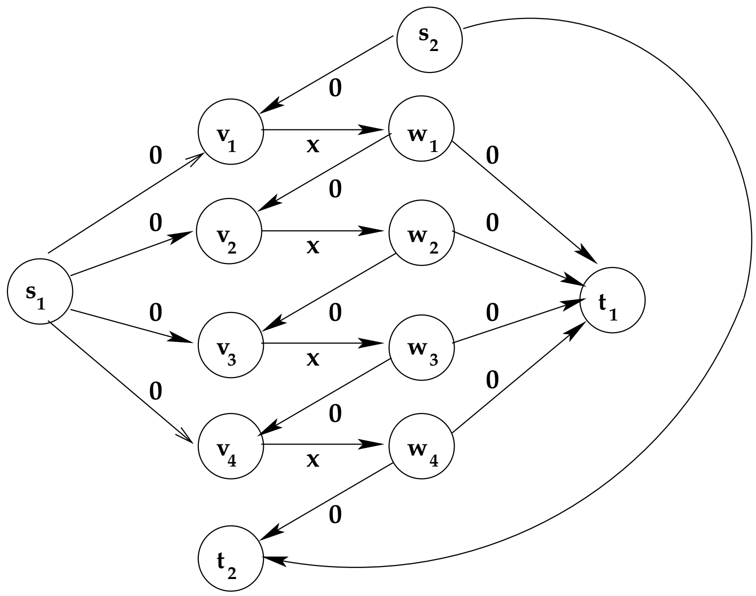

Figure 1: A bad two-commodity example for the ℓ_p norm with $p > 1$.

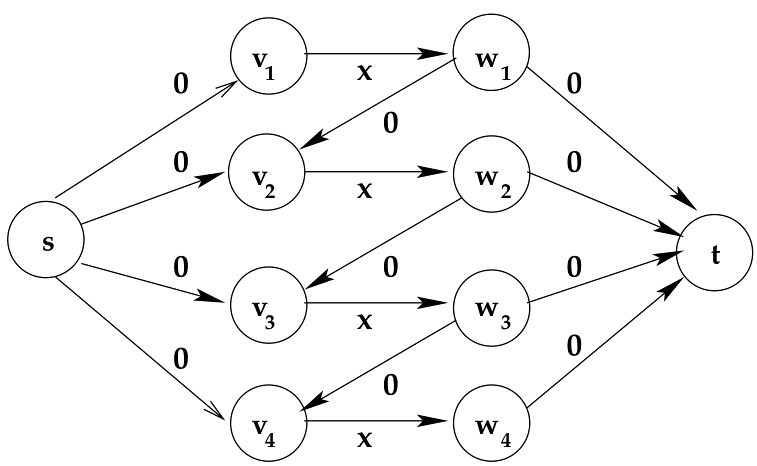

Figure 2: A bad single-commodity example for the ℓ_∞ norm.

B Missing Proofs

B.1 Missing Proofs from Section 3

Proof of Theorem 3.2: Fix an instance (G, Γ, c) with cost functions in a set $\mathcal{C}$, a Nash flow f, and a feasible flow f^*. For each commodity i, let $r_i = \sum_{P \in \mathcal{P}_i} f_P$ and $r_i^* = \sum_{P \in \mathcal{P}_i} f_P^*$. First write

$$\begin{aligned} C(f^*) &= \sum_{e \in E} c_e(f_e^*) f_e^* + \sum_{i=1}^{k} \int_{r_i^*}^{R_i} \Gamma(x) dx \\ &= \sum_{e \in E} c_e(f_e^*) f_e^* + \sum_{i=1}^{k} \int_{r_i^*}^{r_i} \Gamma(x) dx + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x) dx \\ &\geq \sum_{e \in E} c_e(f_e^*) f_e^* + \sum_{i=1}^{k} [r_i - r_i^*] \Gamma(r_i) + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x) dx, \end{aligned}$$

where the first equality follows from a reversal of sums ($\sum_P c_P(f^*) f_P^* = \sum_e c_e(f_e^*) f_e^*$) and the inequality follows from the fact that each function Γ_i is nonincreasing. (This inequality holds both if $r_i^* \geq r_i$ and if $r_i \geq r_i^*$.)

Next, the definition (2) of $\alpha(\mathcal{C})$ yields, for each edge e (with $r = f_e$ and $x = f_e^*$):

$$c_e(f_e^*) f_e^* \geq \frac{c_e(f_e) f_e}{\alpha(\mathcal{C})} + c_e(f_e)[f_e^* - f_e].$$

Hence

$$\begin{aligned} C(f^*) &\geq \sum_{e \in E} \left[\frac{c_e(f_e) f_e}{\alpha(\mathcal{C})} + c_e(f_e)[f_e^* - f_e] \right] + \sum_{i=1}^{k} [r_i - r_i^*] \Gamma(r_i) + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x) dx \\ &= \frac{1}{\alpha(\mathcal{C})} \sum_{e \in E} c_e(f_e) f_e + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x) dx + \sum_{e \in E} c_e(f_e)[f_e^* - f_e] + \sum_{i=1}^{k} [r_i - r_i^*] \Gamma(r_i) \\ &\geq \frac{1}{\alpha(\mathcal{C})} \sum_{e \in E} c_e(f_e) f_e + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x) dx \\ &= \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} c_P(f) f_P + \sum_{i=1}^{k} \int_{r_i}^{R_i} \Gamma(x) dx, \end{aligned} \tag{3}$$

where (3) follows from Proposition 3.1. Since $\alpha(\mathcal{C}) \geq 1$ for every set $\mathcal{C}$, inequality (3) implies that $C(f^*) \geq C(f)/\alpha(\mathcal{C})$, and the proof is complete. ■

B.2 Missing Proofs from Section 4

Missing Proofs from Section 4.2

Proof of Lemma 4.5: Parts (a) and (b) follow from the definitions. Part (c) follows from part (b) and the fact that if all flow paths of f have cost $c(f)$, then $c_e(f_e) \leq c(f)$ for all edges e with $f_e > 0$. For part (d), suppose for contradiction that $e = (v, w) \in \delta^-(S)$ with $f_e > 0$. Let $P \in \mathcal{P}$ be a path with $f_P > 0$ and $e \in P$. Recall that $d(u)$ denotes the minimum cost (w.r.t. f and the ℓ_∞ norm) of an s-u path. By the definition of S, $d(u) < c(f)$ for all $u \in S$; in particular, $d(w) < c(f)$.

Let P' be the s-w subpath of P, which concludes with the edge e. Since $e \in \delta^-(S)$ and $s \in S$, an earlier edge of P' lies in $\delta^+(S)$. By the definition of S, $c_{e'}(f_{e'}) \geq c(f)$ for all $e' \in \delta^+(S)$ (otherwise S could be enlarged), so the ℓ_∞ norm of P' is at least $c(f) > d(w)$. But this contradicts the subpath-optimality of f. ■

Proof of Theorem 4.6: Define the s-t cut S as in Lemma 4.5. We now define an instance on a network of parallel links. Let $V' = \{s', t'\}$ and let E' be a set of parallel edges (all directed from s' to t') in one-to-one correspondence with the edges of $\delta^+(S)$. Edges of E' inherit cost functions c from their counterparts in $\delta^+(S)$.

Let $G' = (V', E')$ and consider the instance (G', r, c). Parts (a) and (d) of Lemma 4.5 imply that f routes precisely r units of flow on the edges of $\delta^+(S)$; it therefore naturally induces (by projection) a flow g feasible for (G', r, c). Moreover, parts (b) and (c) of Lemma 4.5 imply that g is a Nash flow for (G', r, c) with cost $r \cdot c(f)$—the same cost as f in (G, r, c). Note that when we discuss the cost of flows in the network of parallel links (G', r, c), the path norm is irrelevant.

We now discuss f^*, which might route strictly more than r units of flow on the edges of $\delta^+(S)$ (if some flow path of f^* contains more than one edge of $\delta^+(S)$). In this case, we define $g^*_e \leq f^*_e$ for all $e \in \delta^+(S)$ in the following way: a path $P \in \mathcal{P}$ with $f^*_P > 0$ only contributes to the g^*-value of the most expensive (largest value of $c_e(f^*_e)$) edge in $P \cap \delta^+(S)$, with ties broken arbitrarily. (Since P is an s-t path and S is an s-t cut, $P \cap \delta^+(S) \neq \emptyset$.) Then g^* can be viewed as a flow feasible for (G', r, c) satisfying

$$\sum_{e \in E'} c_e(g^*_e) g^*_e \leq \sum_{e \in E'} c_e(f^*_e) g^*_e \sum_{P \in \mathcal{P}} c_P(f^*) f^*_P = C(f^*);$$

thus the cost of g^* in (G', r, c) is at most that of f^* in (G, r, c).

We have established that the price of anarchy in (G, r, c) is at most that in (G', r, c). But since the latter instance is a network of parallel links and can therefore be viewed as basic, its price of anarchy is at most $\alpha(\mathcal{C})$. ■

Proof of Theorem 4.7: The proof of Theorem 4.7 is similar to that Theorem 4.6 and is omitted. ■

Missing Proofs from Section 4.3

We now outline the proofs of Theorems 4.8 and 4.9. The first step is to linearly order the vertices of a network so that the cost of a Nash flow breaks down nicely across several cuts. The following three propositions were previously known for the ℓ_1 norm [30], and the proofs for the general ℓ_p case are similar.

Proposition B.1 *Let (G, r, c) be a single-commodity instance with path norm ℓ_p for some $p \in [1, \infty)$. Let f be a flow feasible for (G, r, c) and for a vertex v, let $d(v)$ denote the minimum-norm of an s-v path with respect to f.*

- *(a) For every edge $e = (v, w)$, $d(w) \leq (d(v)^p + c_e(f_e)^p)^{1/p}$.*
- *(b) The flow f is at Nash equilibrium if and only if $d(w) = (d(v)^p + c_e(f_e)^p)^{1/p}$ whenever $e = (v, w)$ is an edge with $f_e > 0$.*

Proposition B.2 *Let (G, r, c) be a single-commodity instance with path norm ℓ_p for some $p \in [1, \infty)$. If f is a Nash flow for (G, r, c), then there is an acyclic Nash flow $\tilde{f}$ with $C(\tilde{f}) = C(f)$.*

Proposition B.3 *Let (G, r, c) be a single-commodity instance with path norm ℓ_p for some $p \in [1, \infty)$, and let f be an acyclic Nash flow. Define $d(v)$ as in Proposition B.1. Then the vertices of G can be sorted topologically w.r.t. the flow f such that s comes and first and such that the values $d(v)$ are nondecreasing in the ordering.*

Proposition B.3 gives a *monotone ordering* of the vertices that induces $n-1$ cuts (the first vertex, the first two vertices, and so on). Monotone orderings have emerged as a basic tool for analyzing single-commodity selfish routing networks, and have been used previously in the context of Braess's Paradox and related issues [30, 32], as well as for a pricing problem [12]. Monotone orderings were only used in [12, 30, 32] for the ℓ_1 path norm, however, and these prior works used them only to prove inequalities on the relative amounts of flow on different edges (as opposed to the cost of flows, which we are concerned with here). Moreover, for bounding Braess's Paradox and related quantities [30, 32], it sufficed to consider a *single* good cut; this is similar in spirit to our proofs of Theorems 4.6 and 4.7. Here, in a price of anarchy context, we will use these orderings to identify a *sequence* of good cuts, which we will then analyze separately and combine the results. Because we must aggregate the analyses of many cuts with respect to a fairly general path norm, the following analysis will be more involved than the previous applications of monotone orderings [12, 30, 32].

The next definition provides a sequence of cost functions, where the ith set of cost functions is designed to isolate the cost of a flow across the ith cut of the network. This would be easy if each edge of the network participated in at most one such cut, but this of course need not be the case. There are two properties we will require of these cost functions, which work in opposition to each other. First, the ith set of cost functions should be "uniform" in some sense. Second, the cost of an edge should be accurately accounted for over the sequence of cost functions, no matter how many cuts of the network the edge participates in.

Definition B.4 Let (G, r, c) be a single-commodity instance with path norm ℓ_p for some $p \in [1, \infty)$, and let f be an acyclic Nash flow. Define $d(v)$ as in Proposition B.1 and sort the vertices $v_1, \ldots, v_n$ as in Proposition B.3. For an edge e and an integer $i \in \{1, 2, \ldots, n-1\}$, the ith cost function $c_e^{(i)}$ of e is defined as follows:

- if $e \notin \delta^+(\{v_1, v_2, \ldots, v_i\})$, then $c_e^{(i)}$ is zero everywhere;
- if $e \in \delta^+(\{v_1, v_2, \ldots, v_i\})$, then $c_e^{(i)} = \lambda_e^{(i)} c_e$, where $\lambda_e^{(i)}$ is the unique number such that

$$d(v_{i+1}) = (d(v_i)^p + [\lambda_e^{(i)} c_e(f_e)]^p)^{1/p}.$$

For a flow f^* feasible for (G, r, c), we then define $c_P^{(i)}(f^*) = (\sum_{e \in P} [c_e^{(i)}(f_e^*)]^p)^{1/p}$.

The second part of Definition B.4 is not well defined if $c_e(f_e) = 0$; in this case Propositions B.1(a) and B.3 imply that $d(v) = d(w)$, so we can take $\lambda_e^{(i)} = 0$.

The next lemma states that the cost functions $c_e^{(i)}$ do indeed serve the purpose described prior to Definition B.4. We defer the proof to the full version.

Lemma B.5 *With the assumptions and notation of Definition B.4, the following statements hold.*

(a) For each $i = 1, 2, \ldots, n-1$, there is a constant $A_i > 0$ such that $c_P^{(i)}(f) \geq A_i$ for all $P \in \mathcal{P}$, with equality holding whenever $f_P > 0$.

(b) For each $i = 1, 2, \ldots, n-1$, $d(v_{i+1}) = \|A_1, \ldots, A_i\|$, where each A_j is defined as in (a).

(c) For every feasible flow f^* for (G, r, c) and every path $P \in \mathcal{P}$, $c_P(f^*) \geq \|c_P^{(1)}(f^*), \ldots, c_P^{(n-1)}(f^*)\|$.

The proofs of the next two lemmas, which analyze the inefficiency of a Nash flow according to a single set of the cost functions described Definition B.4, are analogous to those for Theorems 4.6 and 4.7, respectively, where Lemma B.5(a) plays the role originally served by Lemma 4.5. In the statements of the lemmas, we will use $C^{(i)}$ to denote the cost of a flow with respect to the cost functions $c^{(i)}$.

Lemma B.6 *Let (G, r, c) be a single-commodity instance with the ℓ_p path norm ($p < \infty$) and cost functions in the set $\mathcal{C}$. Let f and f^* be acyclic Nash and feasible flows for (G, r, c), respectively, and define the cost functions $c_e^{(i)}$ as in Definition B.4. Then for each $i = 1, 2, \ldots, n-1$,*

$$C^{(i)}(f^*) \geq \frac{C^{(i)}(f)}{\alpha(\mathcal{C})}.$$

Lemma B.7 *Let (G, r, c) be a single-commodity instance with the ℓ_p path norm ($p < \infty$). Let f be an acyclic Nash flow for (G, r, c) and f^* a feasible flow for $(G, 2r, c)$, respectively, and define the cost functions $c_e^{(i)}$ as in Definition B.4. Then for each $i = 1, 2, \ldots, n-1$,*

$$C^{(i)}(f^*) \geq C^{(i)}(f).$$

The proof for the ℓ_∞ case was complete at this point. For the ℓ_p path norms with $p < \infty$, the most crucial step lies directly ahead: aggregating the "cut-by-cut" bounds of Lemmas B.6 and B.7 into a bound for the entire network. Next, we use such an aggregation to prove Theorem 4.8.

Proof of Theorem 4.8: First, note that by Proposition B.2 there is no loss of generality in assuming that f is acyclic; we can then define the cost functions $c_e^{(i)}$ as in Definition B.4.

Since the only tool at our disposal for relating the costs of f and f^* is Lemma B.6, which relates their "cut costs", our first goal will be to use Lemma B.5 to express the cost $C(f^*)$ of f^* in terms of these cut costs. Specifically, write

$$\begin{aligned}
C(f^*) &= \sum_{P \in \mathcal{P}} f_P^* c_P(f^*) \\
&\geq \sum_{P \in \mathcal{P}} f_P^* \left\| c_P^{(1)}(f^*), \ldots, c_P^{(n-1)}(f^*) \right\| && (4) \\
&= \sum_{P \in \mathcal{P}} \left\| f_P^* \cdot c_P^{(1)}(f^*), \ldots, f_P^* \cdot c_P^{(n-1)}(f^*) \right\| && (5) \\
&\geq \left\| \sum_{P \in \mathcal{P}} f_P^* \cdot c_P^{(1)}(f^*), \ldots, \sum_{P \in \mathcal{P}} f_P^* \cdot c_P^{(n-1)}(f^*) \right\|, && (6)
\end{aligned}$$

where (4) follows from Lemma B.5(c), and (5) and (6) follow since ℓ_p is a norm (and thus satisfies linearity under scalar multiplication and the Triangle inequality). Applying Lemma B.6 and the monotonicity of $\|\cdot\|$, we then obtain

$$C(f^*) \geq \left\| \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot c_P^{(1)}(f), \ldots, \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot c_P^{(n-1)}(f) \right\|. \tag{7}$$

Finally, we aim to reverse the first argument to recover the cost of the subpath-optimal Nash flow f. Since the Triangle inequality of $\|\cdot\|$ is only useful in one direction, we will rely on the stronger assertions of Lemma B.5, which are tailored for Nash flows, to accomplish this. Precisely, we have, for some nonnegative constants $A_1, \ldots, A_{n-1}$,

$$C(f^*) \geq \left\| \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot A_1, \ldots, \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot A_{n-1} \right\| \tag{8}$$

$$= \frac{\sum_{P \in \mathcal{P}} f_P}{\alpha(\mathcal{C})} \|A_1, \ldots, A_{n-1}\| \tag{9}$$

$$= \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot d(t) \tag{10}$$

$$= \frac{C(f)}{\alpha(\mathcal{C})}, \tag{11}$$

where (8) follows from (7) and Lemma B.5(a), (9) from the linearity of $\|\cdot\|$ under scalar multiplication, (10) from Lemma B.5(b) with $i = n$, and (11) from Definition 2.1 and the definition of $d(t)$. This completes the proof. ■

A similar proof, combined with Lemma B.7, establishes Theorem 4.9.