

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

Richard Cole[†]

Yevgeniy Dodis[‡]

Tim Roughgarden[§]

March 31, 2011

Abstract

We study the price of anarchy of selfish routing with variable traffic rates and when the path cost is a non-additive function of the edge costs. Non-additive path costs are important, for example, in networking applications, where a key performance metric is the achievable throughput along a path, which is controlled by its bottleneck (most congested) edge. We prove the following results.

- In multicommodity networks, the worst-case price of anarchy under the ℓ_p path cost with $1 < p \leq \infty$ can be dramatically larger than under the standard ℓ_1 path cost.
- In single-commodity networks, the worst-case price of anarchy under the ℓ_p path cost with $1 < p < \infty$ is no more than with the standard ℓ_1 path norm. (A matching lower bound follows trivially from known results.) This upper bound also applies to the ℓ_∞ path cost if and only if attention is restricted to the natural subclass of equilibria generated by distributed shortest-path routing protocols.
- For a natural cost-minimization objective function, the price of anarchy with endogenous traffic rates (and under any ℓ_p path cost) is no larger than that in fixed-demand networks. Intuitively, the worst-case inefficiency arising from the “tragedy of the commons” is no more severe than that from routing inefficiencies.

Keywords: Routing; price of anarchy; traffic equilibria; shortest-path protocols

1 Introduction

1.1 The Price of Anarchy and Variable Demand

The *price of anarchy* is the worst-case ratio between the objective function values of a Nash equilibrium and an optimal outcome of a game, and is an important quantitative measure of the inefficiency

*An extended abstract of this paper appeared in the Proceedings of the 17th Annual ACM-SIAM Symposium on Discrete Algorithms, January 2006.

[†]Department of Computer Science, New York University, 251 Mercer Street, New York, NY 10012. This research was supported in part by NSF grants CCF-0515127 and CCF-0830516. Email: cole@cs.nyu.edu.

[‡]Department of Computer Science, New York University, 251 Mercer Street, New York, NY 10012. This research was 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. This research was supported in part by ONR grant N00014-04-1-0725, DARPA grant W911NF-04-9-0001, and an NSF Postdoctoral Fellowship. Email: tim@cs.stanford.edu.

of selfish behavior. Over the past ten years, the price of anarchy has been successfully analyzed in a diverse array of applications, such as scheduling, routing, facility location, network design, and resource allocation (see [18, Chapters 17–21]). The primary research agenda of these works is to identify problem domains and conditions under which the price of anarchy is guaranteed to be close to 1, and hence selfish behavior causes only modest efficiency loss.

Much of this previous work studies 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. We expect that the demand for an uncongested resource will be high, and that this demand will fall as the resource becomes more congested. Allowing variable demand inevitably leads to a trade-off 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 trade-off with a stark, famous example: *the tragedy of the commons* [15].

1.2 The Tragedy of the Commons

The tragedy of the commons refers to a shared resource that is destroyed by overconsumption. In lieu of the traditional bovine example [15], we illustrate the 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 non-negative, 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: then a $1 - \epsilon$ fraction of the population enjoys 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 most previous works on the price of anarchy of “selfish routing”, travel is mandatory for all agents — that is, the amount of network traffic is exogenous (fixed), rather than endogenous as in the above example. This brings us to the first goal of this paper.

- (1) Quantify the inefficiency of selfish routing when traffic rates depend on the network congestion.

1.3 Bottleneck Links and Nonlinear Aggregation Functions

Question (1) illuminates a second issue with previous studies of the price of anarchy of selfish routing: in the standard model, each network edge is given a congestion-dependent cost function, and the cost of a path is defined using the additive aggregation function, as the sum of its edges’ costs. This aggregation function may be the most natural one, but it is not appropriate 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. The studies of Qiu et al. [19] and Akella, Chawla, and Seshan [1] single out the choice of the additive aggregation function over the

bottleneck link metric as a key disconnect between the standard selfish routing model and typical concerns in networking applications.

The bottleneck link metric corresponds to using the ℓ_∞ norm as the aggregation function. Banner and Orda [5] 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 [9]. The ℓ_∞ norm also arises when robustness to bursty traffic [4] or to growing demand [27] is a priority. These types of applications motivate the second goal of this paper.

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

1.4 Our Results: Nonlinear Aggregation Functions

We give several tight bounds on the price of anarchy in single- and multicommodity selfish routing networks, with variable demand and nonlinear aggregation functions. We focus on the ℓ_p norms for $1 \leq p \leq \infty$.

We first describe our matching positive and negative results for the second goal (2). On the negative side, we give examples in Section 3.1 that demonstrate the following.

- For every $1 < p \leq \infty$, there is a family of two-commodity networks with linear cost functions and fixed traffic rates 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 arbitrary networks [23].) The “bicriteria bound” of [23] — asserting that with arbitrary cost functions, the cost of an equilibrium is no more than that of an optimal solution with double the traffic — also fails to hold.
- For the ℓ_∞ norm, there is a family of single-commodity networks with linear cost functions and fixed traffic rates in which the price of anarchy grows polynomially with the network size, and in which the bicriteria bound of [23] does not hold.

We present the following matching positive results for the case of fixed traffic rates in Section 3.

- 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 well-understood $p = 1$ case [11, 20, 23]. For example, it is at most $4/3$ with linear cost functions, $\approx d/\ln d$ with bounded-degree polynomial cost functions, and so on. The bicriteria bound of [23] also holds in such networks.

Because single-commodity networks of parallel links provide tight lower bounds for the ℓ_1 case [11, 20], the same lower bounds apply to the ℓ_p norm for every $p \geq 1$ — for a single-edge path, all of our path norms coincide. Thus, our upper bounds are the best possible.

- 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 [23] also holds for this subclass of equilibria. Again, these upper bounds are the best possible.

We next compare and contrast these contributions with other results in the literature. Our positive result for subpath-optimal equilibria is reminiscent of “price of stability” analyses which concern the best equilibrium of a game [2, 3, 11], but it is much stronger. The best equilibrium of a game typically cannot be reached without centralized intervention, but subpath-optimal equilibria are the natural outcome of decentralized optimization from a networking perspective: if an equilibrium is computed by a distributed shortest-path routing protocol, it is automatically subpath-optimal (see Section 3.2).

The separation that we prove between the worst-case price of anarchy in single- and in multi-commodity networks with the ℓ_p norm with $p > 1$ stands in contrast to the case of the ℓ_1 norm, where there is never such a separation [11, 20]. Indeed, all previously known proof techniques for bounding the price of anarchy of selfish routing (for the ℓ_1 norm) do not refer to the number of commodities of a network, nor to any combinatorial structure whatsoever (as made precise in [22]). These proof techniques for the ℓ_1 case therefore appear incapable of extending to the ℓ_p case with $p > 1$ — our results require a new and fundamentally combinatorial approach.

Finally, several recent works [5, 7, 8, 12, 17] obtain results on the price of anarchy with the ℓ_∞ aggregation function in the conceptually similar but technically different *atomic* selfish routing model, where there is a finite number of players who each control a non-negligible amount of traffic.

1.5 Our Results: Variable Traffic Rates

To address the first goal (1), we augment the basic selfish routing model so that each player has a fixed benefit from 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. This augmented model has been extensively studied in the transportation science literature (e.g. [13]), where the traffic is said to be *elastic*.

With elastic traffic, there are two important quantities: the benefit to the players and the cost incurred. Arguably the most natural way to optimize jointly these two quantities is to maximize the difference between them; this objective has been studied previously for selfish routing with elastic traffic by Chau and Sim [10] and Vetta [26]. However, meaningful approximation guarantees exist for this mixed-sign objective function only under very strong assumptions (as in [10, 26]); this issue is evident already in the single-link example in Section 1.2. To obtain broader insights about the inefficiency of selfish routing with elastic demand, we instead consider the objective of minimizing the travel cost of participating players plus the (lost) benefit of non-participating players. This objective is equivalent to the previous one for exact optimization, but not for relative approximation. It is analogous to the “prize-collecting” objectives widely studied in approximation algorithm design (e.g., [14]).

Our main result for elastic traffic is that *the tragedy of the commons is no worse than the inefficiency of fixed-demand selfish routing*. We establish this by a reduction: for every single- or multicommodity network with elastic traffic, and for every ℓ_p path norm, we exhibit a network with inelastic traffic, essentially the same cost functions, and with at least as large a price of anarchy (under the same path norm). This reduction immediately implies the following upper bounds on the price of anarchy with elastic traffic: for the standard ℓ_1 path norm, all of the well-known upper bounds for inelastic traffic carry over to the elastic traffic case, even in multicommodity networks; for the ℓ_p path norm with $p \in (1, \infty)$, the same upper bounds apply to all single-commodity networks with elastic traffic; and for the ℓ_∞ path norm, these upper bounds apply to all subpath-optimal equilibria in single-commodity networks with elastic traffic. Since selfish routing with elastic traffic

is a strict generalization of the model with inelastic traffic — inelastic traffic is equivalent to infinite benefits — these upper bounds are the best possible.

2 The Model

2.1 Instances

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 $s_1, \dots, s_k \in V$ and sinks $t_1, \dots, t_k \in V$. Second, Γ is a vector of non-increasing, 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 non-increasing amounts to ordering players according to their benefit of participating. Finally, c is a vector of non-negative, continuous, non-decreasing *cost functions*, indexed by E . The function c_e denotes the per flow-unit cost incurred by traffic that uses the edge e , given the amount of traffic on the edge.

2.2 Paths, Flows, and Aggregation Functions

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 i th 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, \dots, e_m and f is a flow, then the per flow-unit *cost* $c_P(f)$ of P with respect to f is

$$c_P(f) = \|c_{e_1}(f_{e_1}), \dots, c_{e_m}(f_{e_m})\|$$

for some aggregation function $\|\cdot\|$. In the traditional selfish routing model, $\|\cdot\|$ is the sum function. In this paper, we allow $\|\cdot\|$ to be any ℓ_p norm $\|\cdot\|_p$ with $1 \leq p \leq \infty$, where, by definition,

$$\|v_1, \dots, v_m\|_p = (v_1^p + \dots + v_m^p)^{1/p}$$

if $p < \infty$ and $\|v_1, \dots, v_m\|_p = \max_i v_i$ if $p = +\infty$ (where the v_i 's are non-negative). We sometimes call such an aggregation function a *path norm*.

2.3 Nash Flows

Intuitively, a flow is at *Nash equilibrium* (or is a *Nash flow*) if no player can do better by changing its mind—by switching paths or by switching whether or not to travel.

Definition 2.1 (Flow at Nash Equilibrium) 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 $f_P > 0$, $c_P(f) \leq c_{P'}(f)$;

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

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 $\Gamma_i(a) \geq \Gamma_i(r_i) = c_i(f)$ 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 $\Gamma_i(a) \leq \Gamma_i(r_i) = c_i(f)$ for all $a \in [r_i, R_i]$).

Existence of a flow at Nash equilibrium can be established in a number of ways; for example, it is a consequence of the very general results of Schmeidler [25].

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

Remark 2.3 (Instances with Inelastic Traffic) Section 3 focuses on instances with *inelastic traffic*. Such instances can be modeled with elastic traffic by defining the functions Γ to be sufficiently large everywhere. Section 3 equivalently defines 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. (If desired, the traffic rates and cost functions can be scaled so that $\sum_{i=1}^k r_i = 1$.) The definition of a Nash flow is then merely part (a) of Definition 2.1.

2.4 The Price of Anarchy and the Pigou Bound

We define the *combined cost* $CC(f)$ of a flow f feasible for an instance (G, Γ, c) with induced traffic rates r as the travel cost plus the lost benefit from non-participants:

$$CC(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P + \sum_{i=1}^k \int_{r_i}^{R_i} \Gamma_i(x) dx. \quad (1)$$

This objective function is simply the joint cost incurred by all of the players in the flow f , where the cost of a single player x of commodity i is $c_P(f)$ if it is routed on the path P and $\Gamma_i(x)$ if it is not routed at all.

An *optimal flow* for an instance is one that minimizes the combined cost over all feasible flows. The *price of anarchy* 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. For instances with inelastic traffic, the second term of (1) vanishes. The *cost* of a feasible flow f is then defined as $C(f) = \sum_{P \in \mathcal{P}} c_P(f) f_P$.

We call instances with inelastic traffic and the additive aggregation function *basic instances*. We only study generalizations of basic instances, so our best-case scenario is to prove upper bounds on the price of anarchy that match those for basic instances. The price of anarchy in such instances depends on the set of allowable cost functions. Define the *Pigou bound* $\alpha(\mathcal{C})$ of a non-empty set of cost functions \mathcal{C} to be:

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

The Pigou bound of \mathcal{C} is essentially the worst-possible price of anarchy in a two-node, two-edge network where one edge has a constant cost function and the other edge has a cost function in \mathcal{C} .

Since all of the ℓ_p norms coincide in networks of parallel edges, $\alpha(\mathcal{C})$ lower bounds the price of anarchy in networks with cost functions in \mathcal{C} with respect to every ℓ_p norm (assuming that \mathcal{C} contains all of the constant cost functions).

The following facts are known for basic instances [11, 20]. First, the price of anarchy of every (multicommodity) instance (G, r, c) with cost functions in a set \mathcal{C} is at most the Pigou bound $\alpha(\mathcal{C})$. 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 non-negative coefficients and degree at most d , then $\alpha(\mathcal{C}) \approx d/\ln d$. (See [20] for more examples.) Qualitatively, these results imply that the price of anarchy is small in basic instances provided cost functions are not “extremely steep.”

3 Nonlinear Path Norms

This section bounds the price of anarchy in selfish routing networks under the ℓ_p norms and with *inelastic* traffic; the next section extends these bounds to networks with elastic traffic. We will discover that the price of anarchy of selfish routing behaves differently in each of the three cases of $p = 1$, $p \in (1, \infty)$, and $p = +\infty$. Section 3.1 exhibits two examples that shape our goals for the rest of the section. Section 3.2 identifies a natural subclass of Nash flows for the ℓ_∞ norm, justifies them from a networking perspective, and proves optimal bounds on their inefficiency in single-commodity networks. Section 3.3 treats the ℓ_p norms with $p \in (1, \infty)$ and proves tight bounds on the inefficiency of arbitrary Nash flows in single-commodity networks.

3.1 Motivating Examples

We now give the two examples promised in the Introduction. The first shows that there are no good bounds on the price of anarchy of selfish routing with the ℓ_p norm in multicommodity networks, even with linear cost functions and inelastic traffic, when $p > 1$.

Example 3.1 (A Lower Bound for the ℓ_p Norm in Multicommodity Networks) Fix an ℓ_p norm with $1 < p \leq \infty$ and consider the two-commodity network shown in Figure 1. For a parameter $k \geq 1$, there are k internally disjoint paths $s_1 \rightarrow v_i \rightarrow w_i \rightarrow t_1$ ($i \in \{1, 2, \dots, k\}$). Edges (v_i, w_i) have the cost function $c(x) = x$; other edges in these paths have zero cost. There are $k - 1$ *cross edges* (w_i, v_{i+1}) (for $i \in \{1, 2, \dots, k - 1\}$), 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 three-hop s_1 - t_1 paths, and routes the second commodity’s traffic on the direct s_2 - t_2 edge. 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 edge, the following flow f is at Nash equilibrium: route the first commodity’s traffic evenly across the three-hop s_1 - t_1 paths and the second commodity’s traffic on the s_2 - t_2 path that contains 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 [23]—stating that a Nash flow is no more expensive than




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

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 for worst-case flows at Nash equilibrium under the ℓ_∞ path norm.

Example 3.2 (A Lower Bound for the ℓ_∞ Norm in Single-Commodity Networks) Suppose we modify the network of Example 3.1 by removing s_2 , t_2 , and the edges incident to them. This yields the network shown in 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 3.1, the bicriteria bound of [23] also fails in this example.

Examples 3.1 and 3.2 justify restricting our attention to single-commodity networks and, for the ℓ_∞ norm, to natural subclasses of equilibria.

3.2 The ℓ_∞ Norm and Subpath-Optimal Nash Flows

We next consider single-commodity networks with the ℓ_∞ norm. Example 3.2 shows that additional restrictions are needed to prove a good bound on the price of anarchy. We require only a modest and natural extra condition on Nash flows, stating that the Nash flow condition holds for all intermediate nodes v and not just the destination t .




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

Definition 3.3 (Subpath-Optimal Flow at Nash Equilibrium) 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 $v = t$. The zigzag Nash flow of Example 3.2 is not subpath-optimal, while the optimal flow is. Notions similar to subpath-optimal equilibria were also proposed, for different purposes, in [1, 5].

Subpath-optimal Nash flows are well motivated. For example, suppose a flow at Nash equilibrium is computed by a Bellman-Ford-type shortest-path algorithm, like a “distance vector protocol” such as OSPF (see e.g. [16]). If such an algorithm uses the cost functions $\{c_e(\cdot)\}$ for edge lengths, and it evaluates path lengths according to an ℓ_p norm, then its fixed points are flows at Nash equilibrium under this norm (see Bertsekas and Tsitsiklis [6]). Such a shortest-path routing protocol computes, by definition, shortest s - v paths for all possible destinations v . Thus, with the ℓ_∞ path norm, the fixed points of such an algorithm automatically satisfy the subpath-optimality property.

We now prove bounds on the inefficiency of subpath-optimal Nash flows. Examples 3.1 and 3.2 show that our proof techniques must make crucial use of both the subpath-optimal assumption and the combinatorial structure of single-commodity networks.

The next lemma identifies a “minimal cut” with respect to a subpath-optimal Nash flow. The plan is to treat the edges crossing this cut as a network of parallel links, enabling bounds on the price of anarchy and also an analogue of the bicriteria bound of [23]. 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 sticking out of (sticking into) the set S .

Lemma 3.4 (Minimal Cut Lemma) *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 $e \in \delta^-(S)$.

Proof: 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 there is an edge $e = (v, w) \in \delta^-(S)$ 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_∞ path 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 part (c), the ℓ_∞ norm of P' is at least $c(f) > d(w)$. This contradicts the subpath-optimality of f . ■

Theorem 3.5 (Price of Anarchy Bound for the ℓ_∞ Norm) *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 $C(f) \leq \alpha(\mathcal{C}) \cdot C(f^*)$.*

Proof: Define the s - t cut S as in the Minimal Cut Lemma (Lemma 3.4). We now define an instance on a network of parallel edges. 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 the Minimal Cut Lemma 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 3.4 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) with the ℓ_∞ path norm. Note that when we discuss the cost of flows in the network of parallel links (G', r, c) , the path norm is irrelevant.

The feasible flow f^* for (G, r, c) 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 (i.e., 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

$$\sum_{e \in E'} c_e(g_e^*)g_e^* \leq \sum_{e \in E'} c_e(f_e^*)g_e^* \leq \sum_{P \in \mathcal{P}} c_P(f^*)f_P^* = C(f^*).$$

Thus the cost of g^* in (G', r, c) is at most that of f^* in (G, r, c) under the ℓ_∞ path norm.

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

As discussed in Section 2.4, the upper bound in Theorem 3.5 is the best possible. Simple examples [23] also show that the following bicriteria bound is optimal.

Theorem 3.6 (Bicriteria Bound for the ℓ_∞ Norm) *Let (G, r, c) be a single-commodity instance and f a subpath-optimal Nash flow under the ℓ_∞ norm. If f^* is feasible for $(G, 2r, c)$, then $C(f) \leq C(f^*)$.*

Proof: Define the cut S , the instance (G', r, c) and the flow g feasible for it, the instance $(G', 2r, c)$ and the flow g^* feasible for it, as in the proof of Theorem 3.5. The bicriteria bound for basic instances [23] implies that the cost of g^* in (G', r, c) is at least that of g in (G', r, c) . Since the cost of g^* in $(G', 2r, c)$ is at most that of f^* in $(G, 2r, c)$ under the ℓ_∞ path norm and the cost of g in (G', r, c) equals that of f in (G, r, c) under the ℓ_∞ path norm, the proof is complete. ■

3.3 The ℓ_p Norms

While the ℓ_1 and ℓ_∞ norms are the best motivated ones, it is also interesting to consider the ℓ_p path norms with $p < \infty$. This section extends Theorems 3.5 and 3.6 to these path norms. The proofs here are more involved as the Minimal Cut Lemma (Lemma 3.4) has only weak analogues for the ℓ_p path norms with $p < \infty$. We instead argue about *many* cuts, and then aggregate the results into a bound on the overall cost of a Nash flow. Since Nash flows under the ℓ_p norm with $p < \infty$ are automatically subpath-optimal, we can bound their cost without any extra restrictions.

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 path norm (e.g. [21]); extending their proofs to the general ℓ_p case requires only cosmetic changes, which we omit here.

Proposition 3.7 *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 cost 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 3.8 *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 3.9 *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 3.7. Then the vertices of G can be ordered, with s first, such that: (i) every edge e with $f_e > 0$ travels forward in the ordering; and (ii) the values $d(v)$ induced by f are non-decreasing in the ordering.*

The next definition provides a sequence of cost functions, where the i th set of cost functions is designed to isolate the cost of a flow across the i th cut of the network. We require two properties. First, the i th 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 of the cuts the edge participates in.

Definition 3.10 (Cut Cost Functions) *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 3.7, sort*

the vertices v_1, \dots, v_n as in Proposition 3.9, and define $S_i = \{v_1, \dots, v_i\}$ for each $i = 1, 2, \dots, n-1$. For an edge e and an integer $i \in \{1, 2, \dots, n-1\}$, the i th cost function $c_e^{(i)}$ of e is defined as follows:

- if $e \notin \delta^+(S_i)$, then $c_e^{(i)}$ is zero everywhere;
- if $e \in \delta^+(S_i)$, 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 3.10 is not well defined if $c_e(f_e) = 0$; in this case Propositions 3.7(a) and 3.9 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 purposes described prior to Definition 3.10.

Lemma 3.11 (Properties of Cut Cost Functions) *With the assumptions and notation of Definition 3.10, the following statements hold.*

- (a) *For every $i = 1, 2, \dots, n-1$, there is a constant $A_i \geq 0$ such that $c_e^{(i)}(f_e) = A_i$ for every $e \in \delta^+(S_i)$ and $c_e^{(i)}$ is identically zero for every other edge. Moreover, every path P with $f_P > 0$ includes either one edge of $\delta^+(S_i)$ (if S_i is an s - t cut) or no edges of $\delta^+(S_i)$ (otherwise); and includes no edge of $\delta^-(S_i)$.*
- (b) *For each $i = 0, 1, \dots, n-1$, $d(v_{i+1}) = \|A_1, \dots, A_i\|$, where each A_j is defined as in (a).*
- (c) *For every feasible flow f^* for (G, r, c) and every path $P \in \mathcal{P}$, $c_P(f^*) \geq \|c_P^{(1)}(f^*), \dots, c_P^{(n-1)}(f^*)\|$.*

Proof: For part (a), fix $i \in \{1, 2, \dots, n-1\}$. The first assertion holds by construction — only edges of $\delta^+(S_i)$ have non-zero cost under the cost functions $c^{(i)}$ — with A_i the unique number satisfying $d(v_{i+1})^p = d(v_i)^p + A_i^p$. The second holds because the vertices are sorted topologically with respect to f .

We next establish part (b) by induction on i . Since $v_1 = s$ and $d(s) = 0$, the base case ($i = 0$) holds. For $i > 0$, we have

$$d(v_{i+1})^p = d(v_i)^p + A_i^p = \sum_{i=1}^i A_i^p = \|A_1, \dots, A_i\|^p,$$

where the first equality follows from the definition of A_i and the second from the inductive hypothesis. The inductive step and part (b) thus hold.

Part (c) requires a non-trivial proof. Fix a feasible flow f^* and an s - t path $P \in \mathcal{P}$. Let P' be the edges of P that travel forward with respect to the given topological ordering of the vertices of G . We first claim that for every edge $e = (v_i, v_j) \in P'$,

$$c_e(f_e^*)^p \geq \sum_{q=i}^{j-1} \left[c_e^{(q)}(f_e^*) \right]^p. \quad (3)$$

To see this, fix an edge $e = (v_i, v_j) \in P'$ with $i < j$. We can assume that $c_e(f_e^*) > 0$. For every $q \in \{i, \dots, j-1\}$, the definitions of $\lambda^{(q)}$ and $c^{(q)}$ ensure that

$$d(v_{q+1})^p - d(v_q)^p = \left(\lambda_e^{(q)} c_e(f_e) \right)^p = \left(\lambda_e^{(q)} c_e(f_e^*) \right)^p \cdot \frac{c_e(f_e)^p}{c_e(f_e^*)^p} = \left(c_e^{(q)}(f_e^*) \right)^p \cdot \frac{c_e(f_e)^p}{c_e(f_e^*)^p}.$$

Summing over all $q \in \{i, \dots, j-1\}$ then gives

$$d(v_j)^p - d(v_i)^p = \left[\sum_{q=i}^{j-1} \left(c_e^{(q)}(f_e^*) \right)^p \right] \cdot \frac{c_e(f_e)^p}{c_e(f_e^*)^p}. \quad (4)$$

Finally, Proposition 3.7(a) implies that

$$d(v_j)^p - d(v_i)^p \leq c_e(f_e)^p = c_e(f_e^*)^p \cdot \frac{c_e(f_e)^p}{c_e(f_e^*)^p}. \quad (5)$$

Combining (4) and (5) then verifies (3).

Next, we can apply (3) to obtain

$$\begin{aligned} (c_P(f^*))^p &= \sum_{e \in P} c_e(f_e^*)^p \\ &\geq \sum_{e \in P'} c_e(f_e^*)^p \\ &\geq \sum_{e=(v_i, v_j) \in P'} \sum_{q=i}^{j-1} [c_e^{(q)}(f_e^*)]^p. \end{aligned}$$

Reversing the order of summation and recalling the definition of $c^{(i)}$ then completes the proof:

$$\begin{aligned} (c_P(f^*))^p &\geq \sum_{i=1}^{n-1} \sum_{e \in P \cap \delta^+(S_i)} [c_e^{(i)}(f_e^*)]^p \\ &= \sum_{i=1}^{n-1} [c_P^{(i)}(f_e^*)]^p \\ &= \|c_P^{(1)}(f^*), \dots, c_P^{(n-1)}(f^*)\|^p. \end{aligned}$$

■

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 in Definition 3.10, are analogous to those for Theorems 3.5 and 3.6, respectively, where Lemma 3.11(a) plays the role originally served by the Minimal Cut Lemma. In the statements of the lemmas, we use $C^{(i)}$ to denote the cost of a flow with respect to the cost functions $c^{(i)}$.

Lemma 3.12 *Let (G, r, c) be a single-commodity instance with the ℓ_p path norm ($p \in [1, \infty)$) 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 3.10. Then for each $i = 1, 2, \dots, n-1$,*

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

Proof: Let the sink t of G appear as the m th vertex in the topological ordering of the vertices with respect to the acyclic flow f . If $i \geq m$, then $C^{(i)}(f) = 0$ and the lemma holds. Otherwise, consider the s - t cut $S_i = \{v_1, \dots, v_i\}$. Define the instance $(G', r, c^{(i)})$ on a network of parallel links as in the proof of Theorem 3.5, with edges of G' corresponding to those of $\delta^+(S_i)$. Restricting f to the edges of $\delta^+(S_i)$ induces a flow g in G' . Using the notation and the assertions in Lemma 3.11(a), g is feasible and at Nash equilibrium for $(G', r, c^{(i)})$, and its cost is $r \cdot A_i = C^{(i)}(f)$. As in the proof of Theorem 3.5, the flow f^* induces a flow g^* feasible for $(G', r, c^{(i)})$ with cost at most that of f^* in $(G, r, c^{(i)})$ under the ℓ_p path norm. The instance $(G', r, c^{(i)})$ can be viewed as a basic instance with cost functions that are scalar multiples of functions in \mathcal{C} . Taking the closure of a set of cost functions under multiplication by positive scalars does not change its Pigou bound [20]. Thus, the cost of g^* in $(G', r, c^{(i)})$, and hence of f^* in $(G, r, c^{(i)})$, is at least $C(g)/\alpha(\mathcal{C}) = C^{(i)}(f)/\alpha(\mathcal{C})$. ■

Modifying the proof of Lemma 3.12 in the obvious way (as in the proof of Theorem 3.6) yields the following.

Lemma 3.13 *Let (G, r, c) be a single-commodity instance with the ℓ_p path norm ($p \in [1, \infty)$). Let f be an acyclic Nash flow for (G, r, c) and f^* a feasible flow for $(G, 2r, c)$, and define the cost functions $c_e^{(i)}$ as in Definition 3.10. Then for each $i = 1, 2, \dots, n-1$,*

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

We now prove bounds on the price of anarchy and a bicriteria bound by aggregating the “cut-by-cut” bounds of Lemmas 3.12 and 3.13 into a bound for the entire network.

Theorem 3.14 (Price of Anarchy Bound for the ℓ_p Norm) *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 $C(f) \leq \alpha(\mathcal{C}) \cdot C(f^*)$.*

Proof: We can assume without loss of generality that f is acyclic (Proposition 3.8). We first use Lemma 3.11 to lower bound the cost of f^* in terms of the corresponding cost functions $c^{(i)}$. Specifically, write

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

$$= \sum_{P \in \mathcal{P}} \left\| f_P^* \cdot c_P^{(1)}(f^*), \dots, f_P^* \cdot c_P^{(n-1)}(f^*) \right\| \tag{7}$$

$$\geq \left\| \sum_{P \in \mathcal{P}} f_P^* \cdot c_P^{(1)}(f^*), \dots, \sum_{P \in \mathcal{P}} f_P^* \cdot c_P^{(n-1)}(f^*) \right\|, \tag{8}$$

where (6) follows from Lemma 3.11(c), and (7) and (8) follow since $\|\cdot\|$ is a norm, and thus is linear under scalar multiplication and satisfies the Triangle inequality. Applying Lemma 3.12 and the monotonicity of $\|\cdot\|$, we obtain

$$C(f^*) \geq \left\| \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot c_P^{(1)}(f), \dots, \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot c_P^{(n-1)}(f) \right\|. \tag{9}$$

We now reverse the argument to recover the cost of the Nash flow f . Since the Triangle inequality is only useful in one direction, we instead use the stronger assertions of Lemma 3.11. Precisely, let the sink t be the m th vertex in the underlying ordering of the vertices of G . Then, for some non-negative constants A_1, \dots, A_{m-1} , we have

$$C(f^*) \geq \left\| \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot A_1, \dots, \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot A_{m-1} \right\| \quad (10)$$

$$= \frac{\sum_{P \in \mathcal{P}} f_P}{\alpha(\mathcal{C})} \|A_1, \dots, A_{m-1}\| \quad (11)$$

$$= \frac{1}{\alpha(\mathcal{C})} \sum_{P \in \mathcal{P}} f_P \cdot d(t) \quad (12)$$

$$= \frac{C(f)}{\alpha(\mathcal{C})}, \quad (13)$$

where (10) follows from (9) and Lemma 3.11(a), (11) from the linearity of $\|\cdot\|$ under scalar multiplication, (12) from Lemma 3.11(b) with $i = m$, and (13) from Definition 2.1 and the definition of $d(t)$. This completes the proof. ■

An entirely analogous proof, combined with Lemma 3.13, establishes a bicriteria bound.

Theorem 3.15 (Bicriteria Bound for the ℓ_p Norm) *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 $C(f) \leq C(f^*)$.*

4 Bounds on the Tragedy of the Commons with Elastic Traffic

This section shows that the worst-possible inefficiency arising from the tragedy of the commons has no greater magnitude than that arising from the routing inefficiencies of selfish traffic. We establish this by “reducing” networks with elastic traffic to those with inelastic traffic in the following sense.

Theorem 4.1 (Reduction from Elastic Traffic to Inelastic Traffic) *For every instance (G, Γ, c) with elastic traffic and cost functions in a set \mathcal{C} , and every ℓ_p path norm with $1 \leq p \leq \infty$, there is an instance $(\hat{G}, \hat{r}, \hat{c})$ with inelastic traffic and cost functions that are either constant or in \mathcal{C} , in which the price of anarchy (under the ℓ_p path norm) is at least that of (G, Γ, c) .*

For the ℓ_∞ path norm, this reduction also applies to the price of anarchy with respect to subpath-optimal Nash flows.

Proof: Consider an instance (G, Γ, c) and an ℓ_p path norm with $1 \leq p \leq \infty$. Let f be a Nash flow and f^* an optimal flow. Let r and r^* denote the corresponding induced vectors of traffic rates. For each commodity i , we can discard all traffic between $\max\{r_i, r_i^*\}$ and R_i : this decreases the combined cost of f and f^* by a common amount and can only increase the price of anarchy. We can therefore assume that for each commodity i , either f or f^* routes all of the traffic of that commodity.

Obtain a network \hat{G} from G as follows. For each commodity i , add new vertices \hat{s}_i and \hat{t}_i , which are the source and sink vertices for commodity i in \hat{G} . Add edges (\hat{s}_i, s_i) and (t_i, \hat{t}_i) with

constant cost $\hat{c}(x) = 0$. Finally, add the edge (\hat{s}_i, \hat{t}_i) with constant cost $\hat{c}(x) = \Gamma_i(r_i)$, and define \hat{r}_i as $\max\{r_i, r_i^*\}$. Observe that the \hat{s}_i - \hat{t}_i paths of \hat{G} , excluding the direct \hat{s}_i - \hat{t}_i edge, enjoy a natural and cost-preserving bijective correspondence with the s_i - t_i paths of G .

We complete the proof by showing that the price of anarchy under the ℓ_p norm in the instance $(\hat{G}, \hat{r}, \hat{c})$ with inelastic traffic is at least that in the original instance (G, Γ, c) . First, observe that f and f^* naturally induce feasible flows \hat{f} and \hat{f}^* in $(\hat{G}, \hat{r}, \hat{c})$, with the traffic of a commodity i that is not routed by a flow in (G, Γ, c) being sent on the direct (\hat{s}_i, \hat{t}_i) edge in $(\hat{G}, \hat{r}, \hat{c})$. Since the cost of the direct \hat{s}_i - \hat{t}_i path is always $\Gamma_i(r_i)$, independent of the path norm, the induced flow \hat{f} is at Nash equilibrium in $(\hat{G}, \hat{r}, \hat{c})$ under the ℓ_p path norm. Also, for the ℓ_∞ path norm, the flow \hat{f} is subpath-optimal in $(\hat{G}, \hat{r}, \hat{c})$ provided f is subpath-optimal in (G, Γ, c) . Since every Γ_i is non-increasing, the cost of \hat{f} in $(\hat{G}, \hat{r}, \hat{c})$ is at least the combined cost of f in (G, Γ, c) . For the same reason, the cost of \hat{f}^* in $(\hat{G}, \hat{r}, \hat{c})$ is at most the combined cost of f^* in (G, Γ, c) . The price of anarchy of $(\hat{G}, \hat{r}, \hat{c})$ is at least $C(\hat{f})/C(\hat{f}^*) \geq CC(f)/CC(f^*)$, which completes the proof. ■

Adding constant cost functions to a set \mathcal{C} does not affect its Pigou bound $\alpha(\mathcal{C})$ (recall Section 2.4). Thus, the following three corollaries follow immediately from Theorem 4.1 and from known results for the ℓ_1 path norm [11, 20] and Theorems 3.5 and 3.14.

Corollary 4.2 *Let (G, Γ, c) be an instance with cost functions in \mathcal{C} , f a Nash flow under the ℓ_1 norm, and f^* a feasible flow. Then $CC(f) \leq \alpha(\mathcal{C}) \cdot CC(f^*)$.*

Corollary 4.3 *Let (G, Γ, 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 $CC(f) \leq \alpha(\mathcal{C}) \cdot CC(f^*)$.*

Corollary 4.4 *Let (G, Γ, c) be a single-commodity instance with cost functions in \mathcal{C} , f a Nash flow under the ℓ_p norm with $p \in (1, \infty)$, and f^* a feasible flow. Then $CC(f) \leq \alpha(\mathcal{C}) \cdot CC(f^*)$.*

Corollary 4.2 also holds more generally for the nonatomic congestion games studied in [24].

5 Conclusions

This paper established bounds on the price of anarchy of selfish routing with nonlinear path aggregation functions and with variable traffic. The worst-case price of anarchy is unbounded in multicommodity networks with nonlinear aggregation functions, and even in single-commodity networks with the ℓ_∞ path norm, even when every edge cost function is linear. On the positive side, the worst-case price of anarchy in single-commodity networks with nonlinear aggregation functions — and for the ℓ_∞ norm, with attention restricted to the subclass of equilibria generated by distributed shortest-path routing protocols — is the same as that for the well-understood ℓ_1 norm. All of these bounds on the price of anarchy extend to networks with elastic traffic, for a natural cost-minimization objective.

References

- [1] A. Akella, S. Chawla, and S. Seshan. Realistic models for selfish routing in the Internet. Unpublished manuscript, 2003. Preliminary version appeared as CMU SCS Technical Report Number CMU-CS-02-163.

- [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. *SIAM Journal on Computing*, 38(4):1602–1623, 2008.
- [3] E. Anshelevich, A. Dasgupta, É. Tardos, and T. Wexler. Near-optimal network design with selfish agents. *Theory of Computing*, 4(1):77–109, 2008.
- [4] A. Banerjee and J. Yoo. Minimizing maximum logical link congestion in packet switched optimal networks. In *Proceedings of the 1997 IEEE International Conference on Communications (ICC)*, volume 3, pages 1298–1302, 1997.
- [5] R. Banner and A. Orda. Bottleneck routing games in communication networks. *IEEE Journal on Selected Areas of Communication*, 25(6):1173–1179, 2007.
- [6] D. P. Bertsekas and J. N. Tsitsiklis. *Parallel and Distributed Computation: Numerical Methods*. Prentice-Hall, 1989. Second Edition, Athena Scientific, 1997.
- [7] C. Busch and M. Magdon-Ismaïl. Atomic routing games on maximum congestion. *Theoretical Computer Science*, 410(36):3337–3347, 2009.
- [8] I. Caragiannis, C. Galdi, and C. Kaklamanis. Network load games. In *Proceedings of the 16th Annual Symposium on Algorithms and Computation (ISAAC)*, volume 3827 of *Lecture Notes in Computer Science*, pages 809–818, 2005.
- [9] J. H. Chang and L. Tassiulas. Energy conserving routing in wireless ad-hoc networks. In *Proceedings of INFOCOM*, pages 22–31, 2000.
- [10] 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.
- [11] 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.
- [12] A. Epstein, M. Feldman, and Y. Mansour. Efficient graph topologies in network routing games. *Games and Economic Behavior*, 66(1):115–125, 2009.
- [13] N. H. Gartner. Optimal traffic assignment with elastic demands: A review; Part I. Analysis framework. *Transportation Science*, 14(2):174–191, 1980.
- [14] 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.
- [15] G. Hardin. The tragedy of the commons. *Science*, 162:1243–1248, December 13 1968.
- [16] S. Keshav. *An Engineering Approach to Computer Networking*. Addison-Wesley, 1997.
- [17] V. Mazalov, B. Monien, F. Schoppmann, and K. Tiemann. Wardrop equilibria and price of stability for bottleneck games with splittable traffic. In *Proceedings of the Second Annual Workshop on Internet and Network Economics (WINE)*, volume 4286 of *Lecture Notes in Computer Science*, pages 331–342, 2006.

- [18] N. Nisan, T. Roughgarden, É. Tardos, and V. Vazirani, editors. *Algorithmic Game Theory*. Cambridge University Press, 2007.
- [19] 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.
- [20] T. Roughgarden. The price of anarchy is independent of the network topology. *Journal of Computer and System Sciences*, 67(2):341–364, 2003.
- [21] T. Roughgarden. *Selfish Routing and the Price of Anarchy*. MIT Press, 2005.
- [22] T. Roughgarden. Intrinsic robustness of the price of anarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing (STOC)*, pages 513–522, 2009.
- [23] T. Roughgarden and É. Tardos. How bad is selfish routing? *Journal of the ACM*, 49(2):236–259, 2002.
- [24] T. Roughgarden and É. Tardos. Bounding the inefficiency of equilibria in nonatomic congestion games. *Games and Economic Behavior*, 49(2):389–403, 2004.
- [25] D. Schmeidler. Equilibrium points of nonatomic games. *Journal of Statistical Physics*, 7(4):295–300, 1973.
- [26] 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.
- [27] 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.