Bertrand competition in networks*

Shuchi Chawla[†] Tim Roughgarden[‡]

Abstract

We study price-of-anarchy type questions in two-sided markets with combinatorial consumers and limited supply sellers. Sellers own edges in a network and sell bandwidth at fixed prices subject to capacity constraints; consumers buy bandwidth between their sources and sinks so as to maximize their value from sending traffic minus the prices they pay to edges. We characterize the price of anarchy and price of stability in these "network pricing" games with respect to two objectives—the social value (social welfare) of the consumers, and the total profit obtained by all the sellers. In single-source single-sink networks we give tight bounds on these quantities based on the degree of competition, specifically the number of monopolistic edges, in the network. In multiple-source single-sink networks, we show that equilibria perform well only under additional assumptions on the network and demand structure.

1 Introduction

The Internet is a unique modern artifact given its sheer size, and the number of its users. Given its (continuing) distributed and ad-hoc evolution, as well as emerging applications, there have been growing concerns about the effectiveness of its current routing protocols in finding good routes and ensuring quality of service. Congestion and QoS based pricing has been suggested as a way of combating the ills of this distributed growth and selfish use of resources (see, e.g., [4, 6, 7, 9, 11]). Unfortunately, the effectiveness of such approaches relies on the cooperation of the multiple entities implementing them, namely the owners of resources on the Internet, or the ISPs. The ISPs' goals do not necessarily align with the social objectives of efficiency and quality of service; their primary objective is to maximize their own market share and profit.

In this paper we consider the following question: given a large combinatorial market such as the Internet, suppose that the owners of resources selfishly price their product so as to maximize their own profit, and consumers selfishly purchase bundles of products to maximize their own utility, how does this effect the functioning of the market as a whole?

We consider a simple model where each edge of the network is owned by a distinct selfish entity, and is subject to capacity constraints. Each consumer is interested in buying bandwidth along a path from its source to its destination, and obtains a fixed value per unit of flow that it can send along this path; consumers are therefore single-parameter agents. The game proceeds by the sellers first picking (per-unit-bandwidth) prices for the edges they own, and then the consumers buying their most-desirable paths (or not buying anything if all the paths are too expensive). An outcome of the game (a collection of prices and the paths bought by consumers) is called a Nash equilibrium if no seller can improve her profit by changing her price single-handedly. Note that the consumers already play a best-response to the prices. We compare the performance of equilibria in this game to that of the best state achievable through coordination, under two metrics—the efficiency or social value of the system, and the total profit earned by all the edges.

Economists have traditionally studied the properties of equilibria that emerge in pricing games with competing firms in single-item markets (see, for example, [14, 15] and references therein). It is well known [10], for example,

*An extended abstract of this paper is to appear in the 1st Annual Symposium on Algorithmic Game Theory (SAGT), 2008.

[†]Computer Sciences Department, University of Wisconsin - Madison. `shuchi@cs.wisc.edu`. Supported in part by NSF CAREER Award CCF-0643763.

[‡]Department of Computer Science, Stanford University. `tim@cs.stanford.edu`. Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, and an Alfred P. Sloan Fellowship.

that in a single-good free market, oligopolies (two or a few competing firms) lead to a socially-optimal equilibrium[1]. On the other hand, a monopoly can cause a lot of inefficiency in the market by selfishly maximizing its own profit. Fortunately the extent of this inefficiency is bounded by a logarithmic factor in the ratio of the maximum consumer utility to the minimum consumer utility, as well as by a logarithmic factor in the number of consumers.

These classical economic models ignore the combinatorial aspects of network pricing, namely that consumers have different geographic sources and destinations for their traffic, and goods (i.e., edges) are not pure substitutes, but rather are a complex mix of substitutes and complements, as defined by the network topology. So a timely and basic research question is: which properties of standard price equlbrium models carry over to network/combinatorial settings? For example, are equilibria still guaranteed to exist? Are equilibria fully efficient? Does the answer depend in an interesting way on the network/demand structure?

The network model captures the classical single-item setting in the form of a single-source single-sink network with a single edge (modeling a monopoly), or multiple parallel edges (modeling an oligopoly). In addition, we investigate these questions in general single-source single-sink networks, as well as multiple-source single-sink networks. Our work can we viewed as a non-trivial first step toward understanding price competition in general combinatorial markets.

Our results

We study the price of anarchy, or the ratio of the performance of the worst Nash equilibrium to that of an optimal state, for the network pricing game with respect to social value and profit. We give matching upper and lower bounds, as a function of the degree of competition in the network, and the ratio $\mathcal{L}$ of the maximum and minimum customer valuations. For instances with a high price of anarchy, a natural question is whether there exist any good equilibria for the instance. We provide a negative answer in most such cases, giving strong lower bounds on the price of stability, which quantifies the ratio of the performance of the *best* Nash equilibrium for the instance to that of an optimal solution.

For single-source single-sink networks, we provide tight upper and lower bounds on the prices of anarchy and stability (see Section 3). Although in a network with a single monopolistic edge, these quantities are $O(\log \mathcal{L})$ for social value, both become worse as the number of monopolies increases. The price of stability, for example, increases exponentially with the number k of monopolies, as $\Theta(\mathcal{L}^{k-1})$ for $k > 1$. The equilibrium prices in these instances are closely related to the min-cut structure of the instances.

With respect to profit, as is expected, networks that contain no monopolies display a large price of anarchy and stability because competition hurts the profits of all the firms, while networks with a single monopoly perform very well. One may suspect that as competition decreases further (the number of monopolies gets larger), collective profit improves. We show instead that the price of stability for profit also increases exponentially with the number of monopolies.

In multiple-source single-sink networks, the behavior of Nash equilibria changes considerably (see Section 4). In particular, equilibria do not always exist even in very simple directed acyclic networks. When they do exist, some instances display a high price of stability (polynomial in $\mathcal{L}$) despite strong competition in the network. In addition to the presence of monopolies, we identify other properties of instances that cause such poor behavior: (1) an uneven distribution of demand across different sources, and (2) congested subnetworks (congestion in one part of the network can get "carried over" to a different part of the network in the form of high prices due to the selfishness of the edges). We show that in a certain class of directed acyclic networks with no monopolies, in which equilibria are guaranteed to exist, the absence of the above two conditions leads to good equilibria. Specifically, the price of stability for social value in such networks is at most $1/\alpha$ where α is the sparsity of the network. Once again, we use the sparse-cut structure of the network to explicitly construct good equilibria.

[1] To be precise, there are two models of competition in an oligopolistic market—Bertrand competition, where the firms compete on prices, and Cournot competition, where they compete on quantity. The former always leads to a socially-optimal equilibrium; the latter may not. In this paper we will focus on the Bertrand model. See Section 5 for a brief discussion of the Cournot model.

Related work

The literature on quantifying the inefficiency of equilibria is too large to survey here; see [13] and the references therein for an introduction.

Recently, several researchers have studied the existence and inefficiency of equilibria in network pricing models where consumers face congestion costs from other traffic sharing the same bandwidth [8, 1, 2, 12, 16]. As in this paper, all of these other works consider network pricing games in which consumers are interested in routing their demands over paths in a network, edges of which are owned by selfish sellers. The routing cost faced by each consumer has two components: the price charged by each edge on the path, and the latency faced by the consumer's flow owing to congestion on the path. In addition to selfish pricing, this congestion-based externality among consumers leads to highly inefficient outcomes even in very simple networks (such as single-source single-sink series-parallel networks [2]). The cost model considered by us is a special case of this latency-based cost function, in which the latency faced by a flow is 0 as long as all capacity constraints along the path are satisfied, and ∞ otherwise. Furthermore, in our model, latency (congestion) costs are paid by edges, rather than by consumers, and therefore force the edges to raise their prices just enough for the capacity constraints to be met. Owing to the generality of the latency functions they consider, these other papers study extremely simple network models. Acemoglu and Ozdaglar [1, 2], for example, assume that all consumers are identical, and have unbounded values (i.e. they simply minimize their total routing cost). They analyze the game in single-source single-sink networks with parallel links. (Some of their results also extend to single-source single-sink series-parallel networks.) Likewise, Hayrapetyan et al. [8] consider single-source single-sink networks with parallel links, but in addition allow different values for different consumers. In contrast, we consider general single-source single-sink as well as multiple-source single-sink topologies with the simpler capacity-based cost model. In effect, our work isolates the impact of selfish pricing on the efficiency of the network in the absence of congestion effects. Although capacity constraints in our model mimic some congestion effects, we see interesting behavior even in the absence of capacity constraints when the market contains monopolies. The instances we consider display a large range of behavior in the performance of equilibria, depending on the network and demand structure.

Another recent work closely related to ours is a network formation model introduced by Anshelevich et al. [3] in which neighboring agents form bilateral agreements to both buy and sell bandwidth simultaneously. The game studied by Anshelevich et al. can be thought of as a meta-level game played by agents when they first enter the network and install capacities based on anticipated demand. Then, once the network is formed, a different game is played between the agents owning edges and consumers. This second game is the one that we analyze. Furthermore, in the model considered by Anshelevich et al. there are no latencies or capacity constraints, instead there is a fixed cost for routing each additional unit of flow.

2 Model & notation

A network pricing game (NPG) is characterized by a directed graph $G = (V, E)$ with edge capacities $\{c_e\}_{e \in E}$, and a set of users (traffic matrix) endowed with values. Each edge is owned by a distinct ISP. (Many of our results can be easily extended to the case where a single ISP owns multiple edges.) The value associated with each chunk of traffic represents the *per-unit monetary value* that the owner of that chunk obtains upon sending this traffic from its source to its destination. User values are represented in the form of *demand curves*[2], $\mathcal{D}_{(s,t)}$, for every source-destination pair (s, t), where for every ℓ, $\mathcal{D}_{(s,t)}(\ell)$ represents the amount of traffic with value at least ℓ. When the network has a single source-sink pair, we drop the subscript (s, t). We use $\mathcal{D}$ to denote the "demand suite", or the collection of these demand curves, one for each source-sink pair. Without loss of generality, the minimum value is 1, that is, $\mathcal{D}_{(s,t)}(1) = \mathbf{F}^{\text{tot}}_{s,t}$, the total flow between s and t, for all pairs (s, t), and we use $\mathcal{L}$ to denote the maximum value—$\mathcal{L} = \sup\{\ell | \mathcal{D}_{(s,t)}(\ell) > 0\}$.

We extend the classic Bertrand model of competition to network pricing. The NPG has two stages. In the first stage, each ISP (edge) e picks a price π_e. In the second stage each user picks paths between its source and destination to send its traffic. We assume that users can split their traffic into infinitesimally small chunks, and spread it across multiple paths, or send fractional amounts of traffic. Each user picks paths to maximize her utility, $u = v - \min_P \sum_{e \in P} \pi_e$, where the minimum is over all paths P from the user's source to its destination, and v is its value (or sends no flow if the minimum total price is larger than its value v). This selection of paths determines the amount of traffic f_e on each

[2]We aggregate these curves over all users with the same source and destination pairs.

edge. ISP e's utility is given by $f_e\pi_e$ if $f_e \le c_e$, and $-\infty$ otherwise. ISPs are selfish and set prices to maximize their utility. (We briefly discuss an alternate model of ISP behavior in Section 5 based on the classic Cournot competition.)

A given state in a game (in this case consisting of a set of prices and traffic pattern) is called a Nash equilibrium if no agent wants to deviate from it unilaterally so as to improve its own utility. In a multi-stage game, the relevant concept is that of a *subgame-perfect Nash equilibrium*, where, given any first stage strategies, the second stage strategies form a Nash equilibrium. In the NPG, users are price-takers, that is, they merely follow a best response to the prices set by ISPs, and the responses of different users are decoupled from each other. Therefore, given the first stage strategies, the second stage strategies always form a Nash equilibrium, and the dynamics of the system is determined primarily by the first stage game.

Note that by sending fractional flow, or splitting their traffic across multiple paths, users effectively mimick randomized strategies. ISPs, on the other hand, always pick a deterministic strategy (committing to a fixed price). Therefore, (pure strategy) equilibria do not always exist in these games (indeed in Section 4 we present an example where there are no pure strategy equilibria). Nevertheless we identify some cases in which equilibria do exist, and characterize their performance in those cases.

Note also that if the flow f resulting from the users' strategies in the second part of the game is such that the capacity constraint on an edge e is violated, users using that edge still obtain their value from routing their flow, while the edge e incurs a large penalty. Intuitively, the edge e is forced to compensate those users that are denied service due to capacity constraints, for not honoring its commitment to serve them at its declared price. This situation of course does not arise at an equilibrium – any edge with a violated capacity can improve its profit by increasing the price charged by it.

We evaluate the Nash equilibria of these games with respect to two objectives—social value and profit. The social value of a state S of the network, $\mathbf{Val}(S)$ is defined to be the total utility of all the agents in the system, specifically, the total value obtained by all the users, minus the prices paid by the users, plus the profits (prices) earned by all the ISPs. Since prices are endogenous to the game, this is equivalent to the total value obtained by all the users, and we will use this latter expression to evaluate it throughout the paper. The worst such value over all Nash equilibria is captured by the price of anarchy: the price of anarchy of the NPG with respect to social value, $\mathbf{POA_{Val}}$, is defined to be the minimum over all Nash equilibria $S \in \mathcal{N}$ of the ratio of the social value of the equilibrium to the optimal achievable value $\mathbf{Val}^*$:

$$\mathbf{POA_{Val}}(G, \mathcal{D}) = \frac{\min_{S\in\mathcal{N}(G,\mathcal{D})} \mathbf{Val}(S)}{\mathbf{Val}^*}$$

Here, $\mathbf{Val}^*$ is the maximum total value achievable while satisfying all the capacity constraints in the network (this can be computed by a simple flow LP). Likewise, $\mathbf{POA_{Pro}}$ denotes the price of anarchy with respect to profit:

$$\mathbf{POA_{Pro}}(G, \mathcal{D}) = \frac{\min_{S\in\mathcal{N}(G,\mathcal{D})} \mathbf{Pro}(S)}{\mathbf{Pro}^*}$$

Here $\mathbf{Pro}(S)$ is the total utility obtained by all the ISPs, or alternately the total payment made by all users. The optimal profit achievable $\mathbf{Pro}^*$ is defined to be the maximum profit over all states in which users are in equilibrium, and capacity constraints are satisfied.

In instances with a large price of anarchy, we also study the performance of the best Nash equilibria and provide lower bounds for it. The price of stability of a game is defined to be the *maximum* over all Nash equilbria in the game of the ratio of the value of the equilibrium to the optimal achievable value. We use $\mathbf{POS_{Val}}$ and $\mathbf{POS_{Pro}}$ to denote the price of stability with respect to social value and profit respectively.

3 The network pricing game in single-source single-sink networks

In this section we study the network pricing game in single commodity networks, that is, instances in which every customer has the same source and sink. As the single-item case suggests, the equilibrium behavior of the NPG depends on whether or not there is competition in the network. However, the extent of competition, specifically the number of monopolies, also plays an important role. In the context of a network (or a general combinatorial market), an edge monopolizes over a consumer if *all* the paths (bundles of items) desired by the customer contain the edge.

Definition 1 *An edge in a given network is called a monopoly if its removal causes the source of a commodity to be disconnected from its sink.*

No monopoly. In the absence of monopolies, the behavior of the network is analogous to competition in single-item markets. Specifically, competition drives down prices and enables higher usage of the network, thereby obtaining good social value but poor profit.

Theorem 1 *In a single commodity network with no monopolies,* $\mathbf{POA}_{\mathbf{Val}} = 1$*. Furthermore, there exist instances with* $\mathbf{POS}_{\mathbf{Pro}} = \Theta(\mathcal{L})$*.*

Proof: We first note that an equilibrium supporting the optimal flow (w.r.t. social value) always exists: consider an optimal flow of amount, say, f in the network; let $p = D^{-1}(f)$ if the flow saturates the network, and 0 otherwise; take an arbitrary distribution over min-cuts in the network, and "distribute" this price over the min-cuts; in particular, the price on any edge is p times the probability that e is in the min-cut. We claim that these prices, along with the flow f form an equilibrium. Note that customers play a best-response—each unit of flow admitted pays a total price of p regardless of the path it takes from the source to the destination, and a chunk of flow is admitted if and only if it has value at least p. On the other hand, edges cannot improve their profits by increasing prices unilaterally, because their customers can switch to a different cheaper path. Finally, edges with non-zero prices are saturated, and cannot gain customers by lowering their price.

For a bound on the price of anarchy, consider any equilibrium in the given instance, and suppose that the network is not saturated. If all the traffic is admitted, then $\mathbf{POA}_{\mathbf{Val}} = 1$. Otherwise, there exists an unsaturated edge with non-zero price that is not carrying all of the admitted flow (if there exists a zero-price unsaturated path, then some users are playing suboptimally). Let the edge be e. Then there is a source-sink path P in the graph carrying flow that doesn't contain the edge e. Note that all flow carrying paths charge an equal total price. Edge e can therefore improve its profit by lowering its price infinitesimally and grabbing some of the flow on path P which is not among the cheapest paths any more. This contradicts the fact that the network is in equilibrium.

For the second part of the theorem, we consider a network with unbounded capacity, and assume without loss of generality that $\mathbf{F}^{\text{tot}}_{s,t} = 1$. Our argument above (that $\mathbf{POA}_{\mathbf{Val}} = 1$) implies that in any equilibrium all the traffic is admitted. Therefore the price charged to each user is at most 1 (the minimum value), and at equilibrium, the total profit of the network is 1. On the other hand, suppose that all but an infinitessimal fraction of the users have value $\mathcal{L}$, then a solution admitting only the high-value set of users (and charging a price of $\mathcal{L}$ to each user) has net profit almost $\mathcal{L}$. ■

Single monopoly. As the following theorem shows, the best-case and worst-case performance of single monopoly networks is identical to that of single-link networks.

Theorem 2 *In a single commodity network with a single monopoly,* $\mathbf{POA}_{\mathbf{Val}} = O(\log \mathcal{L})$ *and* $\mathbf{POA}_{\mathbf{Pro}} = 1$*. Furthermore, there are instances for which* $\mathbf{POS}_{\mathbf{Val}} = \Theta(\log \mathcal{L})$*.*

Proof: We prove the second part of the theorem first. This follows by simply considering the $1/x$ demand curve from 1 to $\mathcal{L}$ in a single link unbounded capacity network. The single link then behaves like a single-item monopolist, and without loss of generality charges a price of $\mathcal{L}$, resulting in a social value of 1. Adding an infinitesimal point mass in the demand curve at $\mathcal{L}$ breaks ties among prices and ensures that this is the only equilibrium. The optimal social value, on the other hand, is the total value of all users $\int_1^{\mathcal{L}} 1/x dx = \log \mathcal{L}$.

For the first part of the theorem, we first note that in a single-link network (i.e. a single-item market), the above example is essentially the worst. Specifically, if at equilibrium an x amount of flow is admitted, and each user pays a price of p, then for each value $q < p$, $\mathcal{D}_{(s,t)}(q) \leq px/q$. Therefore, the total value foregone from not routing flow with value less than p is at most $\int_1^p (px/q - x)dq < px \log p < px \log \mathcal{L}$. With respect to profit, a single-link network is optimal by definition.

Now consider any equilibrium in a general single-commodity single-monopoly network. Suppose that in the equilibrium the network has non-zero residual capacity. Then using an argument analogous to the proof of theorem 1, we can argue that any non-monopoly edge in the network must have a price of 0. The sole monopoly in the network then behaves like a single-link network and obtains optimal profit, as well as a social value that is at least an $O(\log \mathcal{L})$ fraction of the optimal value.

Finally, suppose that at equilibrium the network is saturated. Then it is obvious that the social value of the network is equal to the optimal social value. To conclude, we argue that the network also achieves optimal profit at this equilibrium. In particular, we prove that any flow achieving optimal profit saturates the network, so the total price paid by each user in the two states must be equal. Let f be the amount of flow admitted at equilibrium, and suppose that each admitted user in this flow pays an m amount of money to non-monopoly edges. (Note that this amount is the same for all admitted users.) Then, since the sole monopoly in the network has no incentive to deviate, it must be the case that $f = \text{argmax}_g g(D^{-1}(g) - m)$. Now suppose for the sake of contradiction that the amount of flow at a profit-optimal solution is $f' < f$, that is, $f' = \text{argmax}_g gD^{-1}(g)$. Then, $f(D^{-1}(f) - m) \geq f'(D^{-1}(f') - m) > f'D^{-1}(f') - fm \geq f(D^{-1}(f) - m)$ which leads to a contradiction. ■

Multiple monopolies. The performance of the game with multiple monopolies degrades significantly. We first show that the price of anarchy can be unbounded even with 2 monopolies. The next lemma shows that the best Nash equilibrium behaves slightly better but is still a polynomial factor worse than an optimal solution.

Theorem 3 *For every B, there exists a single-source single-sink instance of the NPG containing 2 monopolies, with $\mathcal{L} = 2$, and $\mathbf{POA_{Val}}, \mathbf{POA_{Pro}} = \Omega(B)$.*

Proof: We construct the instance as follows. The network consists of three nodes s, v and t, and two edges (s, v) and (v, t), both with a capacity of 1 each. All customers want to route their traffic from s to t. The demand curve D is given by $D(\ell) = 0$ for $\ell > 2$, $D(\ell) = 1/B$ for $\ell \in (1, 2]$ and $D(\ell) = 1$ for $\ell \leq 1$. Then, we claim that $\pi_e = 1$ for each of the edges is an equilibrium. This is because in order to get more customers, unilaterally any edge would have to decrease its price to 0. Furthermore, there is no incentive to unilaterally increase price because then no customers would route their flow. The social value and profit of this equilibrium are both $2/B$, whereas the optimal social value (with $\pi_e = 1/2$ for both the edges) is $1 + 1/B$ and the optimal profit is 1. ■

Theorem 4 *There exists a family of single-commodity instances of NPG with $\mathbf{POS_{Val}}, \mathbf{POS_{Pro}} = \Omega(\mathcal{L}^{k-1})$, where k is the nunmber of monopolies. Moreover, in all single-commodity graphs with $k > 1$ monopolies, $\mathbf{POS_{Val}}, \mathbf{POS_{Pro}} = O(\mathcal{L}^{k-1})$.*

Proof: For the first part of the theorem, we consider a graph containing a single source-sink path with k edges and unbounded capacities. There are n users, each endowed with a unit flow. The ith user has value v_i with v_i recursively defined: $v_1 = 2$, $v_2 = (1 - \frac{1}{n})\frac{2k}{2k+1}$, $v_{i+1} = (1 - \frac{1}{n})\frac{ik}{ik+1}v_i$ for $i \in [3, n]$. (That is, $v_{i+1} = (1 - \frac{1}{n})^i \prod_{j \leq i} \frac{kj}{kj+1}$ for $i > 1$.) We first claim that this network contains a single equilibrium, one at which each edge charges a price of $v_1/k = 2/k$, and will then prove that this is a poor equilibrium.

To prove the claim, first note that $\pi_e = 2/k$ is indeed an equilibrium: no edge can gain for raising its price (it would lose all its flow in doing so), and $v_i < (k-1)\pi_e$ for $i, k > 1$ implies that in order to attract more flow any edge must lower its price to below 0. Next we show that there are no other equilibria. Consider any state of the network in which an $m > 1$ amount of flow is admitted. Then the minimum price edge along the path charges a price π, which is at most v_m/k and makes a profit of πm. By increasing its price to $(v_{m-1} - v_m + \pi)$, the edge admits at least $m - 1$ users. We now note that $v_{m-1} - v_m$ is strictly larger than $\frac{v_{m-1}}{1+(m-1)k} > \frac{v_m}{(m-1)k} \geq \pi/(m-1)$. Therefore, the new profit of the edge is strictly larger than $\pi(1 + 1/(m-1))(m-1) = \pi m$, and the edge has incentive to deviate. This concludes the proof of the claim.

Now, since the network has unbounded capacity, the optimal solution (with respect to social value) admits the entire flow. Some algebra shows that $v_n = \Theta(n^{-1/k})$. Therefore, the social value of the optimum is $\sum_i v_i = \Omega(n^{1-1/k}) = \Omega(\mathcal{L}^{k-1})$, as $\mathcal{L} = v_1/v_n = \Theta(n^{1/k})$. The total achievable profit is also at least $nv_n = \Omega(n^{1-1/k}) = \Omega(\mathcal{L}^{k-1})$. On the other hand, the social value of the equilibrium, as well as its profit, is $v_1 \cdot 1 = 2$. This concludes the proof of the first part of the theorem.

For the second part, let $\overline{D}$ denote the inverse-demand curve for the network, that is, for every x, an x amount of flow has value at least $\overline{D}(x)$. Without loss of generality, $\overline{D}(0) = \mathcal{L}$, $\overline{D}(F) = 1$, where $F = \mathbf{F}^{\text{tot}}_{s,t}$ is the total amount of flow (or the capacity of the network, whichever is lesser). Let $x^* = \text{argmax}_{x \leq F}\{x^{1/k}\overline{D}(x)\}$. We claim that the following is a Nash equilibrium: each monopoly charges a price of $p^* = \overline{D}(x^*)/k$, and each non-monopoly charges a price of 0. It is obvious that the non-monopolies do not gain anything from increasing their price. So, for the rest of the proof, we focus on the monopolies.

Suppose that a monopoly wants to deviate and change its price to $p' = p^* - \overline{D}(x^*) + \overline{D}(x') \geq 0$, for some $x' \in [0, F]$. Then, the total price of any source-sink path is $\overline{D}(x')$, and it is immediate that the total amount of flow admitted is no more than x'. Then, the profit of the monopoly goes from p^*x^* to at most $p'x'$. We can express the new profit in terms of the old one as follows:

$$p'x' = \left(\frac{\overline{D}(x^*)}{k} - \overline{D}(x^*) + \overline{D}(x')\right)x' \quad \leq \quad \frac{\overline{D}(x^*)x^*}{k}\left(\frac{x'}{x^*}(1-k) + k\left(\frac{x'}{x^*}\right)^{1-1/k}\right)$$

Using the fact that $(1+\epsilon)^\alpha < 1+\alpha\epsilon$ for all $\epsilon > -1$ and for all $\alpha \in (0,1)$, we get $(\frac{x'}{x^*})^{1-1/k} < 1+(1-1/k)(x'/x^*-1)$. Therefore, the above expression becomes

$$p'x' < \frac{\overline{D}(x^*)x^*}{k}\left(\frac{x'}{x^*}(1-k) + k + (k-1)\frac{x'}{x^*} - (k-1)\right) = p^*x^*$$

This proves that the agent has no incentive to deviate. It remains to show that this equilibrium achieves good social welfare. First, note that $\overline{D}(F)F^{1/k} \leq \overline{D}(x^*)(x^*)^{1/k}$. Therefore, $F \leq x^*(\overline{D}(x^*))^k$. Likewise, for any $y \in [0, F]$, $\overline{D}(y) \leq \overline{D}(x^*)(x^*/y)^{1/k}$. So the total value of flow that is not admitted in the above equilibrium is

$$\begin{aligned}\int_{y=x^*}^{y=F}\overline{D}(y)dy \quad &\leq \quad \int_{y=x^*}^{y=F}\overline{D}(x^*)(x^*/y)^{1/k}dy = \frac{\overline{D}(x^*)(x^*)^{1/k}}{(1-1/k)}(F^{1-1/k} - (x^*)^{1-1/k}) \\ &\leq \quad (1-1/k)^{-1}(\overline{D}(x^*)^k x^* - \overline{D}(x^*)x^*) < 2(\overline{D}(x^*))^k x^*\end{aligned}$$

So, the maximum social welfare achievable is strictly less than $2(\overline{D}(x^*))^k x^*$ plus the social value of the above equilibrium, while the equilibrium achieves at least $\overline{D}(x^*)x^*$. The price of stability is therefore no more than $2(\overline{D}(x^*))^{k-1} + 1 \leq 3\mathcal{L}^{k-1}$. It is easy to see that the same bound holds for profit as well. ■

4 Networks with multiple sources

Next we study the NPG in graphs with more general traffic matrices. Specifically different users have different sources, but a common sink. We assume that the network is a DAG with a single sink, and focus on instances that contain no monopolies[3]. Theorem 1 already shows that the price of stability with respect to profit can be quite large in this case. The main question we address here is whether competition drives down prices and enables a near socially optimal equilibrium just as in the single-commodity case.

The results are surprisingly pessimistic. We find that there are networks that admit no pure equilibria. (To maintain flow, we defer the proofs of Theorems 5 and 6 to Section 4.1.)

Theorem 5 *There exists a multi-source single-sink instance of the NPG with no monopolies that does not admit any pure Nash equilibria.*

In networks that admit pure equilibria, the price of stability for social value can be polynomial in $\mathcal{L}$. This can happen (Theorem 6 below) even when the network in question satisfies a certain strong-competition condition, specifically, (1) there is sufficient path-choice – from every node in the graph, there are at least two edge-disjoint paths to the sink, and (2) no edge dominates over a specific user in terms of the capacity available to that user – removing any single edge reduces the amount of traffic that any user or group of users can route by only a constant fraction. We therefore attempt to isolate conditions that lead to a high price of stability, and find two culprits:

[3] We mainly give strong lower bounds on the price of stability. Naturally, the same bounds hold for instances containing monopolies.

1. Variations in demand curves across users—a very high value low traffic user can pre-empt a low value high traffic user.

2. Congestion in the network—congestion in one part of the network (owing to low capacity), can get "carried over" to a different part of the network (in the form of high prices) due to the ISPs' selfishness.

Each of these conditions by itself can cause the network to have a high price of stability.

Theorem 6 *There exists a family of multiple-source single-sink instances satisfying strong competition and containing uniform demand such that* $\mathbf{POS_{Val}} = \Omega(\textit{poly}\ \mathcal{L}, \textit{poly}\ N)$*, where* N *is the size of the network. There exists a family of multiple-source single-sink instances satisfying strong competition and with sparsity* 1 *such that* $\mathbf{POS_{Val}} = \Omega(\textit{poly}\ \mathcal{L}, \textit{poly}\ N)$*.*

Here uniformity of demand and sparsity defined as follows.

Definition 2 *An instance of the NPG* $(G, \mathcal{D})$ *with multiple commodities and a single sink* t *is said to contain* uniform demand *if there exists a demand curve* D *such that for all nodes* s*, the demand curve* $\mathcal{D}_{(s,t)}$ *is either identically zero, or equal to a scalar* $F_{s,t}$ *times* D*.*

Definition 3 *Given a capacitated graph and a demand matrix, the sparsity of a cut in the graph with respect to the demand is the ratio of the total capacity of the cut to the total demand between all pairs* (s, t) *separated by the cut. The sparsity of the graph is the minimum of these sparsities over all cuts in the graph.*

Fortunately, in the absence of the two conditions given above, the network behaves well. In particular, we consider a certain class of DAGs called traffic-spreaders in which equilibria are guaranteed to exist, and show that when each user has an identical demand curve (in terms of the fraction of traffic with a certain value), but potentially different amounts of traffic, the price of stability with respect to the social value is at most $1/\alpha$, where α is the sparsity of the network. We conjecture that this bound on the price of stability holds for all DAGs that admit pure equilibria.

Definition 4 *A DAG with sink* t *is said to be a* traffic spreader *if for every node* v *in the graph, and every two distinct paths* P_1 *and* P_2 *from* v *to* t*, any maximal common subpath of* P_1 *and* P_2 *is a prefix of both the paths. That is, once the two paths* P_1 *and* P_2 *"diverge" they meet again only at the sink* t*.*

The main theorem of this section bounds the $\mathbf{POS_{Val}}$ in traffic spreaders in terms of the sparsity of the underlying graph.

Theorem 7 *Let* $(G, \mathcal{D})$ *be a uniform-demand instance of the NPG where* G *is a traffic spreader and contains no monopolies, and all sources in the graph are leaves, that is, their in-degree is* 0*. Then* $(G, \mathcal{D})$ *always admits a pure Nash equilibrium, and* $\mathbf{POS_{Val}} \leq 1/\alpha$*, where* α *is the sparsity of* G *with respect to* $\mathcal{D}$*.*

We remark that the networks in the proof of Theorem 6 are traffic spreaders, whereas the one in the proof of Theorem 5 is not. Note also that for the above theorem, we do not require the instance to satisfy strong competition. This indicates that the amount of competition in the network has lesser influence on its performance compared to its traffic distribution.

Proof of Theorem 7. We begin with some notation. Given a graph G and a flow f in G satisfying capacity constraints, $G[f]$ is the residual graph with capacities $c'_e = c_e - f_e$. For a graph $G = (V, E)$, set S of nodes, and set E' of edges, we use $G \setminus S$ to denote $(V \setminus S, E[V \setminus S])$, and $G \setminus E'$ to denote $(V, E \setminus E')$.

Given an instance $(G, \mathcal{D})$, $G = (V, E)$, satisfying the conditions in the theorem, we construct an equilibrium using the algorithm below. Let F_v denote the total traffic at source v, and D be a demand curve defined such that $\mathcal{D}_{v,t} = F_v D$ for all v. The algorithm crucially exploits the sparse-cut structure of the network. In particular, we use as subroutine a procedure for computing the maximum concurrent flow in a graph with some "mandatory" demand. We call this procedure $\mathbf{MCFMD}$ (for Maximum Concurrent Flow with Mandatory Demand).

The procedure $\mathbf{MCFMD}$ takes as input a DAG G with single sink t, a set of sources A with demands F_v at $v \in A$, and a set of mandatory-demand sources B with demands M_v at $v \in B$. The procedure returns a cut C and a flow f. Let V_C denote the set of nodes from which t is not reachable in $G \setminus C$. The cut C minimizes "sparsity with mandatory demand" defined as follows:

$$\alpha_M(C) = \frac{\sum_{e \in C} c_e - \sum_{v \in B \cap V_C} M_v}{\sum_{v \in A \cap V_C} F_v}$$

The flow f routes the entire demand M_v of sources $v \in B$ to t, and routes an $\alpha_M(C)$ fraction of every demand F_v at sources $v \in A$ to t.

The next lemma asserts the correctness of this procedure. Specifically, it states that sparsity is equal to maximum concurrent flow in DAGs, even with mandatory demands, when the commodities share a common sink.

Lemma 8 *Let (G, A, B) denote an instance for procedure* $\mathbf{MCFMD}$*, and $\alpha = \alpha_M(C)$ be the sparsity of the cut C produced by the procedure, as defined above. Then, there exists a flow in G that satisfies all capacity constraints, routes an M_v amount of flow from every $v \in B$ to t, routes an αF_v amount of flow from every $v \in A$ to t, and saturates the cut C.*

Proof: We begin by proving the lemma in the case of $B = \emptyset$. In this case, we are simply claiming that the sparsity of the sparsest cut in a DAG with a single common sink is equal to the maximum concurrent flow in the DAG. We first note that the maximum concurrent flow in the graph can be found by solving the program **LP3** (below left). The dual of this program (**LP4**, below right) is a relaxation of the sparsest cut problem on the same graph.

$$\begin{array}{llll} \text{maximize} & \alpha & & \textbf{(LP3)} \\ \text{subject to} & \sum_{P \in \mathcal{P}_v} f_P \geq \alpha F_v & \forall v \in A & \\ & \sum_{P : e \in P} f_P \leq c_e & \forall e & \\ & f_P \geq 0 & \forall v \in A, P \in \mathcal{P}_v & \\ & \alpha \geq 0 & & \end{array}$$

$$\begin{array}{llll} \text{minimize} & \sum_e \pi_e c_e & & \textbf{(LP4)} \\ \text{subject to} & \sum_{v \in A} F_v s_v \geq 1 & & \\ & \sum_{e \in P} \pi_e \geq s_v & \forall v \in A, P \in \mathcal{P}_v & \\ & \pi_e \geq 0 & \forall e & \\ & s_u \geq 0 & \forall v \in A & \end{array}$$

The lemma now follows from the claim that the integrality gap of the second program is 1. Let the value of the program be α. To prove the claim, consider any optimal fractional solution to the dual program **LP4**, and interpret the π_e values as lengths on edges. Then, s_v denotes the shortest distance from v to t. Now, we modify the standard argument for the integrality gap of the mincut LP – pick β uniformly at random from the range $[0, 1]$. Let T_β denote the set of vertices whose shortest distance to t is less than β, and S_β denote the remaining vertices. The expected size of the cut (S_β, T_β) is no more than α. On the other hand, the expected total demand separated, $E_\beta[\sum_{v \in S_\beta} F_v]$, is at least 1. Therefore, there exists a value β for which the integral cut (S_β, T_β) has sparsity no more than α.

Finally, to prove the lemma for a general B, let $\alpha = \alpha_M(C)$ as in the statement of the lemma, and consider the following instance of sparsest cut: we are given a graph G, and set of sources $A \cup B$ with F_v for $v \in A$ as defined in the original instance, and $F_v = M_v/\alpha$ for $v \in B$. Then the sparsity of this new instance is at least α. So, as argued above, there exists a flow in G for this new instance that routes an α fraction of all the commodities in $A \cup B$ and saturates the cut C. The flow satisfies all the requirements of the lemma. ■

Armed with the procedure $\mathbf{MCFMD}$, our algorithm for constructing an equilibrium is as follows. (Note that we do not care about computational efficiency here.)

1. Set $G_1 = G$, $V_1 = V$, $C = \emptyset$, $B_1 = \emptyset$, $i = 1$. Let $A_1 = A$ be the set of all sources in the instance. Let f denote a partial flow in the graph at any instant; initialize f to 0 at each edge.

2. Repeat until A_i is empty:

   (a) Run the procedure $\mathbf{MCFMD}$ on G_i with demands A_i and mandatory demands B_i. Let C_i be the resulting cut and f'_i be the resulting flow. Let $\alpha_i = \alpha_M(C_i)$, $X_i = A_i \cap V_{C_i}$, $Y_i = B_i \cap V_{C_i}$, and $C = C \cup C_i$. Define V_{i+1} to be the set of nodes with paths to t in $G \setminus C$, and S_i to be the subset of $V \setminus V_{i+1}$ reachable from X_i or Y_i in G.

   (b) We now construct a partial flow from f'_i as follows. Let $B' = \{v : \exists u \text{ with } (u \to v) \in C_i\}$, and for all $v \in B'$ let $M_v = \sum_{u:(u\to v)\in C_i} c_{(u,v)}$. Let f_i be a partial flow of amount $\alpha_i F_v$ from each $v \in X_i$, and amount M_v from each $v \in Y_i$ to B', given by the prefices of some of the flow paths in f'_i. Let $f = f + f_i$, $A_{i+1} = A_i \setminus X_i$, and $B_{i+1} = (B \setminus Y_i) \cup B'$. Set $\ell_i = D^{-1}(\alpha_i)$.

   (c) Let $G_{i+1} = G_i \setminus S_i$; repeat for $i = i + 1$.

3. Route all the flow from B_i to t in G_i satisfying capacity constraints. Call this flow f_i, and set $f = f + f_i$.

4. Assign a "height" to every node v in the graph as follows: if there exists an i such that $v \in S_i$, then $h(v) = \min_{i:v\in S_i}\{\ell_i\}$; if there is no such i, then $h(v) = 0$. Furthermore, $h(t) = 0$ for the sink t.

5. For every edge $e = (u \to v)$, let $\pi_e = \max\{h(u) - h(v), 0\}$.

Let I be the final value of the index i. Recall that V_I is the set of nodes that can reach t in G_I. We will show that (π, f) is a Nash equilibrium. This immediately implies the result, because as we argue below, f admits an $\alpha_i \geq \alpha$ fraction of the most valuable traffic from all sources in X_i. We first state some facts regarding the heights $h(v)$ and the flow f.

Lemma 9 *f is a valid flow and routes an α_i fraction of the traffic from all $v \in X_i$ to t. Furthermore, for every i, $1 < i < I$, in the above construction, $\alpha_i \geq \alpha_{i-1}$, and $\alpha_1 > \alpha$, where α is the sparsity of the graph G.*

Proof: We prove by induction that for every i, $\alpha_i \geq \alpha_{i-1}$, and there exists a flow f'_i in $G[f]$ that routes all the flow from B_i to t, and an α_i fraction of the flow at each source in A_i to t. (Here $f = \sum_{j<i} f_j$.) This immediately implies the result. When $i = 1$, we have $f = 0$ and $B_1 = \emptyset$, and this statements holds because the sparsity of the graph is $\alpha = \alpha_1$. Consider some step i, and assume that the statement holds. We will now prove it for $i + 1$. Note that Lemma 8 implies that f'_i saturates the cut C_i. Now divide f'_i into two partial flows—g_i^1 that routes all demand from Y_i to B' and an α_i fraction of the demand from X_i to B', and $g_i^2 = g_i - g_i^1$. Then, g_i^1 is identical to f_i. Furthermore, g_i^2 is a valid flow for the instance $(G_{i+1}, A_{i+1}, B_{i+1})$ of the $\mathbf{MCFMD}$, satisfies capacity constraints in $G[f + g_i^1]$ and is a certificate of the fact that this instance has sparsity $\alpha_{i+1} \geq \alpha_i$. The existence of flow f'_{i+1} is now guaranteed by Lemma 8. ■

Lemma 10 *$V(G_i) = V_i$ for all $i \leq I$, and $h(v) = 0$ if and only if $v \in V_I$.*

Proof: We prove the first part of the lemma by induction. The base case $i = 1$ is trivial. For the inductive step, recall that $V_{i+1} \subseteq V_i$, $V(G_{i+1}) = V_i \setminus S_i$ and $S_i \cap V_{i+1} = \emptyset$. Note that every node $v \in V_i$ must either be reachable from $X_i \cup Y_i$ or be able to reach t in $G_i \setminus C_i$, otherwise the optimality of C_i for the instance (G_i, A_i, B_i) of $\mathbf{MCFMD}$ is contradicted. Therefore, $V_i \subseteq S_i \cup V_{i+1}$, and the claim follows.

For the second part of the lemma, note that $v \in V_I$ immediately implies that for all i, $v \notin S_i$, therefore, $h(v) = 0$. On the other hand, $h(v) = 0$ implies that there is no index i such that $v \in S_i$ (because $\ell_i > 0$ for all i). Therefore, $v \in V_1$ and $V_i = V_{i-1} \setminus S_{i-1}$ for all $i > 1$, implies that $v \in V_i$ for all i. ■

Lemma 11 *For every pair of nodes u and v with $h(u), h(v) > 0$ such that there is a directed path from u to v in G, $h(u) \geq h(v)$.*

Proof: We use the fact that for any i, j with $j > i$, $\alpha_i \leq \alpha_j$ (Lemma 9), and therefore, $\ell_i \geq \ell_j$. Consider the smallest index i with $v \in S_i$; we claim that for all $j > i$ with $u \in S_j$, we also have $v \in S_j$—for any such j, both u and v are absent from V_{j+1}, and if u is reachable from X_j or Y_j in G, so is v. Now let $i_1 = \text{argmax}_i\{u \in S_i\}$ and $i_2 = \text{argmax}_i\{v \in S_i\}$. Then we have $i_1 \leq i_2$, $h(u) = \ell_{i_1}$ and $h(v) = \ell_{i_2}$, and therefore, $h(v) \leq h(u)$. ■

The next lemma follows from noting that if v is a source, then the in-degree of v is 0, and therefore there is a unique j, namely $j = i$, such that $v \in S_j$.

Lemma 12 *For any source v with $v \in X_i$, $h(v) = \ell_i$.*

Lemma 13 *For every node v with $h(v) > 0$, every path from v to t is fully saturated under the flow f.*

Proof: $h(v) > 0$ implies that v belongs to S_i for some index i. Therefore, $v \notin V_{i+1}$ and the cut C separates v from t. The lemma follows by noting that the final flow f saturates C. ■

Lemma 14 *For every source v with $v \in X_i$, every path from v to t has total price at least ℓ_i. Furthermore, there exist at least two edge-disjoint paths P_1 and P_2 from v to t such that $\sum_{e \in P_1} \pi_e = \sum_{e \in P_2} \pi_e = \ell_i$.*

Proof: For the first part of the lemma, note that $h(v) = \ell_i$ (by Lemma 12), $h(t) = 0$, and for every edge $e = (x \to y)$, $\pi_e \geq h(x) - h(y)$.

Consider any two disjoint paths from v to t, say P_1 and P_2 (these clearly exist, as the network contains no monopolies). Let x_1 and x_2 be the closest points in V_I to v along these paths respectively. Let Q_{11} and Q_{12} denote the prefixes of P_1 and P_2 from v to x_1 and x_2 respectively. Note that t is reachable from x_1 and x_2 in $G[V_I]$; let Q_{21} and Q_{22} denote any paths in $G[V_I]$ from x_1 and x_2 to t respectively. These paths are clearly disjoint (as the network is traffic spreading). Now, consider the paths $Q_1 = Q_{11} \cdot Q_{21}$ and $Q_2 = Q_{12} \cdot Q_{22}$, where "$\cdot$" represents concatenation. Then, Q_1 and Q_2 are disjoint; the lengths of Q_{21} and Q_{22} are 0 under the metric π; and, all nodes in Q_{11} and Q_{12} have non-zero heights, implying that the lengths of Q_{11} and Q_{12} are equal to $h(v) - h(x_1) = h(v) - h(x_2) = \ell_i$ (using Lemma 11). These facts together imply the lemma. ■

Lemma 15 *Let P be a flow carrying path from some source $v \in X_i$ to t. Then $\sum_{e \in P} \pi_e = \ell_i$.*

Proof: Following the proof of the previous lemma, we only need to show that if u is the first node on path P such that when the algorithm terminates, $u \in V_I$, then all subsequent nodes on the path P are in V_I. Let u' be the node preceding u in P. Then the edge $(u' \to u)$ belongs to some cut C_i, and $u \in B_{i+1}$. Then, since $u \in V_I$, it must be the case that $u \in B_j$ for all $j > i$. This implies that the partial flow of amount M_u gets routed to t in the last step of the algorithm, and the flow only uses edges in G_I. ■

Finally, we claim that (π, f) is an equilibrium. First observe that we route an $\alpha_i F_v$ amount of flow for every v in X_i (Lemma 9). Each chunk of traffic originating at v that gets routed has value at least $D^{-1}(\alpha_i) = \ell_i$. Therefore, Lemmas 14 and 15 imply that users follow best response. Next, consider any edge $e = (u \to v)$. Note that edge e has no incentive to increase its price – Lemma 14 ensures that all the traffic on this edge has an alternate path of equal total price; so by increasing its price, e risks losing all of its flow. Finally, if the edge has non-zero price, the edge stands to gain from lowering its price only if in doing so it can increase the traffic carried by it. Let e be the edge $(u \to v)$ and C' be the mincut between u and t. Note that $h(u) > 0$. Lemma 13 implies that the cut C' is saturated. Suppose that e has non-zero residual capacity (i.e. $e \notin C'$) and by lowering its price, the edge gains extra traffic without violating the capacity of the cut C'. This means that the extra traffic on e was previously getting routed along a path that crosses the cut C', and furthermore shares a source with the edge e. This contradicts the fact that the network is a traffic spreader. Therefore, no edge has an incentive to deviate. This concludes the proof of the theorem.

4.1 Multi-commodity networks with bad equilibria

We now restate and prove Theorems 5 and 6.

Theorem 5 *There exists a multi-source single-sink instance of the NPG with no monopolies that does not admit any pure Nash equilibria.*

Proof: Consider the example in Figure 1(a). The capacities on edges are as shown in the figure. All unlabeled edges have unbounded capacity. Let D denote the demand curve $D(3) = 3$, $D(2) = 5$, and $D(1) = 6$. The proof is by case analysis. Consider any pure strategy equilibrium in this instance. Note first that the two edges from node 4 to node 2 must charge a price of zero—if not, then one of them can lower its price and gain all of the flow of the other. Second, the total price paid by all flow getting routed from source 4 must be at least 3, otherwise, one of the edges $(4 \to 1)$ or $(2 \to 1)$ is overloaded and is incentivized to raise its price. Now the basic idea is that the traffic from source 5 gets routed exclusively through node 3, and the edge $(3 \to 1)$ faces no competition from the edge $(2 \to 1)$. This creates a virtual monopoly in edge $(3 \to 1)$. Let us now focus on the edges $(3 \to 1)$ and $(6 \to 1)$. Let the price charged by the former in equilibrium be x and that charged by the latter be y. Note that $x, y \leq 3$. For the rest of the proof for simplicity of exposition, we assume that the edges $(6 \to 3)$ and $(5 \to 3)$ charge a price of 0 each, although the same argument holds even without this assumption. Suppose first that $x > y$. Then all the flow originating at 6 uses the edge $(6 \to 1)$. Given the low capacity of this edge, the edge charges a price of at least 3, contradicting $x > y$. Next, suppose that $x < y$. Then $(6 \to 1)$ carries no flow, and either $y = 0$ (contradicting $x < y$), or the edge has incentive to deviate and lower its price. Finally, suppose that $x = y$. Then if x and y are larger than 1, then one of the two edges $(3 \to 1)$ and $(6 \to 1)$ has non-zero residual capacity, and is incentivized to raise its price and steal the other's flow. Otherwise, $x = y = 1$, but in this case $(3 \to 1)$ has an incentive to raise its price to 2, and obtain a profit of 20 instead of 18. ■

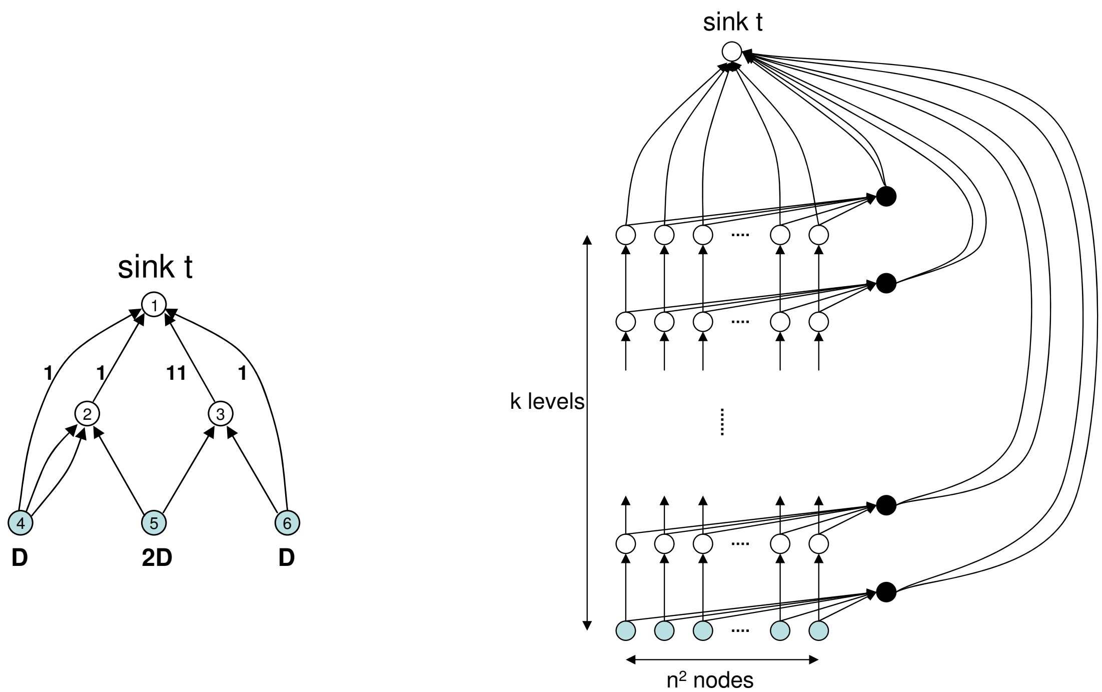

Figure 1: (a) Instance with no pure Nash equilibrium; (b) Instance with high price of stability

Theorem 6 *There exists a family of multiple-source single-sink instances satisfying strong competition and containing uniform demand such that* $\mathbf{POS_{Val}} = \Omega(\textit{poly}\ \mathcal{L}, \textit{poly}\ N)$*, where* N *is the size of the network. There exists a family of multiple-source single-sink instances satisfying strong competition and with sparsity* 1 *such that* $\mathbf{POS_{Val}} = \Omega(\textit{poly}\ \mathcal{L}, \textit{poly}\ N)$.

Proof: For the first part of the theorem, the family of examples, parameterized by integers n and k, $1 < k < n$, is given in Figure 1(b). The network has $N = \Theta(n^2 k)$ nodes. Each edge in the network has capacity n. Let D denote the demand curve used in the proof of Theorem 4, that is, there are n consumers, each with a unit amount of traffic, and the ith user endowed with value $v_i = (1-1/n)^{i-1}\prod_{j<i}\frac{jk}{jk+1}$ for $i > 1$ and $v_1 = 2$. There are $n^2 + k$ different sources. The demand at each light-colored node in the graph (the n^2 nodes at the lowest level) is given by D, while the demand at each dark-colored node (the k nodes to the right) is given by $2nD$. It is immediate that the instance satisfies strong competition and contains uniform demand.

Let us consider an arbitrary equilibrium in this instance. The traffic at each dark-colored nodes faces congestion – at most a $2n$ amount of flow can be admitted from each of these source. Therefore, the price charged by the edges going from these sources to the sink is at least $D^{-1}(1) = 2$. This means that for the traffic at any light-colored node, the path of choice is the k-hop path that goes through each of the k levels in the network, and then to the sink. Each such path then behaves as a network of k monopolies and the results of theorem 4 apply. In particular, any equilibrium admits a total value of $2n^2 + 4nk = O(n^2)$ for $k < n$, whereas the social optimal solution admits a total value of $\Omega(n^2 n^{1-1/k} + 4nk) = \Omega(n^{3-1/k})$. Noting that $\mathcal{L} = \Theta(n^{1/k})$, the price of stability is $\Omega(n^{1-1/k}) = \Omega(\mathcal{L}^{k-1})$.

For the second part of the theorem, the family of examples is obtained by modifying the one used previously in the following way: the demand at every dark-colored node now consists of a $2n$ amount of flow with value 2 (see Figure 1(b)). The equilibria and optimal solutions in the two cases remain the same. However, note that the sparsity in the new instances is 1 (all the flow is admitted in the optimal solution). ■

5 Discussion and Open Questions

In addition to the obvious questions that remain unresolved (e.g. extending Theorem 7 to all instances that admit pure equilibria), we now discuss a few extensions and open problems related to this work.

1. **Multi-parameter users and QoS-based pricing.** In the context of pricing mechanisms for the Internet, one weakness of our work is that it does not take into account quality of service requirements of the users. For example, users may attach different values with different paths. These may be manifested as the differences in preferences consumers may have over different source-sink paths. How does market efficiency and profit depend on the network and traffic structure in this multi-parameter case? One simple way of modeling these preferences is to associate a global QoS parameter with each edge, which assigns an additional per-unit-flow cost to the edge. Several of our results carry over to this setting in a straightforward manner. Notably, Theorem 7 is not necessarily true in this setting. It would be interesting to determine conditions under which the price of stability is small in this new model.

   More generally, different users may attach different costs with the paths. For example, real-time video traffic may prefer a low-jitter path, whereas file transfers may prefer a high bandwidth path. In this setting, in order for the equilibria of the game to behave nicely, we would require a stronger guarantee on competition, namely that no single edge dominates any specific kind of quality of service.

2. **Seller costs.** In a network setting, the marginal costs to sellers for admitting a larger flow are negligible (subject to capacity constraints). In a general two-sided market, sellers may have costs associated with each additional copy of the item that they sell. Again this would modify seller behavior and may lead to better or worse equilibria in the game.

3. **Cournot competition.** As mentioned earlier, an alternate model of competition in two-sided markets is for the sellers to commit to producing (or making available in the market) certain quantities of their product, and then allow market forces to determine the demand and prices. This two-stage game is known as "Cournot competition". A natural question is whether Cournot competition leads to better or worse equilibria compared to Bertrand competition (where sellers commit to prices). Unfortunately, in the case of combinatorial markets, this question is ill posed—in some instances, given the strategies of sellers (the quantities that they commit to), many different sets of prices can arise in the second stage of the game. Whether or not a seller is playing a best response depends on the expected outcome of the second stage. One such instance is a network containing a single source-sink pair, with a single path of length two between the source and the sink. Then, once the

capacities of the two links are determined in the first stage of the game, the total price to be paid by the consumers is uniquely determined; however this price may be split across the two edges in the path in an arbitrary way.

What properties do the equilibria satisfy in instances for which the second stage of the game has a unique solution? One specific class of such instances is single-source single-sink networks with k parallel links. In this case, the prices of anarchy and stability with respect to social value can be as large as $\log \mathcal{L}/k$, but no larger. With respect to profit, the two ratios can be arbitrarily large. We omit the details. Is it possible to define canonical solutions to the second stage of the game such that equilibria in the Cournot game (for some class of instances) are better behaved than those in the Bertrand game?

4. **Market evolution and investment.** An important aspect of studying any market is to determine how the market evolves over time. This is especially interesting in the context of the Internet which has and will continue to evolve in a distributed fashion motivated by economic considerations. Under what conditions do existing ISPs have incentive to invest in more bandwidth?

   As the network evolves, we may expect monopolies to turn into oligopolies, thereby leading to an improvement in social value. Alternately, the price of stability for profit may be high enough in highly competitive markets so as to deter entry of new ISPs (and thereby hurt social value in the long term). Likewise, under what conditions do existing ISPs have incentive to invest in more bandwidth? Weintraub et al. [16] recently studied these questions in the context of latency-based network pricing games.

Acknowledgements

We are grateful to Aditya Akella, Suman Banerjee, and Cristian Estan for bringing to our notice their recent work on a system for QoS-based pricing in the Internet [5], which motivated the network pricing game we study in this paper.

References

[1] D. Acemoglu and A. Ozdaglar. Competition and efficiency in congested markets. *Mathematics of Operations Research*, 32(1):1–31, 2007.

[2] D. Acemoglu and A. Ozdaglar. Competition in parallel-serial networks. *IEEE Journal on Selected Areas in Communications, special issue on Non-cooperative Behavior in Networking*, 25(6):1180–1192, 2007.

[3] E. Anshelevich, B. Shepherd, and G. Wilfong. Strategic network formation through peering and service agreements. In *Proc. 47th IEEE Symp. Foundations of Computer Science*, 2006.

[4] R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 34th ACM Symp. Theory of Computing*, 2003.

[5] C. Estan, A. Akella, and S. Banerjee. Achieving good end-to-end service using bill-pay. In *HOTNETS-V*, 2006.

[6] L. Fleischer, K. Jain, and M. Mahdian. Tolls for heterogeneous selfish users in multicommodity networks and generalized congestion games. In *Proc. 45th IEEE Symp. Foundations of Computer Science*, pages 277–285, 2004.

[7] R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35(12):1969–1985, 1999.

[8] A. Hayrapetyan, É. Tardos, and T. Wexler. A network pricing game for selfish traffic. In *Proc. Symp. Principles of distributed Computing*, pages 284–291, 2005.

[9] J. Mackie-Mason and H. Varian. Pricing congestible network resources. *IEEE J. Selected Areas Commun.*, 13(7):1141–1149, 1995.

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

[11] A. Odlyzko. Paris metro pricing for the internet. In *Proc. 1st ACM Conf. Electronic Commerce*, pages 140–147, 1999.

[12] A. Ozdaglar and R. Srikant. Incentives and pricing in communication networks. In *Algorithmic Game Theory*. Cambridge Press, 2007.

[13] T. Roughgarden and É. Tardos. Introduction to the inefficiency of equilibria. In N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani, editors, *Algorithmic Game Theory*, chapter 17, pages 443–459. Cambridge University Press, 2007.

[14] A. M. Spence. Entry, capacity, investment and oligopolistic pricing. *The Bell Journal of Economics*, 8(2):534–544, 1977.

[15] K. Srinivasan. Multiple market entry, cost signalling and entry deterrence. *Management Science*, 37(12):1539–1555, 1991.

[16] G. Y. Weintraub, R. Johari, and B. Van Roy. Investment and market structure in industries with congestion. Unpublished manuscript, 2007.

A Altruistic ISPs and the optimal social value

Does inefficiency arise in network pricing games due to the selfishness of the ISPs, or the selfishness of the users, or both? In order to answer this question and better understand the game, in this section we consider a version of the game in which ISPs are altruistic, and try to achieve good social efficiency (albeit without centralized collaboration). In particular, each ISP commits to a price π_e, and tries to maximize the amount of traffic it carries subject to capacity constraints—the ISP's utility from carrying an f_e amount of traffic is f_e when $f_e \leq c_e$ and 0 otherwise. Prices on edges are set according to market supply and demand of bandwidth. We use the superscript $\mathcal{A}$ to denote the price of anarchy and stability in this version of the game.

In this section we assume that demand curves are step functions with a finite number of steps. Further, to simplify exposition, we assume that the instance contains a finite number of users, each with a unit amount of flow and uniform value. (The two assumptions are equivalent.) We first characterize the optimal solution to a network pricing game with respect to social value. Then the following linear program computes the optimum:

$$\begin{aligned} \text{maximize} \quad & \sum_u \ell_u \sum_{P \in \mathcal{P}_u} f_P & & \textbf{(LP1)} \\ \text{subject to} \quad & \sum_{P \in \mathcal{P}_u} f_P \leq 1 & & \forall u \\ & \sum_{P : e \in P} f_P \leq c_e & & \forall e \\ & f_P \geq 0 & & \forall u, P \in \mathcal{P}_u \end{aligned}$$

Here u indexes users, and $\mathcal{P}_u$ denotes the set of source-sink paths available to user u, and ℓ_u denotes the user's per-unit value. The dual to this program defines prices supporing the optimal solution:

$$\begin{aligned} \text{minimize} \quad & \sum_e \pi_e c_e + \sum_u s_u & & \textbf{(LP2)} \\ \text{subject to} \quad & \sum_{e \in P} \pi_e \geq \ell_u - s_u & & \forall u, P \in \mathcal{P}_u \\ & \pi_e \geq 0 & & \forall e \\ & s_u \geq 0 & & \forall u \end{aligned}$$

The variables in this dual program can be interpreted as follows: π_e is the price charged by edge e; s_u is the surplus of user u after it has paid the price on a min-price path.

We now show that when ISPs are altruistic, all possible equilibria constitute feasible solutions to the programs **LP1** and **LP2** that together satisfy complementary slackness. The LP duality theorem then implies that all equilibria in this game are optimal with respect to value.

Theorem 16 *In any network, the set of Nash equilibria for NPG*$^{\mathcal{A}}$ *is identical to the set of optimal solutions to the primal and dual programs* **LP1** *and* **LP2** *above, i.e.* $\mathbf{POA}^{\mathcal{A}}_{\mathbf{Val}} = \mathbf{POS}^{\mathcal{A}}_{\mathbf{Val}} = 1$.

Proof: First we note that any pair of optimal solutions to the above primal and dual form a Nash equilibrium—as noted before, users are not motivated to deviate; furthermore, complementary slackness implies that any ISP with non-zero residual capacity charges a price of zero, therefore, no unilateral deviation on part of the ISP can lead to higher flow on its edge.

Next, consider any equilibrium. Then, any edge with non-zero residual capacity and non-zero price can lower its price and potentially improve its usage. Therefore, at equilibrium all edges with non-zero prices are saturated. Furthermore, any $s-t$ path that is not a least cost $s-t$ path has zero flow along it. In other words, flow and prices together satisfy complementary slackness conditions, and are therefore optimal. ■