Bertrand competition in networks

Shuchi Chawla[1]*, Tim Roughgarden[2]**

[1] Computer Sciences Dept., University of Wisconsin - Madison. shuchi@cs.wisc.edu.
[2] Department of Computer Science, Stanford University. tim@cs.stanford.edu.

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., [5, 7, 8, 10, 12]). 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 profit, and consumers selfishly purchase bundles of products to maximize their 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

* Supported by NSF CAREER award CCF-0643763.
** Supported in part by NSF CAREER Award CCF-0448664, an ONR Young Investigator Award, and an Alfred P. Sloan Fellowship.

first picking (per-unit-bandwidth) prices for their edges, and the consumers buying their most-desirable paths (or nothing 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 social value (efficiency) 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, e.g., [15, 16] and references therein). It is well known [11], e.g., that in a single-good free market, oligopolies (two or a few competing firms) lead to a socially-optimal equilibrium[3]. On the other hand, a monopoly can cause an inefficient allocation by selfishly maximizing its own profit. Fortunately the extent of this inefficiency is bounded by a logarithmic factor in the (multiplicative) disparity between consumer values, 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 equilibrium 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 to that of an optimal solution.

[3] 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 the full version [4] of this paper for a brief discussion of the Cournot model.

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.

Related work. The literature on quantifying the inefficiency of equilibria is too large to survey here; see [14] 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 [9, 1, 2, 13, 17]. In these other works, 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, consider single-source single-sink networks with parallel links and assume that all consumers are identical and have unbounded values (i.e. they simply minimize their total routing cost). Hayrapetyan et al. [9] consider the same class of networks 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.

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 in [3] 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. Furthermore, in their model 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*[4], $\mathcal{D}_{(s,t)}$, 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, $\mathcal{D}_{(s,t)}(1) = \mathbf{F}^{\text{tot}}_{s,t}$ 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 edge. ISP e's utility is given by $f_e \pi_e$ if $f_e \leq c_e$, and $-\infty$ otherwise. ISPs are selfish and set prices to maximize their utility.

A given state in a game (in this case consisting of a set of prices and flow) is called a Nash equilibrium if no agent wants to deviate from it unilaterally so as to improve its own utility. Note that in the NPG, users are price-takers, that is,

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

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 the full version of this paper [4] we present an example that admits 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 stage is such that the capacity constraint on an edge e is violated, users using e still obtain their value from routing their flow, while 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 cannot 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 of all the ISPs, or the total payment made by all users. The optimal profit $\mathbf{Pro}^*$ is defined to be the maximum profit over all states in which users are at 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,* $\mathbf{POA_{Val}} = 1$*. Furthermore, there exist instances with* $\mathbf{POS_{Pro}} = \Theta(\mathcal{L})$*.*

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; pick an arbitrary min-cut, and assign a price of p to every edge in the min-cut. These prices, along with the flow f form an equilibrium: edges cannot improve their profits by increasing prices unilaterally, because their customers can switch to a different cheaper path, and, 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 $\mathbf{POA_{Val}} = 1$. Otherwise, there exists an unsaturated edge, say e, with non-zero price that does not carry all of the admitted flow (if there exists a zero-price unsaturated path, then some users are playing suboptimally). Then there is a source-sink path P carrying flow with $e \notin P$. Edge e can then 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, we consider a network with unbounded capacity. Our argument above (that $\mathbf{POA_{Val}} = 1$) 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 the total profit of the network is $\mathbf{F}^{\text{tot}}_{s,t}$. 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 $\mathcal{L}\mathbf{F}^{\text{tot}}_{s,t}$.

Single monopoly. As we show below, 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* 1 *monopoly,* $\mathbf{POA_{Pro}} = 1$ *and* $\mathbf{POA_{Val}} = O(\log \mathcal{L})$*. Moreover, there exist instances with* $\mathbf{POS_{Val}} = \Theta(\log \mathcal{L})$*.*

Proof. The second part follows by 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 monopolist, and w.l.o.g. 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. We omit the straightforward extension to general single commodity networks (see [4]).

Multiple monopolies. The performance of the game with multiple monopolies degrades significantly – the price of anarchy can be unbounded even with 2 monopolies. As we show below, 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. Consider a network with a single source s, a single sink t, an intermediate node v, and two unit-capacity edges (s, v) and (v, t). $\mathbf{F}^{\text{tot}}_{s,t} = 1$; all but a $1/B$ fraction of the traffic has a value of 1; the rest has a value of 2. We claim that $\pi_e = 1$ for each of the edges is an equilibrium: there is no incentive to increase price (and lose all customers), and, in order to get more customers, unilaterally any edge must decrease its price to 0. 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 with* $\mathbf{POS_{Val}}$*,* $\mathbf{POS_{Pro}} = \Omega(\mathcal{L}^{k-1})$*, where* k *is the number 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 $i > 1$.) This network contains a single equilibrium, one at which each edge charges a price of $v_1/k = 2/k$, and admits a single user.

Since the network has unbounded capacity, the optimal solution (for social value) admits the entire flow. Some algebra shows that $v_n = \Theta(n^{-1/k})$. So, 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, i.e., 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 optimal amount of flow. Let $x^* = \text{argmax}_{x \le F}\{x^{1/k}\overline{D}(x)\}$. We claim that the following is an equilibrium: each monopoly charges a price of $p^* = \overline{D}(x^*)/k$, and each non-monopoly charges 0. It is obvious that the non-monopolies have no incentive to increase 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 $p' = p^* - \overline{D}(x^*) + \overline{D}(x') \ge 0$, for some $x' \in [0, F]$. Then, the total price of any source-sink path is $\overline{D}(x')$, and the total amount of flow admitted is no more than x'. The profit of the monopoly goes from p^*x^* to at most $p'x'$, which can be simplified as follows:

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

Here we used $(1+\epsilon)^\alpha < 1 + \alpha\epsilon$ for all $\epsilon > -1$ and for all $\alpha \in (0,1)$. This proves that the agent has no incentive to deviate. It remains to show that this equilibrium achieves good social welfare. First note that $\overline{D}(F)F^{1/k} \le \overline{D}(x^*)(x^*)^{1/k}$. Therefore, $F \le x^*(\overline{D}(x^*))^k$. Likewise, $\forall y \in [0, F]$, $\overline{D}(y) \le \overline{D}(x^*)(x^*/y)^{1/k}$. So the total value of flow not admitted by the equilibrium is

$$\int_{y=x^*}^{y=F} \overline{D}(y)dy \le \int_{y=x^*}^{y=F} \overline{D}(x^*)(x^*/y)^{1/k}dy = \frac{\overline{D}(x^*)(x^*)^{1/k}}{(1-1/k)}(F^{1-1/k} - (x^*)^{1-1/k})$$
$$\le (1-1/k)^{-1}(\overline{D}(x^*)^k x^* - \overline{D}(x^*)x^*) < 2(\overline{D}(x^*))^k x^*$$

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 $2(\overline{D}(x^*))^{k-1} + 1 \le 3\mathcal{L}^{k-1}$. 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 matrix. 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[5]. 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 with no pure equilibria. (See [4] for proofs of the next two theorems.)

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

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:

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 condition alone 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* $\mathbf{POS_{Val}} = \Omega(\text{poly } \mathcal{L}, \text{ poly } N)$*, 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* $\mathbf{POS_{Val}} = \Omega(\text{poly } \mathcal{L}, \text{poly } N)$.

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 s, $\mathcal{D}_{(s,t)}$ is either 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 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 demand is uniform, 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.*

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 for Theorem 7, 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 **MCFMD** (for Maximum Concurrent Flow with Mandatory Demand).

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 $v \in B$. It 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 an $\alpha_M(C)$ fraction of demands F_v at sources $v \in A$ to t. The next lemma asserts the correctness of this procedure (see [4] for a proof): sparsity is equal to maximum concurrent flow in DAGs with a single sink, even with mandatory demands.

Lemma 1. *Let (G, A, B) be an instance for* **MCFMD***, and $\alpha = \alpha_M(C)$ be the sparsity of the cut C produced by the procedure. 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, an αF_v amount of flow from every $v \in A$ to t, and saturates C.*

Armed with this procedure, 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 **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) 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 (see [4] for the proofs of these lemmas).

Lemma 2. *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.*

Lemma 3. *$V(G_i) = V_i$ for all $i \leq I$, and $h(v) = 0$ if and only if $v \in V_I$. For any source v with $v \in X_i$, $h(v) = \ell_i$.*

Lemma 4. *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, $h(u) \geq h(v)$. Furthermore, for every node v with $h(v) > 0$, every path from v to t is fully saturated under the flow f.*

Lemma 5. *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 $\sum_{e \in P_1} \pi_e = \sum_{e \in P_2} \pi_e = \ell_i$.*

Lemma 6. *Let P be a flow carrying path from $v \in X_i$ to t. Then $\sum_{e \in P} \pi_e = \ell_i$.*

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. Each chunk of traffic originating at v that gets routed has value at least $D^{-1}(\alpha_i) = \ell_i$. Therefore, Lemmas 5 and 6 imply that users follow best response. Next, consider any edge $e = (u \to v)$. Note that e has no incentive to increase its price – Lemma 5 ensures that all the traffic on this edge has an alternate path of equal total price. Finally, if the edge has non-zero price, it can gain from lowering its price only if this increases the traffic through it. Let C' be the mincut between u and t. Note that $h(u) > 0$. Lemma 4 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.

5 Discussion and Open Questions

We consider a simplistic model for network pricing. A more realistic model should take into account quality of service requirements of the users, which may be manifested in the form of different values for different paths between the same source-destination pairs. In general combinatorial markets it would also be interesting to consider the effect of production costs on the pricing game, and this may change the behavior of the market considerably. Finally, an alternate model of competition in two-sided markets is for the sellers to commit to producing certain quantities of their product, and allowing market forces to determine the demand and prices. This two-stage game, known as "Cournot competition", may lead to better or worse equilibria compared to Bertrand competition. We include a brief discussion of these extensions in the full version of this paper [4].

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 [6], 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 *Foundations of Computer Science*, 2006.
4. S. Chawla and T. Roughgarden. Bertrand competition in networks. http://www.cs.wisc.edu/~shuchi/papers/Bertrand-competition.pdf.
5. R. Cole, Y. Dodis, and T. Roughgarden. Pricing network edges for heterogeneous selfish users. In *Proc. 34th ACM Symp. Theory of Computing*, 2003.
6. C. Estan, A. Akella, and S. Banerjee. Achieving good end-to-end service using bill-pay. In *HOTNETS-V*, 2006.
7. 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.
8. R. J. Gibbens and F. P. Kelly. Resource pricing and the evolution of congestion control. *Automatica*, 35(12):1969–1985, 1999.
9. 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.
10. J. Mackie-Mason and H. Varian. Pricing congestible network resources. *IEEE J. Selected Areas Commun.*, 13(7):1141–1149, 1995.
11. Mas-Colell, Whinston, and Green. *Microeconomic Theory.* Oxford, 1995.
12. A. Odlyzko. Paris metro pricing for the internet. In *Proc. 1st ACM Conf. Electronic Commerce*, pages 140–147, 1999.
13. A. Ozdaglar and R. Srikant. Incentives and pricing in communication networks. In *Algorithmic Game Theory.* Cambridge Press, 2007.
14. 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.
15. A. M. Spence. Entry, capacity, investment and oligopolistic pricing. *The Bell Journal of Economics*, 8(2):534–544, 1977.
16. K. Srinivasan. Multiple market entry, cost signalling and entry deterrence. *Management Science*, 37(12):1539–1555, 1991.
17. G. Y. Weintraub, R. Johari, and B. Van Roy. Investment and market structure in industries with congestion. Unpublished manuscript, 2007.