

A New Look at Selfish Routing

Christos Papadimitriou* Gregory Valiant†

University of California, Berkeley

christos@cs.berkeley.edu gvaliant@cs.berkeley.edu

Abstract: We revisit price of anarchy in network routing, in a new model in which routing decisions are made by self-interested components of the network, as opposed to by the flows as in [12]. This significant departure from previous work on the problem seeks to model Internet routing more accurately. We propose two models: the *latency model* in which the network edges seek to minimize the average latency of the flow through them on the basis of knowledge of latency conditions in the whole network, and the *pricing model* in which network edges advertise pricing schemes to their neighbors and seek to maximize their profit.

We show two rather surprising results: the price of stability in the latency model is unbounded — $\Omega(n^{\frac{1}{60}})$ — even with linear latencies (as compared with $\frac{4}{3}$ in [12] for the case in which routing decisions are made by the flows themselves). However, in the *pricing model* in which edges advertise pricing schemes — how the price varies as a function of the total amount of flow — we show that, under a condition ruling out monopolistic situations, all Nash equilibria have optimal flows; that is, the price of anarchy in this model is *one*, in the case of linear latencies with no constant term.

Keywords: selfish routing; price of anarchy; network routing

1 Introduction

Computer Science is not like farming. It is about building things right, not watching them grow in hope. It is therefore natural that many computer scientists had been uncomfortable about the Internet and the mysterious ability of its many thousands of autonomous systems to connect, quite efficiently and reliably, millions of nodes through uncoordinated, decentralized, and ultimately selfish decisions. The technical expression of this discomfort was the concept of the price of anarchy (POA) [6], gauging the loss of efficiency due to this lack of design and unity of purpose. Therefore, the result of Roughgarden and Tardos [10, 12] that this only entails a 33% loss of efficiency, was a welcome, and celebrated, reassurance.

The model of [12] deviates significantly from the realities of the Internet (as it happens, the Internet is never explicitly mentioned in [12] as the underlying motivation): they assume that routing decisions are made

- by the flows in the network, not self-interested components of the network; and in fact
- based on information about network conditions far downstream.

These weaknesses of the model of [12] were pointed out very early, for example in [9]. *In this paper we examine what happens when one removes these as-*

sumptions, bringing the game much closer to the real Internet. Briefly, it gets much worse, and then much better.

If routing decisions are made by selfish network components, one immediate question is, *what is the objective of each component?* We examine two possibilities:

1. One possible answer is to postulate that each network component charges a per-unit-of-flow price to its flows, and seeks to maximize its profit: payments received, minus payments made, minus cost. It is rather natural to postulate that the operating costs of each component are proportional to the cumulative amount of time the flows spend within the domain of that component. We call this the *pricing model*.
2. In the absence of monetary charges, another natural assumption is the following: each component strives to enhance its reputation by routing in such a way that its flows arrive at their destination with the least possible delay. We call this the *latency model*.

These are the two models considered in this paper. For simplicity, throughout this paper we assume that each selfish network component is simply an edge (as opposed to the more realistic assumption that the agents can be collections of edges). We note that many of our results hold in the case that the nodes are the routing agents, and we suspect that there are few, if any, qualitative differences in behavior.

*Research supported by NSF grant CCF-0635319.

†Research supported under a National Science Foundation Graduate Research Fellowship.

Let us start with the latter model — that is, when edges split their flow so as to minimize the downstream latency experienced by it. (Notice that this model addresses only our first objection above to the [12] model.) It turns out that, in this case, Braess’s famous paradox network (which is worst-case in the Roughgarden-Tardos model) behaves much better: the loss decreases from 33% to 3%. But there is one important — and ominous — difference: this loss *scales*, in that it persists no matter how much flow traverses the network (while in [12]’s model the price of anarchy is very close to 1 for all but a narrow range of traffic rates); this enables a recursive construction that establishes the POA is *unbounded* in this model.

Theorem: *The price of anarchy of the latency model is $\Omega(n^{\frac{1}{10}})$, where n is the number of nodes in the network.*

In fact, we prove something much stronger: Not just the price of anarchy (the cost of the worst Nash equilibrium) is bad; but so is the *price of stability* (the cost of the *best* Nash equilibrium). We also show (Theorem 3.1) that pure Nash equilibria in this game always exist.

There is an interesting issue in defining Nash equilibria in this and similar games, that is worth discussing here. A pure Nash equilibrium is, of course, a set of choices (routing decisions by edges) such that no edge can gain by changing its decision, *if one assumes that all other edges do not change theirs*. In this particular game, the highlighted sentence is a little ambiguous. If I change my outgoing flows, downstream edges will have different incoming flows, and so they will *have* to change something in their outgoing flows! There are two ways to define what it means that an edge “does not change its decision.” In the *fixed ratio* version of the model Nash equilibrium is defined assuming that all other edges keep allocating their incoming flows to their outgoing edges in the same ratios as before. In the *fully rational* variant, it is assumed that downstream edges readjust optimally to the new situation. This latter model is less realistic (as it suggests that autonomous systems in the Internet anticipate sophisticated behavior by far-away autonomous systems, something that is quite implausible); however, it does correspond more closely to the demanding notions of rationality of mainstream Game Theory. As it turns out, these variants behave quite similarly in terms of our present interests and results (even though we can show, see Proposition 3.2, that they can lead to different equilibria).

There is one ray of hope in the latency model: In a network of parallel links (in which we already have nontrivial price of anarchy in the Roughgarden-

Tardos model), the price of anarchy is one. This is easy to see, but it does raise our hopes, because it can be interpreted to mean that the reason for the abysmal performance of the latency model is its dependence on non-local information; perhaps by addressing this problem (by pursuing the second departure above from the Roughgarden-Tardos model) good performance can be restored. (As a matter of fact, our positive result is more general: The price of anarchy is one even in series parallel networks (Theorem 3.3).)

Which brings us to the non-local information problem. The procrustean way of fixing it, by restricting information about marginal latencies to the immediate neighborhood of each edge, results in essentially random decisions and terrible worst-case performance. On the other extreme, it could be argued that no fixing is needed: the latency model does not really require non-local information, in that what is needed by each routing agent to reach a decision is only the average marginal delay of the immediately downstream edges — the kind of information that could be propagated in the network through local exchanges, and be reported by the downstream neighbors. But this argument misses the important point of *incentive compatibility*: The next edge could lie about its flow’s downstream delays — for example, exaggerating them so that its upstream neighbors avoid the paths it uses, and so its own average latency improves.

By what means can information about latency conditions propagate upstream in a manner that is both efficient and incentive compatible? The answer is *prices* — the economist’s instrument of choice for informational efficiency.¹ In Section 3 we describe a model in which edges make both routing and pricing decisions, all based on local information.

This is not the first time that prices have been suggested in this context, see for example [3–5, 13] and Chapter 22 in [8]. However, in these previous works, while prices are set by the edges as they are here, *routing decisions are still made by flows*, as in the Roughgarden-Tardos model. It is the combination of pricing and routing decisions by the edges that makes the difference.

Suppose that we assume that edges charge per-unit-of-flow prices to upstream neighbors routing flow through them, and compete with other edges for upstream flow by announcing these prices. The utility

¹Economic history studies some astonishingly long chains of exchanges between neighbors, such as the trade routes for tin from Cornwall and Afghanistan to Bronze Age Greece ca. 1500 BC [2, 7], which were apparently created in the absence of any non-local information and solely by the local propagation of price-like information.

of each edge now boils down to the payments received from upstream neighbors, minus the amounts paid to downstream neighbors, *minus the costs incurred due to the flow through the edge*. It is natural to assume that this latter cost is proportional to the total latency suffered by the flow at the edge — since this latency reflects the total amount of computational work, or energy, needed in order to process the flow. And this defines our new game: Strategies are pricing schemes (per unit price as a linear function of the flow routed), announced to upstream neighbors, and decisions about splitting the received flow among the downstream neighbors. The utility of an edge is the sum total of the payments it receives, minus the sum of its payments downstream, minus the total activity taking place at the edge.

As it often happens in the analysis of competition in networks [1], to arrive at a meaningful answer we need to exclude “monopolistic” networks, in which one edge has too much control over the outcome. The condition needed is rather technical: there must be an optimum flow in which, for every node other than the destination, there are two edge-disjoint flow-carrying paths from the node to the destination. In the absence of this, things fall apart. Importantly, however, bad examples in the latency model game can be constructed to satisfy this condition. We assume latency functions of the form $\ell(x) = a \cdot x$. We can prove:

Theorem: *The price of anarchy in monopoly-free networks in which edges announce linear pricing schemes is one.*

The price charged by each edge at equilibrium is flat (constant), and, importantly, coincides with the marginal latency to the destination (which, at optimality, is independent of the path) — in other words, precisely the non-local information the decision-makers are missing to make efficient decisions. That is, prices here are a mechanism for achieving total efficiency — in networks where it can be achieved, of course. Our proof starts by establishing that at equilibrium all pricing schemes are flat. We then show that charging above the marginal latency by any subset of the edges leads to a redirection of flow to the competition. The argument is based on a result in resistive electrical networks (Lemma 4.9) that seems to be new to that field (the connection to resistive networks is this: with latency functions of the form $a \cdot x$, edges are resistors with resistance a and optimum flows are (energy-minimizing) currents.)

One parenthesis on our restricted latency functions: we believe that the result holds for much more general functions (any nondecreasing function, possibly positive at zero), and that the proof is within the reach

of our method, despite the fact that it has eluded us. One may argue that this restriction trivializes our result, or at least renders it unsurprising, since the price of anarchy for such latency functions is one already in the Roughgarden-Tardos model. A cursory inspection of our proof reveals that this conclusion is wrong: the mathematical origins and proof techniques of the two results could not be more different (one comes from optimization, the other from variational analysis of resistive networks). And it is a very delicate result: We show that the price of stability for such latencies is *unbounded* without prices, and similarly for the price of anarchy when the prices are restricted to being flat. Finally, we note that *the price of stability is 1 for arbitrary nondecreasing latency functions*, underscoring the incomparability of our results with previous ones on the strategic flows models.

Naturally, the Economics literature abounds with results showing that prices beget efficiency, most notably the so-called First Theorem of Welfare Economics (stating that a price equilibrium results in an allocation that is Pareto efficient). The nature (and, of course, proof) of our result is very different; we show an extreme form of efficiency, viz. maximization of social welfare. Additionally, the presence of (nonconstant) latencies in our network precludes the application of many Economics tools and results. It would be interesting to investigate whether our result can be generalized to broader contexts in Economic Theory, such as trading networks.

Finally, let us recall that the main stated goal and motivation behind the body of literature on game-theoretic routing is to gain insights into the Internet — insights that may be useful in guiding its evolution. Seen this way, our work can be interpreted as rigorous, if tentative, evidence that:

- Selfish routing may be more inefficient than we thought;
- prices bring efficiency in a subtle way;
- common practices for preventing monopolistic situations (such as double-sourcing, i.e. having agreements with several downstream providers) are necessary for this efficiency to arise; and
- short-term competition between routing agents, informed by congestion conditions — something that is, arguably, absent from the current Internet — is crucial for achieving this efficiency.

Some of these insights were inaccessible with existing models. But of course, much work needs to be done to make these implications less tentative.

In Section 2 we define our models, in Section 3 we study the latency model and prove, among other results, our lower bound for general networks. We then

prove our POA = 1 result for the pricing model in Section 4, while in Section 5 we briefly discuss the many questions left open by our work.

2 The Models

We introduce two distinct models, one for latency minimization and one for pricing.

The Latency Model

A *network* $N = (V, E, s, t, \ell)$ is a directed acyclic graph with nodes V , edges E , source $s \in V$ and a sink $t \in V$, and for each edge $e \in E$ a *latency function* $\ell_e : \mathbb{R}_+ \mapsto \mathbb{R}_+$, assumed to be linear of the form $\ell_e(x) = a_e \cdot x$, where $a_e > 0$. We call all edges of the form (u, t) *terminal edges*. The set of non-terminal edges is denoted E_{nt} . (Actually, all the results in Section 3 hold in the more general setting with edge latencies of the form $\ell_e(x) = a_e x + b_e$, for $a_e, b_e \geq 0$.)

We assume that flow of total value r is injected in s and siphoned off t . Each edge $e = (u, v)$, is a strategic player with strategy space Δ_k , the $k - 1$ -dimensional simplex $\{(\alpha_1, \dots, \alpha_k) : \alpha_i \geq 0, \sum_i \alpha_i = 1\}$, where there are $k > 0$ edges leaving v , e_1, \dots, e_k . If edge $e = (u, v)$ plays strategy $A = (\alpha_{e_1}, \dots, \alpha_{e_k})$, this means that if the total flow of e is f_e , edge e_i will receive $\alpha_{e_i} \cdot f_e$ of this flow (in addition to any other flows received from other edges leading into v). To treat the source s uniformly, we assume that there is a single edge (s, s') leaving s , and this edge decides how the flow injected to s splits.

Thus, a strategy profile $\mathcal{A} = (A^e : e \in E_{nt})$ of all nonterminal edges defines a flow $f^{\mathcal{A}}$ in the network. The utility of edge $e = (u, v)$ is then defined to be the average latency of the flow through e to t ; that is,

$$U_e(\mathcal{A}) = - \sum_{(e_0=e, e_1, \dots, e_p) \in \mathcal{P}_{v,t}} \sum_{i=1}^p \ell_{e_i}(f_{e_i}^{\mathcal{A}}) \prod_{j=1}^i \alpha_{e_j}^{e_j-1}$$

where by $\mathcal{P}_{v,t}$ we denote the set of all paths from e to t .

A (*pure*) *Nash equilibrium* is a strategy profile \mathcal{A} such that, for each edge e , A^e is the distribution that maximizes U_e if all other distributions are kept the same. We shall see (Theorem 1) that such equilibria always exist. This is the equilibrium in the *fixed ratio model*.

An equilibrium in the *fully rational model* is a strategy profile \mathcal{A} such that, for all e , if e changes its distribution A^e in any way, and all downstream edges respond by optimally changing their distributions, the utility U_e does not improve. As we shall see, this model results in potentially different equilibria, but in a way that does not affect our upper and lower bounds.

The Pricing Model

In the pricing model our network $N = (V, E, s, t, \ell)$ is an *undirected* graph, again with a source, sink, and latency functions (and a single edge $[s, s']$ out of s). Our reason for adopting undirected networks is technical convenience, as certain insights from resistive networks that we need are more accessible in this case; we strongly believe that our results hold for dags (as well as for many other extensions), but we have not shown this yet.

A *strategy* for a nonterminal edge $e = (u, v)$ is now an object of the form $S = (A_u, A_v, P_u, P_v)$, where A_u is a distribution of the flow in the (v, u) direction into the other edges incident upon u , A_v a distribution of the (u, v) flow into the other edges incident upon v , P_u is a set of pricing functions (assumed to be of the form $p(x) = a \cdot x + b$ with $a, b \geq 0$), one advertised to each edge incident upon u , and P_v is a similar set of advertised pricing schemes to the edges incident upon v . (At equilibrium, of course, only one direction of the flow will be in effect.) Edge (s, s') 's strategy only has the A_u component.

The utility of edge $e = (u, v)$ under the strategy profile $\mathcal{S} = \{S^{e'}\}$ is defined as the sum of all payments received, minus the sum of all payments made for outgoing flows, minus the total latency at the edge:

$$\begin{aligned} U_e(\mathcal{S}) = & \sum_{e'=(w,u)} p_{u,e'}^e(f_{(w,u)}^{\mathcal{S}}) \cdot A_{u,e'}^e \\ & + \sum_{e'=(v,w)} p_{v,e'}^e(f_{(v,w)}^{\mathcal{S}}) \cdot A_{v,e'}^e \\ & - \sum_{e'=(w,u)} p_{u,e'}^e(f_{(v,u)}^{\mathcal{S}}) \cdot A_{u,e'}^e \\ & - \sum_{e'=(v,w)} p_{v,e'}^e(f_{(u,v)}^{\mathcal{S}}) \cdot A_{v,e'}^e \\ & - [f_{(u,v)}^{\mathcal{S}} + f_{(v,u)}^{\mathcal{S}}] \cdot \ell_e([f_{(u,v)}^{\mathcal{S}} + f_{(v,u)}^{\mathcal{S}}]). \end{aligned}$$

A *Nash equilibrium* is a strategy profile in which each edge plays the best response to the pricing and routing choices of the remaining edges together with the assumption that immediate upstream neighbors will always adjust their routing so as to route optimally given the advertised prices. Additionally, we assume that edges that route zero flow advertise price schemes $p(\cdot)$ such that $p(0)$ is equal to their true cost—that is, edges that don't receive flow advertise the smallest price for which they would be willing to route flow. We believe that this assumption is reasonable, and that it can be justified in several natural ways. One informal justification is that these are network operators who are being undercut by more efficient competitors, so it would be irrational for them to

further exaggerate their costs. Also, this assumption would be obviated by a stronger definition of rationality based on local information beyond immediate neighbors. In the absence of this assumption, information might not effectively propagate, and unrealistic equilibria with high cost can arise.

The pricing model behaves very badly in the absence of some assumption that eliminates monopolies, that is, situations in which a player can demand and get unboundedly high prices. Define a flow in the network to be *optimal* if its total latency is as small as possible among all flows. That is, an optimal flow is the solution of the min-cost problem associated with the network; it turns out to be, in the case of the $a_e \cdot c$ latencies we are considering, the current flowing through this network under voltage difference of r between s and t , where the resistance of edge e is a_e .

Definition 2.1 We say that a given instance of the routing problem is *monopoly free* if in every optimal flow, any edge that routes positive flow is either directly connected to the destination, or has at least two downstream edge-disjoint paths to the destination that carry positive flow.

3 Minimizing Latency

In this section we focus on the model in which edges seek to minimize the latency experienced by their flow. We begin by stating some basic facts about the Nash equilibria for the *fixed ratio* and *fully rational* model dynamics, then prove our main results about the price of anarchy for series parallel networks and the price of stability for arbitrary networks. We note that throughout this section, identical results can be obtained by having the routing agents be the network nodes, rather than the edges.

Theorem 3.1 In both the fixed ratio model and the fully rational model with directed acyclic graphs, pure strategy Nash Equilibria exist.

The proof of Proposition 3.1 is by induction on the edges in topological order, beginning with the edges connected to the sink. We defer details to the full version.

Although the fixed ratio and fully rational models differ, one might hope that the equilibria of the fixed ratio model would be a subset of equilibria of the fully rational model. The following proposition, whose proof is by a simple example (which we defer to the full version) demonstrates that this need not be the case, and highlights the sensitivity of the equilibria to the choices in definition of best responses.

Proposition 3.2 There are instances of strategic routing networks with latencies $\ell_e(x) = a_e x + b_e$ for which the fixed ratio model and the fully rational model have disjoint sets of equilibria.

The above proposition illustrates that the fixed ratio model and the fully rational model yield distinctly different games. The results in the remainder of this section illustrate that, although there are differences in the two models, similar price of anarchy and price of stability results hold for both models. We believe that it would be worthwhile to try to extend these results to the class of all ‘reasonable’ definitions of best responses.

Theorem 3.3 For both the fixed ratio model and the fully rational model, the price of anarchy is one, for series parallel networks.

The proof of the above theorem is by induction on the graph structure, and is included in the Appendix.

Theorem 3.4 For arbitrary, even monopoly-free networks, the price of stability is $\Omega(n^{\log_3(51/50)})$ for linear latency functions.

Our proof of Theorem 3.4 is by construction, and relies on recursively embedding the network given in Figure 2, which we refer to as the *beetle* network. We begin by characterizing the unique equilibrium of the beetle network.




Figure 1: The beetle network. For any traffic rate, $POA = 51/50$.

ÐÇÆÚËÄ, ÊöþÔÅ 3, 2009 at 4:52 pm

Lemma 3.5 The price of stability of the beetle network is $51/50$, for any positive traffic rate entering the network.

We defer the proof of the above lemma to the Appendix.

Remark 3.6 *The price of stability of the beetle network is independent of the amount of (nonzero) traffic entering the network; this is a fundamental difference in behaviors of our models and the Roughgarden and Tardos’s strategic flows model, for which any given network has a relatively small range of traffic rates which induce significant POA.*

Proof of Theorem 3.4: We exploit the independence of the price of stability and traffic rate in the beetle network by recursively replacing edges (a, t) , (b, t) , and (c, t) with copies of the entire beetle network. From Lemma 3.5, at equilibrium the beetle with input traffic rate r behaves like a single edges connected to the sink, with latency $\ell(x) = \frac{51}{50}x$, and thus, after $k - 1$ recursions, the unique equilibrium of the network will resemble that of a single link of latency $\ell(x) = \left(\frac{51}{50.3}\right)^k x$. The optimal flow has cost equivalent to that of a single link of latency $\ell(x) = \frac{x}{3^k}$. Finally, note that the network has size $n = 2 + 3^k$ nodes, and thus, in terms of the network size, the price of stability is $\Omega(n^{\log_3(51/50)})$, as desired.

Note that although the beetle network does not satisfy the monopoly-free condition, by splitting the (c, t) edge into two edges of latency $\ell(x) = 2x$, and replacing all latency 0 edges with edges of latency $\ell(x) = \epsilon x$ for any small $\epsilon > 0$, the network will be monopoly-free, and the price of anarchy approaches that for the beetle network as $\epsilon \rightarrow 0$.

4 Maximizing Profit

We now consider the setting in which edges advertise prices to their neighbors. The hope, which proves to be well-founded, is that the prices allow information to propagate in a local fashion in such a manner as to ensure societally ‘good’ equilibria. As a motivation for considering advertised *pricing schemes*—a price-per-unit-flow that is a nondecreasing function of the amount of traffic—the following proposition summarizes what happens if edges can only advertise constant functions.

Proposition 4.1 *If edges can only advertise constant per-unit-flow prices, then the price of stability is one in monopoly-free networks with arbitrary nondecreasing latencies, but worst-case price of anarchy is at least $n - 2$ even for linear latency functions.*

For the $\text{POS} = 1$ result, it is not hard to verify that the following is an equilibrium: The flow is an optimal flow, and each edge advertises a constant price c_e equal to the marginal cost of its internal latency plus the minimum price advertised by its downstream

neighbors. For the price of anarchy result, consider a network of n parallel links of latency $\ell(x) = x$. We defer the details of these proofs to the full version.

Intuitively, the inefficient equilibria caused by requiring that prices be constants arise because the strategy-space is too restricted; an edge has no way of indicating that it would like to route a small amount of flow at a modest price. The problem is that an edge cannot control how much flow it receives—it receives either zero flow, or all the flow from its upstream neighbors depending on whether its price is greater or less than the price of its competitors. As we shall see, this problem can be remedied by expanding the strategy space to allow linear nondecreasing functions as pricing schemes. At a high level, this added power enables edges to control the amount of flow they receive from each upstream neighbor, allowing them to be more specific in expressing their selfishness; somewhat remarkably, this increased expressive power results in optimal flows at equilibrium.

Theorem 4.2 *In monopoly free instances with edge latency functions of the form $\ell(x) = a_e x$ with $a_e > 0$, the price of anarchy is one if nodes can advertise linear, nondecreasing price schemes.*

To prove Theorem 4.2, we’ll argue that in every equilibrium, any edge that is competing with another edge will end up advertising a constant pricing schemes. Then, we show that the value of this constant must be equal to the edge’s true marginal cost for routing the amount of flow it receives, and thus it is the threat of advertising a nonconstant pricing scheme that forces the upstream edges to split their outgoing flow in the desired fashion. Finally, we’ll require a technical lemma that relies on properties of electrical resistive networks to show that in any equilibrium it must be the case that all edges that receive positive flow split their outgoing flow to at least two edges, from which our claim follows.

First, we observe that the optimal equilibrium described in the sketch of the proof of Proposition 4.1 remains an equilibrium in this more general setting in which edges may advertise nondecreasing linear prices, yielding the following proposition:

Proposition 4.3 *The following instance is a Nash equilibrium: fix an optimal flow f , and let each edge $e = (u, v)$ advertises a constant price c_e equal to the marginal cost of all $u - t$ paths that carry positive flow.*

We now show that in all equilibria, edges that compete over flow must advertise constant pricing schemes by demonstrating that an edge that competes

with an adjacent edge for flow from a common upstream neighboring edge has an improving deviation unless it is advertising a constant pricing scheme. We begin with an easy technical lemma, whose proof we defer to the full version.

Lemma 4.4 *Let $p(x) = ax + b$, with some fixed $a > 0$ be the pricing scheme of edge $e = (u, v)$, that is competing with edge $e' = (u, v')$ for flow from an upstream neighbor. Given a fixed non-decreasing price scheme $p'(x) = a'x + b'$ of edge e' and total flow f_u to be routed through u , the amount of flow f_e routed through e is a continuous nonincreasing function of b .*

Remark 4.5 *It is precisely in the case that both the derivative of p and p' are zero that Lemma 4.4 breaks. At first glance, this seems troubling given our characterization of equilibria as exclusively having constant prices. Nevertheless, as noted above, the threat of either e or e' using a nonconstant price is enough to ensure that flow through u splits in the desirable fashion.*

Lemma 4.6 *In a pure strategy Nash equilibrium, for any edge e and an upstream edge u sharing an endpoint, either e receives all of u 's flow, or e advertises a constant pricing scheme to u .*

Proof: For clarity of exposition, we give the proof in the case that edge e is competing with a single edge e' for flow from u . The proof in the general case in which e is competing with multiple edges is similar though slightly more involved, and we defer the details to the full version.

Let $p(x) = ax + b$, with some fixed $a > 0$ be the pricing scheme of edge $e = (u, v)$. Fix a non-decreasing price scheme $p'(x) = a'x + b'$ of competing edge $e' = (u, v')$ and total flow $f_u > 0$ to be routed through u by an upstream neighbor. If e receives flow f_e with $0 < f_e < f_u$, then e can strictly improve its utility by advertising price $p(x) = ax/2 + b'$ for some $b' \geq b$.

Consider the pricing scheme $p_0(x) = ax/2 + b_0$ where $b_0 = b + af_e/2$ is chosen so that $p_0(f_e) = p(f_e)$. We claim that the pricing scheme $p_0(\cdot)$, induces an increased flow to e . By assumption the upstream neighbor chose the routing $f_e, f_u - f_e$ to edges e, e' respectively so as to equalize the marginal costs to e and e' . With pricing scheme p_0 , the marginal cost of routing to e has been decreased, and thus with price p_0 , e will receive flow $f_e^0 > f_e$.

There is some value $b_1 > b$ such that e would receive zero flow were it to advertise price $p(x) = ax/2 + b_1$. By Lemma 4.4, there must be some constant β , with

$b_0 < \beta < b_1$ such that pricing scheme $p(x) = ax/2 + \beta$ induces the upstream neighbor to split the flow in the same manner as was done with pricing scheme $p(x) = ax + b$. To complete the proof, observe that given this new pricing scheme, e receives the same amount of flow as for its original pricing scheme, yet receives a strictly higher per-unit-flow payment.

Given that edges actively competing for flow will advertise constant pricing schemes at equilibrium, the following lemma states that either the flow at equilibrium is the optimal flow, or there are edges who have ‘local monopolies’, meaning that they receive all of an upstream neighbor’s flow.

Lemma 4.7 *In any equilibrium in which all edges that receive nonzero flow split this flow between at least two downstream edges (aside from edges connected directly to the sink), the routing must be an optimal flow.*

At a high level, the proof of the above lemma follows from noting that, by Lemma 4.6, since at such an equilibrium all edges will advertise constant pricing schemes, it must advertise the same constant, c to all its upstream neighbors. Then, for any $\epsilon > 0$, we show that an edge can augment its advertised pricing scheme so as to receive any amount of flow in some range independent of ϵ , at price-per-unit $c - \epsilon$. It then follows that each edge (that doesn’t receive all of an upstream neighbor’s flow) must receive its preferred amount of flow, given its constant advertised price, which corresponds to its true marginal cost. The details can be found in the Appendix.

To complete the proof of Theorem 4.2, we must show that given the monopoly-free condition, there are no equilibria in which some edges receive all of an upstream neighbor’s flow. Intuitively, such a local monopoly seems at odds with our monopoly-free condition—if no monopolies exist when all edges advertise their true marginal costs, then if some edge were to exaggerate his costs, it seems like the flow should penalize him rather than rewarding him with the entire flow from an upstream neighbor. This intuition turns out to be correct. We will show that if some set of edges S advertises prices more than their true marginal costs, at equilibrium it must be the case that at least one of them receives flow from a neighbor who also sends positive flow to a different edge, and thus by Lemma 4.6 we have a contradiction. Our proof relies on viewing the equilibrium flow as an optimal flow in a related network, and then using ideas from electrical circuit analysis to characterize this optimal flow.

Lemma 4.8 *Given a network N at equilibrium, the flow is an optimal flow for a related network N' , which is a network with identical structure to that of N but with latency functions modified as follows. For every edge that does not receive the entire flow of an upstream neighbors, the edge latency in N' is the same as in N . For edges that receive the entire flow from an upstream edge, the latency in N' is $l_{N'}(x) = (a + \alpha)x$ for some $\alpha \geq 0$ where the latency of the edge in N was $l(x) = ax$.*

The proof of the above lemma is intuitively clear, and we defer details to the full version. Now, to conclude our proof of Theorem 4.2 it suffices to consider possible equilibria in which only edges that receive nonzero flow at the optimal flow advertise exaggerated marginal costs. This is true because at least one of these edges must receive nonzero flow in the optimal flow, otherwise the routing would still be the optimal routing. Furthermore, for each edge e that doesn't receive flow at the optimal flow who exaggerates its marginal cost at an equilibrium, we can construct a related instance of the game in which everything is identical, except where edge e 's true marginal cost is equal to its advertised cost at equilibrium (by increasing its latency function, which obviously does not change the optimal flow, and doesn't change the fact that it is at equilibrium).

Lemma 4.9 *Consider a network N where the monopoly-free property holds, and some set of edges e_1, \dots, e_k with the property that no two edges share a source in the optimal flow. Let F denote the optimal flow in N and consider increasing the coefficients of the latency functions of edges e_1, \dots, e_k to yield network N' . Let F' denote an optimal flow for N' . If all edges e_1, \dots, e_k carry nonzero flow in F' , then there is at least one $e_i = (u, v)$ such that some nonzero flow is routed along another edge $e' = (u, v')$.*

Before proving the above lemma, we state a fact adapted from circuit analysis. Consider a network N with edge latencies $l_e(x) = a_e x$ and an optimal flow F . Choose an edge $e = (u, v)$ carrying flow f_e at the optimal flow, and consider increasing its latency function from $l(x) = a_e x$ to $l(x) = (a_e + \delta)x$ to get network N' . Define the associated network N_m to be identical to N , but with an extra node v' inserted into edge e . Let the latency of (u, v') be $l(x) = (a_e + \delta)x$, remove the edge connecting v' and v . Let F_m denote the optimal flow in N_m with source v' and sink v given a traffic rate that induces latency δf_e between v' and v . Note that such a traffic rate exists because the optimal flow scales directly with the input traffic

rate. We now have the following fact:

Fact 4.10 *In the above setup,*

$$F_m + F = F',$$

where we view F_m as a circulation in network N' by merging vertex v' and v .

Proof: First observe that $F_m + F$ is a valid flow for network N' , since F is a valid flow for N' , and F_m is a circulation. Next, to see that $F_m + F$ is optimal for N' , it suffices to check that the sum of the marginal costs around any cycle is zero (where 'backwards' oriented flows contribute negatively). First consider a cycle that contains edge (u, v) ; let f_1, \dots, f_k denote the flows of F around the cycle with f_1 being the flow on edge (u, v) and let f_1^m, \dots, f_k^m denote the flows on these edges in F_m . We have $\sum_i \frac{f_i}{|f_i|} (a_i |f_i|) = 0$, and $(a_i + \delta) f_1^m + \sum_{i=2}^k a_i f_i^m = \delta f_1$, and thus

$$(a_1 + \delta)(f_1^m + f_1) + b_1 + \sum_{i=2}^k \frac{f_i}{|f_i|} (a_i |f_i + f_i^m|) = 0,$$

as desired. For a cycle that does not include edge (u, v) , the argument is similar.

We are now ready to prove Lemma 4.9, completing our proof of Theorem 4.2.

Proof of Lemma 4.9: First note that Fact 4.10 generalizes to a set of edges $S = (u_1, v_1), \dots, (u_k, v_k)$ who increase their latencies, by viewing the network N_m as a resistive network identical to N' , but with voltage sources inserted in each edge of increased latency. Consider the set of nodes $U \subset \{u_1, \dots, u_k\}$ of highest voltage potential. It must be the case that the voltage potential of all nodes in U is higher than any other node in network N_m , aside from the added nodes v' . For each $u_i \in U$, let (u_i, w_i) be an edge that receives positive flow in F , and note that since $(u_i, w_i) \notin S$ by assumption, it must be the case that w_i has strictly lower voltage than u_i . Thus in F_m , no current travels from w_i to u_i , and thus the total flow on edge (u_i, w_i) in F' must be at least that in F , which is positive as claimed.

5 Discussion and Open Problems

We have revisited the classical problem of the price of anarchy in routing, and showed that the price of anarchy becomes unbearable when decisions are made by autonomous systems and are based on knowledge of global traffic conditions, but becomes exquisite in

the presence of prices (and absence of monopolistic situations). One immediate question is, how robust are these results? Do they extend to models in which latencies are more general? (We conjecture that they do extend, by a simple variant of our proof technique that is still eluding us, to arbitrary nondecreasing continuous functions—for example, we can show that any non-optimal equilibrium must involve at least three edges who advertise prices higher than their true marginal costs.) When more elaborate pricing schemes are allowed? When the networks are directed? Or when the latencies, as well as the pricing and routing decisions, are at the nodes? When autonomous systems are modeled, more accurately, as hyperedges? Or when there are multiple sources and multiple sinks? We believe that our results extend to many of these models, but we can see technical obstacles for each extension.

One important question is, what do these result tell us about the Internet, as it is today and as it may evolve in the near future? In the real Internet edges/autonomous systems exchange traffic through the BGP protocol, which implements existing, and more-or-less long-term, agreements between them — agreements that often involve prices. In view of our results, both the positive one and the negative one, we may want to ask: *Which of these current practices are impeding the Internet's ideal efficiency?*

Acknowledgements

We thank Yaron Singer for helping formulate the routing models, and for many valuable discussions throughout.

References

- [1] A. Archer, É. Tardos “Frugal path mechanisms,” ACM Transactions on Algorithms (TALG), 2007.
- [2] K. Arrow, personal communication, October 2007.
- [3] R. Cole, Y. Dodis, T. Roughgarden “Pricing network edges for heterogeneous selfish users,” STOC 2003.
- [4] A. Hayrapetyan, E. Tardos and T. Wexler “A network pricing game for selfish traffic” Distributed Computing 19(4): pp. 255-266, 2007.
- [5] X. Huang, A. Ozdaglar, D. Acemoglu “Efficiency and Braess’ Paradox under Pricing in General Networks,” IEEE Journal on Selected Areas in Communication, special issue on Price-Based Access Control and Economics, vol. 24, no. 5, pp. 977-991, May 2006.
- [6] E. Koutsoupias, C. H. Papadimitriou “Worst-case Equilibria,” STACS 1999.
- [7] R. Maddin, T. S. Wheeler, J. D. Muhly “Tin in the ancient Near East: old questions and new finds,” *Expedition*, 1977.
- [8] N. Nisan, T. Roughgarden, É. Tardos and V. Vazirani (editors), *Algorithmic Game Theory*, Cambridge University Press, 2006.
- [9] C. H. Papadimitriou “Algorithms, Games, and the Internet,” STOC 2001.
- [10] T. Roughgarden “Selfish Routing,” PhD dissertation, Cornell, 2002.
- [11] T. Roughgarden “The price of anarchy is independent of the network topology,” Journal of Computer and System Sciences, 2003.
- [12] T. Roughgarden, Éva Tardos “How bad is selfish routing?” JACM 2002.
- [13] M. J. Smith “The marginal cost taxation of a transportation network,” Transportation Research, issue 13, pp. 237-242, 1979.

Appendix

Proof of Theorem 3.3: We proceed by induction on the graph structure. The claim is trivially satisfied for the base case—a single link. Assuming the claim holds for series parallel networks α, β , we consider the two possible types of compositions. The claim holds trivially in the case that we compose α and β in series, via associating the sink t_α with the source s_β , and letting the new source and new sink be s_α, t_β , respectively.

For the case where we compose α and β in parallel by associating s_α with s_β , and t_α with t_β , in both the fractional and fully rational dynamics, all edges other than the source edge will still route optimally at equilibrium because after the composition, the utility of these nodes for routing in a given manner is identical to their utilities given such a routing in the subnetwork α , or β . Thus it suffices to consider the routing decisions of the source under the two dynamics. Under the fully rational dynamics, the source will route as in an optimal flow because it assumes the downstream nodes will route optimally for their flow, which, by our inductive hypothesis, means optimally for the network as a whole, and optimal routings of series parallel networks consist of optimal routings in each component. Under the fixed ratio dynamics at equilibrium, by our inductive hypothesis, the source’s evaluation of the marginal cost of sending flow to either subnetwork α , or β is the true marginal cost, because at optimum the marginal costs along all paths are equal, and thus the source will have a deviation unless these marginal costs are equal, completing our proof.

Proof of Lemma 3.5: Assume a traffic rate of r entering the network. At equilibrium, edges e and e' will each equalize the marginal cost experienced by *their* traffic along the two possible routes available to each

node. Thus we have that

$$\begin{aligned} 2f_{(a,c)} + f_{(b,c)} &= 2f_{(a,t)} \\ 2f_{(b,c)} + f_{(a,c)} &= 2f_{(b,t)}. \end{aligned}$$

For the fully rational model, using the above characterization of the behavior of edges e and e' together with the conservation of flow constraints, the cost of the source as a function of $f_{(s,a)}$ is proportional to $89r^2 - 50rf_{(s,a)} + 50f_{(s,a)}^2$, which is minimized by setting $f_{(s,a)} = f_{(s,b)} = r/2$. Similarly for the fixed ratio model, given the above characterization of the equilibrium behaviors of edges e , e' , and the constraints imposed by conservation of flow, one can express the source's best-response move as a function of $f_{(s,a)}$. Requiring that there is no best-response move yields that $f_{(s,a)} = f_{(s,b)} = r/2$, as in the case of the fully rational model.

To conclude, note that $f_{(s,a)} = f_{(s,b)} = r/2$ and the characterization of how edges e, e' route at equilibrium imply that $f_{(a,c)} = f_{(b,c)} = r/5$, and $f_{(b,t)} = f_{(a,t)} = 3r/10$, and thus the total cost of the flow is $\frac{17r^2}{50}$. The optimal flow routes $r/3$ flow on each of the edges incident to the sink, and thus has cost $\frac{r^2}{3}$.

Proof of Lemma 4.7: From Lemma 4.6, we know that at such an equilibrium all edges will advertise constant pricing schemes. Consider an edge e that is competing with edges e_1, \dots, e_k for flow from some edge u . Let edge e_i advertise (necessarily constant) pricing scheme p_i . For any $f_0 \in (0, f_u]$, and any $\epsilon > 0$, there is a pricing scheme $p_{f_0}(\cdot)$ which will induce u to route exactly f_0 units of flow to v_0 at price-per-unit $c - \epsilon$, where $c := \min_i(p_i)$. Indeed, it is easily verified that the pricing scheme

$$p_{f_0}(x) := \frac{\epsilon}{f_0}x + c - 2\epsilon$$

induces such a routing.

Now, we show that under the conditions of our lemma, edge e must advertise the same constant to each of its upstream neighboring edges. Let c_1, \dots, c_k be the constant prices it advertises to its k upstream neighboring edges that route positive flows $f_{1,e}, \dots, f_{k,e}$, respectively to e . Let f_1, \dots, f_k be the total amount of flow arriving at each upstream neighbor, and by assumption we have $0 < f_{i,e} < f_i$. First, observe that $c_1 = c_2 = \dots = c_k$; if this were not the case, assume without loss of generality that $c_1 > c_2$, and from above, for any $\epsilon > 0$, e could augmenting its advertised prices so as to receive flow $f'_{1,e} = \min(f_1, f_{1,e} + f_{2,e}) > f_{1,e}$ at price $c_1 - \epsilon$ and flow $f'_{2,e} = f_{2,e} - (f'_{1,e} - f_{1,e}) < f_{2,e}$ at price $c_2 - \epsilon$. Such a change would strictly increase the utility of e .

To see that the flow is optimal, observe that e will have an improving deviation unless it receives the amount of flow f_e^* that maximizes its utility given its advertised constant price-per-unit flow c . By definition,

$$f_e^* := (cx - xC(x)),$$

where $C(x)$ is the private cost per unit flow of edge e as a function the total flow it receives, encapsulating both the internal latency of e , and the prices of the downstream edges. Setting the derivative of the above expression equal to zero, we see that $c = \frac{d[xC(x)]}{dx}$, which is precisely the marginal cost of node e evaluated at the amount of flow it receives.

A symmetric argument applies to the competitors of e , and thus in order for the instance to be at equilibrium, one of the following conditions must hold:

- A single edge e_i receives the entire flow f_u at cost c and all other edges e_j have cost-per-unit flow strictly larger than c .
- $m \geq 2$ edges e_{i_1}, \dots, e_{i_m} receive positive flow from u , and each advertises the same constant price c , and receives flows f_1, \dots, f_m , respectively, such that for all $j \in \{1, \dots, m\}$,

$$c = \frac{d[xC_j(x)]}{dx}(f_j),$$

where $C_j(x)$ is the private cost per unit flow for edge v_{i_j} to route x units of flow. Furthermore, every edge v_i that receives zero flow from u has cost-per-unit flow at least c for routing the flow they receive.

Given an instance at equilibrium in which every edge that receives nonzero flow routes nonzero flow to at least two distinct edges, the above requirements of equilibria guarantee that the routing induces a socially optimal flow. To see why, observe that every edge that routes nonzero flow must satisfy the second condition above, and thus all source-sink paths that carry nonzero flow will have identical costs c^* , equal to their marginal costs, and all source-sink paths that carry zero flow have marginal cost at least c^* ; this is precisely the characterization of optimal flows.