

# Selfish Routers and the Price of Anarchy

Christos Papadimitriou \* and Gregory Valiant\*\*

**Abstract.** We revisit Roughgarden and Tardos's price of anarchy in network routing, in a new model in which routing decisions are made by the *edges* as opposed to the flows. We propose two models: the *latency model* in which 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 edges advertise pricing schemes to their neighbors and seek to maximize their profit. We show the counterintuitive result that the price of stability in the latency model is  $\Omega(n^{\frac{1}{60}})$ , even with linear latencies (as compared with  $\frac{4}{3}$  for the case in which routing decisions are made by the flows themselves). However, in the *pricing model* in which edges advertise pricing schemes — functions dictating how the price varies as a function of the total amount of flow — we show the surprising result that, under a condition ruling out monopolistic situations, all Nash equilibria have societally optimal flows; that is, the price of anarchy in this model is *one*.

---

\* University of California, Berkeley. Email: christos@cs.berkeley.edu.

\*\* University of California, Berkeley. Work supported under a National Science Foundation Graduate Research Fellowship. Email: gvaliant@cs.berkeley.edu.

## 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 area of *the price of anarchy* [5], gauging the cost of this lack of strategic homogeneity, is the technical expression of the field's discomfort.

Roughgarden's thesis [9–11] is by far the most influential work in this subject to date. The main result of Roughgarden and Tardos's famous paper "*How bad is selfish routing?*" [11] is the wonderfully reassuring statement that, under a linear latency assumption, the price of anarchy in routing is upper bounded by about 33% — not insignificant, but certainly neither the disaster one could have feared.

In this paper we revisit and revise the Roughgarden-Tardos model to we obtain two strong and rather surprising results:

- If routing decisions are made not by users as in the Roughgarden-Tardos model but by the edges of the network, aiming to improve average latency of their flows to the destination, and based on information about congestion conditions downstream, then the price of anarchy is *unbounded*.
- But if the routing decisions of an edge are based on *prices* announced and charged by the downstream edges, then the price of anarchy becomes *one*.

The authors of [11] never mention the Internet as the context and motivation of their work, but their results have been interpreted this way in the literature. For example, [8] has these comments after reciting the Roughgarden-Tardos result on the price of anarchy in routing:

“But, of course, in today's Internet flows do not choose paths. In the Internet, routers direct traffic based on local information [...] *What is the price of the Internet architecture?*”

(original emphasis.) *In the present paper we take on this question.*

Notice that, in the above quotation from [8], two distinct aspects of Roughgarden and Tardos's model are identified to be in disagreement with Internet reality:

- (1) The decision makers are not the flows, but the routers (by this, of course, it is meant not the individual routing machines, but the autonomous systems doing the routing, modeled by edges in this papers).
- (2) The decisions are based on local information, for example without full knowledge of downstream congestion.

This paper is an account of what happens if one tries to address these two issues, one by one and in this order.


Recall the classical example known as *Braess's paradox* (Figure 1), in which the optimum flow has cost  $\frac{3}{2}$ , and selfish decisions by individual flows lead to a cost of 2, thus achieving the  $\frac{4}{3}$  worst case. To address (1) above is to ask, *what happens if the routing decisions are made by edges instead of flows?* That is, let us assume that each edges splits the flow through it among the edges adjacent to its downstream endpoint in such a way as to *minimize the average latency to the destination*. We call this the *strategic edges model*, as opposed to Roughgarden and Tardos's *strategic flows model*.

It is not an unreasonable assumption that edges strive to minimize latency to the destination. Edges model autonomous systems, and, intuitively, delivering packets efficiently is these entities' *raison d'être*. Presumably, there are also flows originating at each edge (the flows of its customers, not appearing in this simple example and model), and the edge wants to please them, keep them, and expand them. Besides, if the latency is unsatisfactory, flows will be routed elsewhere, this may be reflected in the network topology as it evolves at a different time scale, and the autonomous system will be out of business.<sup>1</sup>

Coming back to Braess's paradox in the strategic edges model, let us call  $y$  the fraction of the unit flow that the source edge directs to the edge  $(s, a)$  and  $z$  the amount edge  $(s, a)$  directs to edge  $(a, b)$ . It is then trivial to optimize edge  $(s, a)$ 's choice of  $z$  as a function of  $y$  (the cost to be minimized is  $y - z + z \cdot (1 - y + z)$ ), and, given that, the source edge's choice of  $y$ . The result of this calculation is the Nash equilibrium of the new game: The price of anarchy descends from  $\frac{4}{3}$  to  $\frac{31}{30}$ !

This seems very exciting. Could it be that this more realistic model yields a much better price of anarchy than the original one? Our result here is devastating:

<sup>1</sup> At this point, of course, our model contains no reason why an edge may want to stay in business; all in good time.



**Fig. 1.** Braess's network, for which  $\text{POA} = 4/3$  in the strategic flows model but  $\text{POA} = 31/30$  in the strategic edges model.

**Theorem:** *The price of anarchy of the strategic edges model is  $\Omega(n^{\frac{1}{60}})$ , where  $n$  is the number of nodes in the network.*

The proof is by a recursive construction; we do not know any satisfactory upper bound. 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 1) that pure Nash equilibria in this game always exist. One important technical reason why the situation is so far worse than the Roughgarden/Tardos model is this: In Roughgarden-Tardos, the 33% price of anarchy kicks in only when the traffic from source to destination is at its peak of 1. In contrast, the much more modest 3% deviation in our new model persists at low traffic rates; it is this feature that enables our recursive construction.

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*. Note that, 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 1, that they can lead to different equilibria).

There is one ray of hope in the strategic edges 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 strategic edges model is its dependence on non-local information; perhaps by addressing this problem (by pursuing departure (2) above from the Roughgarden-Tardos model) good performance can be restored. (As a matter of fact, our result is more general: The price of anarchy is one even in series parallel networks (Theorem 2).)

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 strategic edges model does not really require non-local information, in that what is needed by each edge 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 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.<sup>2</sup> 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, 4, 12] and Chapter 22 in [7]. However, in these previous works prices are set by the edges, but *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 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 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, either there are two edge-disjoint flow-carrying paths from the node to the destination, or a direct flow-carrying edge to the destination. In the absence of this, things fall apart. Importantly, however, the bad examples in the strategic edges game 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 constant, and, perhaps not surprisingly, 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 edges are missing in this model. 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 loss of flow to the competition. This is a subtle argument based on a sophisticated analysis of electrical networks.

In Section 2 we define our models, in Section 3 we study the strategic edges 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 questions left open by our work.

## 2 The Models

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

### The Latency Minimization 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  $\Delta_k$ , 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.

<sup>2</sup> Economic 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, 6], which were apparently created in the absence of any non-local information and solely by the propagation of price-like information.

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,

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

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 strategic 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  $S = \{S^{e'}\}$  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:

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

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.

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  $t$ , and the resistance of edge  $e$  is  $a_e$ .

**Definition 1.** We say that a given instance of the routing problem is *monopoly free* if in every optimal flow, any node that routes positive flow is either directly connected to the destination via an edge carrying positive flow, or has at least two 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.

**Theorem 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 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 1.** *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 2.** *For both the fixed ratio model and the fully rational model, the  $POA = 1$  for series parallel networks.*

*Proof.* 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  $\alpha, \beta$ , we consider the two possible types of compositions. The claim holds trivially in the case that we compose  $\alpha$  and  $\beta$  in series, via associating the sink  $t_\alpha$  with the source  $s_\beta$ , and letting the new source and new sink be  $s_\alpha, t_\beta$ , respectively.

For the case where we compose  $\alpha$  and  $\beta$  in parallel by associating  $s_\alpha$  with  $s_\beta$ , and  $t_\alpha$  with  $t_\beta$ , 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  $\alpha$ , or  $\beta$ . 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  $\alpha$ , or  $\beta$  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.

**Theorem 3.** *For arbitrary, even monopoly-free networks, the worst-case  $POS \geq \Omega(n^{\log_3(51/50)})$  for linear latency functions.*

Our proof of Theorem 3 is by construction, and relies on recursively embedding the network given in Figure 3, which we refer to as the *beetle* network. We begin by characterizing the unique equilibrium of the beetle.



**Fig. 2.** The *beetle* network. For any traffic rate,  $POA = 51/50$ .

**Lemma 1.** *The POS of the beetle network is  $51/50$ , for any positive traffic rate entering the network.*

*Proof.* 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  $\frac{r^2}{3}$ .

*Remark 1.* The POS 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:* To complete our proof, we exploit the independence of the POS 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 1, 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  $\ell(x) = \left(\frac{51}{50 \cdot 3}\right)^k x$ . 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 replacing all latency 0 edges with edges of latency  $l(x) = \epsilon x$  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 non-decreasing function of the amount of traffic—we first consider what happens if edges can only advertise constant functions.

**Proposition 2.** *If edges can only advertise constant per-unit-flow prices, then  $POS = 1$  in monopoly-free networks with arbitrary nondecreasing latencies, but worst-case  $POA \geq n - 2$  even for linear latency functions.*

*Proof.* Fix a societally optimal flow  $f$ , and assume each edge  $e = (u, v)$  advertises a constant price  $c_e$  equal to the marginal cost of its internal latency plus the minimum price advertised by its downstream neighbors. Because  $f$  is optimal, all  $v - t$  paths carrying positive flow have the same marginal latencies, and thus if  $e$  routes positive flow,  $c_e = \frac{d\ell_e(x)}{dx}(f_e) + c_{e'}$ , where  $c_{e'}$  is the common price-per-unit-flow advertised by each of the downstream neighbors of  $e$  that carry flow in  $f$ . Thus  $c_e$  is simply the marginal cost of all  $e - t$  paths that carry flow in  $f$ . To see that this is an equilibrium, note that since the latencies are nondecreasing, the marginal cost is at least the latency, and thus every edge has nonnegative utility. By the monopoly-free condition, if any edge raises its advertised price, it will receive zero flow and thus have zero utility. If an edge  $e$  lowers its price from  $c_e$  to  $c'_e$ , it will receive some amount of flow

$f'_e \geq f_e$ . Now, note that  $\frac{dU_e(f_e, c_e)}{df_e} = c_e - \frac{d\ell_v(x)}{dx}(f_e) - c_{e'}$ , and thus even at price  $c_e$ , edge  $e$  maximizes its utility by receiving  $f_e$  units of flow, and  $U_e(f_e, c_e) \geq U_e(f'_e, c_e) \geq U_e(f'_e, c'_e)$ .

For the price of anarchy, consider a network of  $n + 2$  nodes consisting of  $n$  parallel links between a source and sink with each link having latency  $\ell(x) = x$ , and traffic rate  $r = 1$ . The instance where one edge receives all the traffic, and advertises a price  $c_e = 1$  is clearly an equilibrium, and has  $\text{POA} = n$ .

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 societally optimal flows at equilibrium.

**Theorem 4.** *In monopoly free instances with edge latency functions of the form  $\ell(x) = a_e x$  with  $a_e > 0$ ,  $\text{POA} = 1$  if nodes can advertise linear, nondecreasing price schemes.*

To prove Theorem 4, 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 will be equal to the edge's true marginal cost for routing the amount of flow it receives. 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 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 proof of Proposition 4 remains an equilibrium in this more general setting in which edges may advertise nondecreasing linear prices. The arguments used in the first part of the proof of Proposition 4 hold, and yield the following proposition:

**Proposition 3.** *The following instance is a Nash equilibrium: fix a societally 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 compete over flow must advertise constant pricing schemes by demonstrating that a competing edge routing positive flow has an improving move unless it is advertising a constant pricing scheme. To show this, we'll begin by considering the simple setting in which an edge  $e$  is competing with one other edge,  $e'$ , for the traffic from a single upstream edge  $u$ . We then extend this analysis to the scenario in which  $e'$  represents a collection of edges.

We begin with an easy technical lemma.

**Lemma 2.** *Let  $p(x) = ax + b$ , with some fixed  $a > 0$  be the pricing scheme of edge  $e$ . Given a fixed non-decreasing price scheme  $p'(x) = a'x + b'$  of edge  $e'$  and total flow  $f_u$  to be routed by  $u$ , the amount of flow  $f_{(u,e)}$  routed to  $e$  is a continuous nonincreasing function of  $b$ .*

*Proof.* For ease of notation, we'll let  $f_{(u,e)}(b)$  denote the flow routed to  $e$  given  $p(x) = ax + b$ . The component of edge  $u$ 's cost from edges  $e, e'$  is

$$(f_u - f_{(u,e)}(b))p'(f_u - f_{(u,e)}(b)) + f_{(u,e)}(b)p(f_{(u,e)}(b)).$$

Because both  $p(\cdot)$ , and  $p'(\cdot)$  are nondecreasing, this function is convex with minimum at

$$f_{(u,e)}(b) = \frac{-b + b' + 2a'f_u}{2(a + a')}.$$

Given that  $a > 0$ ,  $f_{(u,e)}(b)$  is well-defined, and is a continuous decreasing function of  $b$ . In the case that the minimal solution is at some value outside of  $[0, f_u]$ , then  $f_{(u,e)}(b) = 0$ ,  $f_{(u,e)}(b) = 1$  accordingly, which preserves the continuity of  $f_{(u,e)}(b)$ , as desired.

*Remark 2.* It is precisely in the case that both the derivative of  $p$  and  $p'$  are zero that Lemma 2 breaks. At first glance, this seems troubling given our characterization of equilibria as exclusively having constant prices. Nevertheless, the *threat* of either  $e$  or  $e'$  using a nonconstant price is enough to ensure that  $u$  splits flow in the desirable fashion.

**Lemma 3.** *Let  $p(x) = ax + b$ , with some fixed  $a > 0$  be the pricing scheme of edge  $e$ . Fix a non-decreasing price scheme  $p'(x) = a'x + b'$  of competing edge  $e'$  and total flow  $f_u > 0$  to be routed by  $u$ . If  $e$  receives flow  $f_{(u,e)}$  with  $0 < f_{(u,e)} < f_u$ , then  $e$  can strictly improve its utility by advertising price  $p(x) = ax/2 + b'$  for some  $b' \geq b$ .*

*Proof.* Consider the pricing scheme  $p_0(x) = ax/2 + b_0$  where  $b_0 = b + af_{(u,e)}/2$  is chosen so that  $p_0(f_{(u,e)}) = p(f_{(u,e)})$ . We claim that the pricing scheme  $p_0(\cdot)$ , induces an increased flow to  $e$ . By assumption  $u$  chose the routing  $f_{(u,e)}, f_u - f_{(u,e)}$  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  $f_{(u,e)}^0 > f_{(u,e)}$ .

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 2, there must be some constant  $\beta$ , with  $b_0 < \beta < b_1$  such that pricing scheme  $p(x) = ax/2 + \beta$  induces  $u$  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.

The following lemma extends the above arguments to the case that edge  $e$  is competing with more than one other edge, and has several upstream edges routing flow.

**Lemma 4.** *In a pure strategy Nash equilibrium, for any edge  $e$  and an upstream edge  $u_i$  sharing an endpoint, either  $e$  receives all of  $u_i$ 's flow, or  $e$  advertises a constant pricing scheme to  $u_i$ .*

*Proof.* We start by considering the case when  $e$  is competing with edges  $e_1, \dots, e_k$ , for flow from upstream neighboring edge  $u$ . Given differentiable nondecreasing pricing schemes  $p_1(\cdot), \dots, p_k(\cdot)$  advertised by the competitors, for the purpose of evaluating whether  $e$  should deviate from a given pricing scheme, it can model the routing decisions of  $u$ , by replacing the  $e_1, \dots, e_k$  with a single competitor  $e'$  that has pricing scheme  $p'(\cdot)$ , where  $p'(x)$  is defined to be the price-per-unit-flow of optimally routing flow  $x$  among  $e_1, \dots, e_k$ . Since each  $p_i(\cdot)$  is nondecreasing, it follows that  $p'(\cdot)$  is nondecreasing and continuous. Even if each  $p_i(\cdot)$  is linear,  $p'(\cdot)$  will not necessarily be linear. Nevertheless the proof of Lemma 2 does not rely in a fundamental way on the linearity of  $p'(\cdot)$ , and can be easily modified so as to apply to arbitrary nondecreasing competitor's strategies. Given that Lemma 2 holds, the proof given above for Lemma 3 applies without modification to general continuous nondecreasing price functions  $p'(\cdot)$ . Thus our lemma holds for the case of competition with multiple nodes. In the case of multiple upstream neighbors, if  $e$  advertises a nonconstant pricing scheme to  $u_i$ , then by the above extension to Lemma 3 it can alter the price it advertises to  $u_i$  so as to receive the same amount of flow from  $u_i$  at a strictly higher price, thereby increasing its utility and thus the instance could not have been at equilibrium.

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 5.** *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.*

*Proof.* From Lemma 4, 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 societally optimal flows.

To complete the proof of Theorem 4, 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 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 6.** *Given a network  $N$  at equilibrium, the flow is an optimal flow for the network  $N'$ , which is a network with identical structure to  $N$  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  $l_{N'}(x) = (a_N + \alpha)x + b_N$  for some  $\alpha \geq 0$  where the latency of the edge in  $N$  was  $l(x) = ax + b$ .*

*Proof.* Fixing the flow and advertised prices of our equilibrium of  $N$ , we start from the sink and work our way towards the source in topological order, computing the edge latencies of  $N'$  that would induce the optimal flow to be the observed flow. To illustrate, consider an edge connected to the sink with latency to the sink in  $N$   $l_N(x) = ax + b$  receiving flow  $f$ , and advertising a pricing scheme  $p(\cdot)$ , we set  $l_{N'} = \frac{p(f)}{2f} + b$ . Since we are at equilibrium  $\frac{p(f)-b}{2f} = a$  if the node in question is competing with other edges by Lemma 4. Additionally,  $\frac{p(f)-b}{2f} \geq a$  otherwise the edge would have an improving deviation by raising its price given flow  $f$ .

Now, observe that 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 7.** 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 from  $u$  along another edge  $e'$ .

We begin with a fact adapted from circuit analysis. Consider a  $N$  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  $l(x) = (a_e + \delta)x$ , and we 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  $\delta f_e$  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 5** 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  $F$  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 7, completing our proof of Theorem 4.

*Proof of Lemma 7:* Note that Fact 5 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 Roughgarden-Tardos framework 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 (linear with a constant term, or concave)? 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.

## Acknowledgements

We thank Yaron Singer for helping formulate the models for strategic edges, 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. 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.
5. E. Koutsoupias, C. H. Papadimitriou “Worst-case Equilibria,” STACS 1999.
6. R. Maddin, T. S. Wheeler, J. D. Muhly “Tin in the ancient Near East: old questions and new finds,” *Expedition*, 1977.
7. N. Nisan, T. Roughgarden, É. Tardos and V. Vazirani (editors), *Algorithmic Game Theory*, Cambridge University Press, 2006.
8. C. H. Papadimitriou “Algorithms, Games, and the Internet,” STOC 2001.
9. T. Roughgarden “Selfish Routing,” PhD dissertation, Cornell, 2002.
10. T. Roughgarden “The price of anarchy is independent of the network topology,” Journal of Computer and System Sciences, 2003.
11. T. Roughgarden, Éva Tardos “How bad is selfish routing?” JACM 2002.
12. M. J. Smith “The marginal cost taxation of a transportation network,” Transportation Research, issue 13, pp. 237-242, 1979.