

Strategic Pricing in Next-hop Routing with Elastic Demands*

Elliot Anshelevich[†] Ameya Hate* Koushik Kar[‡]

Abstract

We consider a model of next-hop routing by self-interested agents. In this model, nodes in a graph (representing ISPs, Autonomous Systems, etc.) make pricing decisions of how much to charge for forwarding traffic from each of their upstream neighbors, and routing decisions of which downstream neighbors to forward traffic to (i.e., choosing the next hop). Traffic originates at a subset of these nodes that derive a utility when the traffic is routed to its destination node; the traffic demand is elastic and the utility derived from it can be different for different source nodes. Our next-hop routing and pricing model is in sharp contrast with the more common source routing and pricing models, in which the source of traffic determines the entire route from source to destination. For our model, we begin by showing sufficient conditions for prices to result in a Nash equilibrium, and in fact give an efficient algorithm to compute a Nash equilibrium which is as good as the centralized optimum, thus proving that the price of stability is 1. When only a single source node exists, then the price of anarchy is 1 as well, as long as some minor assumptions on player behavior is made. The above results hold for arbitrary convex pricing functions, but with the assumption that the utilities derived from getting traffic to its destination are linear. When utilities can be non-linear functions, we show that Nash equilibrium may not exist, even with simple discrete pricing models.

1 Introduction

The ubiquitous impact of the Internet on modern life is a testimony to its growth in the past two decades. One of the principal factors behind this growth has been the decentralization of control, which also allows it to be modeled naturally as a system of interacting but independent, self-interested agents. More specifically, the Internet can be viewed as a collection of ISPs or ASes (Autonomous Systems) that are interested in routing and pricing traffic to maximize their individual revenues [2, 7, 15]. Similar frameworks have also been applied to the study of relaying/routing of traffic in wireless ad-hoc networks [19, 20]. The study of large decentralized networks of self-interested agents, with regard to their efficiency, has sparked an enormous amount of interest, as the insight thus earned can be used to extract maximum utility from existing infrastructure, as well as to make good policy decisions.

We consider the interactions of self-interested agents in a network at a very abstract level, where each agent is modeled as a self-interested node in a graph. Traffic originates

*This work was supported in part by NSF grants CNS-1017932 and CCF-0914782.

[†]Department of Computer Science, Rensselaer Polytechnic Institute. {eanshel,hatea}@cs.rpi.edu

[‡]Department of Electrical, Computer & Systems Engineering, Rensselaer Polytechnic Institute.

at a subset of these nodes that derive a utility when the traffic is routed to its destination node, which may be many hops away. We consider next-hop routing, where each node on the path of the traffic individually determines which node(s) the traffic should be forwarded to (i.e., chooses the “next hop”). Nodes are allowed to *charge* their upstream neighbors for the traffic that they are asked to forward, as is typically done in contracts formed by neighboring Autonomous Systems in the Internet. Thus, a node obtains payments from its upstream neighbors for accepting their traffic for forwarding, and must in turn pay its downstream neighbors for receiving its traffic (i.e., the traffic it has accepted to forward, plus the traffic that it originates). Moreover, the nodes, being self-interested, will always choose to send traffic to the downstream neighbors with the cheapest price. As remarked in [3], for example, there is an interplay between setting the price to forward traffic, and choosing the routing policies of a node, since both decisions can change the profit/cost of a node (Autonomous System). In this paper, we consider both decisions to be under the control of each node, and study the properties of the equilibrium solutions of this game, which we believe captures the fundamental aspects of next-hop routing by self-interested agents.

Our next-hop routing and pricing model is in sharp contrast with the more common “source routing” and pricing models (see e.g., [6, 10, 14, 17]). In the latter models, the source node of the traffic determines its entire route from the source to the destination. Next-hop models provide a better representation of the routing protocols and pricing practices in the current Internet, as well as those that are likely to dominate the future multi-hop wireless networks [15]. In the Internet, traffic flow and service pricing negotiations occur at the inter-domain level, between an ISP and its neighboring ISPs (i.e., ISPs with which it shares a POP (Point-Of-Presence) and has a customer/provider/peering relationship) [11]. Inter-domain routing follows the BGP protocol, where hops at the AS level are determined one at a time [16]. Even though BGP determines this next hop based on information on the entire AS-level path, the benefits of making it strictly next-hop has been argued recently [18]. Our next-hop routing and pricing model also closely captures the Path Vector Contract Switching framework proposed for the future Internet [21], where neighboring ISPs establish contracts (on the amount of flow and its pricing) towards forwarding traffic for a specific destination. Source routing requires knowledge of the entire path at the source node, and this practical limitation has restricted the use of source routing in the Internet, while next-hop routing involves decision making by agents that is much more local and distributed.

In addition to its focus on next-hop routing, our model differs from most existing models in several other aspects as well. We assume that links in our network have fixed capacities, which represents the constraints associated with routing somewhat better than having linear cost functions that depend on how much traffic is being routed. An important feature of our model is the existence of multiple sources that have elastic demands with non-uniform utilities. See the Related Work section on further contrast with existing models.

Model Summary We now give a brief outline of our model. A more detailed description is given in Section 2.

We are given a directed acyclic graph $G = (V, E)$ containing a special sink node t , and edge capacities c_e . As commonly done when analyzing competition in networks [4, 6, 15], we assume that all edges of this DAG, except the ones that are incident on the sink node, have a special non-monopolistic property. For our model, this property essentially ensures that enough capacity exists that no node could charge an infinitely high price for forwarding traffic, and yet have other nodes pay this price because they have no alternative.

The players of our game are all vertices of G except the sink node. An edge $e = (u, v) \in E$ with capacity c_e denotes that player u has the capacity to send a flow of size c_e to player v . Additionally, every player v has an associated *source utility* λ_v . This means that if a player v sends f_v amount of its own flow (flow originating from vertex v) to the sink t , then the player will obtain a utility of $\lambda_v f_v$. Thus, the player demands are *elastic*, since each player can choose an amount of traffic to send in order to maximize its utility. We consider an extension of this model in Section 5 where source utilities are allowed to be non-linear.

Players choose prices on their incoming edges. For every edge $e = (u, v)$ player v chooses a price p_e such that if u sends a flow of size f_e on edge e then u pays an amount $p_e f_e$ to v . Players route flow on outgoing edges such that this minimizes their cost, but are obligated to forward all flow that they receive. Finally, the utility of a player is the total amount of money it receives from upstream players and the utility obtained by sending its own flow minus the amount of money paid to the downstream nodes.

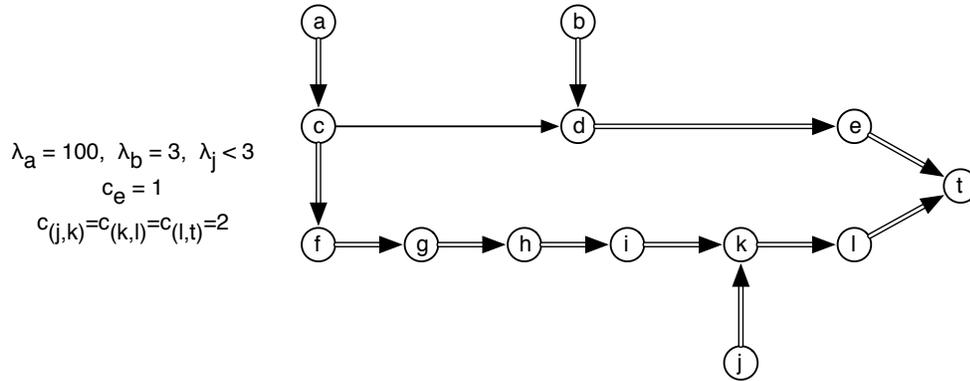


Figure 1: An example illustrating our model, and the “long-range” effects that node utilities can have on edge prices.

Example. To illustrate some key consequences of our model, consider the example in Figure 1. The utility values of all nodes except the ones mentioned are 0. Edge (j, k) , (k, l) , (l, t) have capacity 2 and all other edges have capacity 1. Not all the edges in the graph are pictured: for every edge (u, v) shown in the figure that does not satisfy the non-monopolistic property, there also exist (non-pictured) edges (u, w) , (w, t) of capacities 1 such that w has a high source utility (say 1000). In any optimal solution edge (u, w) will not have any flow on it whereas (w, t) will be saturated. Now it is not difficult to see that any optimal solution will consist of the flow indicated in Figure 1, where the double walled edges are saturated with flow, except edge (j, k) has a flow of size 1 whereas its capacity is 2.

A more careful analysis shows that any Nash equilibrium strategy that yields this optimal flow will have the following edge prices: $p_{(j,k)} = p_{(k,l)} = \lambda_j$, $p_{(b,d)} = 3$. Also $p_{(c,d)} \geq p_{(c,f)} \geq p_{(f,g)} \geq p_{(g,h)} \geq p_{(h,i)} \geq p_{(i,k)} \geq \lambda_j$. This means that the price of edge (c, d) depends on the source utility value of node j , and thus there are “long-range” interactions between source utilities and edge prices.

Results Our main goal involves understanding the properties of stable solutions in this pricing game: specifically we focus on pure Nash equilibrium. In Section 3 we give an efficient

algorithm that constructs a Nash equilibrium strategy that is as good as the optimal solution. In other words, we show that the price of stability is 1, and thus it is always possible to implement traffic pricing that maximizes social welfare. This holds for an arbitrary number of sources with elastic demands of heterogeneous value. We also show in Section 4 that in case of a single source, under some reasonable behavioral assumptions for the players, the price of anarchy is 1, and in fact player prices at equilibrium are unique.

Until this point, the source utilities λ_v and the allowed price functions p_e were considered to be linear. In Section 5, we instead consider the more general case where the source utility can be an arbitrary concave function $\Lambda_v(f_v)$, and the prices can be arbitrary convex functions $\Pi_e(f_e)$. We show that the above results still hold if arbitrary convex prices are allowed, and thus allowing non-linear prices does not impact the quality of equilibrium solutions. On the other hand, if source utilities can be non-linear functions, then we show that pure Nash equilibrium may no longer exist, even for discrete pricing models.

Related work Selfish routing and pricing games have been studied in many contexts (see e.g., [1–4, 9, 10, 12, 13, 17], and the many references in [8, 14]). As mentioned before, most of the work in this area has been done using source routing, where a source of traffic chooses the entire path that the traffic takes. In one such model, [5, 6] consider a game where there are two sets of players. One set of players own edges of the network (edges have finite capacities) and sell capacity to players of the second set. The second set of players obtain utility for routing a unit amount of flow from its source to destination and hence buys capacity on edges along the route if it is profitable to do so. The essential difference between source routing models like the one in [5, 6] and ours is that in our model when a node changes its price on an incoming edge, only its upstream neighbors are immediately affected and may change their routes; whereas in the source routing model, any change in prices by the first set of players is seen by all players of the second set, and can immediately result in a globally different routing. Thus next hop routing operates much more on local knowledge [18].

Perhaps the most relevant paper to ours is by Papadimitriou and Valiant [15], in which they define a next hop routing model where players are edges of a network, and argue for the importance of next-hop routing models. The strategy of players in their game is similar to our model: players charge neighbors for processing and forwarding flow. Unlike our model where edges have capacities, their model has edges with (linear) latencies, and deals only with a single source with a fixed demand. Although their results can most likely be extended to networks with multiple sources, the crucial complications in our model arise from the fact that we consider sources with elastic demands and non-uniform utilities. This leads to complex interaction of prices as illustrated by the example in Figure 1, and prevents us from using the methods from [15] for analysis. To further illustrate the differences between our model and the one from [15], notice that in our model, the price of anarchy for single source games is 1, while in [15], the price of anarchy can be large.

Another recent next-hop routing model is discussed by Xi and Yeh [20]. As in [15], their model only considers a single source with a fixed amount of traffic demand. The links in [20] have latency functions instead of capacities. For these reasons, just as with [15], equilibrium solutions in [20] are very different from the ones in our model, and have a very different structure. In essence, the complexity in our game arises from the interplay between different source utilities and edge capacities, while in [15] and [20] it arises due to the presence of latency functions. Finally, [3] considers a somewhat general routing game that can also include next-hop routing as a special case, but uses a very different pricing mechanism from

the one considered here.

2 Model

We are given a directed acyclic graph $G = (V, E)$ containing a special sink node t , and edge capacities c_e . To arrive at a meaningful model of price competition, we assume that all edges of this DAG, except the ones that are incident on the sink node, have a special non-monopolistic property. This property says that for any edge $e = (u, v)$, even if this edge is removed from the graph, the total capacity of outgoing edges of node u will be greater than the total capacity of its incoming edges. The rationale behind this property will be explained once the model has been illustrated in more detail.

The set of players in our game consists of all vertices of G , except the sink node. An edge $e = (u, v) \in E$ with capacity c_e denotes that player u has the capacity to send a flow of size c_e to player v . Additionally, every player v has an associated *source utility* λ_v . This means that if a player v sends f_v amount of its own flow (flow originating from vertex v) to the sink t , then the player will obtain a utility of $\lambda_v \cdot f_v$. We consider an extension of this model in Section 5 where source utilities are allowed to be non-linear.

Player Strategy The core behavior of our model is that a vertex can charge a price for the flow being sent through it. For example, since edge $e = (u, v)$ is incident on v , node v can set a price p_e on this edge. If vertex u sends a flow f_e on this edge then vertex u has to pay an amount equalling $f_e \cdot p_e$ to vertex v . (We will consider a model where the price per packet changes with the amount of flow being sent on the edge in Section 5.)

So apart from gaining utility by sending their own flow, vertices also gain utility for receiving flows and lose utility by paying the next-hop vertices that receive their flow. Also, observe that for the flow to reach the sink, it is essential that intermediate nodes forward the flow reliably. Hence every player is required to forward all incoming flow (alternatively, we can think of there being a very large penalty for accepting payment for incoming flow that the player has no intention of forwarding). The forwarding of all flow is always possible since the non-monopolistic property ensures that for every node the total outgoing capacity is always greater than the total incoming capacity.

Imagine a situation where in order to route all its incoming flow, node u always has to forward some flow to vertex v . Since it is obligatory for u to forward all its incoming flow, v can charge an exorbitantly high price on edge (u, v) , and u would have to pay it. In other words, v can act as a monopoly. Similarly to [15], existence of such a structure in G may lead to no meaningful equilibrium and the non-monopolistic property for edge (u, v) obviates precisely this situation. Since the sink node is not a player, it does not set prices for edges that are incident on it (we assume that price is fixed at 0, and that wlog these edges are always saturated); neither do these edges need to satisfy the non-monopolistic property, as long as the outgoing capacity of each node is at least as large as the incoming capacity.

In order to provide a more formal definition of the player strategies and the resulting flow, we define the following terminology. Let E_v^{in} and E_v^{out} be the set of incoming and outgoing edges for node v respectively. The vector of flows on the incoming edges of vertex v is denoted by f_v^{in} and on the outgoing edges is denoted by f_v^{out} . Similarly, the vector of prices on the incoming edges is denoted by p_v^{in} and on the outgoing edges is denoted by p_v^{out} . Let f_v be the amount of own flow (flow originating at vertex v) sent by vertex v to sink t .

We assume that a vertex always sends or forwards flow by choosing the outgoing edges that have the lowest price and have free capacity. Also, if there exists an outgoing edge with free capacity and has $p_e < \lambda_v$ then the node will always send its own flow on such an edge, and will never send its own flow on edges with $p_e > \lambda_v$. To make this precise, we define a notion of valid flows, which are flows where every vertex forwards flow in order to maximize its utility. Specifically, given the prices p_v^{out} and flows f_v^{in} , we define the set of valid resulting flows f_v^{out} to be $\mathcal{F}_v(f_v^{in}, p_v^{out})$, which are all flows satisfying the following conditions:

- $\forall e \in E_v^{out} : f_e \leq c_e$ (usual capacity constraint);
 $f_v = \sum_{e \in E_v^{in}} f_e - \sum_{e \in E_v^{out}} f_e \geq 0$ (usual flow conservation)
- $\forall e \in E_v^{out} : f_e > 0$ only if for every $e' \in E_v^{out} \setminus e$ with $p_{e'} < p_e$, e' is saturated (send on cheapest edges first),
- $\forall e \in E_v^{out} : p_e < \lambda_v$ implies that e is saturated (send own flow if profitable), and $p_e > \lambda_v$ and $f_e > 0$ imply that $f_v = 0$ (don't send own flow if unprofitable).

Any way of forwarding flow to maximize v 's utility obeys these conditions. The last condition holds since if v is sending its own flow, but $f_e > 0$ for some edge $e \in E_v^{out}$ with $p_e > \lambda_v$, then v could re-distribute its flow so that it is sending its own flow on edge e , and then improve its utility by sending less of its own flow.

When all prices in p_v^{out} are distinct from each other and λ_v , then $\mathcal{F}_v(f_v^{in}, p_v^{out})$ is a unique flow resulting from forwarding flow from all from all incoming edges, beginning with the edges of least cost, and then sending its own flow on remaining edges with free capacity and cost $p_e < \lambda_v$. Now consider instead a situation where two or more outgoing edges have the same price. So long as they have the same price, the utility of player v is not affected by the choice of edge on which it sends a flow. Similarly, when there exists an outgoing edge with free capacity and $p_e = \lambda_v$, the vertex is indifferent towards the choice of sending its own flow on the edge. In this model we assume that both these tie-breaking choices are left up to the players and are part of their strategy. More formally, since each valid flow in $\mathcal{F}_v(f_v^{in}, p_v^{out})$ corresponds to a tie-breaking rule selected by player v , we associate these tie-breaking rules with a flow generation function γ_v which, given the incoming flows and outgoing prices, produces an outgoing flow. The set of these flow generation functions is denoted by Γ_v :

$$\Gamma_v = \{\gamma_v | \forall f_v^{in}, p_v^{out} : \gamma_v(f_v^{in}, p_v^{out}) = f_v^{out} \in \mathcal{F}_v(f_v^{in}, p_v^{out})\}$$

In other words, Γ_v contains all functions that generate only valid out-flows. Hence the strategy set of each player v is $\mathcal{R}_+^{|E_v^{in}|} \times \Gamma_v$, and a strategy of the player is given by the tuple $\{p_v^{in}, \gamma_v\}$ where $p_v^{in} \in \mathcal{R}_+^{|E_v^{in}|}$ and $\gamma_v \in \Gamma_v$. We denote the collective strategy of all players by $\{P, \gamma\}$.

Outcome Each flow generating function γ_v needs incoming flow and prices on outgoing edges in order to compute the resulting flow. Given a strategy $\{P, \gamma\}$, the prices are already known. The algorithm to produce the resulting flow is then simply: Iterate over $v \in V$ in topologically sorted order (recall that our graph is a DAG), and set $f_v^{out} = \gamma_v(f_v^{in}, p_v^{out})$. We denote the resulting flow by $f(P, \gamma)$: this is the outcome of the strategy $\{P, \gamma\}$.

Utility and Best Response Given the output flow $f(P, \gamma)$, the utility of player v is given by the following expression:

$$utility_v(P, \gamma) = \sum_{e \in E_v^{in}} f_e \cdot p_e - \sum_{e \in E_v^{out}} f_e \cdot p_e + \lambda_v \cdot f_v$$

Consider a player v who is computing its best response to a strategy $\{P, \gamma\}$. Notice that by changing its prices p_v^{in} , the resulting flow f_v^{in} may become completely different from $f(P, \gamma)$. If this were not the case, then v could always raise its incoming price, knowing that this would increase its utility since the flow would remain the same. In essence, players in this game anticipate changes in flow that result from price changes, but myopically assume that the prices of all other nodes remain the same when computing their own best response. Such behavior is reasonable in ISP routing settings, for example, since price setting takes place on a much slower time scale than routing.

3 Uniform Nash Equilibrium and Price of Stability

We first prove useful sufficient conditions for a strategy to be a Nash equilibrium.

Theorem 1 *If flow $f(P, \gamma)$ and prices P satisfy the following conditions, then strategy $\{P, \gamma\}$ is a Nash equilibrium:*

- (a) *For every node u , the price on all edges of E_u^{out} , except edge (u, t) if it exists, is the same. Let this price be denoted by y_u .*
- (b) *If $f_u > 0$ then $y_u = \lambda_u$; if $f_u = 0$ then $y_u \geq \lambda_u$.*
- (c) *For $\forall v \neq t$, if edge (u, v) has a positive flow on it, then $y_u \geq y_v$.*
- (d) *For $\forall v \neq t$, if edge $e = (u, v)$ is unsaturated ($f_e < c_e$), then $y_u \leq y_v$.*

We first prove the following lemmas, which state that if a player deviation includes increasing its price on some incoming edge, then no flow is sent on this edge after such a deviation.

Lemma 1 *Let $\{P, \gamma\}$ satisfy the conditions of Theorem 1. Consider a unilateral deviation by player v where v changes its strategy from (p_v^{in}, γ_v) to $(p_v^{in'}, \gamma_v')$ such that for some edge $e = (u, v)$, $p_e' > p_e$. Let $\{P', \gamma'\}$ denote the resulting collective strategy of all players. Then the flow on edge e in $f(P', \gamma')$ equals 0.*

Proof. We will let f denote the flow $f(P, \gamma)$, and f' denote the flow $f(P', \gamma')$. Consider first the case where node u is not sending its own flow in $f(P', \gamma')$, i.e., $f'_u = 0$. Since $p_e' > p_e$, condition (a) of Theorem 1 says that $\forall e' \in E_u^{out} \setminus e$, $p_{e'} < p_e'$. Also because of the non-monopolistic property we know that total capacity on the set of edges $E_u^{out} \setminus e$ is at least as much as the capacity on edges E_u^{in} . Therefore the flow generating algorithm will route all incoming flow to u on edges of $E_u^{out} \setminus e$ in $f(P', \gamma')$.

Now suppose instead that $f'_u > 0$. By condition (b), we know that $y_u \geq \lambda_u$, and thus that $p_e' > \lambda_u$. Thus no valid flow $f(P', \gamma')$ will send any flow on edge e (since node u loses utility by sending flow on edge e). Hence in $f(P', \gamma')$ there would be no flow on edge e . ■

We are now ready to prove Theorem 1.

Proof. Let $\{P, \gamma\}$ satisfy the conditions of Theorem 1, and suppose to the contrary that for some player v , there exists a strictly improving unilateral deviation from $\{P, \gamma\}$, where v changes its strategy from (p_v^{in}, γ_v) to $(p_v'^{in}, \gamma_v')$. Denote by $\{P', \gamma'\}$ the collective player strategy after v 's deviation, and let $f = f(P, \gamma)$, $f' = f(P', \gamma')$.

Let the set of incoming edges of v be partitioned into the following:

- $E_{inc} \subseteq E_v^{in}$ such that $\forall e \in E_{inc} : p_e' > p_e$
- $E_{dec} \subseteq E_v^{in}$ such that $\forall e \in E_{dec} : p_e' < p_e$
- $E_{free} \subseteq E_v^{in}$ such that $\forall e \in E_{free} : p_e' = p_e$ and $f_e < c_e$
- $E_{sat} \subseteq E_v^{in}$ such that $\forall e \in E_{sat} : p_e' = p_e$ and $f_e = c_e$

Note that $E_{inc} \cup E_{dec} \cup E_{free} \cup E_{sat} = E_v^{in}$. Now, utility of node v in strategy (P, γ) can be stated as follows:

$$utility_v(P, \gamma) = \sum_{e \in E_{inc} \cup E_{dec} \cup E_{free} \cup E_{sat}} f_e \cdot (p_e - y_v) + f_v \cdot (\lambda_v - y_v) + c_{(v,t)} \cdot y_v. \quad (1)$$

This is because for every unit of flow that node v is forwarding or sending, it must pay y_v to an outgoing edge, except for $c_{(v,t)}$ units of flow, which can be sent to t with price 0 (we will use the convention that $c_{(v,t)} = 0$ is edge (v, t) does not exist).

Notice that each term $f_e(p_e - y_v)$ is nonnegative due to condition (c), and the term $f_v(\lambda_v - y_v)$ equals zero due to condition (b). Similarly, the utility of v in strategy $\{P', \gamma'\}$ is the same as in Equation 1 with f' replacing f and p_e' replacing p_e . Note that, since we assumed that every edge (v, t) is saturated (since node v can always forward or send its own flow to fill up edge (v, t) to capacity without decreasing its utility), then the term $c_{(v,t)} \cdot y_v$ is present both before and after the deviation. We now analyze the effect of price changes from p_e to p_e' on the sum in Equation 1 for each of the four subsets of E_v^{in} .

- E_{inc} : According to Lemma 1, all edges with price increase will not have any flow in $f(P', \gamma')$. Hence the term $\sum_{e \in E_{inc}} f_e' \cdot (p_e' - y_v)$ equals $0 \leq \sum_{e \in E_{inc}} f_e \cdot (p_e - y_v)$.
- E_{dec} : Since $p_e' < p_e$ for edges in E_{dec} , then the only way that it is possible for $f_e' \cdot (p_e' - y_v)$ be greater than $f_e \cdot (p_e - y_v)$ is if $f_e' > f_e$. This implies that edge e is not saturated in f , i.e, $f_e < c_e$. According to condition (d), edges in E_{dec} that have free capacity in $f(P, \gamma)$ have $p_e \leq y_v$. Thus $(p_e' - y_v) < (p_e - y_v) \leq 0$, and so it is not possible that $f_e' \cdot (p_e' - y_v) > f_e \cdot (p_e - y_v)$ because each term $f_e \cdot (p_e - y_v)$ is nonnegative. Therefore, $\sum_{e \in E_{dec}} f_e' \cdot (p_e' - y_v) \leq \sum_{e \in E_{dec}} f_e \cdot (p_e - y_v)$.
- E_{free} : For an edge $e \in E_{free}$, we know by condition (d) that $p_e' = p_e \leq y_v$. Since each term $f_e(p_e - y_v)$ is nonnegative, this implies that $0 = f_e(p_e - y_v) \geq f_e'(p_e' - y_v)$. Thus $\sum_{e \in E_{free}} f_e' \cdot (p_e' - y_v) \leq \sum_{e \in E_{free}} f_e \cdot (p_e - y_v)$.
- E_{sat} : Since these edges are saturated in $f(P, \gamma)$, then $f_e' \leq f_e$. Also from condition (c) we know that $p_e' = p_e \geq y_v$. It follows that $\sum_{e \in E_{sat}} f_e' \cdot (p_e' - y_v) \leq \sum_{e \in E_{sat}} f_e \cdot (p_e - y_v)$.
- f_v : As mentioned above, we know because of condition (b) that $f_v \cdot (\lambda_v - y_v) = 0$. The only way that $f_v' \cdot (\lambda_v - y_v) > 0$ is if $\lambda_v - y_v > 0$, which contradicts condition (b). Thus, this term equals 0 both before and after v 's deviation.

Hence the change in node v 's strategy can only result in a drop in utility of node v . Thus the collective strategy $\{P, \gamma\}$ is a Nash equilibrium. ■

Theorem 1 gives sufficient conditions for a strategy to be a Nash equilibrium. We will call such strategies *uniform*, since all the outgoing prices are the same for every node in such a solution. As we will show below, good uniform Nash equilibria always exist, and can be efficiently computed.

Definition 1 *Uniform Nash equilibrium: Any Nash equilibrium strategy that satisfies the conditions of Theorem 1 is a **uniform** Nash equilibrium.*

3.1 Computing a Nash Equilibrium As Good as the Optimal Solution

By an optimal solution to this game, we will mean one in which the sum of the utilities of all players is maximized. Since the price paid by players to each other cancels out in the sum, optimal solutions are ones in which $\sum_v \lambda_v f_v$ is maximized. We will call a flow f^* *socially optimal* if $\sum_v \lambda_v f_v^*$ is maximum over all flows that obey capacity constraints and where f_v^* flow originates at node v , with all flow ending at the sink t (it is easy to see that, without loss of generality, all edges incident on t are saturated). Clearly, $\sum_v \lambda_v f_v^*$ is the social welfare in an optimal solution, since if all prices are set to 0, and γ is such that f^* is the resulting flow, then this results in social welfare of $\sum_v \lambda_v f_v^*$. We now prove that flow f^* can also be achieved by a Nash equilibrium solution, i.e., that the price of stability of this game is 1.

Theorem 2 *Given a socially optimal flow f^* , there exists a collective strategy $\{P, \gamma\}$ such that $f(P, \gamma) = f^*$ and $\{P, \gamma\}$ is a uniform Nash Equilibrium. In other words, the price of stability is 1.*

Proof. We will prove this theorem by constructing a strategy $\{P, \gamma\}$ such that $f(P, \gamma) = f^*$, and then showing that P and $f(P, \gamma)$ satisfy the conditions of Theorem 1.

We first set the prices P . Let V_S be the set of nodes that send their own flow in f^* (i.e., $f_v^* > 0$), and V_{NS} be the set of nodes that do not send their own flow in f^* . Since our goal is to form prices that satisfy condition (a) of Theorem 1, we will set all prices on outgoing edges in P to be the same for each node: denote these prices as y_v for node v . For every vertex $v \in V_S$, set $y_v = \lambda_v$, thus satisfying condition (b) as well. We now need to assign prices to the vertices of V_{NS} .

Let G^r be the residual graph of f^* , i.e., it contains an edge (v, u) for every edge $e = (u, v)$ with $f_e^* > 0$ (call these *backward* edges), and an edge (u, v) for every edge $e = (u, v)$ with $f_e^* < c_e$ (call these *forward* edges). The following lemma is a simple consequence of the fact that f^* is socially optimal.

Lemma 2 *If $u \in V_S$ is reachable from v in G^r , then $\lambda_u \geq \lambda_v$.*

Proof. To see that $\lambda_u \geq \lambda_v$, consider the path from v to u in G^r . This path is an augmenting path of the flow f^* . Since $u \in V_S$, we know that $f_u^* > 0$. We can thus use a flow path from u to t together with the augmenting path from v to u in order to augment the flow f^* by decreasing the flow f_u^* by $\varepsilon > 0$, and increasing the flow f_v^* by ε , while still satisfying all capacity constraints. If $\lambda_u < \lambda_v$, this would result in a flow with higher social welfare than f^* , giving us a contradiction. Thus, $\lambda_u \geq \lambda_v$. This is the only place where we use the fact that f^* is a socially optimal flow. ■

To assign prices to nodes of V_{NS} , we now proceed to obtain upper and lower bounds on the y_v values of nodes $v \in V_{NS}$.

Lemma 3 *In every uniform Nash equilibrium strategy $\{P^*, \gamma^*\}$ with $f(P^*, \gamma^*) = f^*$ and prices y_v , we have that for all $u \in V_S$ and $v \in V$ with u being reachable from v in G^r , it must be that $\lambda_u \geq y_v$.*

Proof. We can show this inductively on the distance from v to u in G^r . Since this is a uniform Nash equilibrium and $u \in V_S$, then by condition (b) we know that $\lambda_u = y_u$. Now consider a node v such that there is an edge (v, w) in G^r with u being reachable from w in G^r by a path of length at most k . By the inductive hypothesis, $\lambda_u \geq y_w$. If (v, w) is a forward edge, then by condition (d), we know that $y_v \leq y_w \leq \lambda_u$, as desired. If instead (w, v) is a backward edge, then by condition (c), we similarly know that $y_v \leq y_w \leq \lambda_u$. ■

Let R_v^h be the set of nodes $u \in V_S$ such that u is reachable from v in G^r . Then, using Lemma 3 we know that $h_v = \min_{u \in R_v^h} \lambda_u$ is an upper bound for y_v in any uniform Nash equilibrium that results in flow f^* .

Lemma 4 *In every uniform Nash equilibrium strategy $\{P^*, \gamma^*\}$ with $f(P^*, \gamma^*) = f^*$ and prices y_v , we have that for all $u \in V_S$ and $v \in V$ with v being reachable from u in G^r , it must be that $\lambda_u \leq y_v$.*

Proof. The proof is the same as the proof of Lemma 3, with conditions (c) and (d) reversed. ■

Let R_v^l be the set of nodes $u \in V_S$ such that v is reachable from u in G^r . Then, using Lemma 4 we know that $l_v = \max_{u \in R_v^l} \lambda_u$ is a lower bound for y_v in any uniform Nash equilibrium that results in flow f^* .

We now set the y_v values for $v \in V_{NS}$ in the following manner: For every vertex v , $y_v = h_v$.

Now that we have fully defined the prices P , we can choose flow generation functions $\gamma_v \in \Gamma_v$ for all v such that $f(P, \gamma) = f^*$. To do this, simply set γ_v to be such that $\gamma_v(f_v^{in*}, \vec{y}_v) = f_v^{out*}$, where f_v^{in*} and f_v^{out*} are the flows into and out of v in flow f^* , and \vec{y}_v is a vector of prices which equals y_v on every edge of E_v^{out} . To show that this is a valid flow generation function, we just need to prove that $f_v^{out*} \in \mathcal{F}_v(f_v^{in*}, \vec{y}_v)$, so we check all the necessary conditions listed in Section 2. Clearly, the capacity and conservation conditions are satisfied since f^* satisfies them. All edges have the same outgoing price, so v 's utility remains the same no matter which edges its flow is forwarded on. Finally, notice that for all v , we have that $y_v \geq \lambda_v$, since if $v \in V_S$ then $y_v = \lambda_v$, and if $v \in V_{NS}$, then this also holds due to Lemma 2, which implies that $h_v \geq \lambda_v$. Thus, with these prices it is never strictly profitable for a node v to send its own flow, and $p_e = y_v$ can only be strictly greater than λ_v when $v \in V_{NS}$, i.e., when $f_v^* = 0$, as desired. Thus, we have fully defined a strategy $\{P, \gamma\}$ such that $f(P, \gamma) = f^*$. All that is left to prove is that $\{P, \gamma\}$ is a uniform Nash equilibrium:

Condition (a). We know that condition (a) is satisfied since all edges of E_v^{out} were given the same price y_v .

Condition (b). Condition (b) is satisfied by the prices generated by our algorithm, since we set prices of all $v \in V_S$ to λ_v , and by Lemma 2, $y_v = h_v \geq \lambda_v$ for all $v \in V_{NS}$.

Condition (c). Consider an edge $e = (u, v)$ with $f_e^* > 0$. Then, edge (v, u) is a backward edge in G^r . If $u \in V_S$ and $v \in V_{NS}$, then Lemma 3 immediately implies that $y_u = \lambda_u \geq h_v = y_v$.

Since G^r has an edge (v, u) , then all nodes reachable from u in G^r are also reachable from v . Thus by Lemma 3, if $u, v \in V_{NS}$, then $h_u \geq h_v$, and so $y_u \geq y_v$, as desired.

If $u \in V_{NS}$ and $v \in V_S$, then let w be the node of V_S reachable from u in G^r such that $h_u = \lambda_w$. Since w is reachable from u , then w is also reachable from v in G^r . Therefore, by Lemma 2, we know that $y_u = \lambda_w \geq \lambda_v = y_v$.

Finally, if both u and v are in V_S , then Lemma 2 immediately implies that $y_u = \lambda_u \geq \lambda_v = y_v$, as desired. Hence condition (c) is satisfied.

Condition (d). Consider an edge $e = (u, v)$ with $f_e^* < c_e$. Then, edge (u, v) is a forward edge in G^r . If $v \in V_S$ and $u \in V_{NS}$, then Lemma 3 immediately implies that $y_u = h_u \leq \lambda_v = y_v$.

Since G^r has an edge (u, v) , then all nodes reachable from v in G^r are also reachable from u . Thus by Lemma 3, if $u, v \in V_{NS}$, then $h_u \leq h_v$, and so $y_u \leq y_v$, as desired.

If $v \in V_{NS}$ and $u \in V_S$, then let w be the node of V_S reachable from v in G^r such that $h_v = \lambda_w$. Since w is reachable from v , then w is also reachable from u in G^r . Therefore, by Lemma 2, we know that $y_v = \lambda_w \geq \lambda_u = y_u$.

Finally, if both u and v are in V_S , then Lemma 2 immediately implies that $y_u = \lambda_u \leq \lambda_v = y_v$, as desired. Hence condition (d) is satisfied.

Hence we have shown that a collective strategy $\{P, \gamma\}$ can be computed in polynomial time such that it is a uniform Nash equilibrium and $f(P, \gamma) = f^*$. ■

4 Reasonable Assumptions and Price of Anarchy

Consider a game where all players with non-zero λ_v value are not neighbors of the sink. Now consider a strategy for this game where every vertex charges a very high price (say bigger than the highest λ_v value). Given this pricing strategy, no vertex will send its own flow and still every vertex will have no incentive to deviate, i.e., the strategy will be in Nash equilibrium. This is because no vertex would unilaterally reduce the prices of its incoming edges, given that they will have to pay a large amount to forward any flow sent to them. Nodes that have edges incident to the sink will not change their prices as there is no hope of obtaining any flow and hence, any profit. In this Nash equilibrium strategy the total utility of players is 0 and hence the price of anarchy is unbounded. These “bad equilibria” cannot be eliminated even after introducing pairwise deviations.

In order to eliminate such unrealistic solutions from consideration, work dealing with similar scenarios made some reasonable assumptions about player behavior. For example, [15] assumes that if a player does not receive any flow on its incoming edge, then she never charge an unnecessarily large price for this edge. In this section, we make the same assumption on the players’ pricing strategy:

Property 1 *If a vertex v does not receive any flow on edge (u, v) , then it sets $p_{(u,v)}$ to be the price of the cheapest unsaturated outgoing edge of v , if one exists.*

We call pricing strategies that satisfy this property *reasonable*. This property simply says that given an edge (u, v) that has no flow on it, node v will charge the minimum price such that potential flow on this edge will not result in loss of utility for v . Below we show that,

at least for single-source games (i.e., games where only one node has a non-zero λ value), this additional assumption on player behavior causes all equilibria to become as good as the optimum solution.

Theorem 3 *For a single source game where players form reasonable pricing strategies, the price of anarchy is 1.*

Proof. Let the source node be s . Any node that is not reachable from s can never receive a flow and hence the strategy of such a node will not have consequence on the outcome of the game. Hence we assume w.l.o.g. that no such node exists. Note that if there exists an edge from the source to sink then it will always be saturated. Hence we also assume w.l.o.g. that no such edge exists.

Due to the non-monopolistic property of the network, the capacity of the minimum cut, and hence the size of the largest $s - t$ flow, will be the sum of the capacities of the outgoing edges of s . So in the socially optimal solution, all outgoing edges of s are saturated.

In order to prove the theorem we make use of the following notation. Let N_s be the set of neighboring vertices of the source s . If c_e is the capacity of the edge e then the size of the socially optimal flow is $\sum_{v \in N_s} c_{(s,v)}$. Let N_t be the set of nodes adjacent to the sink t .

Lemma 5 *In any Nash equilibrium for a single source game where players form reasonable pricing strategies, for all $(u, v) \in E$ such that $u \neq s$, $p_{(u,v)} = 0$.*

Proof. We prove this lemma by induction. Consider a topological sort order of the network, with the first element in the order being node s followed by the nodes belonging to N_s . The last element will be the sink node t whereas the penultimate element will be a node v such that $v \in N_t$. According to the topological ordering, all outgoing edges of v will be incident to t and hence all outgoing edges of v will have price $p_e = 0$ in any Nash equilibrium strategy, say $\{P, \gamma\}$. Let this be the base case for induction.

Now consider an arbitrary node v_i in the ordering $s, v_1, v_2, \dots, v_k, t$. We assume that the prices of all outgoing edges of $v_{i+1}, v_{i+2}, \dots, v_k$ in $\{P, \gamma\}$ are 0. Suppose to the contrary that in strategy $\{P, \gamma\}$, the price of edge (v_i, v_j) , where $(i < j \leq k)$, is non-zero, and assume wlog that (v_i, v_j) is such an edge with largest price. Then consider the following two cases:

- Case 1. Under strategy $\{P, \gamma\}$, there does not exist any flow on the edge (v_i, v_j) . In this case, Property 1 tells us that node v_j would set the price of (v_i, v_j) to 0. This contradicts our assumption that (v_i, v_j) has non-zero price.
- Case 2. Under strategy $\{P, \gamma\}$, there does exist a non-zero flow on edge $e = (v_i, v_j)$. This implies that node v_i is not sending its own flow, by the last property of the flow generation function (i.e., it would be unprofitable for v_i to send its own flow since its λ value equals 0, and it sends positive flow on edge e with price greater than 0). From the non-monopolistic property of the network, and since v_i is not sending its own flow, we know that there exists another node $v_{j'}$ such that the edge $e' = (v_i, v_{j'})$ is not saturated. Also from the inductive hypothesis we know that the price of all outgoing edges of $v_{j'}$ is 0. Let $f = f(P, \gamma)$.

In this case, consider a deviation by node $v_{j'}$ in which it changes the price on edge e' to $p_e - \epsilon$ while keeping the rest of its strategy the same, where $p_e > 0$ is the price of edge e , and

$$\epsilon < \frac{p_e \min\{f_e, c_{e'} - f_{e'}\}}{c_{e'}}.$$

Call this new collective strategy $\{P', \gamma\}$. We claim that the utility of $v_{j'}$ strictly increases after this deviation. Consider the flow generation of $f(P', \gamma)$. All outgoing prices and incoming flow vectors are the same as in $\{P, \gamma\}$ until node v_i , so the incoming flow into node v_i is the same in both $f(P, \gamma)$ and $f(P', \gamma)$.

First consider the change in flow on e' . By setting its price to be smaller than p_e , $v_{j'}$ has guaranteed that it receives at least $\min\{f_e, c_{e'} - f_{e'}\}$ extra flow on edge e' , since flow is sent on edges with smallest price first. Thus the utility of $v_{j'}$ due to edge e' changes from at most $p_e \cdot f_{e'}$ (we assumed that p_e is the largest price of all edges leaving v_i), to at least $(p_e - \epsilon)(f_{e'} + \min\{f_e, c_{e'} - f_{e'}\})$. The utility due to this edge strictly increases, since $\epsilon(f_{e'} + \min\{f_e, c_{e'} - f_{e'}\}) \leq \epsilon c_{e'} < p_e \min\{f_e, c_{e'} - f_{e'}\}$.

Now consider the utility of $v_{j'}$ due to other incoming edges. All edges from nodes earlier than v_i in the topological ordering are sending the same amount of flow to $v_{j'}$ as before the deviation. All edges from nodes after v_i in the topological ordering have price 0, and so flow on them does not change the utility of node $v_{j'}$. Combined with the fact that all outgoing edges from node $v_{j'}$ have price 0, we know that the utility of $v_{j'}$ strictly increases after the deviation. This contradicts our assumption that strategy $\{P, \gamma\}$ is a Nash equilibrium.

Hence by induction we have shown that there does not exist any Nash equilibrium strategy where $p_{(u,v)} > 0$ for any $u \neq s$. This proves the lemma. ■

Using the result of Lemma 5, we will now show that in any Nash equilibrium strategy, the outgoing edges of s will be saturated. This will imply that every Nash equilibrium is socially optimum, since for a single source s , the social welfare of a solution with flow f is simply $\lambda_s \cdot f_s$.

Suppose to the contrary that in some Nash equilibrium strategy, the edge (s, v) is not saturated. In this case node v can price edge (s, v) just below λ_s so that (s, v) is saturated by the flow algorithm. This will always be a beneficial deviation for v since Lemma 5 tells us that the price on all outgoing edges of v is 0. This is clearly a contradiction and hence no such edge exists. ■

The above proof implies that prices at Nash equilibrium are unique. It also fairly easy to see that a strategy (P, γ) where all neighbors v of s price the edges (s, v) at λ_s is a Nash equilibrium satisfying Property 1, and it results in the socially optimum flow given the appropriate choice of functions γ .

5 Non-linear Utility and Price Functions

In previous sections we analyzed the case where the utility of sending one unit of own flow (will also be referred to as ‘per packet’) was constant for the player. We will now study the case where the utility is a concave function of the amount of flow sent. This mirrors the fact that sending more flow usually has diminishing returns for the player. We denote this utility function as $\Lambda_v(f_v)$ where f_v is the total amount of own flow sent by node v and Λ_v is continuously differentiable, concave, and non-decreasing. Additionally, denote the derivative of Λ_v by λ_v : in the old model this was a constant, but now it is a non-increasing function.

Similarly when players receive flow on an incoming edge, the processing cost for each unit of flow generally increases with the total amount of flow. Hence we look at the case

where the price charged by each player for incoming flow is a convex function. We denote this by $\Pi_e(f_e)$ where f_e is the flow on edge e and Π_e is a continuously differentiable convex non-decreasing function. Let π_e be the derivative of Π_e : in the old model this was called p_e and was a constant; now it is a non-decreasing function.

This more general model is formally defined as follows. Note that our assumptions of the graph being acyclic and non-monopolistic still hold.

Strategy: Once the price functions Π_e for every outgoing edge E_v^{out} and quantity of incoming flow f_v^{in} are set, player v is assumed to forward the flow and send its own flow in order to maximize its utility. In the earlier model where price per packet π_e was always constant, a node would start with sending flow on an edge with the cheapest price per packet and saturate it before moving on to other edges. In this model however, the price per packet on an edge may increase with the amount of flow present on the edge. Hence, if we ignore capacities, a node will always send flow on outgoing edges such that $\pi_e(f_e)$ is equal on all outgoing edges that have flow on them. It is easy to show that such a choice will minimize the node's cost.

Similarly, a node will only send its own flow if the utility of sending a packet is at least as much as the cost of sending it, otherwise sending own flow would result in lowering the player's utility. Hence a player would send its own flow until the utility of the marginal packet $\lambda_v(f_v)$ is equal to its price.

To state the above observations formally, consider a node v that has incoming edges E_v^{in} and outgoing edges E_v^{out} . Given the set of price functions Π_v^{out} and incoming flows f_v^{in} , we define the set of valid resulting flows f_v^{out} to be $\mathcal{F}_v(f_v^{in}, \Pi_v^{out})$, which are all flows satisfying the following conditions:

1. $\forall e \in E_v^{out} : f_e \leq c_e$ and $\sum_{e \in E_v^{in}} f_e + f_v = \sum_{e \in E_v^{out}} f_e$
2. $\forall e \in E_v^{out}$, if $f_e > 0$ then $\forall e' \in E_v^{out} \setminus e : \pi_e(f_e) \leq \pi_{e'}(f_{e'})$ or $f_{e'} = c_{e'}$
3. $\forall e \in E_v^{out} : \pi_e(f_e) < \lambda_v(f_v)$ implies that e is saturated, and if $f_v > 0$ then $\forall e \in E_v^{out}, f_e > 0 : \pi_e(f_e) \leq \lambda_v(f_v)$

The first condition just says that the generated flow should obey flow conservation and edge capacities. The second condition says that flow is first sent to the edges with smallest marginal price. Finally, the third condition says that traffic is sent by node v if the marginal utility of sending a packet is less than the cost of sending it, and is not sent by v if the marginal utility of sending a packet is more than the cost of sending it. It is easy to show (through first order conditions for maximizing utility) that satisfying the above conditions maximizes utility for players. Note that these subsume the conditions in Section 2, replacing p_e with $\pi_e(f_e)$ and λ_v with $\lambda_v(f_v)$: hence this model is a strict generalization.

As in the previous model, each of these valid flows correspond to a tie-breaking rule selected by player v . We associate these tie-breaking rules with a flow generation function γ_v , which, given the incoming flows and outgoing prices, produces an outgoing flow. The set of these flow generation functions is denoted by Γ_v , defined in exactly the same manner as in Section 2. Hence the strategy of a player v is given by the tuple $\{\Pi_v^{in}, \gamma_v\}$. We denote the collective strategy of all players by $\{\Pi, \gamma\}$.

Outcome and Utility: Also like the previous model, the flow outcome can be easily determined if the nodes are considered in the topological sort order and is denoted by $f(\Pi, \gamma)$.

Given the resulting flow $f(\Pi, \gamma)$, utility of player v is given by the following expression:

$$utility_v(\Pi, \gamma) = \sum_{e \in E_v^{in}} \Pi_e(f_e) - \sum_{e \in E_v^{out}} \Pi_e(f_e) + \Lambda_v(f_v)$$

5.1 Non-Linear Prices

In this section we assume that utility functions Λ_v are linear, and show that all our results for linear price functions also hold for arbitrary convex price functions, thus showing that allowing players to set non-linear prices does not make the system any worse. To do this, we prove analogues of Theorems 1 and 3. By proving that the conditions from Theorem 1 imply that a solution is a Nash equilibrium, even when changing your strategy to an arbitrary price function is allowed, we immediately get the consequence that the price of stability is 1, since we already showed how to create an optimal solution satisfying these conditions in Theorem 2.

Theorem 4 *For instances with linear utility functions and non-decreasing, convex price functions: if flow $f(\Pi, \gamma)$ and pricing strategy Π satisfy the following conditions, then strategy $\{\Pi, \gamma\}$ is a Nash equilibrium:*

- (a) *For every node u , the price function on all edges of E_u^{out} , except edge (u, t) if it exists, is identical and linear, i.e., $\pi_{(u,v)}(x) = p_e$ for all x and constant p_e . Let the constant price per packet p_e be denoted by y_u .*
- (b) *If $f_u > 0$ then $y_u = \lambda_u$; if $f_u = 0$ then $y_u \geq \lambda_u$.*
- (c) *For $\forall v \neq t$, if edge (u, v) has a positive flow on it, then $y_u \geq y_v$.*
- (d) *For $\forall v \neq t$, if edge $e = (u, v)$ is unsaturated ($f_e < c_e$), then $y_u \leq y_v$.*

Proof. Let $\{\Pi, \gamma\}$ satisfy the conditions of Theorem 4, and suppose to the contrary that for some player v , there exists a strictly improving unilateral deviation from $\{\Pi, \gamma\}$, where v changes its strategy from (Π_v^{in}, γ_v) to $(\Pi_v^{in'}, \gamma_v')$. Denote by $\{\Pi', \gamma'\}$ the collective player strategy after v 's deviation, and let $f = f(\Pi, \gamma)$, $f' = f(\Pi', \gamma')$.

For any edge $e = (u, v) \in E_v^{in}$ which has constant price per packet p_e before deviation, let Π'_e be the price function after deviation. We define the following terms:

- Let f_e^l be the flow of f'_e that “costs at most p_e ”. Formally, let $L = \{x | \pi'_e(x) \leq p_e\}$ be the set of flow amounts that cost less than p_e per packet, and let \bar{L} be the supremum of L . Then, $f_e^l = \min(\bar{L}, f'_e)$.
- Let $f_e^h = \max(0, f'_e - f_e^l)$ be the rest of the flow of f'_e .

Note that $f_e^l + f_e^h = f'_e$. Since for all e , all functions $\pi_e(x) = p_e$ are constant by condition (a), then just as in the proof of Theorem 1, the utility of node v in strategy $\{\Pi, \gamma\}$ can be stated as follows:

$$utility_v(\Pi, \gamma) = \sum_{e \in E_v^{in}} f_e \cdot (p_e - y_v) + f_v \cdot (\lambda_v - y_v) + c_{(v,t)} \cdot y_v \quad (2)$$

Notice that each term $f_e(p_e - y_v)$ is nonnegative due to condition (c), and the term $f_v(\lambda_v - y_v)$ equals zero due to condition (b).

Similarly, the utility of v in strategy $\{\Pi', \gamma'\}$ is the same as in Equation 2 except for the fact that the edge prices can now be arbitrary non-decreasing convex functions. It can be stated as follows:

$$utility_v(\Pi', \gamma') = \sum_{e \in E_v^{in}} [\Pi'_e(f'_e) - f'_e \cdot y_v] + f'_v \cdot (\lambda_v - y_v) + c_{(v,t)} \cdot y_v \quad (3)$$

Note that, since we assumed that every edge (v, t) is saturated (since node v can always forward or send its own flow to fill up edge (v, t) to capacity without decreasing its utility), then the term $c_{(v,t)} \cdot y_v$ is present both before and after the deviation.

We now analyze the effect of price function changes from $\Pi_e(x) = p_e \cdot x$ to $\Pi'_e(\cdot)$ on the sum in Equations 2 and 3 for each edge in E_v^{in} . For any edge $e \in E_v^{in}$, the first two terms of Equation 3 can be represented as follows:

$$\Pi'_e(f'_e) - f'_e \cdot y_v = \int_0^{f'_e} \pi'_e(x) dx - f'_e \cdot y_v + \int_{f'_e}^{f_e^h} \pi'_e(x) dx - f_e^h \cdot y_v \quad (4)$$

We first claim that for $\forall e \in E_v^{in}$, $f_e^h = 0$. Suppose to the contrary that there exists an edge $e = (u, v) \in E_v^{in}$ such that $f_e^h > 0$. Condition (b) of Theorem 4 says that $\lambda_u \leq p_e < \pi'_e(f'_e)$. Then from condition (3) of definition of flow generating function we know that node u does not send its own flow in f' , i.e., node u only forwards incoming flow in f' . Since $f'_u = 0$, the non-monopolistic property for edge e requires that either $f'_e = 0$ or $\exists e' \in E_u^{out} \setminus e$ such that e' is not saturated in f' . Also, condition (a) of the theorem says that $\forall e' \in E_u^{out} \setminus e$, $\pi'_{e'} = p_e$. Now, since we have assumed $f_e^h > 0$, it has to be that $\exists e' \in E_u^{out} \setminus e$ such that $\pi'_{e'}(f'_{e'}) = p_e < \pi'_e(f'_e)$ where edge e' is not saturated. But this violates condition (2) of definition of the flow generating function.

Hence $f_e^h = 0$ and the term $\int_{f'_e}^{f_e^h} \pi'_e(x) dx - f_e^h \cdot y_v = 0$.

In order to analyze the term $\int_0^{f'_e} \pi'_e(x) dx - f'_e \cdot y_v$, we look at two cases.

- e is saturated in f'_e : In this case $f'_e \leq f_e = c_e$. Also, $\pi'_e(x) \leq p_e$ for $0 \leq x \leq f'_e$ (by definition of f'_e). Combined with the fact that $f_e \cdot (p_e - y_v) \geq 0$, it implies that $\int_0^{f'_e} \pi'_e(x) dx - f'_e \cdot y_v \leq f'_e(p_e - y_v) \leq f_e \cdot (p_e - y_v)$.
- e is unsaturated in f'_e : In this case condition (d) of the theorem says that $\pi'_e(x) \leq p_e \leq y_v$ for $0 \leq x \leq f'_e$. Since the term $f_e(p_e - y_v)$ is nonnegative, this implies that $0 = f_e(p_e - y_v) \geq \int_0^{f'_e} \pi'_e(x) dx - f'_e \cdot y_v$.

Hence for every edge $e \in E_v^{in}$, $f_e \cdot (p_e - y_v) \geq \Pi'_e(f'_e) - f'_e \cdot y_v$.

Also we know because of condition (b) that $f_v \cdot (\lambda_v - y_v) = 0$. The only way that $f'_v \cdot (\lambda_v - y_v) > 0$ is if $\lambda_v - y_v > 0$, which contradicts condition (b). Thus, this term equals 0 both before and after v 's deviation.

Hence the change in node v 's strategy can only result in a drop in utility of node v . Thus the collective strategy $\{\Pi, \gamma\}$ is a Nash equilibrium. ■

In section 4 we showed that when prices have to be linear, the price of anarchy is 1 for networks with a single source under the mild assumption that players do not set large prices without a good reason (Property 1). In this section, we show that the same result holds if prices are allowed to be convex non-decreasing functions. Note that utilities are still

linear. Since edge prices are allowed to be functions, we call pricing strategies that satisfy the following property as reasonable:

Property 2 *If a vertex v does not receive any flow on edge (u, v) , then it sets $\Pi_{(u,v)}(x) = p_{(u,v)}x$ where $p_{(u,v)}$ is the cheapest marginal price of all unsaturated outgoing edges of v , if one exists.*

We now prove the theorem.

Theorem 5 *If node utilities are linear and edge prices are allowed to be convex non-decreasing functions then, for a single source game where players form reasonable pricing strategies, the price of anarchy is 1.*

Proof. Let the source node be s . Any node that is not reachable from s can never receive a flow and hence the strategy of such a node will not have consequence on the outcome of the game. Hence we assume w.l.o.g. that no such node exists. Note that if there exists an edge from the source to sink then it will always be saturated. Hence we also assume w.l.o.g. that no such edge exists.

Due to the non-monopolistic property of the network, the capacity of the minimum cut, and hence the size of the largest $s - t$ flow, will be the sum of the capacities of the outgoing edges of s . So in the socially optimal solution, all outgoing edges of s are saturated.

In order to prove the theorem we make use of the following notation. Let N_s be the set of neighboring vertices of the source s . If c_e is the capacity of the edge e then size of the socially optimal solution will be given by $\sum_{v \in N_s} c_{(s,v)}$. Let N_t be the set of nodes adjacent to the sink t .

Lemma 6 *If node utilities are linear and edge prices are allowed to be convex non-decreasing functions then, in any Nash equilibrium for a single source game where players form reasonable pricing strategies, for all $(u, v) \in E$ such that $u \neq s$, $\Pi_{(u,v)} = 0$.*

Proof. We prove this lemma by induction. Consider a topological sort order of the network, with the first element in the order being node s followed by the nodes belonging to N_s . The last element will be the sink node t whereas the penultimate element will be a node v such that $v \in N_t$. According to the topological ordering, all outgoing edges of v will be incident to t and hence all outgoing edges of v will have price $\Pi_e = 0$ in any Nash equilibrium strategy, say $\{\Pi, \gamma\}$. Let this be the base case for induction.

Now consider an arbitrary node v_i in the topological ordering $s, v_1, v_2, \dots, v_k, t$. We assume that the prices of all outgoing edges of $v_{i+1}, v_{i+2}, \dots, v_k$ in $\{\Pi, \gamma\}$ are 0. We will first show that all outgoing edges of v_i satisfy the following property:

Property 3 *If f_e is the flow on edge e then $\Pi_e(x) = p_e x$ for $0 \leq x \leq f_e$ for some constant $p_e \geq 0$.*

Suppose to the contrary that in strategy $\{\Pi, \gamma\}$, the price of edge (v_i, v_j) , where $(i < j \leq k)$, does not satisfy Property 3. Let edge (v_i, v_j) have a concave non-decreasing price function $\Pi_e(x)$, and a flow of size f_e (possibly zero).

Then consider the following two cases:

- Case 1. Under strategy $\{\Pi, \gamma\}$, there does not exist any flow on the edge (v_i, v_j) , i.e., $f_e = 0$. In this case, Property 2 tells us that node v_j would set the price of (v_i, v_j) to 0. This contradicts our assumption that (v_i, v_j) does not satisfy Property 3.

Case 2. Under strategy $\{\Pi, \gamma\}$, there does exist a non-zero flow on edge $e = (v_i, v_j)$, i.e., $f_e > 0$. Now consider a deviation to a new price function for edge e : $\Pi'_e(x) = p'_e x$ for $0 \leq x \leq f_e$ and $\Pi'_e(x) = \Pi_e(x)$ for $x > f_e$ such that $p'_e < \pi_e(f_e)$ and $f_e \cdot p'_e > \int_0^{f_e} \pi_e(x) dx$. Note that such a constant p'_e exists since $\pi_e(x)$ is a monotonic non-decreasing function and $\pi_e(x) < \pi_e(f_e)$ for some $x \leq f_e$.

Call this new collective strategy $\{\Pi', \gamma\}$. We claim that the utility of v_j strictly increases after this deviation. Consider the flow generation of $f(\Pi', \gamma)$. All outgoing prices and incoming flow vectors are the same as in $\{\Pi, \gamma\}$ until node v_i , so the incoming flow into node v_i is the same in both $f(\Pi, \gamma)$ and $f(\Pi', \gamma)$.

First consider the change in the flow on e . We know that node v_i is only forwarding flow and not sending any of its own (since it is sending flow on an edge with non-zero price). Also, since edge e received a flow of size f_e in $f(\Pi, \gamma)$, condition (2.) of definition of flow generating function says that for all $e' \in E_{v_i}^{out} \setminus e$, either $\pi_{e'}(f_{e'}) \geq \pi_e(f_e) > p'_e$ or e' was saturated. Hence e would receive a flow of size at least f_e in $f(\Pi', \gamma)$. For node v_j , the utility due to this edge strictly increases, since $f_e \cdot p'_e > \int_0^{f_e} \pi_e(x) dx$.

Now consider the utility of v_j due to other incoming edges. All edges from nodes earlier than v_i in the topological ordering are sending the same amount of flow to v_j as before the deviation. All edges from nodes after v_i in the topological ordering have outgoing price 0, and so flow on them does not change the utility of node v_j . Combined with the fact that all outgoing edges from node v_j have price 0 and that v_j is not sending its own flow, we know that the utility of v_j strictly increases after the deviation. This contradicts our assumption the strategy $\{\Pi, \gamma\}$ is a Nash equilibrium. Therefore in every Nash equilibrium strategy, (v_i, v_j) must satisfy Property 3.

Hence all outgoing edges of v_i will satisfy Property 3. We will use this observation to show that all outgoing edges of v_i in fact have price functions $\Pi_e = 0$.

Suppose to the contrary that in strategy $\{\Pi, \gamma\}$, the price of edge (v_i, v_j) , where $(i < j \leq k)$, is non-zero. We have already shown that this edge will satisfy Property 3. Assume w.l.o.g. that (v_i, v_j) is such an edge with the largest p_e as defined by Property 3. The argument that this is not possible is essentially the same as in Lemma 5. Hence by induction we have shown that there does not exist any Nash equilibrium strategy where $\Pi_{(u,v)} \neq 0$ for any $u \neq s$. This proves the lemma. ■

Using the result of Lemma 6, we will now show that in any Nash equilibrium strategy, the outgoing edges of s will be saturated. This will imply that every Nash equilibrium is socially optimum, since for a single source s , the social welfare of a solution with flow f is simply $\lambda_s \cdot f_s$.

Suppose to the contrary that in some Nash equilibrium strategy, the edge (s, v) is not saturated. In this case node v can price edge (s, v) just below λ_s so that (s, v) is saturated by the flow algorithm. This will always be a beneficial deviation for v since Lemma 6 tells us that the price on all outgoing edges of v is 0. This is clearly a contradiction and hence no such edge exists. ■

5.2 Non-linear Utilities Can Cause Non-Existence of Equilibrium

In this section we show that, if utility functions Λ_v can be non-linear, then there may not exist a pure Nash equilibrium.

Theorem 6 *If the player utilities Λ_v are concave non-decreasing functions, then pure Nash equilibrium may not exist.*

Proof. Consider the example in Figure 2.

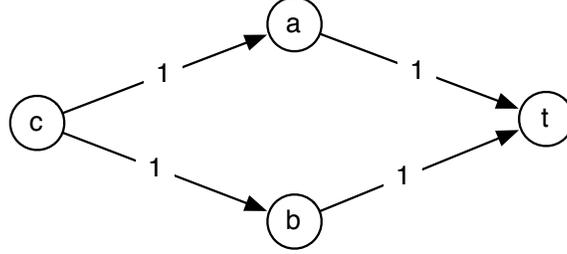


Figure 2: $\Lambda_c(f) = -9f^2 + 37f$ for $f \leq 2$. $\Lambda_a = \Lambda_b = 0$. All edges have capacity 1.

Lemma 7 *For any Nash equilibrium strategy for example in Figure 2, there exists another equilibrium strategy where the price functions for edges (c, a) and (c, b) are linear.*

Proof. Note that in all Nash Equilibrium solutions, there will always be some flow on edge (c, a) . If there is no flow on edge (c, a) then node a 's utility is 0, whereas by lowering the price $\Pi_{(c,a)}$ sufficiently, a can ensure that it receives some flow on the edge, which provides positive utility. The same argument symmetrically holds for edge (c, b) .

We first show that in Nash equilibrium strategy $\{\Pi, \gamma\}$ for this example, all edges must satisfy Property 3.

If this is not the case then let $\{\Pi, \gamma\}$ be a Nash equilibrium strategy where we assume wlog that edge (c, a) has a concave non-decreasing price function $\Pi_e(x)$, and a positive flow of size f_e and (c, a) does not satisfy Property 3. The utility of node a can be stated as :

$$utility_a = \Pi_e(f_e) = \int_0^{f_e} \pi_e(x) dx \quad (5)$$

Now consider a deviation to a new price function for edge (c, a) : $\Pi'_e(x) = p'_e x$ for $0 \leq x \leq f_e$ and $\Pi'_e(x) = \Pi_e(x)$ for $x > f_e$ such that $p'_e < \pi_e(f_e)$ and $f_e \cdot p'_e > \int_0^{f_e} \pi_e(x) dx$. Note that such a constant p'_e exists since $\pi_e(x)$ is a monotonic non-decreasing function and $\pi_e(x) < \pi_e(f_e)$ for some $x \leq f_e$ (since (c, a) does not satisfy Property 3).

Note that in $\{\Pi', \gamma\}$, the price function of edge (c, b) is the same whereas the marginal price for edge (c, a) has decreased. By condition (3.) of definition of the flow generating function we know that:

$$\lambda_c(f_c) \geq \pi_e(f_e) > p'_e \quad (6)$$

This means that node c would send a flow of size at least f_c after deviation. If edge (c, b) was saturated in the flow $f(\Pi, \gamma)$ then edge (c, a) would receive a flow a size at least f_e after deviation. If (c, b) was not saturated in flow $f(\Pi, \gamma)$ then from condition (2.) of definition of flow generating function we know that $\pi_{(c,b)}(f_{(c,b)}) \geq \pi_e(f_e) > p'_e$. Therefore, in this case too, edge (c, a) would receive a flow a size at least f_e after deviation.

And since $f_e \cdot p'_e > \int_0^{f_e} \pi_e(x) dx$, utility of node a as given by Equation 5, will increase. This contradicts the assumption that strategy $\{\Pi, \gamma\}$ was in Nash equilibrium. Therefore in all Nash equilibrium pricing strategies, edges will satisfy Property 3.

Now consider a new pricing strategy $\bar{\Pi}$ constructed from a Nash equilibrium strategy $\{\Pi, \gamma\}$ as follows: for every edge e with flow f_e on it and $\pi_e(f_e) = p_e$ let $\bar{\Pi}_e(f) = f \cdot p_e$ for $f > 0$. We have already shown that in $f(\Pi, \gamma)$, $\forall e \in E$: for $0 < f \leq f_e$, $\Pi_e(f) = f \cdot p_e$.

Note that the only difference between the pricing strategies Π and $\bar{\Pi}$ is that $\forall e \in E$, $\pi_e(f) \geq p_e$ for $f > f_e$ where as $\bar{\pi}_e(f) = p_e$ for $f > f_e$. Therefore there exists a tie-breaking strategy $\bar{\gamma}$ such that $f(\bar{\Pi}, \bar{\gamma}) = f(\Pi, \gamma)$. Since the utility of players a and b depend only on the flow received from node c , their utility is identical in both strategies.

Now, if there exists a profitable deviation for say node a in strategy $\{\bar{\Pi}, \bar{\gamma}\}$ then the same deviation will be profitable in strategy $\{\Pi, \gamma\}$. For example since the utility of node a depends only on flow from node c , it could reduce price of edge (c, a) to obtain more flow. But this deviation would also be valid in $\{\Pi, \gamma\}$. Since we assumed $\{\Pi, \gamma\}$ to be a Nash equilibrium strategy, no such deviation exists. Hence $\{\bar{\Pi}, \bar{\gamma}\}$ is also a Nash equilibrium strategy, thus proving statement of the Lemma. ■

Since we have shown in Lemma 7 that for every Nash equilibrium strategy, there exists one with linear prices, we need only show that such Nash equilibria do not exist.

Suppose such an equilibrium exists for this example: let p_a be the price per packet of edge (c, a) , and p_b be the price per packet of edge (c, b) in this equilibrium: the prices of the other two edges must be 0 since they are edges to the sink.

Note that in all Nash Equilibrium solutions, there will always be some flow on edge (c, a) . If there is no flow on edge (c, a) then node a 's utility is 0, whereas by lowering the price $p_{(c,a)}$ sufficiently, a can ensure that it receives some flow on the edge, which provides positive utility. The same argument symmetrically holds for edge (c, b) .

Lemma 8 *In any Nash equilibrium strategy for the example in Figure 2, edges (c, a) and (c, b) will have the same price.*

Proof. Suppose to the contrary that $p_a > p_b$. Then, by increasing its price by $\epsilon < p_a - p_b$, node b will still receive the same amount of flow on edge (c, b) according to the flow generating algorithm, since (c, a) has positive flow on it at price p_a , and so (c, b) will be saturated at both price p_b and $p_b + \epsilon$. This deviation increases the utility of node b , and thus this cannot be a Nash equilibrium. ■

Lemma 9 *In any Nash equilibrium strategy for the example in Figure 2, edges (c, a) and (c, b) will be saturated.*

Proof. By Lemma 8 we know that in any Nash equilibrium, $p_a = p_b$. Now consider a Nash equilibrium strategy with resulting flow f , where edge (c, a) is not saturated. We know that $f_{(c,a)} > 0$ and $f_{(c,b)} > 0$. Consider a deviation by node a where it sets its incoming price to $p_a - \epsilon$ for $\epsilon < p_a \min\{1 - f_{(c,a)}, f_{(c,b)}\}$. Now edge (c, a) becomes the preferred edge and hence receives more flow (flow that was being routed to (c, b)). Specifically, edge (c, a) must now receive at least $\min\{1 - f_{(c,a)}, f_{(c,b)}\}$ more flow, and thus node a now receives at least $(p_a - \epsilon) \min\{1, f_{(c,a)} + f_{(c,b)}\}$ utility. This is strictly larger than the utility a received before the deviation, since $p_a f_{(c,a)} + \epsilon \min\{1, f_{(c,a)} + f_{(c,b)}\} < p_a \min\{1, f_{(c,a)} + f_{(c,b)}\}$. This contradicts our assumption that the strategy is in Nash equilibrium, and the same argument symmetrically holds for node b . ■

Now consider the flow amount f_c sent by node c in a Nash equilibrium in Example 2. We know that if $p = p_a = p_b$ is the price per packet set on both edges out of c , then node c will send a flow of size at most f_c such that $\lambda_c(f_c) = p$ (since any more flow would be unprofitable). Hence, we can define the inverse function such that $\lambda_c^{-1}(p) = f_c$. Specifically, here $\Lambda_c(f) = -9f^2 + 37f$ and hence $\lambda_c(f) = -18f + 37$ and $\lambda_c^{-1}(p) = (37 - p)/18$. Lemma 9 says that in any Nash equilibrium both edges (c, a) and (c, b) have to be saturated. For that to be true, node c has to send 2 units of its own flow. This means that the price per packet on the outgoing edges of c cannot be more than 1, and thus the utility of node a in any Nash equilibrium can be no more than 1.

Now consider a deviation where node a increases its incoming price p_a to 10. In this situation, a will still receive a flow of size $\lambda_c^{-1}(10) - 1 = 0.5$ giving it a utility of $10 * 0.5 = 5$. This means that for all strategies that satisfy the conditions laid down by Lemma 9, there exists an improving deviation for player a . Hence there does not exist any Nash equilibrium solution for this example. ■

Discrete Prices The non-existence of Nash equilibrium in the previous example relied heavily on the fact that the prices could be changed by an infinitesimal amount. We now show that even if the prices were to be discrete, the existence of Nash equilibrium is not guaranteed. Let us assume that players are allowed to pick discrete prices and say the smallest unit is δ (e.g., we can say that players can pick integer prices, and thus $\delta = 1$).

Using this Lemma we show that for the example in Figure 2, there does not exist a Nash equilibrium strategy.

Theorem 7 *If utility functions are non-linear then there exists an instance that does not have a pure Nash equilibrium strategy, even with discrete price structures.*

Proof. Consider the example in Figure 2 and let δ be the smallest unit of discrete price that players are allowed to pick. For this proof, we will assume that δ is much smaller than 1: we can handle the general case with simple re-scaling of the function Λ_c and the edge capacities. Let $e = (c, a)$ and $e' = (c, b)$. We can still assume that prices are linear, so let p_e be the price per packet on edge e , and $p_{e'}$ be the price per packet on edge e' . We first prove the following lemma.

Lemma 10 *If $f_e > 0$ and $f_{e'} > 0$ in a Nash equilibrium solution, then $|p_e - p_{e'}| \leq \delta$.*

Proof. Suppose to the contrary that $p_e < p_{e'} - \delta$. According to the flow generating algorithm, e will be saturated since a higher-priced edge, e' , has flow on it. Now by increasing the price of edge e by δ , node a can receive the same amount of flow but at a higher price thus making it a profitable deviation. This is a contradiction since the strategy was supposed to be in Nash equilibrium. ■

We first claim that there does not exist a Nash equilibrium strategy in which node c emits a flow of size less than or equal to 1. In order to prove this claim, suppose to the contrary that ≤ 1 flow is being emitted. Then, without loss of generality, let edge (c, a) be the edge that receives less flow (can be equal) out of the two. This implies that $p_e \geq p_{e'}$. Recall that the function $\lambda_c^{-1}(p) = (37 - p)/18$ gives the total amount of flow emitted by c when the price of both of its outgoing edges is p . The flow on edge e is at most $\lambda_c^{-1}(p_e)/2$ when $p_e = p_{e'}$. When $p_e = p_{e'} + \delta$ instead, then it is still true that the flow on edge e is at most $\lambda_c^{-1}(p_e)/2$. To see this, notice that this is trivially true if $\lambda_c^{-1}(p_e) \geq 2$, since the

capacity of edge e is 1. Otherwise, the flow on edge e is at most $\lambda_c^{-1}(p_e) - 1 < \lambda_c^{-1}(p_e)/2$, since we can assume that $\lambda_c^{-1}(p_e) < 2$. Thus we can give an upper bound on the utility of node a as follows:

$$utility_a \leq \frac{p_e \lambda_c^{-1}(p_e)}{2} = \frac{p_e \cdot (37 - p_e)}{2 \cdot 18}$$

Given this upper bound for utility of a , it can be shown through simple differentiation that the maximum occurs at $p_e = 18.5$. Note that since we are assuming that $f_c \leq 1$, then $p_e \geq 19$. Therefore, $utility_a \leq \frac{19 \cdot (37 - 19)}{2 \cdot 18} = 9.5$. But in this case node a can improve its utility by reducing the price of edge (c, a) to 18 and thereby getting completely saturated and obtaining a utility of 18. This is clearly a contradiction since we assumed the strategy to be in Nash equilibrium.

This means that if there exists an equilibrium, then $f_c > 1$, and thus price $p_e < 19$. Consider such an equilibrium strategy. From Lemma 10 we know that the price of the two edges (c, a) and (c, b) can differ by at most δ . Let us suppose first that the prices are the same.

Now, if the amount of flow on edge (c, a) is f_a then the utility of node a is $f_a p_e$. If a deviates by reducing its price by δ , then it will be saturated (since $p_e < 19$, and thus c sends at least 1 unit of flow). Since we are assuming that the current strategy is Nash equilibrium, this deviation should not be profitable. Hence:

$$f_a p_e \geq p_e - \delta$$

The same holds for node b :

$$f_b p_e \geq p_e - \delta$$

Adding the two inequalities gives the following:

$$\begin{aligned} p_e(f_a + f_b) &\geq 2(p_e - \delta) \\ \implies p_e \lambda_c^{-1}(p_e) &\geq 2p_e - 2\delta \\ \implies p_e \frac{(37 - p_e)}{18} &\geq 2p_e - 2\delta \\ \implies p_e^2 - p_e - 36\delta &\leq 0 \\ \implies p_e &\leq \frac{1 + \sqrt{1 + 144\delta}}{2} \end{aligned}$$

Let us assume that δ is sufficiently small so that $p_e \leq 3$. At this price, even if node a receives all the flow it can, its utility will be at most 3. But this means that it can profitably deviate by increasing its price to 10, giving it a flow of size 0.5 and thus utility 5.

This contradicts the assumption that the strategy is in Nash equilibrium. A similar analysis works if we assume that the price of edges (c, a) and (c, b) differ by δ . ■

References

- [1] D. Acemoglu, R. Johari, and A.E. Ozdaglar. Partially optimal routing. *IEEE Journal on Selected Areas in Communications* 25 (6): 11481160, 2007.
- [2] Elliot Anshelevich, Bruce Shepherd, and Gordon Wilfong. Strategic Network Formation through Peering and Service Agreements. *Games and Economic Behavior*, doi:10.1016/j.geb.2011.01.002.

- [3] Elliot Anshelevich and Gordon Wilfong. Network Formation and Routing by Strategic Agents using Local Contracts. *Proc. 4th International Workshop On Internet And Network Economics (WINE 2008)*.
- [4] A. Archer, É. Tardos. Frugal path mechanisms. *ACM Transactions on Algorithms (TALG)*, 2007.
- [5] Shuchi Chawla and Feng Niu. The Price of Anarchy in Bertrand Games. *EC 2009*.
- [6] Shuchi Chawla and Tim Roughgarden. Bertrand Competition in Networks. *In Proceedings of the 1st International Symposium on Algorithmic Game Theory (SAGT '08)*, 2008.
- [7] J. Feigenbaum, C. H. Papadimitriou, R. Sami, and S. Shenker, “A BGP-based mechanism for lowest-cost routing,” *Distributed Computing*, 18(1):6172, 2005.
- [8] Joan Feigenbaum, Michael Schapira, and Scott Shenker. Distributed Algorithmic Mechanism Design. Chapter 14 in: N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani (eds.), *Algorithmic Game Theory*, Cambridge University Press.
- [9] Alexander Hall, Evdokia Nikolova, and Christos Papadimitriou. Incentive-Compatible Interdomain Routing with Linear Utilities. *Internet Mathematics*. Volume 5, Number 4 (2008), 395-410. (Special Issue for Selected papers from WINE'07.)
- [10] A. Hayrapetyan, É. Tardos, and T. Wexler. A network pricing game for selfish traffic. In *Distributed Computing*, March 2007.
- [11] G. Huston, “Interconnection, peering, and settlements,” *Proceedings of the Internet Global Summit*, The Internet Society, San Jose, CA, June 1999.
- [12] R. Johari, S. Mannor, and J.N. Tsitsiklis. A contract-based model for directed network formation. *Games and Economic Behavior* 56 (2): 201224, 2005.
- [13] Hagay Levin, Michael Schapira and Aviv Zohar. Interdomain Routing and Games. *SIAM Journal on Computing (SICOMP)*, Special Issue on Selected Papers from STOC 2008.
- [14] Asuman Ozdaglar and R. Srikant. Incentives and Pricing in Communication Networks. Chapter 22 in: N. Nisan, T. Roughgarden, É. Tardos, and V. V. Vazirani (eds.), *Algorithmic Game Theory*, Cambridge University Press.
- [15] Christos Papadimitriou and Gregory Valiant. A New Look at Selfish Routing. *Innovations in Computer Science*, 2010.
- [16] Y. Rekhter and T. Li, “A Border Gateway Protocol 4 (BGP-4)”, RFC 4271, January 2006.
- [17] T. Roughgarden and É. Tardos. How Bad is Selfish Routing? *Journal of the ACM* 2002.
- [18] M. Schapira, Y. Zhu, and J. Rexford, “Putting BGP on the right path: A case for next-hop routing,” *Proceedings of HotNets-IX*, Monterey, CA, October 2010.

- [19] V. Srinivasan, P. Nuggehalli, C.F. Chiasserini, and R. R. Rao, "An Analytical Approach to the Study of Cooperation in Wireless Ad Hoc Networks," *IEEE Transactions on Wireless Communications*, 4(2):722-733, 2005.
- [20] Yufang Xi and Edmund M. Yeh. Pricing, competition, and routing in relay networks. *In Proceedings of the 47th annual Allerton conference on Communication, control, and computing (Allerton'09)*, 2009.
- [21] M. Yuksel, A. Gupta, K. Kar, and S. Kalyanaraman, "Contract-Switching for Managing Inter-Domain Dynamics," In *Next-Generation Internet Architectures and Protocols*, Eds. B. Ramamurthy, G. N. Rouskas, K. M. Sivalingam, Cambridge University Press, pp.136-153, December 2010.