Pricing, Competition, and Routing for Selfish and Strategic Nodes in Multi-hop Relay Networks

Yufang Xi and Edmund M. Yeh

Department of Electrical Engineering

Yale University

New Haven, CT 06520, USA

Email: {yufang.xi, edmund.yeh}@yale.edu

Abstract

We study a pricing game in multi-hop relay networks where nodes price their services and route their traffic selfishly and strategically. In this game, each node (1) announces pricing functions which specify the payments it demands from its respective customers depending on the amount of traffic they route to it and (2) allocates the total traffic it receives to its service providers. The profit of a node is the difference between the revenue earned from servicing others and the cost of using others' services. We show that the socially optimal routing of such a game can always be induced by an equilibrium where no node can increase its profit by unilaterally changing its pricing functions or routing decision. On the other hand, there may also exist inefficient equilibria. We characterize the loss of efficiency by deriving the price of anarchy at inefficient equilibria. We show that the price of anarchy is finite for oligopolies with concave marginal cost functions, while it is infinite for general topologies and cost functions.

I. INTRODUCTION

It has been widely recognized that cooperation in networks formed by autonomous and selfish nodes cannot be achieved unless sufficient incentives are provided to the nodes. Such incentives normally take the form of payment or reward to the nodes if they help forward other nodes' traffic [1]–[4]. A node is usually willing to participate in routing only if it can charge more than the cost of servicing the transit traffic. While a selfish node always prices its service with the ultimate aim of maximizing its profit, it has to do so strategically since the customers it courts may potentially buy services from other nodes. Thus, there exists a trade-off in each node's pricing decision. That is, higher charges yield larger profit margins but risk losing market share to its competitors.

In this work, we study the game that arises from the selfish and strategic pricing behavior of relay nodes in a unicast multi-hop relay network consisting of one source and one destination. A node is selfish in the sense that maximizing its own profit is its sole objective. Being strategic means that a node is able to optimally design its pricing based on the anticipation of its competitors' best response to its action. Specifically, in this game each node is a service provider to a group of nodes (its customers), and when it needs to forward the traffic received from its customers, the node itself becomes a customer that uses the services of some other group of nodes. As a service provider, the node announces pricing functions which specify the payments it demands from its respective customers depending on the amount of traffic the customers route to the node. As a customer, the node allocates the total traffic it receives to its service providers in a way that minimizes the sum of its own transmission costs and the payments made to the service providers. Such a game can exist in both wireline and wireless networks, where communications consume resources and nodes are often selfish agents. When a network, especially a wireless network, is formed in an ad hoc manner, a node is typically aware of its neighbors only. A rational node thus always bases its pricing and routing decisions on the strategies adopted by its neighbors. We will show that such a game always has equilibria where no node can increase its profit by unilaterally changing its pricing functions or routing decision. Furthermore, depending on the network topology and the nodes' response strategy, the global routing configuration at an equilibrium may or may not be socially efficient. We characterize the loss of efficiency by deriving the price of anarchy at inefficient equilibria. It is found that the price of anarchy is finite for some link cost functions and topologies, while it is infinite for others.

Pricing schemes were introduced into network resource allocation problems first as a means of decomposing a global optimization into sub-problems solved by individual agents [5]. In addition to being a facilitating device, pricing serves as an essential mechanism for inducing social efficiency when users (source nodes) selfishly choose their routes [6]. It is well known that without appropriate pricing, e.g. marginal cost pricing, selfish routing inevitably results in loss of efficiency, which in general can be arbitrarily large [7], [8].

When service providers are also mindful of their self interest, they will use pricing to their own advantage rather than to heed any social mission. With both users and service providers behaving selfishly, the network increasingly approximates a free market, where prices can assume a variety of functions and lead to direct or indirect competition among service providers. For example, pricing network services according to their quality helps to match each type of service with the customers that value it the most [9], [10]. By modelling the interaction between the service provider and the users as a Stackelberg game, [11] shows

that when the service provider always adopts the profit-maximizing price, its revenue per unit bandwidth and the net utility of each user both improve with the number of users. When multiple service providers are present in a network, price competition inevitably ensues [12]–[14]. It is demonstrated in [12], [13] that cooperation in pricing is in the best interest of service providers who jointly serve the same customers. The dire consequence of non-cooperation is explicitly analyzed in [14], which shows that price competition in parallel-serial networks can result in arbitrarily large efficiency loss.

In this paper, we analyze the pricing game in multi-hop relay networks where a node can compete for traffic from multiple nodes and can allocate its received traffic to multiple nodes. Thus, in general, a node is both a service provider and a customer. Another distinctive feature of the game we consider is that the bid from each service provider to a targeted customer is a (possibly nonlinear) *pricing function*, which specifies the price contingent on the amount of service provided. Previous work on pricing games almost exclusively assume a constant unit price from every service provider, which in our terms means restricting pricing functions to be linear. It turns out that the generalization from linear to nonlinear pricing allows for a much richer set of possibilities in pricing games. Even in economics literature, the issue of nonlinear pricing is quite new and challenging [15].[1] Equilibria derived from such a general framework represent the most fundamental outcomes of pricing games in multi-hop networks.

We show that the socially optimal routing can always be induced by an equilibrium of the routing/pricing game where no node can increase its profit by unilaterally changing its pricing functions or routing decision. On the other hand, there also exist inefficient equilibria. In particular, we show that in an oligopoly routing/pricing game, inefficient equilibria are always monopolistic, i.e., a dominant relay carries all the flow from the source. We prove that the price of anarchy at such inefficient equilibria is equal to the number of relays in an oligopoly if marginal cost functions are concave. In this case, the worst inefficient equilibria arise with linear marginal cost functions. When marginal cost functions are convex, however, the price of anarchy can be arbitrarily large. Unlike the case of oligopolies, inefficiency in general multi-hop relay networks stems not only from dominant relays exercising monopolistic pricing power, but also from the myopia of dominant relays. We demonstrate that the inability of a node to gauge the impact of its pricing beyond its local neighborhood can lead to an infinitely large price of anarchy.

[1]The nonlinear pricing game we study can be seen as a generalized menu auction [16] where each bidder offers a continuum of options along with their prices.

II. NETWORK MODEL AND PROBLEM FORMULATION

A. Network Traffic and Multi-hop Routing

We consider a relay network represented by a directed graph $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with one *source* s and one *destination* w, and a set of *relays* which can be used to forward traffic in a multi-hop fashion from s to w. The source s needs to send traffic of a fixed rate R_s to w,[2] which can be carried through links in $\mathcal{E}$.

We assume that there is no direct link between s and w. That is, traffic from s has to be routed to w via relays in a multi-hop fashion. To make matters simple, we assume $\mathcal{G}$ contains only nodes and links which are on the paths from s to w. Since route discovery is not a main concern of this work, we assume such a $\mathcal{G}$ is given a priori, and is loop-free. Each node, however, needs only to be aware of its neighbors (predecessors, siblings, and offsprings) as specified below.

For node i, h is a *predecessor* if $(h, i) \in \mathcal{E}$. Denote the set of i's predecessors by $\mathcal{P}_i$. For any $h \in \mathcal{P}_i$, define $\mathcal{S}_i^h \triangleq \{j \neq i : h \in \mathcal{P}_j\}$. That is, $\mathcal{S}_i^h$ is the set of nodes which share the common predecessor h with i. These are the nodes who compete with i for h's traffic in the pricing game to be introduced later. We will refer to them as *siblings* of i with respect to h. Finally, i is said to be an *offspring* of h if $(h, i) \in \mathcal{E}$. Let the set of h's offsprings be denoted by $\mathcal{O}_h$. The above notation is illustrated in Figure 1. We make a

$$\mathcal{P}_i = \{g, h\}$$
$$\mathcal{O}_g = \{k, i\}$$
$$\mathcal{O}_h = \{i, j\}$$
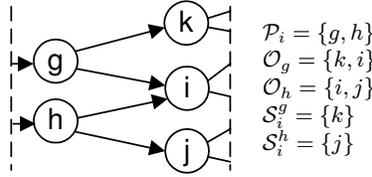$$\mathcal{S}_i^g = \{k\}$$
$$\mathcal{S}_i^h = \{j\}$$

Fig. 1.   Illustration of predecessors, siblings and offsprings.

simplifying yet plausible assumption that $\bigcup_{h \in \mathcal{P}_i} \mathcal{S}_i^h \cap \mathcal{P}_i = \emptyset$, i.e., no node can be both a sibling and a predecessor of any other node.

By our assumption on $\mathcal{G}$, s is the only node without any predecessor while w is the only node without any offspring. Since the pricing game to be studied can arise only if there are multiple relays competing for the traffic from their common predecessor, we assume in $\mathcal{G}$ that every node i except w has multiple offsprings unless $w \in \mathcal{O}_i$.

Denote the rate of flow on $(i, j) \in \mathcal{E}$ by f_{ij}. A link flow vector $\boldsymbol{f} \triangleq (f_{ij})_{(i,j) \in \mathcal{E}}$ is a *routing* of the session traffic if it satisfies the flow conservation constraint: $\sum_{h \in \mathcal{O}_s} f_{sh} = R_s$, $\sum_{k \in \mathcal{P}_w} f_{kw} = R_s$, and for

---
[2]We will discuss the problem involving an elastic session in Section V.

each relay i,

$$\sum_{h \in \mathcal{P}_i} f_{hi} = \sum_{j \in \mathcal{O}_i} f_{ij} \triangleq r_i,$$

where r_i denotes the incoming flow rate at i.

B. Link Cost and Pricing Functions

Each link has a strictly increasing and strictly convex cost function $D_{ij}(f_{ij})$, which is *private information* to i and j only. For example, $D_{ij}(f_{ij})$ can represent the queuing delay incurred on (i, j) with arrival rate f_{ij}, e.g. the average occupancy function $f_{ij}/(c_{ij} - f_{ij})$ of an M/M/1 queue with service rate c_{ij}. As another example, if the links are wireless, $D_{ij}(f_{ij})$ can measure the transmission power required for achieving rate f_{ij}. For example, if the link transmission rate f_{ij} is determined by transmission power P_{ij} as $f_{ij} = W \log(1 + K P_{ij})$ for some constants $W, K > 0$,[3] then $P_{ij} = \frac{1}{K}\left(2^{f_{ij}/W} - 1\right) \triangleq D_{ij}(f_{ij})$, which is strictly increasing and convex in f_{ij}. As suggested by the examples, the analytical framework presented above applies to both wireline and wireless networks.

For analytical purposes, we further assume that $D_{ij}(\cdot)$ is continuously differentiable with derivative $d_{ij}(\cdot)$. By previous assumptions, $d_{ij}(\cdot)$ is positive and strictly increasing. The *socially optimal* routing is the routing that minimizes the network cost $\sum_{(i,j)} D_{ij}(f_{ij})$. Because link costs are strictly convex, the socially optimal routing is uniquely characterized by the condition that every path from s to w with positive flow has the minimum marginal cost among all paths. For otherwise, one can reduce the total cost by shifting an infinitesimal amount of flow from a path with non-minimum marginal cost to a minimum-marginal-cost path.

We model the source and relays as selfish agents who must pay for the costs on their outgoing links. While the source has to send all its traffic out, it strives to do this with the minimum cost. On the other hand, a relay has an incentive to forward traffic for its predecessors only if it is adequately rewarded for its service in the form of payment by its predecessors. The amount of payment is determined as follows.

Suppose a node h has incoming flow of rate $r_h > 0$. Each $i \in \mathcal{O}_h$ announces a pricing function $P_i^h(\cdot)$ which specifies the payment $P_i^h(f_{hi})$ it demands should h forward traffic of rate f_{hi} to it.[4] For analytical purposes, we assume that $P_i^h(t)$ is continuously differentiable with derivative $p_i^h(t)$. Note that $P_i^h(\cdot)$

---

[3]Assume that with proper time or frequency scheduling, transmission on different links are non-interfering.

[4]The domain of $P_i^h(\cdot)$ must contain the interval $[0, r_h]$. And as we will see, a selfish and strategic relay i always designs $P_i^h(\cdot)$ tailored to the total traffic r_h. However, for simplicity we do not express such dependence in the notation.

provides h a *continuum* of options, namely the rate-price pairs $(f_{hi}, P_i^h(f_{hi}))$.[5] After learning $(P_i^h(\cdot))_{i \in \mathcal{O}_h}$, h decides on the allocation of r_h and makes payments to its offsprings accordingly.

C. Pricing Game

We assume every node is selfish and strategic.[6] The source s thus always allocates the total flow R_s to the nodes in $\mathcal{O}_s$ so as to minimize its total cost, which includes the costs on its outgoing links and the payments to its offsprings. Specifically, given the pricing functions $P_i^s(\cdot)$ of $i \in \mathcal{O}_s$, the optimal allocation of R_s from the perspective of s is any

$$(f_{si}^*)_{i \in \mathcal{O}_s} \in \underset{(f_{si}) \in \mathcal{F}_s(R_s)}{\arg\min} \sum_{i \in \mathcal{O}_s} D_{si}(f_{si}) + P_i^s(f_{si}), \tag{1}$$

where $\mathcal{F}_i(r)$ is defined as $\{(f_{ik})_{k \in \mathcal{O}_i} \geq \mathbf{0} : \sum_k f_{ik} = r\}$.

A relay i is a predecessor to some nodes, and is an offspring to some other nodes. As a predecessor, it acts just like s. That is, it allocates the total incoming flow in the most cost efficient way from its own perspective. Thus, the traffic allocation adopted by i when it has incoming flow r_i is any

$$(f_{ik}^*)_{k \in \mathcal{O}_i} \in \underset{(f_{ik}) \in \mathcal{F}_i(r_i)}{\arg\min} \sum_{k \in \mathcal{O}_i} D_{ik}(f_{ik}) + P_k^i(f_{ik}). \tag{2}$$

Denote the minimum value in (2) by $D_i(r_i)$. Note that $D_i(r_i)$ represents the minimum cost to i for forwarding flow of rate r_i.[7] It is easy to show that $D_i(\cdot)$ is continuous and increasing with piecewise continuous derivative denoted by $d_i(\cdot)$.

As an offspring, i designs $P_i^h(\cdot)$ for every $h \in \mathcal{P}_i$ with the aim of maximizing its profit in competition with its siblings (discussed in depth later). It does this with the assumption that h always allocates in the most cost efficient way, and that r_h for each $h \in \mathcal{P}_i$ stays constant at the current value irrespective of its choice of $P_i^h(\cdot)$. While the first assumption is very reasonable, the second one requires some justification.

Theoretically, since r_h of $h \in \mathcal{P}_i$ is the outcome of the optimal allocation by h's predecessors, it in general cannot stay constant if h changes its pricing functions. However, h's pricing functions presumably are tied to i's choice of $P_i^h(\cdot)$ as the total cost to h is partly leveraged by the price charged by i. Once

---

[5]If i has multiple predecessors, $P_i^h(t)$ for one $h \in \mathcal{P}_i$ is an agreement exclusively between i and h, independent of the flow rates allocated to i by other predecessors. Presumably, however, i designs $P_i^h(t)$ for all $h \in \mathcal{P}_i$ jointly because $r_i = \sum_{h \in \mathcal{P}_i} f_{hi}$, and i has to pay its offsprings to get r_i forwarded.

[6]The destination w is the only node that plays no active role in the pricing game described below. It passively accepts the flow assigned by its predecessors, who treat it as an offspring using a uniformly-zero pricing function. Because $\mathcal{G}$ is assumed to be loop-free and provide directed path(s) to w from every other node, the total flow arriving at w must be equal to R_s.

[7]Although not explicit from the notation, $D_i(\cdot)$ depends on $(P_k^i(\cdot))_{k \in \mathcal{O}_i}$.

h reacts to the change in $P_i^h(\cdot)$ by updating its own pricing functions, r_h is inevitably adjusted by h's predecessors. On the other hand, relays are usually too myopic to note this chain reaction since they have very limited knowledge of the network. Recall that we made a practical assumption that a relay is aware of only its predecessors, siblings and offsprings. As a result, it can at best predict the impact of its strategy on the traffic allocation by its predecessors, but not nodes further upstream. It is therefore reasonable for i to consider only the competition with its siblings for the flow their common predecessors currently have.

We now formally define the (static) pricing game (PG) as having the following components:

- The set of players $\mathcal{I} = \mathcal{N}\backslash\{s, w\}$: relays in $\mathcal{G}$.
- Strategy of player i: continuously differentiable pricing functions $P_i^h(\cdot)$ for all $h \in \mathcal{P}_i$.
- Payoff to player i: the profit made by servicing $(f_{hi}^*)_{h\in\mathcal{P}_i}$:

$$\sum_{h\in\mathcal{P}_i} P_i^h(f_{hi}^*) - D_i\left(\sum_{h\in\mathcal{P}_i} f_{hi}^*\right), \tag{3}$$

where the routing $(f_{ij}^*)_{(i,j)\in\mathcal{E}}$ is most cost efficient from the perspective of every node, i.e., (1)-(2) hold for s and every relay i where $r_i = \sum_{h\in\mathcal{P}_i} f_{hi}^*$.

A pricing game is fully characterized by the tuple $(\mathcal{G}, (D_{ij}(\cdot)), R_s)$. In the rest of the paper, we will study the outcome of the PG with *myopic* players as described above. The focus of our work is to investigate whether the PG has an equilibrium where no relay can increase its profit by unilaterally changing its pricing functions, and when an equilibrium exists, how the resulting routing compares to the socially optimal one.

D. Best Response and Equilibrium

Each player i in the PG is assumed to be myopic in the sense that it knows the pricing functions of all its competitors as well as its downstream nodes. Based on its local information $\boldsymbol{L}_i \triangleq ((r_h, (P_j^h(\cdot))_{j\in\mathcal{S}_i^h})_{h\in\mathcal{P}_i}, (P_k^i(\cdot))_{k\in\mathcal{O}_i})$, player i anticipates a payoff $\Gamma_i(\boldsymbol{P}_i; \boldsymbol{L}_i)$ when adopting $\boldsymbol{P}_i \triangleq (P_i^h(\cdot))_{h\in\mathcal{P}_i}$, where

$$\Gamma_i(\boldsymbol{P}_i; \boldsymbol{L}_i) \triangleq \sum_{h\in\mathcal{P}_i} P_i^h(f_{hi}^*) - \sum_{k\in\mathcal{O}_i}\left[D_{ik}(f_{ik}^*) + P_k^i(f_{ik}^*)\right],$$

$(f_{hj}^*)_{j\in\mathcal{O}_h}$ is the optimal allocation of r_h by $h \in \mathcal{P}_i$ given $(P_j^h(\cdot))_{j\in\mathcal{O}_h}$, and $(f_{ik}^*)_{k\in\mathcal{O}_i}$ is the optimal allocation of $r_i^* = \sum_{h\in\mathcal{P}_i} f_{hi}^*$ given $(P_k^i(\cdot))_{k\in\mathcal{O}_i}$.

Definition 1: A pricing function profile $(P_i^h(\cdot))_{h\in\mathcal{P}_i}$ is a best response to local information $\boldsymbol{L}_i$ if

$$\Gamma_i(\boldsymbol{P}_i; \boldsymbol{L}_i) = \max_{\boldsymbol{Q}_i \text{feasible}} \Gamma_i(\boldsymbol{Q}_i; \boldsymbol{L}_i),$$

where $\boldsymbol{Q}_i = (Q_i^h(\cdot))_{h \in \mathcal{P}_i}$ is feasible if every component is a continuously differentiable function.

Denote the set of best responses to $\boldsymbol{L}_i$ by $\mathcal{B}_i(\boldsymbol{L}_i)$. Before we proceed, we must prove that $\mathcal{B}_i(\boldsymbol{L}_i)$ is non-empty.

Lemma 1: For any $\boldsymbol{L}_i$, the set $\mathcal{B}_i(\boldsymbol{L}_i)$ is non-empty. Furthermore, $\boldsymbol{P}_i \in \mathcal{B}_i(\boldsymbol{L}_i)$ if and only if for all $h \in \mathcal{P}_i$ and all $t \in [0, r_h]$,

$$B_i^h(t) \triangleq D_{hi}(t) + P_i^h(t) \geq B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - t), \tag{4}$$

and

$$B_i^h(\tilde{f}_{hi}) = B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - \tilde{f}_{hi}), \tag{5}$$

where

$$B_{\hat{i}}^h(r) \triangleq \min_{\substack{(f_{hj}) \in \mathcal{F}_h(r) \\ f_{hi} = 0}} \sum_{j \in \mathcal{S}_i^h} B_j^h(f_{hj}) \tag{6}$$

and $(\tilde{f}_{hi})_{h \in \mathcal{P}_i}$ is a vector that maximizes

$$\bar{\Gamma}_i(\boldsymbol{f}_i; \boldsymbol{L}_i) \triangleq \sum_{h \in \mathcal{P}_i} \left[B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - f_{hi}) - D_{hi}(f_{hi}) \right] - D_i \left(\sum_{h \in \mathcal{P}_i} f_{hi} \right)$$

over all the $\boldsymbol{f}_i = (f_{hi})$ such that $0 \leq f_{hi} \leq r_h$ for all $h \in \mathcal{P}_i$.

Before giving the proof, we first provide some intuitive explanations for the lemma. The function $B_i^h(t)$ gives the total cost h spends on routing traffic of rate t to i. Since $D_{hi}(\cdot)$ is fixed and known to i, it is equivalent to treat $B_i^h(\cdot)$ as the pricing function i uses to charge h. With this view, h makes a lump-sum payment to i determined by $B_i^h(\cdot)$, and lets i pay for the cost $D_{hi}(\cdot)$ on link (h, i). For convenience, from now on we assume each relay i announces $B_i^h(\cdot)$ to $h \in \mathcal{P}_i$ and siblings $j \in \mathcal{S}_i^h$. By (6), $B_{\hat{i}}^h(r)$ represents the minimum cost h can achieve by forwarding traffic of rate r to offsprings other than i. It will become evident in the next proof that from i's viewpoint, the competition from all $j \in \mathcal{S}_i^h$ can be aggregated into a virtual competitor $\hat{i}^h$ using pricing function $B_{\hat{i}}^h(\cdot)$. Thus, it is as if i were competing with one relay $\hat{i}^h$ in each "market" $h \in \mathcal{P}_i$. The vector $(\tilde{f}_{hi})_{h \in \mathcal{P}_i}$ represents the "market shares" that jointly yield the maximum (anticipated) profit to i. Pricing functions $B_i^h(\cdot)$ which satisfy the conditions in Lemma 1 induce $h \in \mathcal{P}_i$ to allocate the ideal "market share" $\tilde{f}_{hi}$ to i and give i the maximum profit. This is because (5) implies that allocating $\tilde{f}_{hi}$ to i and the rest to other relays yields the same cost to h as allocating all the traffic to other relays. So conditions (4) and (5) combined imply that no other allocation costs less than the above two schemes.

Proof of Lemma 1: First notice that for fixed $\boldsymbol{L}_i$, the profit of i when it adopts $(B_i^h(\cdot))_{h\in\mathcal{P}_i}$ is upper bounded as follows:

$$
\begin{aligned}
\Gamma_i((B_i^h(\cdot)); \boldsymbol{L}_i) &= \sum_{h\in\mathcal{P}_i} \left[B_i^h(f_{hi}) - D_{hi}(f_{hi}) \right] - D_i\left(\sum_{h\in\mathcal{P}_i} f_{hi} \right) \\
&\leq \sum_{h\in\mathcal{P}_i} \left[B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - f_{hi}) - D_{hi}(f_{hi}) \right] - D_i\left(\sum_{h\in\mathcal{P}_i} f_{hi} \right) \\
&= \bar{\Gamma}_i(\boldsymbol{f}_i; \boldsymbol{L}_i) \leq \bar{\Gamma}_i(\tilde{\boldsymbol{f}}_i; \boldsymbol{L}_i),
\end{aligned}
$$

where $\boldsymbol{f}_i \triangleq (f_{hi})_{h\in\mathcal{P}_i}$ is the optimal amount of traffic allocated to i when i uses $(B_i^h(\cdot))$. The first inequality holds because for every $h \in \mathcal{P}_i$, $B_i^h(f_{hi}) \leq B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - f_{hi})$. For otherwise, h would find the cost of allocating r_h exclusively to all $j \in \mathcal{S}_i^h$ ($B_{\hat{i}}^h(r_h)$) strictly less than the cost of allocating f_{hi} to i and $r_h - f_{hi}$ to $j \in \mathcal{S}_i^h$ ($B_i^h(f_{hi}) + B_{\hat{i}}^h(r_h - f_{hi})$), a contradiction. The second inequality follows from the definition of $\tilde{\boldsymbol{f}}_i$.

Notice that the upper bound $\bar{\Gamma}_i(\tilde{\boldsymbol{f}}_i; \boldsymbol{L}_i)$ is independent of $(B_i^h(\cdot))$. It is tight if and only if both inequalities hold with equality. To make the second inequality tight, it is necessary and sufficient to have the traffic allocation $\boldsymbol{f}_i$ induced by $(B_i^h(\cdot))$ equal to $\tilde{\boldsymbol{f}}_i$. Given that, the first inequality is tight if (4)-(5) hold, since they guarantee that allocating $\tilde{f}_{hi}$ to i and $r_h - \tilde{f}_{hi}$ to $j \in \mathcal{S}_i^h$ is in h's best interest. They are also necessary because (5) is prerequisite for the first inequality to be tight, and consequently (4) cannot be violated either. For example, if for some $h \in \mathcal{P}_i$ and some $t \in [0, r_h]$, $B_i^h(t) < B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - t)$. Then allocating t to i and $r_h - t$ to $j \in \mathcal{S}_i^h$ would incur strictly less cost to h than $B_{\hat{i}}^h(r_h) = B_i^h(t)(\tilde{f}_{hi}) + B_{\hat{i}}^h(r_h - \tilde{f}_{hi})$. Thus, h would not have allocated $\tilde{f}_{hi}$ to i, which is a contradiction. $\square$

Note that best response $(B_i^h(\cdot))$ always exists because, for instance, $B_i^h(t) = B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - t)$ for $t \in [0, r_h]$ satisfies (4)-(5). The (pure-strategy) Nash equilibria are defined as the fixed points of the best response mapping.

Definition 2: Pricing function profiles $\boldsymbol{P}_i, i \in \mathcal{I}$ constitute an equilibrium if for all $i \in \mathcal{I}$, $\boldsymbol{P}_i \in \mathcal{B}_i(\boldsymbol{L}_i)$ where the incoming flow vector $(r_h)_{h\in\mathcal{P}_i}$ contained in $\boldsymbol{L}_i$ results from the routing $(f_{ij}^*)_{(i,j)\in\mathcal{E}}$ that is most cost efficient from the perspective of every individual node, i.e., (1)-(2) hold for s and all $i \in \mathcal{I}$,

$$
r_h = \sum_{g\in\mathcal{P}_h} f_{gh}^*, \quad \text{if } h \neq s
$$

and $r_h = R_s$ if $h = s$.

It is easy to see that if $(\boldsymbol{P}_i)_{i \in \mathcal{I}}$ constitutes an equilibrium, $\Gamma_i(\boldsymbol{P}_i; \boldsymbol{L}_i)$ must coincide with the actual payoff of i.

Definition 3: An equilibrium $(\boldsymbol{P}_i)_{i \in \mathcal{I}}$ is *efficient* if it induces the socially optimal routing. In this case, $(\boldsymbol{P}_i)_{i \in \mathcal{I}}$ is said to induce the social optimum.

Before proving the existence of equilibria and analyzing their efficiency in the general setting, we first study pricing games under some simple network topologies. The analysis of these games not only provide valuable insight into the general problem but also have significant implications in their own right.

III. EQUILIBRIA IN OLIGOPOLY

The simplest topologies within our framework are those including a single layer of relays, e.g. the one in Figure 2. Here, N relays each have a direct link from s and a direct link to w. They compete for the
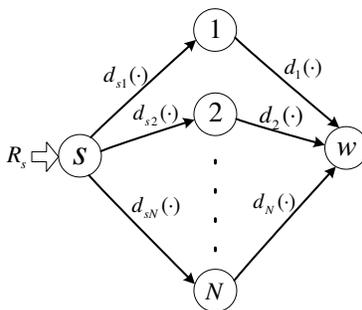

Fig. 2.   Oligopoly with N relays.

total flow R_s by advertising their pricing functions $\beta_i(\cdot) = d_{si}(\cdot) + p_i(\cdot), i = 1, \cdots, N$. From now on, we will more often refer to the derivatives $d_{ij}(\cdot)$ and $p_i(\cdot)$ as link cost and pricing functions, since they appear to be more convenient for marginal cost analysis at equilibria. Also we will use simplified notation whenever appropriate, e.g. the superscript is omitted from $p_i(\cdot), \beta_i(\cdot)$ as s is the only predecessor to every relay. We refer to a pricing game under such a topology as an *oligopoly*. Define $\lambda_i(t) \triangleq d_{si}(t) + d_i(t)$. An oligopoly PG is fully characterized by the tuple $(N, (\lambda_i(\cdot))_{i=1}^N, R_s)$.

Because all $\lambda_i(\cdot)$ are strictly increasing, the socially optimal routing $(r_i^*)_{i=1}^N$ is unique and is given by

$$\lambda_i(r_i^*) = \min_{j=1,\cdots,N} \lambda_j(r_j^*)$$

if $r_i^* > 0$.

We now analyze the routing established by the oligopoly PG. Given $(\beta_i(\cdot))_{i \in \mathcal{O}_s}$, the self-interest of s leads it to adopt the most cost efficient routing $(f^*_{si})_{i=1}^N$ such that

$$\beta_i(f^*_{si}) = \min_{j=1,\cdots,N} \beta_j(f^*_{sj}) \tag{7}$$

whenever $f^*_{si} > 0$. Whether $(f^*_{si}) = (r^*_i)$ or not depends on how $(\beta_i(\cdot))$ are chosen by the individual relays.

A. Best Response and Existence of Equilibria

We apply Lemma 1 to the oligopoly PG. Here,

$$B_i(t) = \int_0^t \beta_i(r) \ dr,$$

$$B_{\hat{i}}(t) = \min_{\sum_{j \neq i} f_{sj} = t} \sum_{j \neq i} \int_0^{f_{sj}} \beta_j(r) \ dr.$$

It is easy to show that $B_{\hat{i}}(t)$ is continuous and increasing. Its derivative, denoted by $\beta_{\hat{i}}(t)$, is in general piecewise continuous. For $t \in (0, R_s)$, let the left and right limits of $\beta_{\hat{i}}(t)$ be denoted by $\beta_{\hat{i}}(t)^-$ and $\beta_{\hat{i}}(t)^+$.[8] By Lemma 1, the best response of i given $\beta_{\hat{i}}(\cdot)$ can be simply characterized by

$$\int_0^t \beta_i(r) \ dr \begin{cases} \geq \int_0^t \beta_{\hat{i}}(R_s - r) \ dr, & 0 \leq t \leq R_s \\ = \int_0^t \beta_{\hat{i}}(R_s - r) \ dr, & t = f^*_i, \end{cases} \tag{8}$$

where

$$f^*_i \in \arg\max_{0 \leq f_i \leq R_s} \int_0^{f_i} \beta_{\hat{i}}(R_s - r) - \lambda_i(r) \ dr. \tag{9}$$

To gain an intuitive idea of the above conditions, suppose $\beta_{\hat{i}}(R_s - r)$ and $\lambda_i(r)$ are given by the dashed and solid curves in Figure 3. A typical best response $\beta_i(r)$ is shown as the dotted curve. In particular,
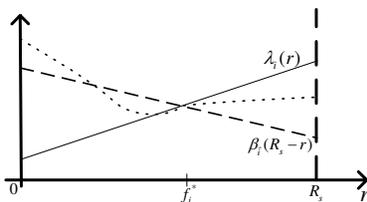

Fig. 3. Typical best response curve in oligopoly.

one can let $\beta_i(r)$ coincide with $\beta_{\hat{i}}(R_s - r)$ on $[0, f^*_i]$ and let $\beta_i(r) \geq \beta_{\hat{i}}(R_s - r)$ on $(f^*_i, R_s]$. Such a best response will be referred to as a *replicating* response. As we will show, oligopoly equilibria induced by replicating responses are always efficient while equilibria induced by other best responses are not necessarily efficient.

[8]It is understood that $\beta_{\hat{i}}(0)$ has only a right limit and that $\beta_{\hat{i}}(R_s)$ has only a left limit.

B. Efficient Equilibria

Theorem 1: The socially optimal routing of an oligopoly can always be induced by an equilibrium.

Proof: We prove the theorem by constructing an equilibrium that induces the socially optimal routing (r_i^*). Define $\lambda^* \triangleq \min_{j=1,\cdots,N} \lambda_j(r_j^*)$. Let $\beta_i(r) \equiv \lambda^*$ for all i. Then $\beta_i(r) = \beta_{\hat{i}}(R_s - r) = \lambda^*$ is a best response for all i with $f_i^* = r_i^*$. Thus, $(\beta_i(\cdot))$ constitutes an equilibrium which results in the routing (r_i^*). $\square$

Because the socially optimal routing always exists, we can conclude that there *always* exists an *efficient* equilibrium for any oligopoly pricing game.

Although we used constant $(\beta_i(\cdot))$ (or linear pricing functions $(B_i(\cdot))$) to construct an efficient equilibrium in the proof, efficient equilibria can be established by nonlinear pricing functions as well. For instance, Figure 4 depicts an equilibrium in a duopoly PG where the two relays adopt $\beta_1(\cdot), \beta_2(\cdot)$ of a more general shape. Notice that in a duopoly, $\beta_{\hat{1}}(t) = \beta_2(t)$ and $\beta_{\hat{2}}(t) = \beta_1(t)$.
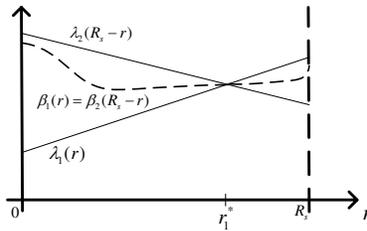

Fig. 4.   General (focal) equilibrium in duopoly.

To derive a simple criterion for checking the efficiency of an equilibrium, we need to make the following distinction. A routing $(f_i)_{i=1}^N$ is said to be *monopolistic* if $f_m = R_s$ for some relay m and $f_j = 0$ for all $j \neq m$. In this case, m is called the *dominant* relay. An equilibrium is *monopolistic* if it induces a monopolistic routing. A routing is said to be *competitive* if there are at least two relays i, j such that $f_i > 0, f_j > 0$. An equilibrium is *competitive* if it induces a competitive routing.

Theorem 2: If an oligopoly equilibrium is competitive, it must be efficient.

We will need the next lemma to prove Theorem 2.

Lemma 2: At an oligopoly equilibrium $(\beta_i(\cdot))$ which induces routing (f_i^*), if $0 < f_i^* \leq R_s$, then for all $j \neq i$,

$$\beta_{\hat{i}}(R_s - f_i^*)^+ \leq \beta_j(f_j^*).$$

If $0 \leq f_i^* < R_s$, then for all $j \neq i$ such that $f_j^* > 0$,

$$\beta_{\hat{i}}(R_s - f_i^*)^- \geq \beta_j(f_j^*).$$

Proof: By definition, if $t < R_s$, $\beta_{\hat{i}}(t)^+ = \lim_{\Delta \to 0^+}(B_{\hat{i}}(t + \Delta) - B_{\hat{i}}(t))/\Delta$. Therefore when $f_i^* > 0$,

$$
\begin{aligned}
\beta_{\hat{i}}(R_s - f_i^*)^+ &= \lim_{\Delta \to 0^+} \frac{1}{\Delta} \{B_{\hat{i}}(R_s - f_i^* + \Delta) - B_{\hat{i}}(R_s - f_i^*)\} \\
&\overset{(a)}{=} \lim_{\Delta \to 0^+} \frac{1}{\Delta} \left\{ \min_{\sum_{j \neq i} f_j = R_s - f_i^* + \Delta} \sum_{j \neq i} \int_0^{f_j} \beta_j(r) \ dr - \sum_{j \neq i} \int_0^{f_j^*} \beta_j(r) \ dr \right\} \\
&\overset{(b)}{\leq} \lim_{\Delta \to 0^+} \frac{1}{\Delta} \left\{ \int_0^{f_k^* + \Delta} \beta_k(r) \ dr + \sum_{j \neq i, k} \int_0^{f_j^*} \beta_j(r) \ dr - \sum_{j \neq i} \int_0^{f_j^*} \beta_j(r) \ dr \right\} \\
&= \lim_{\Delta \to 0^+} \frac{1}{\Delta} \int_{f_k^*}^{f_k^* + \Delta} \beta_k(r) \ dr = \beta_k(f_k^*).
\end{aligned}
$$

Here, equation (a) follows from the fact that $(f_j^*)_{j=1}^N$ is the equilibrium routing of R_s induced by $(\beta_j(\cdot))_{j=1}^N$. Inequality (b) is obtained by substituting the minimum-cost routing of $R_s - f_i^* + \Delta$ by an arbitrary routing, namely $f_k^* + \Delta$ allocated to k and f_j^* to each $j \neq i, k$. The second inequality in the lemma can be proved in a similar manner. $\qquad\square$

Proof of Theorem 2: Let (f_i^*) be the routing induced by a competitive equilibrium $(\beta_i(\cdot))$. Let m, n be any two relays such that $f_m^* > 0, f_n^* > 0$. It is enough to show that $\lambda_m(f_m^*) = \lambda_n(f_n^*)$ and that $\lambda_m(f_m^*) \leq \lambda_j(f_j^*)$ for any j with $f_j^* = 0$. By (7), $\beta_m(f_m^*) = \beta_n(f_n^*)$. The best response condition (9) implies that $\beta_{\hat{m}}(R_s - f_m^*)^- \leq \lambda_m(f_m^*) \leq \beta_{\hat{m}}(R_s - f_m^*)^+$. By Lemma 2, $\beta_n(f_n^*) \leq \beta_{\hat{m}}(R_s - f_m^*)^-$ and $\beta_{\hat{m}}(R_s - f_m^*)^+ \leq \beta_n(f_n^*)$. In conclusion, $\beta_n(f_n^*) = \lambda_m(f_m^*)$. By symmetry, we can show that $\beta_m(f_m^*) = \lambda_n(f_n^*)$. Therefore, $\lambda_m(f_m^*) = \lambda_n(f_n^*)$. Now suppose $f_j^* = 0$. By (9) and Lemma 2, $\lambda_j(0) \geq \beta_{\hat{j}}(R_s)^- \geq \beta_n(f_n^*) = \lambda_m(f_m^*)$. So the proof is complete. $\qquad\square$

C. Inefficient Equilibria

Theorem 2 does not rule out the possibility of inefficient equilibria. In fact, an equilibrium may be inefficient if it is monopolistic. For example, the socially optimal routing of the duopoly PG represented by Figure 5 is $(r_1^*, R_s - r_1^*)$ whereas the equilibrium depicted leads to a monopolistic routing $(R_s, 0)$. In this example, relay 2 adopts a pricing function $\beta_2(\cdot)$ such that $\int_0^{R_s} \beta_2(r) \ dr = \int_0^{R_s} \lambda_2(r) \ dr$ and $\beta_2(R_s - r) > \lambda_1(r)$ for all $r \in [0, R_s]$. Given such a $\beta_2(\cdot)$, relay 1 would want to acquire all the flow (cf. (9)) by using $\beta_1(\cdot)$ such that $\int_0^t \beta_1(R_s - r) \ dr < \int_0^t \lambda_2(r) \ dr$ and $\int_0^t \beta_1(r) \ dr > \int_0^t \beta_2(R_s - r) \ dr$ for
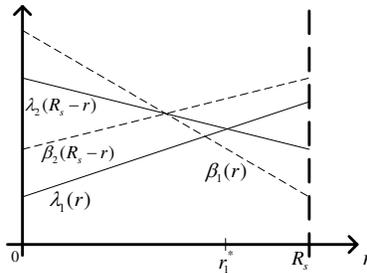
Fig. 5. Inefficient equilibrium in duopoly.

all $t \in (0, R_s)$ and $\int_0^{R_s} \beta_1(R_s - r) \ dr = \int_0^{R_s} \lambda_2(r) \ dr = \int_0^{R_s} \beta_2(r) \ dr$. Thus, it satisfies (8) and leaves relay 2 no incentive to acquire any traffic. So the monopolistic equilibrium holds.

In general, monopolistic equilibria in an oligopoly PG have the following property.

Theorem 3: If an oligopoly equilibrium is monopolistic with dominant relay m, we must have

$$\int_0^{R_s} \lambda_m(r) \ dr \leq \int_0^{R_s} \lambda_j(r) \ dr$$

for any other relay j.

Proof: Consider any j with $f_j^* = 0$ in a monopolistic equilibrium. The condition (9) implies that

$$\int_0^{R_s} \lambda_j(r) \ dr \geq \int_0^{R_s} \beta_{\hat{j}}(R_s - r) \ dr.$$

On the other hand,

$$\int_0^{R_s} \beta_{\hat{j}}(R_s - r) \ dr = \int_0^{R_s} \beta_{\hat{j}}(r) \ dr = \int_0^{R_s} \beta_m(r) \ dr,$$

since from the perspective of s, the optimal allocation of R_s to all the relays except j still assigns all the traffic to m. It follows from m's best response conditions (8)-(9) that

$$\int_0^{R_s} \beta_m(r) \ dr = \int_0^{R_s} \beta_{\hat{m}}(R_s - r) \ dr \geq \int_0^{R_s} \lambda_m(r) \ dr.$$

Thus the proof is complete. □

The next conclusion easily follows from Theorem 3.

Corollary 1: If the socially optimal routing of an oligopoly is monopolistic, then every equilibrium of the oligopoly is monopolistic and efficient.

Proof: It can deduced from the uniqueness of the socially optimal routing and Theorem 2 that every equilibrium of such an oligopoly must be monopolistic. By Theorem 3, the dominant relay m of such an

equilibrium has the minimum $\int_0^{R_s} \lambda_m(r)\, dr$ among all relays. But such an m must be the dominant relay in the socially optimal routing. □

It is shown next that there always exists a monopolistic equilibrium in an oligopoly. Thus, we have the following conclusion.

Corollary 2: If the socially optimal routing of an oligopoly is competitive, then there exists an inefficient (monopolistic) equilibrium.

Proof: We need only show that there exists a monopolistic equilibrium in such an oligopoly. Let all $\beta_j(\cdot)$ be the same strictly decreasing function $\beta(\cdot)$ such that $\int_0^t \beta(R_s - r)dr \leq \int_0^t \lambda_j(r)dr$ for all j and $t \in [0, R_s)$ but $\int_0^{R_s} \beta(R_s - r)dr = \int_0^{R_s} \lambda_m(r)dr$ where $m \in \arg\min_j \int_0^{R_s} \lambda_j(r)dr$. Since $\beta(\cdot)$ is strictly decreasing, $\beta_{\hat{j}}(r) = \beta(r)$ for all j. By construction, $f_j^* = 0$ is an ideal flow to $j \neq m$ (cf. (9)) whereas $\beta_m(\cdot) = \beta(\cdot)$ and $f_m^* = R_s$ jointly satisfy m's best response conditions (8)-(9). So the monopolistic equilibrium is established. □

When an oligopoly has inefficient equilibria, it is of interest to compare the worst-case network cost under an inefficient equilibrium to the optimal cost.

D. Price of Anarchy

The price of anarchy, as a measure of loss of social efficiency due to selfish behavior of individual agents, was studied in the literature on selfish routing [7], [8]. In this work, the price of anarchy of a general PG is defined as follows.

Definition 4: The price of anarchy $\rho(\mathcal{G}, (D_{ij}(\cdot)), R_s)$ of a pricing game $(\mathcal{G}, (D_{ij}(\cdot)), R_s)$ is the ratio of the maximum cost at an equilibrium to the socially optimal cost, i.e.,

$$\rho(\mathcal{G}, (D_{ij}(\cdot)), R_s) \triangleq \frac{\max_{(f_{ij}) \in \mathcal{F}^E} \sum_{(i,j) \in \mathcal{E}} D_{ij}(f_{ij})}{\sum_{(i,j) \in \mathcal{E}} D_{ij}(f_{ij}^*)},$$

where $\mathcal{F}^E$ is the collection of all routings that can be induced by an equilibrium of $(\mathcal{G}, (D_{ij}(\cdot)), R_s)$ and $(f_{ij}^*)_{(i,j) \in \mathcal{E}}$ is the socially optimal routing of the game.

In this section, we study the price of anarchy specifically for oligopolies. As we will show, $\rho(N, (\lambda_i(\cdot)), R_s)$ is equal to N when marginal cost functions are concave, e.g. when cost functions are quadratic. However, the price of anarchy can be arbitrarily large when when marginal cost functions are convex, as is the case for the cost functions discussed in Section II-B.

Theorem 4: If the cost derivatives $(\lambda_i(\cdot))$ are concave, $\rho(N, (\lambda_i(\cdot)), R_s)$ of an oligopoly pricing game is upper bounded by the number of relays N. The upper bound is achieved when the cost derivatives are linear.

Proof: Let the socially optimal routing be $(r_i^*)_{i=1}^N = (\alpha_i R_s)_{i=1}^N$ where the coefficients (α_i) are nonnegative and sum to one. The optimal cost then is

$$D^* = \sum_{i=1}^N \int_0^{\alpha_i R_s} \lambda_i(r) \ dr.$$

Since $\lambda_i(r)$ is concave, it can be shown that $\int_0^{\alpha_i R_s} \lambda_i(r) dr \geq \alpha_i^2 \int_0^{R_s} \lambda_i(r) dr$ where equality holds when $\lambda_i(r)$ is linear. Therefore, D^* is lower bounded as

$$D^* \geq \sum_{i=1}^N \alpha_i^2 \int_0^{R_s} \lambda_i(r) \ dr.$$

Recall that inefficient equilibria in an oligopoly are monopolistic such that the dominant relay m satisfies Theorem 3. The price of anarchy, which is the ratio of the cost at any monopolistic equilibrium (ME) to D^*, is upper bounded as

$$\frac{D^{ME}}{D^*} \overset{(a)}{\leq} \frac{\int_0^{R_s} \lambda_m(r) \ dr}{\sum_{i=1}^N \alpha_i^2 \int_0^{R_s} \lambda_i(r) \ dr} \overset{(b)}{\leq} \frac{\int_0^{R_s} \lambda_m(r) \ dr}{\sum_{i=1}^N \alpha_i^2 \int_0^{R_s} \lambda_m(r) \ dr} = \frac{1}{\sum_{i=1}^N \alpha_i^2} \overset{(c)}{\leq} N.$$

Next we specify the condition under which the upper bound is achieved. Notice that (a) holds with equality if and only if $\int_0^{\alpha_i R_s} \lambda_i(r) dr = \alpha_i^2 \int_0^{R_s} \lambda_i(r) dr$ for all i. This requires each $\lambda_i(r)$ to be a linear function. Inequality (b) is tight when $\int_0^{R_s} \lambda_i(r) dr = \int_0^{R_s} \lambda_m(r) dr$ for every $i \neq m$. Hence, all the relays must have the same linear $\lambda_i(r)$. Thus, (r_i^*) must be the uniform, hence competitive, allocation, i.e., $\alpha_i = 1/N$ for all i. This is exactly what is needed to make (c) tight.

Now it remains to find the pricing functions which can induce the monopolistic equilibrium attaining the upper bound. Let $\beta_i(r) = \lambda(R_s - r) \triangleq \beta(r)$ for every i.[9] Since $\beta_i(r)$ is strictly decreasing for every i, $\beta_{\hat{i}}(r) = \beta(r) = \lambda(R_s - r)$. Since $\beta_{\hat{i}}(R_s - r) = \lambda(r)$, every relay is indifferent to having any amount of flow. Thus, the monopolistic equilibrium can be sustained. $\square$

Unlike the selfish routing games considered in [7], [8], for which the price of anarchy is independent of the topology [17], Theorem 4 indicates that $\rho(N, (\lambda_i(\cdot)), R_s)$ of an oligopoly PG explicitly depends on topology through N. Such a conclusion implies that the more intensive (larger N) the competition is,

---

[9]Here we have omitted the subscript of $\lambda_i(\cdot)$ in light of the symmetry.

the more inefficient the market becomes if it is monopolized. The situation is even worse if the relays in an oligopoly have convex $\lambda_i(\cdot)$. In this case, the price of anarchy can be arbitrarily large.

Theorem 5: For a fixed number $N \geq 2$ of relays and for any $M > 0$, there exists an oligopoly $(N, (\lambda_i(\cdot))_{i=1}^{N}, R_s)$ with convex $(\lambda_i(\cdot))$ such that $\rho(N, (\lambda_i(\cdot)), R_s) \geq M$.

Sketch of Proof: We can construct an oligopoly with N relays such that the socially optimal routing is competitive. By Corollary 2, inefficient monopolistic equilibria exist. However, within the class of convex functions, $\lambda_m(\cdot)$ of the dominant relay m can be designed so that $\int_0^{R_s} \lambda_m(\cdot) \geq M D^*$, where D^* is the optimal cost. $\qquad\square$

E. Focal Equilibria

Although possible, inefficient equilibria in an oligopoly are very unlikely to happen. The example in Figure 5 represents a highly pathological situation. Such an equilibrium is reached only if the subtle relationships between $\beta_2(\cdot)$ and $\lambda_1(\cdot)$ and between $\beta_1(\cdot)$ and $\lambda_2(\cdot)$ are satisfied. These relationships, however, can be established only by coincidence, since relay 2 cannot observe $\lambda_1(\cdot)$ and relay 1 cannot observe $\lambda_2(\cdot)$. In a general PG, it is arguably most rational for a relay to use a replicating response as described in Section III-A.

Definition 5: A focal equilibrium of a general pricing game is an equilibrium where every relay i adopts the replicating response to its local information $\boldsymbol{L}_i = ((r_h, (B_j^h(\cdot))_{j \in \mathcal{S}_i^h})_{h \in \mathcal{P}_i}, (B_k^i(\cdot))_{k \in \mathcal{O}_i})$, i.e., for all $h \in \mathcal{P}_i$,

$$B_i^h(t) \begin{cases} = B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - t), & t \in [0, f_{hi}^*] \\ \geq B_{\hat{i}}^h(r_h) - B_{\hat{i}}^h(r_h - t), & t \in (f_{hi}^*, r_h] \end{cases}, \tag{10}$$

where $(B_{\hat{i}}^h(\cdot), f_{hi}^*)_{h \in \mathcal{P}_i}$ are as specified in Lemma 1.[10]

In this section, we investigate focal equilibria in oligopolies. Such equilibria are not only reasonable for implementation, but also, more importantly, *always* efficient. The next theorem establishes the existence of focal equilibria in an oligopoly.

Theorem 6: The socially optimal routing of an oligopoly is always induced by a focal equilibrium.

---

[10]Since the derivative $\beta_{\hat{i}}^h(\cdot)$ of $B_{\hat{i}}^h(\cdot)$ is in general piecewise continuous, we henceforth allow the derivative $\beta_i^h(\cdot)$ of $B_i^h(\cdot)$ to be piecewise continuous. Let $\beta_i^h(r)^-$ and $\beta_i^h(r)^+$ denote the left and right limits of $\beta_i^h(\cdot)$ at r.

Proof: Note that the equilibrium constructed in the proof of Theorem 1 is a focal equilibrium. $\qquad\square$

Figure 4 illustrates a focal equilibrium in a duopoly induced by nonlinear pricing functions $B_1(\cdot), B_2(\cdot)$. The linear pricing equilibrium used in the proof is a special case of Figure 4 such that the curves of $\beta_1(r)$ and $\beta_2(r)$ are the same horizontal line that passes the point where $\lambda_1(r)$ and $\lambda_2(R_s - r)$ intersect. Notice that the focal equilibria encompassed by the above example have a common property. That is, the curves $\beta_i(\cdot)$ and $\lambda_i(\cdot)$ of all i intersect at the point that corresponds to the social optimum. The next theorem states that such a phenomenon is no coincidence.

Theorem 7: Every focal equilibrium of an oligopoly is efficient.

Proof: In light of Theorem 2, we only need to show that any monopolistic focal equilibrium is efficient. Let $(\beta_i(\cdot))$ be the pricing functions that induce such an equilibrium where m is the dominant relay. By the best response condition (9) and the fact that $\beta_m(\cdot)$ is a replicating response to $\beta_{\hat{m}}(\cdot)$, at $f_m^* = R_s$,

$$\lambda_m(R_s) \leq \beta_{\hat{m}}(0)^+ = \beta_m(R_s)^-.$$

Applying (9) and Lemma 2 to any $j \neq m$, we have

$$\lambda_j(0) \geq \beta_{\hat{j}}(R_s)^- \geq \beta_m(R_s)^-.$$

Therefore, $\lambda_m(R_s) \leq \lambda_j(0)$ for any $j \neq m$, which implies that the monopolistic routing is efficient. $\qquad\square$

To summarize, as the most reasonable outcomes of a pricing game, focal equilibria always exist and are always efficient in oligopolies. We have yet to find out whether these properties hold in general PGs. In the remainder of the paper, we focus on the class of focal equilibria when we study pricing games in multi-hop networks.[11] For brevity, we will drop the qualifier "focal" henceforth.

IV. EQUILIBRIA IN GENERAL PRICING GAME

In this section, we consider a general multi-hop relay network with one source-destination pair as described in Section II. As in Section III, we assume that in every local competition, a relay i declares $\beta_i^h(\cdot) = p_i^h(\cdot) + d_{hi}(\cdot)$ to an $h \in \mathcal{P}_i$ and all $j \in \mathcal{S}_i^h$.

Notice that if $h \in \mathcal{P}_i$ has $r_h = 0$, then technically any $\beta_i^h(\cdot)$ is a best response since the local competition involving i and $j \in \mathcal{S}_i^h$ is vacuous. To prevent absurd equilibria resulting from such arbitrary pricing,

---

[11]We deliberately ignore the type of inefficient equilibria discussed in Section III-C because there is no new discovery we can make about them in the general PG. They are inefficient in oligopolies, and therefore inefficient in general PGs.

however, we assume that i uses *honest pricing* $\beta_i^h(t) = d_{hi}(t) + d_i(t + \sum_{h' \in \mathcal{P}_i \setminus h} f_{h'i}^*)$ when $r_h = 0$. Here, $f_{h'i}^*$ is the flow i intends to acquire from $h' \in \mathcal{P}_i \setminus h$,[12] $d_i(t + \sum_{h' \in \mathcal{P}_i \setminus h} f_{h'i}^*)$ is the derivative of the minimum cost incurred by i in forwarding traffic to its offsprings. Thus, $\beta_i^h(t)$ exactly matches the actual cost of i for forwarding flow from h. Honest pricing, though restrictive, is in line with i's self interest. Being honest with its own cost, i maximally alleviates the burden of h, whose cost is partly leveraged by $\beta_i^h(\cdot)$. Therefore, honest pricing can be seen as the best effort by i to help improve the competitiveness of h, in the hope of earning profit from h should h later receive positive r_h from its predecessors.

We will frequently use the following terms. A *path* is a concatenation of links from s to w, while a *sub-path* is a contiguous segment of a path from a relay to w. Given a routing $(f_{ij})_{(i,j) \in \mathcal{E}}$, a path/sub-path is said to have positive flow if $f_{mn} > 0$ for every (m, n) in that path/sub-path. Otherwise, the path/sub-path is said to have zero flow. The *marginal cost* on a path/sub-path is the sum of $d_{mn}(f_{mn})$ over all (m, n) on that path/sub-path.

Theorem 8: The socially optimal routing of a general PG can be induced by an equilibrium.

Theorem 8 is proved by constructing pricing functions such that every $\beta_i^h(\cdot)$ is constant and is equal to the minimum marginal cost on any sub-path from h to w. Showing that the pricing functions constitute an equilibrium is not a trivial generalization of the proof for Theorem 1. For details, please refer to [18].

A. Inefficient Equilibria and Price of Anarchy

Unlike the oligopoly case, in a general PG, not all focal equilibria are efficient. The inefficiency of an equilibrium in general PG's is caused not only by the manipulative behavior of dominant relays but also by the *myopia* of nodes. We illustrate this point by the game depicted in Figure 6. The derivative of link
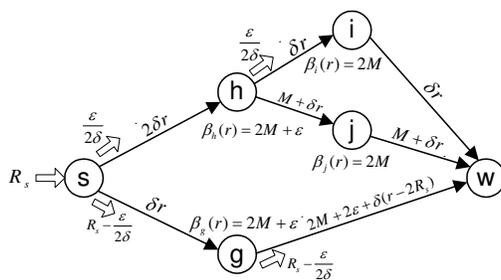

Fig. 6.   Arbitrarily bad equilibrium of a general PG.

cost functions are marked above each link, e.g. $d_{gw}(r) = 2M + 2\varepsilon + \delta(r - 2R_s)$, where M, ε and δ are

---

[12]Node i need not use honest pricing for $h' \neq h$ if h' has positive incoming flow.

positive constants such that $M \gg \varepsilon R_s$ and $M \gg \delta R_s$. The pricing function of each node is marked above the node. There are three paths from s to w, of which (s, h, i, w) has the smallest marginal cost $4\delta r$ even when $r = R_s$. So the socially optimal routing should allocate R_s entirely to the path (s, h, i, w). However, the equilibrium shown in Figure 6 leads to only $\varepsilon/(2\delta)$ being routed on (s, h, i, w) while the rest is routed on (s, g, w). In fact, s is indifferent among all allocations of R_s to h and g since $\beta_h(\cdot) = \beta_g(\cdot) \equiv 2M + \varepsilon$. Figure 7 explains why such $\beta_h(\cdot)$ and $\beta_g(\cdot)$ are h and g's best responses in their competition. Notice that
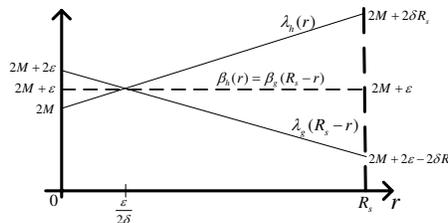

Fig. 7. Competition between h and g.

h is able to win only $\varepsilon/(2\delta)$ of the total flow because it has cost $\lambda_h(r) = d_{sh}(r) + d_h(r) = 2M + 2\delta r$. This inflated cost is a consequence of i's myopic pricing. Since i has superior cost $(\lambda_i(r) = 2\delta r)$ relative to j $(\lambda_j(r) = 2M + 2\delta r)$, i can afford to match j's pricing function $\beta_j(\cdot) \equiv 2M$. Neither i nor j has any incentive to deviate from $2M$ since i has acquired all the flow $\varepsilon/(2\delta)$ while making the maximum possible profit while j would suffer a loss if it tried to win a positive share by bidding lower than $2M$. Although i could have made more profit if it cut its price, thereby making h more competitive, it is unable to discover this opportunity as it lacks "global vision".

To conclude, although focal equilibria of a general PG rule out the manipulative pricing by a superior relay (cf. Sec.III-C), the equilibria are susceptible to the inefficiency caused by the dominant relays' myopic pricing. Such a source of inefficiency is intrinsic to networks consisting of selfish nodes who are aware of their neighbors only. The price of anarchy caused by myopic inefficiency can be arbitrarily large. For the example in Figure 6, the equilibrium holds for all large enough M and results in a total cost

$$D^E = \int_0^{\frac{\varepsilon}{2\delta}} 4\delta r \; dr + \int_0^{R_s - \frac{\varepsilon}{2\delta}} 2(M + \varepsilon - \delta R_s) + 2\delta r \; dr = \frac{\varepsilon^2}{2\delta} + 2(M + \varepsilon - \delta R_s)\left(R_s - \frac{\varepsilon}{2\delta}\right) + \delta\left(R_s - \frac{\varepsilon}{2\delta}\right)^2,$$

whereas the optimal cost is

$$D^* = \int_0^{R_s} 4\delta r \; dr = 2\delta R_s^2.$$

Therefore, the price of anarchy is at least D^E/D^*, which can be made arbitrarily large by increasing M.

Although in a multi-hop network equilibria can be arbitrarily inefficient, we will show in the following that there is a class of equilibria which are always efficient.

B. Everywhere Competitive Equilibria

Definition 6: An equilibrium of a general PG is everywhere competitive if it induces a routing $(f_{ij})_{(i,j)\in\mathcal{E}}$ such that $f_{hi} > 0$ for at least two $i \in \mathcal{O}_h$ whenever $r_h > 0$ and $w \notin \mathcal{O}_h$.[13]

Notice that the equilibrium in Figure 6 is not everywhere competitive as i is dominant to j. One would expect that when no relay is unrivalled, mistakes such as the one made by i could be avoided. The next theorem validates this intuition.

Theorem 9: If an equilibrium of a general PG is everywhere competitive, it must be efficient.

Limited by space, we explain the intuitive reasoning behind the theorem. Similar to the situation at a competitive oligopoly equilibrium, we can show that any relay that gets positive flow from its predecessor must be competitive in the sense that its marginal price is equal to its marginal cost at the equilibrium (cf. Proof of Theorem 2). Moreover, if a relay has a predecessor with zero incoming flow, its marginal price still must equal its marginal cost, as required by the honest pricing principle. Finally notice that at the equilibrium, a path/sub-path has positive flow only if it has a marginal price less than or equal to alternative paths/sub-paths. Therefore, such a routing must be socially optimal as marginal prices at an everywhere competitive equilibrium coincide with marginal costs.

V. PRICING GAME WITH ELASTIC SOURCE

So far we have assumed that s has a fixed, inelastic demand. In this section, we show that pricing games with an elastic source can be studied within the same framework we have developed for the inelastic case.

We consider a source s with elastic traffic demand. The source's preference over different admitted rates r_s is measured by a utility function $U_s(r_s)$ such that $U_s(r_s) = U_s(R_s)$ for all $r_s \geq R_s$. In other words, R_s is the maximum desired service rate of s. In the interval $[0, R_s]$, $U_s(\cdot)$ is assumed to be strictly increasing, concave with continuous derivative $u_s(\cdot)$. Taking the approach of [19], we define the *overflow rate* as $f_{sw} \triangleq R_s - r_s$. Thus, at s we have

$$\sum_{i\in\mathcal{O}_s} f_{si} + f_{sw} = R_s. \tag{11}$$

Let $D_{sw}(f_{sw}) \triangleq U_s(R_s) - U_s(r_s)$ denote the utility loss to s resulting from having a rate of f_{sw} rejected from the network. Equivalently, if we imagine that the blocked flow f_{sw} is routed on a *virtual overflow link* directly from s to w [19], then $D_{sw}(f_{sw})$ can simply be interpreted as the cost incurred on the overflow link when its flow rate is f_{sw}. Moreover, as defined, $D_{sw}(f_{sw})$ is strictly increasing,

---

[13]Recall that we have assumed in Section II-A that $|\mathcal{O}_h| \geq 2$ if $w \notin \mathcal{O}_h$.

continuously differentiable, and convex in f_{sw} on $[0, R_s]$. Denote the derivative of $D_{sw}(\cdot)$ by $d_{sw}(\cdot)$. Thus, we can treat the pricing game with an elastic source as one with an inelastic source and an overflow link (s, w). An oligopoly pricing game with an overflow link is illustrated in Figure 8, where the overflow link (s, w) is represented by a dashed arrow. Such an oligopoly is essentially the same as those studied
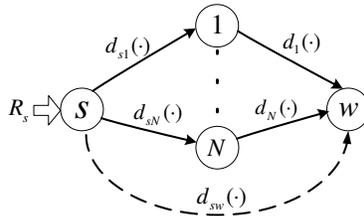

Fig. 8. Oligopoly with an overflow link.

in Section III with the exception that s now has the additional option of sending traffic on link (s, w). From a pricing perspective, we can think of w as directly competing with relays by using a *uniformly-zero pricing function*.

In a general pricing game, the introduction of the overflow link affects only the local competition faced by $i \in \mathcal{O}_s$. Now w becomes a new competitor to all $i \in \mathcal{O}_s$ whose presence changes each i's perception of the competition. Specifically, the pricing function of i's virtual competitor is derived as

$$B_{\hat{i}}^s(r) = \min \sum_{j \in \mathcal{S}_i^s} B_j^h(f_{sj}) + D_{sw}(f_{sw}),$$

where the minimization is taken with respect to nonnegative $((f_{sj})_{j \in \mathcal{S}_i^s}, f_{sw})$ such that $\sum_j f_{sj} + f_{sw} = r$. The conclusions for pricing games with an elastic source are almost verbatim to those for inelastic pricing games. Limited by space, we do not elaborate further.

VI. CONCLUSION

This work presented a game-theoretic analysis of price competition in multi-hop relay networks. The introduction of possibly nonlinear pricing functions to the game enabled us to develop a much richer set of results than if we allowed only constant unit prices. While the socially optimal routing can always be induced by an equilibrium, the game may have inefficient equilibria as well. Furthermore, the existence of competition turns out to be a two-sided coin. On the one side, any competitive equilibrium in oligopoly pricing games and any everywhere competitive equilibrium in general pricing games must be efficient. On the other side, the conclusion that the price of anarchy of an oligopoly is equal to the number of competitors seems to suggest that more intense competition only makes inefficient (monopolistic) equilibria

even worse. Unlike the case of oligopolies, the inefficiency of equilibria in a general pricing game can be attributed not only to manipulative pricing by dominant relays, but also more fundamentally, to the myopia of dominant relays. We showed that the price of anarchy attributed to both the monopolistic and myopic effects is unbounded.

REFERENCES

[1] O. Ileri, S. Mau, and N. Mandayam, "Pricing for enabling forwarding in self-configuring ad hoc networks," *IEEE Journal on Selected Areas in Communications*, vol. 23, no. 1, pp. 151–162, 2005.

[2] A. Blanc, Y. Liu, and A. Vahdat, "Designing incentives for peer-to-peer routing," in *Proceedings of IEEE INFOCOM 2005*, vol. 1, Mar. 2005.

[3] J. Crowcroft, R. Gibbens, F. Kelly, and S. Östring, "Modelling incentives for collaboration in mobile ad hoc networks," *Performance Evaluation*, vol. 57, no. 4, pp. 427–439, 2004.

[4] P. Marbach and Y. Qiu, "Cooperation in wireless ad hoc networks: a market-based approach," *IEEE/ACM Transactions on Networking*, vol. 13, pp. 1325–1338, Dec. 2005.

[5] F. Kelly, A. Maulloo, and D. Tan, "Rate control in communication networks: shadow prices, proportional fairness and stability," *Journal of the Operational Research Society*, vol. 49, 1998.

[6] R. Cole, Y. Dodis, and T. Roughgarden, "Pricing network edges for heterogeneous selfish users.," in *Proceedings of the 35th Annual ACM Symposium on Theory of Computing*, pp. 521–530, Jun. 2003.

[7] T. Roughgarden, E. Tardos "How bad is selfish routing," *Journal of the ACM*, vol. 49, no. 2, pp. 236-259, 2002.

[8] T. Roughgarden, "Selfish routing with atomic players," in *Proceedings of the sixteenth annual ACM-SIAM symposium on Discrete algorithms*, pp. 1184–1185, 2005.

[9] J. Shu and P. Varaiya, "Pricing network services," in *Proceedings of IEEE INFOCOM 2003*, vol. 2, Mar. 2003.

[10] L. He and J. Walrand, "Pricing differentiated internet services," in *Proceedings of IEEE INFOCOM 2005*, vol. 1, Mar. 2005.

[11] T. Basar and R. Srikant, "Revenue-maximizing pricing and capacity expansion in a many-users regime," in *Proceedings of IEEE INFOCOM 2002*, vol. 1, 2002.

[12] L. He and J. Walrand, "Pricing and revenue sharing strategies for internet service providers," in *Proceedings of IEEE INFOCOM 2005*, vol. 1, Mar. 2005.

[13] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple isps," in *Proceedings of IEEE INFOCOM 2005*, vol. 1, Mar. 2005.

[14] D. Acemoglu and A. Ozdaglar, "Competition in parallel-serial networks." To appear in *IEEE Journal of Selected Areas of Communication, Special Issue on Non-Cooperative Behavior in Networking*, 2006.

[15] R. Wilson, *Nonlinear Pricing*. Oxford University Press Inc, USA, 1993.

[16] B. Bernheim and M. Whinston, "Menu auctions, resource allocation, and economic influence," *Quarterly Journal of Economics*, vol. 101, no. 1, pp. 1–32, 1986.

[17] T. Roughgarden, "The price of anarchy is independent of the network topology," *Journal of Computer and System Sciences*, vol. 67, no. 2, pp. 341–364, 2003.

[18] Y. Xi and E. Yeh, "Routing game with selfish and strategic pricing," technical report, Dept. of Electrical Engineering, Yale University, New Haven, CT, July 2007.

[19] D. P. Bertsekas and R. Gallager, *Data Networks*. Prentice Hall, second ed., 1992.