Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI 2015)

Strategic Network Formation through an Intermediary*

Elliot Anshelevich and **Onkar Bhardwaj** and **Koushik Kar**
Rensselaer Polytechnic Institute
110 8th Street, Troy NY 12180
eanshel@cs.rpi.edu, bhardo@rpi.edu, koushik@ecse.rpi.edu

Abstract

Settings in which independent self-interested agents form connections with each other are extremely common, and are usually modeled using network formation games. We study a natural extension of network formation games in which the nodes cannot form connections themselves, but instead must do it through an *intermediary*, and must pay the intermediary to form these connections. The price charged by the intermediary is assumed to be determined by its operating costs, which in turn depend on the total amount of connections it facilitates. We investigate the existence and worst-case efficiency (price of anarchy) of stable solutions in these settings, and especially when the intermediary uses common pricing schemes like proportional pricing or marginal cost pricing. For both these pricing schemes we prove existence of stable solutions and completely characterize their structure, as well as generalize these results to a large class of pricing schemes. Our main results are on bounding the price of anarchy in such settings: we show that while marginal cost pricing leads to an upper bound of only 2, i.e., stable solutions are always close to optimal, proportional pricing also performs reasonably well as long as the operating costs of the intermediary are not too convex.

1 Introduction

Settings in which independent self-interested agents form connections with each other are extremely common, and range from computer networks to social networks to economic networks. Such settings are usually studied using network formation games: a rich category of games which studies properties of networks resulting from agents (nodes) forming relationships (edges) in a strategic manner to maximize their utility (or minimize their cost). The extensive body of work on network formation games (see [Jackson, 2005] for a survey) looks at many different settings, different notions of node utility (e.g. [Jackson and Wolinsky, 1996; Fabrikant *et al.*, 2003]), different concepts of network stability (e.g., [Epstein *et al.*, 2007; Calvó-Armengol and İlkılıç, 2009]), convergence dynamics (e.g., [Dutta *et al.*, 2005; Derks *et al.*, 2008]), etc. For example, the classic work of [Jackson and Wolinsky, 1996] considers an abstract setting in which both agents must agree to form a connection or link between them, but a single agent can break this connection; this work was later extended in many works such as [Bala and Goyal, 2000; Watts, 2003] and forms the foundation of the network formation game which we consider in this work.

The above work on network formation games assumes that agents can form links (connections, relationships, etc.) themselves. The key question that we study in our work, however, is what happens when the agents cannot form connections themselves, and must instead do it by paying an *intermediary.* This occurs, for example, when traders exchanging goods need to pay fees to the freight companies, and in many other settings (see below) when agents are paying some centralized service for their connections, instead of forming connections themselves. If the intermediary charges exorbitant fees which exceed the benefits of forming connections, then nodes would rather not establish connections at all. On the contrary, the work mentioned above can be interpreted as the intermediary being absent, or connections having fixed costs. We are interested in the natural case when the fees charged by the intermediary are determined by its operating costs.

A major motivation for our work comes from Internet Service Providers (ISPs) forming connections via Internet Exchange Points (IXPs). In most basic terms, IXPs are data centers with huge network switches through which ISPs form connections with each other in order to exchange traffic. In return, IXPs recover their operating costs (which depend on the size of infrastructure needed) by charging fees to each member ISP. The recurring component of these fees typically depends on the capacity of the port(s) allocated to each ISP, see e.g., [LINX, 2015; DEC-IX, 2015]. This interaction between ISPs and IXPs is extremely important for the modern Internet: a considerable portion of Internet traffic [Ager *et al.*, 2012; Cardona Restrepo and Stanojevic, 2012] passes through an IXP, and thus requires two ISPs to form a connection with each other by paying an IXP to forward their traffic. In essence an IXP can be thought of as a marketplace in which the seller (the IXP) sells capacity for forming connections to its clients (the ISPs): this is exactly the setting of

*The work was partially supported by NSF awards CCF-1101495 and CNS-1218374.

network formation through an intermediary that we study in this work (see Related Work for more details).

Model Summary. We now outline the model for our network formation game with an intermediary. We assume that the cost of each agent (i.e., node) consists of two components: the connection cost incurred due to all the connections it forms with its neighbors, and the payment made to the intermediary. For any pair of nodes we assume that as the strength of their connection increases, the connection cost incurred by each endpoint due to this connection decreases (e.g., as a pair of ISPs exchange traffic at a higher rate, the losses due to end-customer dissatisfaction decrease; as people form a stronger relationship, they get more benefit from this relationship, etc). Our concept of a stable solution in this setting is a natural extension of the notion of pairwise stability from [Jackson and Wolinsky, 1996]: Using y_{ij} to denote the connection strength between a pair of nodes (ij), we say that a solution is stable if no pair of nodes (ij) is able to increase the strength y_{ij} of their connection in a mutually beneficial way (i.e., to lower the cost for both of them), and no node benefits by reducing the strength of any of its connections. Note that how changing the strength of a connection affects the cost of a node depends not only on the change in the connection cost, but also on the change in the payment made to the intermediary.

We are particularly interested in the case where the payments made by each node to the intermediary are determined by the operating costs of the intermediary. Motivated partly by ISP-IXP interactions, where the operating cost of an IXP depends on the size of the infrastructure needed to carry the traffic of all ISPs, we assume the operating cost of an intermediary to be a function $c(y)$, where y is the total strength of connections formed by all the nodes. Given this, we will focus on the following two popular notions of pricing:

- *Proportional pricing*: Let y_i denote the total strength of the connections formed by a node i. In proportional pricing, the operating cost $c(y)$ of the intermediary is split among the nodes such that the payment made by a node i to the intermediary is proportional to y_i, i.e., node i pays $\frac{y_i}{y}c(y)$.
- *Marginal cost pricing*: In marginal cost pricing, node i makes a payment of $c'(y) \cdot y_i$ to the intermediary.

We will also briefly discuss *Equal pricing* where each node makes a fixed payment $c(y)/n$ to the intermediary (where n is the number of nodes), regardless of connection strengths.

Our goal in this work is to investigate the existence and quality of stable solutions in which payments to the intermediary conform to the above pricing schemes. Towards measuring the quality of a solution, we define the *social cost* of a solution as the sum of costs of all the nodes and the intermediary (the cost of the intermediary is its operating cost minus the total payment it receives). Thus the optimal solution for a given instance is the solution which minimizes the social cost (note that since payments cancel out, the optimal solution does not depend on the pricing schemes). We use the popular notion of *Price of Anarchy* (PoA) to measure the worst-case efficiency of a pricing scheme: the PoA of a pricing scheme is the maximum possible ratio of the social cost of a stable solution (with this pricing scheme) to the social cost of the optimal solution.

Our Contributions. We define network formation games with an intermediary, which are a natural generalization of the basic model in [Jackson and Wolinsky, 1996] as discussed in Related Work. We investigate the existence and PoA of stable solutions in which payments made to the intermediary satisfy well-known notions of proportional pricing and marginal cost pricing. We show that despite requiring the payments to satisfy these pricing schemes, stable solutions always exist for both pricing schemes. In fact, we show a stronger result that stable solutions exist for any pricing scheme in which the payments made by any node u are of the form $p(y) \cdot y_u$, where $p(y)$ is any continuous and increasing function. We also provide a characterization of all stable solutions for any such pricing scheme, in which we show that any stable solution can be generated recursively from other stable solutions. Our main results are on bounding the worst-case efficiency (PoA) of the above pricing schemes. We show that under marginal cost pricing PoA is upper bounded by 2, and thus all stable networks of connections are close to optimum in cost. We also show that under proportional pricing PoA is upper bounded by $\max(2, \sup_{y \geq 0} \frac{yc'(y)}{c(y)})$. For example, under a common assumption that the operating cost $c(y)$ of the intermediary is a convex polynomial of degree at most d, this bound evaluates to at most d, and thus proportional pricing becomes very reasonable as long as the operating costs of the intermediary are not too convex.

2 Related Work

Network formation games have been extensively studied in the literature under many different settings, see [Jackson, 2005; Page and Resende, 2013; Tardos and Wexler, 2007] and the references therein.

Our work is most heavily influenced by the classic network formation game introduced in [Jackson and Wolinsky, 1996]. This work considers a setting in which two nodes can form links with mutual consent (if it benefits both of them) but any node can sever the connection if it hurts its utility. Our network formation game with intermediary (including the notion of stability) can be viewed as a natural extension of this abstract setting in the presence of an intermediary. Note that we also allow connections with fractional strengths instead of just integral strengths $\{0, 1\}$, which is more natural for both networks of Internet traffic and social networks. Whereas [Jackson and Wolinsky, 1996] investigate the structure of the stable solutions and the conditions for which the stable solutions are also optimal, we give bounds on inefficiency of stable solutions with intermediary prices satisfying the well-known notions of proportional and marginal cost pricing. The notions of pairwise stability and price of anarchy have also been used in many works related to network formation games such as [Fabrikant *et al.*, 2003; Calvó-Armengol and İlkılıç, 2009; Corbo and Parkes, 2005; Anshelevich and Hoefer, 2012].

A major motivation for our work is how internet service providers (ISPs) form connections through internet exchange points (IXPs). A typical IXP architecture corresponds to an ethernet switch through which member ISPs can exchange traffic [Augustin *et al.*, 2009; Ager *et al.*, 2012]. In return, each member ISP makes a payment to the IXP of which the recurring charges are determined by the port capacities requested by ISPs: see the following sample fee structures [LINX, 2015; AMS-IX, 2015]. Then each member ISP can form connections with a subset of other ISPs in order to exchange traffic [Ager *et al.*, 2012]. A large number of IXPs, especially in Europe, are operated on a non-profit basis [Ryan and Gerson, 2012] and the operating costs of an IXP are largely determined by the cost of the infrastructure needed for traffic exchange [EuroIX, 2015].

Our setting is generally part of the enormous body of work on markets and pricing (see e.g., [Johari, 2007; Vazirani, 2007] and the reference therein). Unlike this work, however, we consider a seller who is selling pairwise connections instead of items to individual buyers, and thus must analyze pairwise equilibrium solutions with all of the complexities that this entails.

3 Model and Preliminaries

Let there be n nodes. Denote by y_{ij} the strength of the connection between nodes i and j. Let $y_i = \sum_{(ij)\ni i} y_{ij}$ denote the total strength of connections formed by node i. In our primary motivation of the ISP-IXP example, y_{ij} is analogous to the traffic between two ISPs and y_i is analogous to the total traffic exchanged by an ISP. The cost of each node is assumed to be the sum of the following two components:

- Connection cost: As discussed above, we assume that for any pair (ij) of nodes as y_{ij} increases, the connection cost incurred by each endpoint due to this connection decreases. We model the connection cost incurred due to connection (ij) as $\lambda_{ij}(B_{ij} - y_{ij})$, which is a linear function of the connection strength y_{ij}. This connection cost is incurred by both end points. Here B_{ij} is the budget or maximum allowed connection strength for the pair (ij). For the ISP-IXP setting, $\lambda_{ij} \cdot B_{ij}$ is analogous to the benefit obtained by each ISP after exchanging a total traffic B_{ij} committed to the end-users. Then $\lambda_{ij}(B_{ij} - y_{ij})$ is analogous to the losses incurred by each ISP for failing to deliver the committed amount of traffic.
- Payment to the intermediary: We will focus on proportional pricing and marginal cost pricing as discussed in the Introduction. Both these pricing schemes can be expressed as $\gamma \cdot y_i$ where $\gamma = c(y)/y$ for proportional pricing and $\gamma = c'(y)$ for marginal pricing (recall that $c(y)$ denotes the operating cost of the intermediary for a total connection strength of y). This can be also interpreted as each node paying a price γ per unit connection strength. This thinking is consistent with the observation that typically an IXP charges an ISP based on the capacity of allocated port(s), and thus the charges are proportional to y_i. We will call γ as the **price** of the intermediary.

Let $\vec{y}$ denote the vector of connection strengths y_{uv} of all the pairs, which we will also refer to as *allocation vector*. We will call a tuple $(\vec{y}, \gamma)$ of an allocation vector $\vec{y}$ along with an intermediary price γ as a *solution*. Based on the above discussion, the cost of a node i in a solution $(\vec{y}, \gamma)$ is given by

$$C_i(\vec{y}, \gamma) \quad = \quad \gamma \cdot \sum_{(ij)\ni i} y_{ij} + \sum_{(ij)\ni i} \lambda_{ij} \cdot (B_{ij} - y_{ij}) \qquad (1)$$

We assume the operating cost $c(y)$ of the intermediary to be a strictly convex function of total connection strength $y = \sum_{i \neq j} y_{ij}$, with $c'(0) = c(0) = 0$. Taking into account the total payments from the nodes, the cost of the intermediary is then given by

$$X(\vec{y}, \gamma) = c(y) - \gamma \cdot \sum_i \sum_{(ij)\ni i} y_{ij} \qquad (2)$$

We define the social cost of an allocation vector $\vec{y}$ as the sum of the costs of the nodes and the intermediary. Note that since the payments cancel out, this summation is a function of only $\vec{y}$. Thus the social cost of an allocation vector $\vec{y}$ is given by

$$SC(\vec{y}) \quad = \quad c(y) + \sum_i \sum_{(ij)\ni i} \lambda_{ij} \cdot (B_{ij} - y_{ij}) \qquad (3)$$

Drawing a parallel to the capacity of the physical link connecting ISP i to the IXP, we will assume a capacity of L_i for each node, i.e., that $y_i \leq L_i$ for a node i. In more general networks, the capacity L_i corresponds to the total amount of connections that an agent can possibly make and maintain. Thus an instance of the network formation game with an intermediary is completely specified by $(\vec{\lambda}, \vec{B}, \vec{L}, c)$ where $\vec{\lambda}$ is the vector of variables λ_{ij}, $\vec{B}$ is the vector of budget variables B_{ij}, $\vec{L}$ is the vector of node capacity constraints L_i and $c(y)$ is the cost incurred by the intermediary for maintaining a total connection strength of y.

We now describe the notion of a stable solution. In the network formation game analyzed by Jackson and Wolinsky [Jackson and Wolinsky, 1996], a stable network was defined as a network where no pair of nodes could add an edge to lower the cost of both, and no node could remove an adjoining edge to lower its own cost. Analogous to this definition (as well as similar definitions of stability in other works [Corbo and Parkes, 2005; Calvó-Armengol and İlkılıç, 2009]), we consider a solution to be stable in our setting when no (ij) pair of nodes can increase y_{ij} to lower the cost (of both of them), and no node can reduce y_{ij} to lower its own cost. We formally define below the notion of a stable solution. We use $\tau^i_{min}(\vec{y})$ to denote $\min\{\lambda_{ij} : (ij) \ni i \text{ and } y_{ij} > 0\}$, i.e., the minimum λ_{ij} among the pairs adjoining node i having a connection with positive strength.

Definition 3.1 *A solution $(\vec{y}, \gamma)$ formed by the allocation vector $\vec{y}$ and intermediary price γ is said to be a* stable solution *whenever every pair of nodes (ij) satisfies the following conditions:*

1. $y_{ij} > 0$ *implies* $\lambda_{ij} \geq \gamma$
2. $y_{ij} < B_{ij}$ *implies either* $\lambda_{ij} \leq \gamma$ *OR at least one of the endpoints, say i, satisfies* $y_i = L_i$ *and* $\lambda_{ij} \leq \tau^i_{min}(\vec{y})$.

To understand above definition, let us first consider the simpler scenario when all L_i are infinite. In this scenario, there is no limit on how much total connection strength any node is allowed to have, thus the only limiting constraint for a pair of nodes is the intermediary price and their own budget B_{ij}. Consider the following equation obtained from Equation 1:

$$\frac{\partial C_i(\vec{y},\gamma)}{\partial y_{ij}} = \gamma - \lambda_{ij} \tag{4}$$

Thus the relation between γ and λ_{ij} dictates whether increasing or decreasing the connection strength y_{ij} would benefit a pair by lowering the cost of both endpoints: if $\lambda_{ij} < \gamma$ then a pair would benefit by decreasing y_{ij} whereas if $\lambda_{ij} > \gamma$ then a pair would benefit by increasing y_{ij} (See Equation 4). Thus in a stable solution $y_{ij} > 0$ implies $\lambda_{ij} \geq \gamma$ because if $\lambda_{ij} < \gamma$ then both the endpoints would decrease y_{ij} to lower their cost. Similarly, we obtain that $y_{ij} < B_{ij}$ implies $\lambda_{ij} \leq \gamma$ for the simplified scenario of all L_i being infinite.

Introducing finite node capacity constraints L_i does not change the first condition for stability (i.e., $y_{ij} > 0$ implies $\lambda_{ij} \geq \gamma$). However with finite capacity constraints, even when increasing y_{ij} can benefit a pair (i.e., $\lambda_{ij} > \gamma$), one of the endpoints, say i, may not agree on doing so if it is using all its capacity for more beneficial connections (i.e., $y_i = L_i$ and $\lambda_{ij} \leq \tau^i_{min}(\vec{y})$). This implies that if $y_{ij} < B_{ij}$ then either $\lambda_{ij} \leq \gamma$ (same as the case with no capacity constraints) or one of the endpoints, say i, satisfies $y_i = L_i$ and $\lambda_{ij} \leq \tau^i_{min}(\vec{y})$.

To measure the quality of stable solutions, we will use the popular measure of Price of Anarchy (PoA) (see for example [Papadimitriou, 2001; Koutsoupias and Papadimitriou, 1999]). Let $\vec{y^*}$ be the allocation vector which minimizes the social cost (see Equation 3) given an instance I. Then the price of anarchy for marginal cost pricing is given by

$$\text{PoA} = \sup_{I} \sup_{\substack{(\vec{y},\gamma) \text{ is stable} \\ \text{s.t. } \gamma = c'(y)}} \frac{SC(\vec{y})}{SC(\vec{y^*})} \tag{5}$$

In other words, the PoA (for marginal cost pricing) is the maximum factor by which the social cost of a stable solution (with marginal cost pricing) can differ from the social cost of an optimal solution. We can define PoA similarly for other pricing schemes. Now we proceed to characterizing the structure of stable solutions, and later bounding their PoA.

4 Characterization and Existence of Stable Solutions

In this section, we investigate the existence and provide a characterization of stable solutions in which payments made to the intermediary satisfy well-known notions of proportional pricing and marginal cost pricing. We show that despite requiring the payments to satisfy these pricing schemes, stable solutions always exist for both pricing schemes. In fact, we show a stronger result in Theorem 4.3: that stable solutions exist for any pricing scheme in which the payments made by any node i are of the form $p(y) \cdot y_i$, where $p(y)$ is any continuous and increasing function. In Theorem 4.1, we provide a characterization of all stable solutions, which shows that any stable solution can be generated recursively from other stable solutions.

Let us introduce some notation to begin with. We will use λ^k to denote k^{th} biggest value from the set of λ_{ij} variables for a given instance. We will also use λ^0 to denote any arbitrary value greater than λ^1. For example, if there are three pairs in an instance with λ_{ij} values $\{16, 16, 15\}$ then λ^0 is any value greater than 16, $\lambda^1 = 16$ and $\lambda^2 = 15$.

By $T(\alpha)$ we will denote the set of the pairs $\{(ij) : \lambda_{ij} = \alpha\}$, e.g., $T(\lambda^k)$ is the set of pairs having $\lambda_{ij} = \lambda^k$. For rest of the text, we say that a pair (ij) is **tight** whenever either $y_{ij} = B_{ij}$ or at least one of the endpoints, say i, satisfies $y_i = L_i$. Thus a pair is tight whenever its connection strength cannot be increased being constrained by either connection budget or node capacities.

Before we characterize the recursive structure of stable solutions in Theorem 4.1, we give some intuition behind it. Suppose we are given a stable solution $(\vec{f}, \beta)$ for intermediary price β such that $\lambda^2 < \beta < \lambda^1$. Now using this solution, let us construct a stable solution $(\vec{y}, \gamma)$ for a lower intermediary price γ such that $\lambda^3 < \gamma < \lambda^2$ as follows: Initially, let $\vec{y} = \vec{0}$. Now let us set $y_{ij} = f_{ij}$ for all the pairs in $T(\lambda^1)$. This ensures that all the pairs in $T(\lambda^1)$ satisfy the conditions for stability (by virtue of $(\vec{f}, \beta)$ being stable). Now let us iteratively increase y_{ij} for the pairs in $T(\lambda^2)$ until all such pairs are tight. We claim that this makes $(\vec{y}, \gamma)$ a stable solution. To see this, observe that none of the pairs in $T(\lambda^2)$ can violate the first condition for stability. Now let us check for the second condition of stability: All the pairs in $T(\lambda^2)$ have $\lambda_{ij} = \lambda^2 > \gamma$. Also, by making these pairs tight, we have ensured that for every pair (ij) in $T(\lambda^2)$ has at least one endpoint, say i, satisfying $y_i = L_i$ (note that i also satisfies $\lambda_{ij} \leq \tau^i_{min}(\vec{y})$ by construction).

Thus all the pairs with $\lambda_{ij} \in \{\lambda^1, \lambda^2\}$ meet both the conditions for stability, and having $y_{ij} = 0$ for the pairs with $\lambda_{ij} < \gamma < \lambda^2$ ensures that no other pair violates stability conditions. This proves that $(\vec{y}, \gamma)$ is a stable solution. It turns out that if $\gamma = \lambda^2$ then we need not ensure the tightness of all the pairs in $T(\lambda^2)$ while constructing $\vec{y}$.

It can be shown that the above approach of constructing stable solutions can be applied recursively to construct stable solutions for any price γ and vice versa every stable solution can be constructed in such a manner. This gives us the following theorem (we omit the proof for the sake of brevity).

Theorem 4.1 (Characterization) *A solution $(\vec{y}, \gamma)$ where $\gamma \in (\lambda^{k+1}, \lambda^k]$ is a stable solution if and only if there exists a stable solution $(\vec{f}, \beta)$ with $\beta \in (\lambda^k, \lambda^{k-1})$ such that*

(a) If $\gamma = \lambda^k$ then the allocation vector $\vec{y}$ is obtained by increasing f_{ij} for all the pairs in $T(\lambda^k)$ in any arbitrary manner without violating the budget and capacity constraints.

(b) If $\gamma \in (\lambda^{k+1}, \lambda^k)$ then then the allocation vector $\vec{y}$ is obtained by increasing f_{ij} for all the pairs in $T(\lambda^k)$ until all such pairs are tight.

Notice that $(\vec{0}, \lambda^0)$ is the unique stable solution for a price $\gamma = \lambda^0$ where λ^0 denotes any arbitrary value greater than λ^1. Starting with $(\vec{0}, \lambda^0)$, we can keep constructing stable solutions recursively for lower prices using Theorem 4.1, to give us the following:

Theorem 4.2 *There exists a stable solution $(\vec{y}, \gamma)$ for every constant price $\gamma \geq 0$. Moreover, whenever all λ_{ij} are distinct, there exists a unique stable solution for any constant price γ in the open interval $(\lambda^{k+1}, \lambda^k)$ for all $k \geq 1$.*

Note that Theorem 4.2 guarantees the existence of a stable solution for *any constant price* γ independent of any other parameters (with the node payments to the intermediary being of the form $\gamma \cdot y_i$). However, Theorem 4.2 *does not* guarantee the existence of a stable solution for any particular pricing scheme. For example, it could be that for a pricing scheme in which the price is a function $p(y)$ of total connection strength y (which is the case for proportional pricing with $p(y) = c(y)/y$ and marginal cost pricing with $p(y) = c'(y)$), the only stable solutions guaranteed by the above theorem may have the form $(\vec{y}, \gamma)$ with $\gamma \neq p(y)$. Fortunately, we can prove Theorem 4.3, which guarantees the existence of stable solutions for a large class of pricing schemes.

Theorem 4.3 (Existence) *There exists a stable solution of the form $(\vec{y}, p(y))$ for any continuous non-decreasing price function $p(\cdot)$ with $p(0) = 0$.*

The proof of the above theorem involves beginning with the stable solution $(\vec{0}, \lambda^0)$ (where λ^0 is any value greater than λ^1), which has a total connection strength of 0. Then we recursively construct solutions as described in Theorem 4.1, slowly increasing the total connection strength of the allocation vector in the process until the total connection strength and intermediary price satisfy the relation prescribed by the price function $p(\cdot)$. Now we describe the proof in details:

Proof. [Theorem 4.3] Recall that λ^0 denotes some arbitrary value greater than λ^1. Let λ^m denote $\min_{ij} \lambda_{ij}$ and let $\lambda^{m+1} = 0$. Let $\vec{f_0}$ denote the zero vector $\vec{0}$. Consider a sequence of solutions $\{(\vec{f_k}, \gamma^k)\}_{k=0}^m$ where $\gamma^k \in [\lambda^k, \lambda^{k+1})$ and the allocation vector $\vec{f_i}$ is obtained recursively from $\vec{f_{i-1}}$ as specified Theorem 4.1(b) by making all the pairs in $T(\lambda^k)$ tight. Observe that all these solutions are tight for $\gamma^k \in [\lambda^k, \lambda^{k+1})$ as any solution obtained in part(b) of Theorem 4.1 is also stable for intermediary price of λ^k in Theorem 4.1 part (a). We break the further analysis into two cases:
Case $p(f_m) < \lambda^m$: As $(\vec{f_m}, \gamma^m)$ is a stable solution for $\gamma^m \in [\lambda^m, \lambda^{m+1})$ where $\lambda^{m+1} = 0$, we conclude that $(\vec{f_m}, p(f_m))$ is a stable solution.
Case $p(f_m) \geq \lambda^m$: This together with $p(0) = 0$ tells us that there exists a λ^k such that $p(f_{k-1}) \leq \lambda^{k-1}$ and $p(f_k) \geq \lambda^k$. Thus we know that there exists an $\alpha \in [0, 1]$ such that $p(f_{k-1} + \alpha \cdot (f_k - f_{k-1})) = \lambda^k$ using continuity of $p(\cdot)$. This gives us $p(\vec{g}) = \lambda^k$ where $\vec{g} = \vec{f}_{k-1} + \alpha(\vec{f_k} - \vec{f}_{k-1})$. Applying Theorem 4.1 (part (a)), we get that $(\vec{g}, \lambda^k)$ is a stable solution. This together with the already known fact $p(\vec{g}) = \lambda^k$ completes the proof for this remaining case. ∎

Since we assume $c(y)$ is strictly convex with $c'(0) = c(0) = 0$, the marginal cost pricing scheme satisfies Theorem 4.3. It can also be shown that under these conditions $c(y)/y$ is increasing function of y. Thus setting $c(0)/0 = 0$ makes Theorem 4.3 hold for proportional pricing.

5 Price of Anarchy

In this section, we provide our main results on bounding the worst-case efficiency (PoA) of the above pricing schemes. Hence having a small PoA (See Equation 5) implies that the quality of all stable solutions will be close to the social optimum despite of nodes behaving in a self-interested manner.

Indeed, we show that under marginal cost pricing PoA is upper bounded by 2, and thus all stable networks of connections are close to optimum in cost. We also show that under proportional pricing PoA is upper bounded by $\max(2, \sup_{y \geq 0} \frac{yc'(y)}{c(y)})$. For example, under a common assumption that the operating cost $c(y)$ of the intermediary is a convex polynomial of degree at most d, this bound becomes at most d, making proportional pricing very reasonable as long as the operating cost of the intermediary is not too convex.

The critical step for proving both these PoA results is to prove Lemma 5.3 which states that to bound PoA, we need to consider only those instances with all pairs having equal λ_{ij} values. However, before that we will first state some elementary observations, Propositions 5.1-5.2 and Lemma 5.1-5.2 which will be used later to prove our main results.

Observation 5.1 *If a, b, c, d are positive numbers, with $b - d > 0$ and $a/b > c/d$ then $\frac{a-c}{b-d} > \frac{a}{b}$.*

Observation 5.2 *If a, b, c, d, x, y are positive numbers such that $x > y$ and $a/b > c/d$ then $\frac{a+yc}{b+yd} > \frac{a+xc}{b+xd}$*

The following propositions 5.1 and 5.2 allow us to perform minor perturbations in λ_{ij} values of a set of pairs while maintaining stability of a solution $(\vec{y}, \gamma)$. Before we state them, recall that $T(\alpha)$ denotes the set of pairs $\{(ij) : \lambda_{ij} = \alpha\}$.

Proposition 5.1 *If $(\vec{y}, \gamma)$ is a stable solution with $\gamma \leq \lambda^k$, then it is also stable in the instance obtained by changing λ_{ij} for every $(ij) \in T(\lambda^{k-1})$ to some fixed $\Lambda \in [\lambda^k, \lambda^{k-2})$.*

Proposition 5.2 *If $(\vec{y}, \gamma)$ is a stable solution with $\gamma > \lambda^k$, then it is also stable in the instance obtained by changing λ_{ij} for all $(ij) \in T(\lambda^k)$ to some fixed $\Lambda \in (0, \gamma)$.*

Lemma 5.1 *If $(\vec{y}, \gamma)$ is a stable solution with $\gamma < \lambda^1$ and $\vec{g}$ is any other allocation vector then*

$$\sum_i \sum_{\substack{(ij) \ni i \\ \text{s.t. } (ij) \in T(\lambda^1)}} B_{ij} \geq \sum_i \sum_{\substack{(ij) \ni i \\ \text{s.t. } (ij) \in T(\lambda^1)}} (2 \cdot g_{ij} - y_{ij})$$

The following lemma says that budgets B_{ij} have a simplified form in an instance that achieves the worst-case PoA.

Lemma 5.2 *Consider an instance that achieves the worst-case PoA with intermediary price function $p(\cdot)$. Then we have $B_{ij} = \max\{f_{ij}, g_{ij}\}$ where $\vec{g}$ denotes the allocation vector corresponding to the minimum social cost and $(\vec{f}, p(f))$ is a stable solution of maximum social cost in this instance.*

The following critical lemma says that worst-case PoA is achieved in an instance where all pairs have equal λ_{ij}.

Lemma 5.3 *Consider an instance that achieves the worst-case PoA with price function $p(\cdot)$ and suppose that PoA > 2. Then for all pairs (ij) we have $\lambda_{ij} = \lambda$ for some $\lambda > 0$.*

[Proof Outline]: Let $\tau_{\vec{y}}$ denote $\min\{\lambda_{ij} : y_{ij} > 0\}$ for a given allocation vector $\vec{y}$, i.e., $\tau_{\vec{y}}$ is the smallest λ_{ij} that a pair with positive connection strength have in $\vec{y}$. In the instance that achives the worst-case PoA, let $\vec{g}$ be an allocation vector corresponding to the minimum social cost and let $(\vec{f}, p(f))$ be a stable solution of maximum social cost. This means that the ratio $SC(\vec{f})/SC(\vec{g})$ equals the worst-case PoA in the instance under consideration.

The proof of the above lemma involves a series of steps, each requiring careful analysis. The first step involves proving that for the instance under consideration that achieves the worst-case PoA, $\tau_{\vec{g}} \geq \tau_{\vec{f}}$. The second step involves showing that $\lambda_{ij} \in \{\tau_{\vec{g}}, \tau_{\vec{f}}\}$ for every pair (ij). Given this, the final step involves proving $\tau_{\vec{g}} = \tau_{\vec{f}}$.

The theme of all these steps is to make use of Proposition 5.1-5.2 to create an instance by perturbing λ_{ij} values for some intelligently chosen class of pairs, while maintaining the stability of the solution $(\vec{f}, p(f))$. Later carefully applying Lemma 5.1-5.2 we show that in the newly created instance, the ratio $SC(\vec{f})/SC(\vec{g})$ strictly increases. This contradicts our initial assumption that the original instance achieves the worst-case PoA. □

Theorem 5.1 (PoA of proportional pricing) *With proportional pricing, PoA* $\leq \max\left(2, \sup_{f \geq 0} \frac{f \cdot c'(f)}{c(f)}\right)$ *where we define* $c'(0)/c(0) = 0$.

Proof. Assume PoA > 2 as otherwise there is nothing to prove. Consider an instance achieving worst-case PoA with proportional pricing. Let $\vec{g}$ be an allocation vector corresponding to the minimum social cost and let $(\vec{f}, p(f))$ be a stable solution of maximum social cost in this instance. Since we have assumed PoA > 2, this implies that Lemma 5.3 holds. This gives us λ_{ij} is the same for all pairs (ij) and furthermore $\lambda_{ij} = \tau_{\vec{g}} = \tau_{\vec{f}}$ holds for each pair. Thus we get the following equation (where we denote $\sum_i \sum_{(ij) \ni i} B_{ij}$ by B):

$$\frac{SC(\vec{f})}{SC(\vec{g})} = \frac{c(f) + \tau_{\vec{g}} \cdot (B - f)}{c(g) + \tau_{\vec{g}} \cdot (B - g)} \tag{6}$$

We break the further analysis into two cases:

(a) Suppose $g > f$: This gives us $c(f)/c(g) < 1$. This lets us apply Observation 5.1 to Equation 6 gives us

$$\frac{SC(\vec{f})}{SC(\vec{g})} = \frac{c(f) + \tau_{\vec{g}} \cdot (B - f)}{c(g) + \tau_{\vec{g}} \cdot (B - g)} \leq \frac{\tau_{\vec{g}} \cdot (B - f)}{\tau_{\vec{g}} \cdot (B - g)} = \frac{B - f}{B - g} \tag{7}$$

We now claim that whenever $g > f$, we have $\lambda^1 > \gamma$ (where $\gamma = c(f)/f$ for proportional pricing). We know that for the instance under consideration, all λ_{ij} values are equal. Furthermore, we have $\tau_{\vec{g}} \geq c'(g)$ from the optimality of $\vec{g}$ and using $g > f$ with convexity of $c(\cdot)$ also gives us $c'(g) > c'(f) > c(f)/f$. This gives us $\lambda^1 = \tau_{\vec{g}} > \gamma$. Since all the λ_{ij} are equal for the instance under consideration, this gives us $\lambda^1 = \tau_{\vec{f}} = \tau_{\vec{g}} > c(f)/f = \gamma$. Having $\lambda^1 > \gamma$ lets us apply Lemma 5.1 to give us $B \geq 2g - f$. Then we can apply Observation 5.1 to subtract $(B - (2g - f))$ from both numerator and denominator on the right hand side of Equation 7 to obtain $SC(\vec{f})/SC(\vec{g}) \leq 2$. This contradicts our initial assumption that PoA > 2.

(b) Suppose $f \geq g$: This implies $(B - f) \leq (B - g)$. This lets us use Observation 5.1 to subtract $\tau_{\vec{g}} \cdot (B - f)$ from the numerator and the denominator of the ratio in Equation 6 to give us

$$\begin{aligned} \frac{SC(\vec{f})}{SC(\vec{g})} &\leq \frac{c(f)}{c(g) + \tau_{\vec{g}} \cdot (f - g)} \\ &\leq \frac{c(g) + c'(f) \cdot (f - g)}{c(g) + \tau_{\vec{f}} \cdot (f - g)} \quad \text{... by convexity} \\ &\leq \frac{c'(f)}{\tau_{\vec{f}}} \quad \text{... by Observation 5.1} \end{aligned}$$

From the stability conditions and the definition of proportional pricing, we get $\tau_{\vec{f}} \geq \gamma = c(f)/f$. Substituting it in the above expression completes the proof of the theorem. ■

In particular, for a common assumption of $c(y)$ being a polynomial of degree d (without a constant term as we assume $p(0) = 0$), Theorem 5.1 gives us the following:

Theorem 5.2 *PoA* $\leq d$ *with proportional cost sharing whenever the operating cost of intermediary* $c(y)$ *is a polynomial of maximum degree* d *without a constant term.*

Theorem 5.3 (PoA of marginal cost pricing) *For marginal cost pricing, PoA* ≤ 2.

The proof of Theorem 5.3 (omitted for the sake of brevity) is almost similar to that of Theorem 5.1. It is worth noting that the above bound is tight, i.e., there exist stable solutions under marginal pricing which are a factor of 2 away from optimum.

We would like to mention in passing that equal pricing where each node makes a payment of $c(y)/n$ to the intermediary has an unbounded worst-case efficiency. This happens because the connection costs of pairs with high λ_{ij} can get subsidized because of the fees paid by the pairs with less beneficial connections (i.e., small $\lambda_{ij} B_{ij}$). As a result, in stable solutions the pairs with high λ_{ij} can end up setting connections of very high strength to increase the operating costs $c(y)$ in a disproportionate manner. This does not happen in proportional pricing or marginal cost pricing as the payment $p(y) \cdot y_i$ made by a node i depends on total connection strength y *as well as* the strength of the connections y_i formed by a node. We omit the actual example for space constraints.

6 Conclusion

Although different kinds of network formation games have been extensively studied in the literature, the natural setting where nodes need to pay an intermediary in order to form connections has not been explored. Based on the classic network formation game introduced in [Jackson and Wolinsky, 1996], and inspired especially by IXP pricing, we formulated a model of network formation games with intermediaries. For this model, we showed how to efficiently determine stable prices and allocations (and the fact that such prices exist) for a natural class of pricing schemes. Moreover, we proved that all stable solutions are close to optimal for marginal cost pricing (factor of 2 away), and the same is true for proportional pricing as long as the intermediary costs are not too convex.

References

[Ager *et al.*, 2012] Bernhard Ager, Nikolaos Chatzis, Anja Feldmann, Nadi Sarrar, Steve Uhlig, and Walter Willinger. Anatomy of a large european IXP. In *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, pages 163–174. ACM, 2012.

[AMS-IX, 2015] AMS-IX. Amsterdam internet exchange pricing. https://ams-ix.net/services-pricing/pricing, 2015. Retrieved on February 8, 2015.

[Anshelevich and Hoefer, 2012] Elliot Anshelevich and Martin Hoefer. Contribution games in networks. *Algorithmica*, 63(1-2):51–90, 2012.

[Augustin *et al.*, 2009] Brice Augustin, Balachander Krishnamurthy, and Walter Willinger. IXPs: mapped? In *Proceedings of the 9th ACM SIGCOMM conference on Internet measurement conference*, pages 336–349. ACM, 2009.

[Bala and Goyal, 2000] Venkatesh Bala and Sanjeev Goyal. A noncooperative model of network formation. *Econometrica*, 68(5):1181–1229, 2000.

[Calvó-Armengol and İlkılıç, 2009] Antoni Calvó-Armengol and Rahmi İlkılıç. Pairwise-stability and nash equilibria in network formation. *International Journal of Game Theory*, 38(1):51–79, 2009.

[Cardona Restrepo and Stanojevic, 2012] Juan Camilo Cardona Restrepo and Rade Stanojevic. IXP traffic: a macroscopic view. In *Proceedings of the 7th Latin American Networking Conference*, pages 1–8. ACM, 2012.

[Corbo and Parkes, 2005] Jacomo Corbo and David Parkes. The price of selfish behavior in bilateral network formation. In *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, pages 99–107. ACM, 2005.

[DEC-IX, 2015] DEC-IX. German internet exchange pricing. https://www.de-cix.net/products-services/pricing/, 2015. Retrieved on February 8, 2015.

[Derks *et al.*, 2008] Jean Derks, Jeroen Kuipers, Martijn Tennekes, and Frank Thuijsman. Local dynamics in network formation. In *Proc. Third World Congress of The Game Theory Society*, 2008.

[Dutta *et al.*, 2005] Bhaskar Dutta, Sayantan Ghosal, and Debraj Ray. Farsighted network formation. *Journal of Economic Theory*, 122(2):143–164, 2005.

[Epstein *et al.*, 2007] Amir Epstein, Michal Feldman, and Yishay Mansour. Strong equilibrium in cost sharing connection games. In *Proceedings of the 8th ACM conference on Electronic commerce*, pages 84–92. ACM, 2007.

[EuroIX, 2015] EuroIX. European internet exchange association: Starting an IXP – infrastructure and equipment. https://www.euro-ix.net/starting-an-ixp, 2015. Retrieved on February 8, 2015.

[Fabrikant *et al.*, 2003] Alex Fabrikant, Ankur Luthra, Elitza Maneva, Christos H Papadimitriou, and Scott Shenker. On a network creation game. In *Proceedings of the twenty-second annual symposium on Principles of distributed computing*, pages 347–351. ACM, 2003.

[Jackson and Wolinsky, 1996] Matthew O Jackson and Asher Wolinsky. A strategic model of social and economic networks. *Journal of economic theory*, 71(1):44–74, 1996.

[Jackson, 2005] Matthew O Jackson. A survey of network formation models: stability and efficiency. *Group Formation in Economics: Networks, Clubs, and Coalitions*, pages 11–49, 2005.

[Johari, 2007] Ramesh Johari. The price of anarchy and the design of scalable resource allocation mechanisms. *Algorithmic Game Theory*, pages 543–568, 2007.

[Koutsoupias and Papadimitriou, 1999] Elias Koutsoupias and Christos Papadimitriou. Worst-case equilibria. In *STACS 99*, pages 404–413. Springer, 1999.

[LINX, 2015] LINX. London internet exchange service fees. https://www.linx.net/service/servicefees.html, 2015. Retrieved on February 8, 2015.

[Page and Resende, 2013] Frank Page and Joana Resende. Network formation games. http://kenticoqa.uwb.edu/getattachment/business/about/research-series/network-formation-games-uwb.pdf, 2013.

[Papadimitriou, 2001] Christos Papadimitriou. Algorithms, games, and the internet. In *Proceedings of the thirty-third annual ACM symposium on Theory of computing*, pages 749–753. ACM, 2001.

[Ryan and Gerson, 2012] Patrick S Ryan and Jason Gerson. A primer on internet exchange points for policymakers and non-engineers. http://ssrn.com/abstract=2128103, August 2012.

[Tardos and Wexler, 2007] Eva Tardos and Tom Wexler. Network formation games and the potential function method. *Algorithmic Game Theory*, pages 487–516, 2007.

[Vazirani, 2007] Vijay V Vazirani. Combinatorial algorithms for market equilibria. *Algorithmic Game Theory*, pages 103–134, 2007.

[Watts, 2003] Alison Watts. A dynamic model of network formation. In *Networks and Groups*, pages 337–345. Springer, 2003.