Progress on Pricing with Peering

Eui-woong Lee*, David Buchfuhrer*, Lachlan L. H. Andrew*, Ao Tang†, Steven H. Low*
*Department of Computer Science, California Institute of Technology, Pasadena, CA 91125, USA
†Cornell University, Ithaca, NY 14853, USA

***Abstract*—This paper examines a simple model of how a provider ISP charges customer ISPs by assuming the provider ISP wants to maximize its revenue when customer ISPs have the possibility of setting up peering connections. It is shown that finding the optimal pricing is NP-complete, and APX-complete. Customers can respond to price in many ways, including throttling traffic as well as peering. An algorithm is studied which obtains a 1/4 approximation for a wide range of customer responses.**

I. Introduction

The Internet is an aggregation of a large number of networks owned by competing entities, which we generically call Internet service providers (ISPs). It is well known that economic considerations frequently override technical factors in the Internet traffic routing. In general, ISPs seek to maximize their own profit. The border gateway protocol (BGP) provides routing information, but the choice of route is determined by financial considerations as well as engineering factors.

The routes available depend on which ISPs are connected by direct links. When two unconnected ISPs want to exchange data, they must each pay a “provider” ISP to relay their data. Topology measurement results [15] show that, among 4200 or so transit ISPs in the Internet today, only 15 of them are more or less fully connected. This set of ISPs are commonly referred to as tier-1 ISPs, and act as global providers. The remaining transit ISPs either directly or indirectly rely on the tier-1 ISPs for global reachability.

Much attention has recently been paid to ecomonic aspects of interdomain routing [1], [3], [10], [14]. This paper concerns the optimal prices a (tier-1) provider ISP can charge its (tier-2) customers, taking into account the fact that its customers have the option of “peering” with other tier-2 networks by operating their own links to carry some or all of their traffic. In other words, the provider ISP tries to balance the trade-off between high prices and the number of subscribers.

Tier-2 networks are selfish, and do not necessarily peer simply because it is their combined interest to peer; it must be in the interest of both to peer. This work extends previous work [12] by considering a range of conditions under which tier-2 ISPs may decide to peer based on the prices they are charged. It also models the fact that increased prices will cause ISPs to generate less traffic, either by passing costs on to their own customers or by explicitly throttling traffic.

This problem is of interest in the networking community as a simple model that describes the basic relations among ISPs. It serves as a starting point for future extensions to the case of competing tier-1 ISPs. As we will soon see, it also has a rich structure which is of interest to theoretical computer science.

The paper is organized as follows. Section II describes our mathematical model, which captures the optimization problem faced by a provider ISP in deciding what to charge its customers. This problem is shown in Section III to be NP hard. Moreover, it is shown that no polynomial time algorithm exists to achieve an arbitrarily good approximation; that is, it is APX-hard. Section IV develops an algorithm which is guaranteed to achieve a at least 1/4 of the maximum revenue, which generalize the algorithms previously presented in [12]

II. Model and notation

Let $\mathbb{R}^+$ be the set of positive reals, and $\overline{\mathbb{R}}^+ = \mathbb{R}^+ \cup \{0\}$ be the set of non-negative reals, and let $|X|$ denote the cardinality of any set X.

With this notation, consider transit ISPs in two adjacent tiers, modeled as follows. Assume that all ISPs in the top tier cooperate to avoid a price war [14], so that they act as a single provider. This is modeled as a simple undirected graph $G = (V, E)$. Elements of $V = \{v_1, v_2, \dots, v_{|V|}\}$ are called nodes, vertices or customers, and elements of $E \subset V \times V$ are called links or edges. We also assume that the provider is directly connected to every customer $v \in V$. Let $\mu : V \to \overline{\mathbb{R}}^+$ be a price function, such that $\mu(v)$ is the price customer v pays the provider when they exchange one unit of traffic, in either direction.

When customer ISPs are charged more to send data, they typically throttle their own customers' traffic to reduce the load. This paper considers the case that the total traffic volume (per unit time) in both directions between customers u and customer v depends only on the sum of the prices they are charged. Let that total volume be $W(\mu(u)+\mu(v))$. If u and v do not peer, then the traffic goes through their providers and they have to pay a total of $(\mu(u)+\mu(v))W(\mu(u)+\mu(v))$.

If the provider charges too much, then customers will form peering arrangements, and no longer send traffic through the provider. In previous work [12], it was assumed that, for each pair u and v, customers u and v would peer if the sum of the prices charged by the provider, $\mu(u)+\mu(v)$, exceeded a particular amount, $c(u, v)$. This was interpreted as the cost per unit time of peering. However, selfish customers will be more motivated by their own costs than the joint costs to both customers. In this paper, we consider a function $\phi(a, b)$ such that users will peer if $\phi(\mu(u), \mu(v)) \geq c(u, v)$. Note that we still assume that the decision to peer is based purely on the prices, and that the provider knows the value of $c(u, v)$. Thus the revenue to the provider from the edge (u, v) is $(\mu(u) + \mu(v))W(\mu(u) + \mu(v))$ if $\phi(\mu(u), \mu(v)) \leq c(u, v)$, and zero otherwise.

The revenue of the provider is then given by the total, over all pairs of customers who do not peer, of the product of the sum of their prices times the total traffic between them. For any subset of edges $E' \in E$, let

$$\nu(\mu; E') \equiv \sum_{\substack{(u,v)\in E' \\ \phi(\mu(u),\mu(v)) \le c(u,v)}} (\mu(u)+\mu(v))W(\mu(u)+\mu(v)) \tag{1}$$

giving the total revenue from $G = (V, E)$ as

$$\nu(\mu) \equiv \nu(\mu; E). \tag{2}$$

The following mild conditions are placed on the forms ϕ and W may have

1) ϕ and W are left-continuous
2) $\phi \ge 0$, $W \ge 0$
3) $\phi(a, b) = \phi(b, a)$
4) $b \le c \Rightarrow \phi(a, b) \le \phi(a, c)$
5) $x \le y \Rightarrow W(x) \ge W(y)$
6) For any $e_k \in E$, $I_k \equiv \{a|\phi(a, 0) \le c(e_k)\}$ is computable in polynomial time.
7) $xW(x)$ has a finite number of local maxima, and we can compute the locations in polynomial time
8) $B = \sup_x xW(x)$ is known, although possibly infinite.

The following lemma, proved in the appendix, shows that assumption 6 is equivalent to being able to compute $b_k \equiv \sup(I_k) \in \overline{\mathbb{R}}^+ \cup \{-\infty, \infty\}$ in polynomial time.

Lemma 1. *For any $K \in \mathbb{R}$, the set of solutions of*

$$\phi(0, a) \le K \tag{3}$$

is the empty set, $\overline{\mathbb{R}}^+$, or $[0, b]$ for some $b \in \overline{\mathbb{R}}^+$.

Remark:

1) The “price to peer”, c, could equally well model the total cost to u and v of forgoing their connectivity.
2) This model assumes that the costs of peering are primarily on-going costs.
3) In addition to gaining revenue, the provider incurs additional cost for traffic it carries. This is not included in the current model.

□

The problem we seek to solve is to maximize $\nu(\mu)$ for a graph G:

$$\max_{\mu: V \to \overline{\mathbb{R}}^+} \quad \nu(\mu) \tag{P1}$$

with optimal value denoted $\text{opt}(G)$.

This notation is shown for a small example network in Figure 1. The circles represent ISPs, and are labeled with their prices, $\mu(v)$, while the edges are labeled with the peering costs. In this example, if $\phi(a, b) = 2\min(a, b)$, the provider obtains revenue from edges (2,3) and (3,4), but not from (1,3) since $\phi(\mu(1), \mu(3)) > c(1, 3)$. Edge (1,2), where $\phi(\mu(1), \mu(2)) = c(1, 2)$, yields revenue, but shows the discontinuity which gives this problem its combinatorial difficulty. There is negligible traffic between ISPs 1 and 4.

Problem (P1) is related to the problem of finding the maximum cut of a graph. A *cut* of a graph $G = (V, E)$ is a partition of the nodes $\{X, \overline{X}\}$ with $X \cap \overline{X} = \emptyset$, $X \cup \overline{X} = V$. The “size” of the cut of a graph with edge weights c is weighted sum of the “cut edges” with one vertex in each side of the partition, $\sum_{\{(u,v): u\in X, v\in \overline{X}\}} c(u, v)$. This paper uses two NP-complete problems related to maximum cuts. The “simple max-cut” [7], used in Section III, finds the cut of maximum size of an unweighted simple graph ($c(e) = 1$ for all $e \in E$). The “weighted max-cut”, used in Section IV-B, finds the maximum cut of an arbitrary weighted graph.

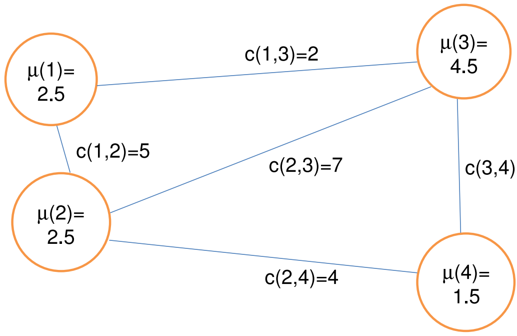

Fig. 1. An example with four ISPs

Note also that the problem can be solved by solving convex programs in the case that ϕ and $xW(x)$ are concave (or linear programs if they are also piecewise linear). As in [12], this is done by considering the maximum revenue achieved enforcing the requirement that $\phi(\mu(u), \mu(v)) < c(u, v)$ for specific subsets of edges. Once again, the number of convex programs that we need to solve grows exponentially as a function of the number of possible links. This combinatoric structure hints that the problem could be NP-hard, which will be formally shown in the next section.

III. NP-completeness

Before presenting an algorithm to solve (P1) to within a constant time, we demonstrate that it is not possible to do better, unless $P = NP$. That is, (P1) is APX-complete.

APX is the set of optimization problems which are approximable to within some constant factor in polynomial time. Under the same assumption that NP-complete problems are difficult ($P \neq NP$), there cannot even exist a polynomial time approximation scheme (PTAS) for APX-complete problems; that is, they have some constant factor within which they cannot be approximated in polynomial time unless $P = NP$. We show here APX-completeness, which implies the more usual NP-completeness but is a stronger result.

An APX-completeness proof consists of showing that the problem is in APX, and providing a polynomial time reduction from some APX-hard problem P to a target problem T, and some function $f : (0, 1] \to (0, 1]$, for which it is possible to find an α-approximation to the instance of P directly by finding an $f(\alpha)$-approximation to T.

The general problem (P1) can be shown to be APX-hard by the following reduction from MAX-CUT to the particular case when $W(x) = 1$ and $\phi(a, b) = a + b$. Note that MAX-CUT

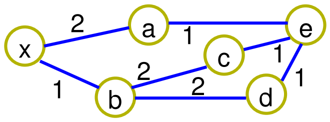

Fig. 2. The construction used in this reduction.

on 3-regular graphs is known to be APX-hard [2], whence the general MAX-CUT also is.

This result is similar to the one presented in [9] for single-minded customers buying pairs of items. In the terms of the present paper, the difference is that [9] allowed multiple edges between pairs of nodes, each with a different cost $c(e)$, while (P1) does not.

Let $d(x)$ denote the degree of node x in the MAX-CUT instance. Attach $d(x)$ copies of the construction in Figure 2 to each node x, and give weight 1 to each edge from the original graph, which we call "original edges".

The following three lemmas are proved in the appendix.

Lemma 2. *The maximum revenue achievable from the graph in Figure 2 given a fixed price for x, $0 \le \mu(x) \le 1$ is $\max(9 - \mu(x), 8 + \mu(x))$.*

Note that these bounds are achievable with $\mu(a) = \mu(c) = \mu(d) = 1$, $\mu(e) = 0$ and either $\mu(b) = 1$ or $\mu(b) = 1 - \mu(x)$. Furthermore, if $\mu(x) > 1$, then edge (x, b) cannot contribute any revenue, whence the maximum revenue is still at most 9. Furthermore, if $\mu(x) > 1$, this prevents revenue from all $d(x)$ original edges. Thus, we only consider the case $\mu(x) \le 1$.

Lemma 3. *If the optimal solution has $0 \le \mu(x) \le 1/2$, then the solution with $\mu(x)$ replaced by 0 is also optimal.*

Lemma 4. *If the optimal solution has $1/2 < \mu(x) \le 1$ and no nodes have prices $\mu(y) \in (0, 1/2]$, then the solution with $\mu(x)$ replace by 1 is also optimal.*

Combining these lemmas, from any optimal prices μ, optimal prices with $\mu(x) \in \{0, 1\}$ can be found. For such an assignment of prices, the revenue from the constructions sums up to $2|E| \cdot 9$ (where $|E|$ is the number of original edges), and the revenue from the original edges is the number of edges between a node of weight 1 and one of weight 0. Since this is the weight of a cut, the original MAX-CUT instance has a maximum cut weight of k if and only if the created graph has a maximum revenue of $18|E| + k$.

Note that if the created graph can give revenue s, the MAX-CUT instance has a cut of size at least $s - 18|E|$. Since every graph has a cut of size at least $|E|/2$, $k \ge |E|/2$. Applying this inequality, if there is a $\frac{36+\alpha}{37}$ approximation to the pricing problem, applying it and using the result to find a MAX-CUT result on a 3-regular graph gives an α-approximation. This shows that (P1) is APX-hard. APX-completeness will follow from the main result of this paper, Theorem 7, which establishes (P1) can be approximated to within a factor of 1/4 in polynomial time.

IV. Approximation algorithms

Since (P1) is APX-hard, it is useful to investigate constant-ratio approximation algorithms. This section presents an algorithm which obtains a 1/4 approximation by applying an approximate max-cut algorithm to a related graph.

A. Bounding the maximum revenue

Following [6] and [12], bounds are found in terms of an upper bound $F(V)$, defined as follows. For any $v \in V$ and any $t \in \overline{\mathbb{R}}^+$, define $\lambda_{v,t} : V \to \mathbb{R}$, by $\lambda_{v,t}(v) = t$ and $\lambda_{v,t}(u) = 0$ for all $u \ne v$. Then let $f(v)$ be the supremum revenue obtainable from ISP v, and $g(v)$ be the corresponding price charged to v. That is,

$$\begin{aligned} f(v) &= \sup\{\nu(\lambda_{v,t}) \mid t \in \overline{\mathbb{R}}^+\} && (4)\\ g(v) &= \inf\{t \in \overline{\mathbb{R}}^+ \mid \nu(\lambda_{v,\tau}) \to f(v) \text{ as } \tau \to t\}. && (5)\end{aligned}$$

It is possible that $f(v)$ or $g(v)$ is infinite, although this would not correspond to realistic functions ϕ and W. An algorithm to compute $f(v)$ and $g(v)$, and to determine if $\operatorname{opt}(G)$ is finite, is deferred until Section V. Nonetheless, the results of this section are valid, using arithmetic on $\overline{\mathbb{R}}^+ \cup \{\infty\}$.

Denote the edges from which v could get revenue when its price is $\mu(v)$ by

$$E_\mu(v) = \{(u, v) \in E \mid \phi(0, \mu(v)) \le c(u, v)\} \qquad (6)$$

and note that, taking $\mu = g$,

$$f(v) = g(v)\, W(g(v))\, |E_g(v)|\,. \qquad (7)$$

For any subset $U \subseteq V$, let

$$F(U) = \sum_{v \in U} f(v). \qquad (8)$$

Lemma 5. *For any graph $G = (V, E)$,*

$$\operatorname{opt}(G) \le F(V). \qquad (9)$$

Proof: Let μ be a function such that $\nu(\mu) = \operatorname{opt}(G)$. As ϕ is monotone increasing and W is monotone decreasing,

$$\begin{aligned} \operatorname{opt}(G) &= \sum_{v \in V} \sum_{\substack{(u,v) \in E \\ \phi(\mu(u),\mu(v)) \le c(u,v)}} \mu(v)\, W(\mu(v) + \mu(u)) \\ &\le \sum_{v \in V} \sum_{\substack{(u,v) \in E \\ \phi(0,\mu(v)) \le c(u,v)}} \mu(v)\, W(\mu(v)) \\ &\le \sum_{v \in V} |E_\mu(v)|\, \mu(v)\, W(\mu(v)) \\ &= F(V) \end{aligned}$$

■

B. Max-cut algorithm

A tighter bound can be obtained using an approximate weighted max-cut algorithm. First, consider how prices μ could be allocated if a suitable cut were known.

Given a cut $\{X, \overline{X}\}$ of G, which cuts edges $E' = \{(u, v) \in E \mid u \in X, v \in \overline{X}\}$, define a price function μ_X which sets price to 0 for nodes on one side of the cut, and $g(v)$ for nodes on the other, as follows:

Let $\mathbf{1}_A = 1$ if A is true and 0 if A is false. If

$$\sum_{v\in X} \sum_{\substack{(u,v)\in E' \\ \phi(0,g(v)) \leq c(u,v)}} g(v)W(g(v)) \geq \sum_{v\in\overline{X}} \sum_{\substack{(u,v)\in E' \\ \phi(0,g(v)) \leq c(u,v)}} g(v)W(g(v)) \tag{10}$$

then,

$$\mu_X(v) \leftarrow g(v)\mathbf{1}_{v\in X} \tag{11}$$

otherwise

$$\mu_X(v) \leftarrow g(v)\mathbf{1}_{v\in \overline{X}} \tag{12}$$

Also, define $c' : E \to \overline{\mathbb{R}}^+ \cup \{\infty\}$ by

$$c'(e) = \sum_{v\in e,\, \phi(0,g(v))\leq c(e)} g(v)W(g(v)), \tag{13}$$

which is either 0, $\min\big(g(u)W(g(u)), g(v)W(g(v))\big)$ or $g(u)W(g(u)) + g(v)W(g(v))$.

Lemma 6. *For an arbitrary cut $\{X, \overline{X}\}$ with cutting edges $E' \subseteq E$, the revenue generated by μ_X is related to the c'-weight of the cut by*

$$\nu(\mu_X) \geq \frac{1}{2} \sum_{e\in E'} c'(e) \tag{14}$$

Proof: Without loss of generality, consider the case that more revenue can be gained from nodes in X than in $\overline{X}$, in the sense of (10). Then, considering revenue only from the cut edges in E',

$$\nu(\mu_X) \geq \sum_{\substack{(v\in X, u\in\overline{X})\in E' \\ \phi(\mu_X(u), \mu_X(v)) \leq c(u,v)}} (\mu(u) + \mu(v))W(\mu(u) + \mu(v)). \tag{15}$$

By (10), μ is given by (11), whence then summand in (15) becomes $g(v)$, giving

$$\begin{aligned} \nu(\mu_X) &\geq \sum_{\substack{(v\in X, u\in\overline{X})\in E' \\ \phi(0,g(v))\leq c(u,v)}} g(v)W(g(v)) \\ &= \sum_{v\in X} \sum_{\substack{(u,v)\in E' \\ \phi(0,g(v))\leq c(u,v)}} g(v)W(g(v)) \\ &\geq \frac{1}{2} \sum_{v\in X\cup\overline{X}} \sum_{\substack{(u,v)\in E' \\ \phi(0,g(v))\leq c(u,v)}} g(v)W(g(v)) \end{aligned}$$

where the last step follows from (10), since the sum of two terms is at most twice the maximum term. Thus, since X and $\overline{X}$ partition V,

$$\begin{aligned} 2\nu(\mu_X) &\geq \sum_{v\in V} \sum_{\substack{(u,v)\in E' \\ c(u,v)\geq\phi((0,g(v))}} g(v)W(g(v)) \\ &= \sum_{e\in E'} \sum_{\substack{v\in e \\ c(u,v)\geq\phi(0,g(v))}} g(v)W(g(v)) \end{aligned}$$

and the result follows by the definition of c' in (13). ■

Algorithm 1 finds a suitable cut, not by performing a max-cut on the graph with the original weights c, but instead with the weights c', which allows the performance to be bounded.

Algorithm 1 Max-cut-based algorithm

```
Given a graph $G = (V, E)$ and peering costs $c : E \to \mathbb{R}^+$:
Compute $f(v)$ and $g(v)$ of (4) and (5) by Algorithm 2
Compute $c'$ according to (13).
Use the approximate weighted max-cut algorithm of [13] to get a cut $X$, $\overline{X}$ of $G$, with edge weights $c'$, of weight at least $\sum_{e\in E} c'(e)/2$.
Let $\mu \leftarrow \mu_X$ given by (10)–(12).
```

The main result of this section is

Theorem 7. *The μ generated by Algorithm 1 satisfies*

$$\nu(\mu) \geq \frac{1}{4}\mathrm{opt}(G) \tag{16}$$

This is a straightforward consequence of Lemma 6 and the following two lemmas.

Lemma 8. *The weight function c' defined by (13) satisfies*

$$F(V) = \sum_{v\in V} f(v) = \sum_{e\in E} c'(e) \tag{17}$$

Proof: By the definitions of f, g and ν and λ,

$$\begin{aligned} F(V) &= \sum_{v\in V} \nu(\lambda_{v,g(v)}) \\ &= \sum_{v\in V} \sum_{\substack{(u,v)\in E \\ \phi(0,g(v))\leq c(u,v)}} g(v)W(g(v)) \\ &= \sum_{e\in E} \sum_{\substack{v\in e \\ \phi(0,g(v))\leq c(u,v)}} g(v)W(g(v)) \end{aligned}$$

and the result follows by the definition of c' in (13). ■

Lemma 9. *The weight of the max-cut satisfies*

$$\sum_{e\in E'} c'(e) \geq \frac{1}{2}F(V) \tag{18}$$

Proof: Since E' is generated using the algorithm of [13], it is guaranteed that

$$\sum_{e\in E'} c'(e) \geq \frac{1}{2}\sum_{e\in E} c'(e) \tag{19}$$

The result then follows from Lemma 8. ■

Note that Algorithm 1 uses a max-cut algorithm which guarantees 1/2 of *sum of edge weights* [13]. It is tempting to think that this result could be tightened by using the celebrated 0.878 approximation algorithm of [8]. However, this is unsuitable here since it guarantees 0.878 of the *true max-cut*, which may be less than 1/2 of the sum of the edge weights.

V. Evaluating $f(v)$ and $g(v)$

An algorithm will now be presented which shows how to compute $f(v)$ and $g(v)$ and determine that the objective function is unbounded. Let

$$h_i(x) = (d - i + 1)xW(x). \tag{20}$$

Algorithm 2 Compute $f(v)$ and $g(v)$

For vertex $v \in V$ and set E_v of edges incident on v:

```
for $k \in E_v$ do
    Compute $b_k$ /* Polynomial time by assumption 6 */
end for
Sort $E_v = \{e_1, e_2, ..., e_d\}$, such that $b_1 \leq b_2 \leq ... \leq b_d$
for $i \leftarrow 1$ to $d$ do
    /* Find max revenue when $\{e_i, ..., e_d\}$ do not peer */
    if $b_i = -\infty$ then
        continue for next $i$.
    else if $b_i \in \overline{\mathbb{R}}^+$ then
        $x_i \leftarrow \text{argmax}_{x \in I_i} xW(x)$
                                /* Polynomial time by Lemma 10 */
        $f_i \leftarrow h_i(x_i)$
    else if $b_i = \infty$ then
        if $B = \sup_x xW(x) < \infty$ then
            if $x_* W(x_*) = B$ for some local maximum $x_*$ then
                $f_i \leftarrow h_i(x_*)$, $x_i \leftarrow x_*$
            else /* $xW(x) \to B$ as $x \to \infty$ by Lemma 11 */
                $f_i \leftarrow (d - i + 1)B$, $x_i \leftarrow \infty$
            end if
        else /* $h_i(x)$ is unbounded, so $\text{opt}(G) = \infty$ */
            $f_i \leftarrow \infty$, $x_i \leftarrow \infty$
        end if
    end if
end for
$f(v) = \max\{f_i | 1 \leq i \leq d\}$
$g(v) = x_j$, where $j = \min\{i | f_i = f(v)\}$
```

Algorithm 2 computes $f(v)$ and $g(v)$ in polynomial time using the following results, proved in the appendix.

Lemma 10. *For all $0 \leq a < b$, $xW(x)$ has a maximum on $I_k = [a, b]$, whose location can be found in polynomial time.*

Lemma 11. *If $B = \sup_x xW(x) < \infty$ and $xW(x) < B$ for all $x \in \overline{\mathbb{R}}^+$, then $xW(x) \to B$ as $x \to \infty$.*

VI. Boundedness of price and revenue

It is important to know under what conditions on ϕ, W and c the problem is well defined, and under what conditions the output of Algorithms 1 and 2 is well defined.

Theorem 12. *Let μ be the price function computed by Algorithm 1 and $f(v)$, $g(v)$ be the values computed by Algorithm 2. Then*

i $f(v) = \infty$ *if and only if* $xW(x)$ *is unbounded and* $b_d = \infty$

ii $g(v) = \infty$ *only if* $b_d = \infty$

iii $\text{opt}(G) = \infty$ *if and only if* $f(v) = \infty$ *for some* v

iv $\mu(v) = \infty$ *only if* $g(v) = \infty$

v $\nu(\mu) = \infty$ *if and only if* $\text{opt}(G) = \infty$.

Proof: By Algorithm 2, $f(v) = \infty$ if and only if for some i, both $h_i(x)$ is unbounded and $b_i = \infty$. These hold if and only if $xW(x)$ is unbounded and $b_d = \infty$, establishing claim (i).

Claim (ii) holds since $g(v) = \infty$ implies that $x_i = \infty$ for some i, which implies $b_d = \infty$.

Note also that $F(V) = \sum_{v \in V} f(v) = \infty$ if and only if $f(v) = \infty$ for some v. Claim (iii) follows since $\sup_v f(v) \leq \text{opt}(G) \leq F(V)$.

Claim (iv) holds since $\mu(v)$ is either 0 or $g(v)$ by (11).

By claim (iii) and (8), $\text{opt}(G) = \infty$ if and only if $F(V) = \infty$. Thus Theorem 7 establishes the "if" of claim (v), and Lemma 5 establishes the "only if". ■

Let μ^* be the optimal price function and μ' be that obtained by Algorithm 1. We now show that the finiteness of μ^* does not imply that of μ', nor the converse. If multiple price functions achieve $\text{opt}(G)$, then μ^* is selected to be one which assigns ∞ to the minimum number of nodes.

Let

$$\phi(a, b) = \min(a + b, 4 + \epsilon) \tag{21}$$

$$xW(x) = \begin{cases} x & x \leq 4 \\ 4 + \epsilon(1 - e^{-(x-4)}) & x > 4 \end{cases} \tag{22}$$

$$V = \{a, b, c, d\} \tag{23}$$

$$E = \{(a, b), (b, c), (c, d)\} \tag{24}$$

Note that edges with cost at least $4+\epsilon$ yield revenue however high the price charged to their nodes is, even though the extra revenue by charging more than 4 is at most ϵ. Manipulating just the cost function c yields examples where Algorithm 1 assigns infinite price to a node but the optimum does not, and vice versa. Note that for any c, the max-cut of G above is $\{\{a, b\}, \{c, d\}\}$.

Example 1. *There is an input where $\mu^*(v)$ is finite for all v, but $\mu'(v)$ is infinite for some v.*

Proof: Let $c(a, b) = 2$, $c(b, c) = 4 + \epsilon$, and $c(c, d) = 2$. Then $\text{opt}(G) = 8$ with $\mu^*(a) = \mu^*(d) = 0$, $\mu^*(b) = \mu^*(c) = 2$. However, $f(b) = f(c) = 4+\epsilon$ with $g(b) = g(c) = \infty$. Since b and c are in different partitions in any 1/2 approximation to the max-cut, Algorithm 1 yields $\mu'(b) = \infty$ or $\mu'(c) = \infty$. ■

Example 2. *There is an input where $\mu'(v)$ is finite for all v, but $\mu^*(v)$ is infinite for some v.*

Proof: Let $c(a, b) = 4 + \epsilon$, $c(b, c) = 4$, and $c(c, d) = 2$. Then $\text{opt}(G) = 10+\epsilon$ with $\mu^*(a) = \infty$, $\mu^*(b) = 4$, $\mu^*(c) = 0$, $\mu^*(d) = 2$. However, $f(b) = 8$, $f(d) = 2$ from $g(b) = 4$, $g(d) = 2$, and their sum exceeds $f(a) + f(c) = 8 + \epsilon$. Therefore, Algorithm 1 chooses b, d, which means $\mu'(a) = \mu'(c) = 0, \mu'(b) = 4, \mu'(d) = 2$, which are all finite. ■

VII. Conclusion

The problem of a tier-1 ISP finding prices to charge its customers to maximize its revenue is NP-hard and APX-hard when customers may form peering arrangements. However, this paper proposed an algorithm which yields a 1/4 approximation for a class of peering criteria. The algorithm is based on an existing weighted max-cut algorithm. Conditions were found under which a given instance of this model yields meaningful (finite) revenue and prices.

Acknowledgments

This material is based upon work supported by the National Science Foundation under Grant No. CNS-0520349.

Appendix

Proof (Lemma 1): Let $\phi_2(a) = \phi(0, a)$, which is left-continuous. Let I be the set of solutions of (3). If I is empty, the result holds. Otherwise, $0 \in I$ since $\phi_2(0) \leq \phi_2(a)$ for all $a \in \overline{\mathbb{R}}^+$. Furthermore, $\phi_2(b) \leq K \Rightarrow \phi_2(a) \leq K$ for $a \leq b$, which implies I is connected and hence an interval. Since $\phi(a, 0)$ is defined only for $a \geq 0$, that implies either I is $\overline{\mathbb{R}}^+$, or for some $b \in \overline{\mathbb{R}}^+$ it is $[0, b]$ or $[0, b)$. However, the last cannot hold since it implies that $\phi_2(b) = K' > K$ whence, by left-continuity, a left-neighborhood of b would not be in I. ■

Proof (Lemma 2): First, consider the maximum possible revenue from edges (x, a) and (a, e). If $\mu(a) > 1$, then there's no revenue from (a, e), so the max revenue is 2. However, if $\mu(a) \leq 1$, the revenue is at most 1 from (a, e) and $\mu(a) + \mu(x) \leq 1 + \mu(x)$ for (x, a). Thus, the revenue from these two edges is at most $2 + \mu(x)$.

Next consider the revenue from the other edges. If no revenue is achieved from the edge (x, b), this is bounded by 6. Otherwise, $\mu(b) \leq 1 - \mu(x)$.

By symmetry, the maximum revenue from the edges incident on c is equal to that of those incident on d, and can be bounded as follows. Using the same arguments as applied to node a, each of these is at most $2 + \mu(b) \leq 3 - \mu(x)$.

So in this case, the total revenue from the five edges under consideration is at most 1 from (x, b), plus $2(3 - \mu(x))$ from the others, for a total of $7 - 2\mu(x)$.

Thus, the maximum revenue from these other edges is at most $\max(6, 7 - 2\mu(x))$. Adding in the maximum revenue from (x, a) and (a, e) of $2 + \mu(x)$, we have a maximum total revenue of $\max(8 + \mu(x), 9 - \mu(x))$. ■

Proof (Lemma 3): Since $0 \leq \mu(x) \leq 1/2$, the maximum revenue from the $d(x)$ constructions is $d(x)(9 - \mu(x))$. Moreover x contributes at most $d(x)\mu(x)$ revenue to the original edges to which it is adjacent. Thus, the combined revenue from x and the nodes in constructions attached to it is at most $9d(x)$. This is also obtained by setting $\mu(x) = 0$. Since setting $\mu(x) = 0$ maximizes the revenue obtainable by the nodes connected to x by original edges, the revenue obtained with $\mu(x) = 0$ is at least as much as obtained with the original $\mu(x)$. ■

Proof (Lemma 4): If $\mu(x) \in (1/2, 1)$, the revenue from the $d(x)$ constructions is $d(x)(8 + \mu(x))$. If we increase $\mu(x)$ to 1, this too will increase. Any node connected to x by an original edge either has weight 0 or weight greater than $1/2$. If the weight is 0, increasing $\mu(x)$ will increase the revenue from that edge. Otherwise, the total cost of the two nodes already exceeds the weight of the edge, and the revenue is unchanged at 0. Thus, increasing $\mu(x)$ increases the total revenue. Since the revenue is left-continuous in $\mu(x)$, the maximum revenue is obtained by replacing $\mu(x) = 1$. ■

Proof (Lemma 10): Because $W(x)$ is left continuous and decreasing, and $a \geq 0$ $xW(x)$ is upper semi-continuous on $[a, b]$. Thus, the extreme value theorem implies that it is bounded above and attains its supremum. The maximum must occur either at a local maximum of $xW(x)$, of which there are a constant (finite) number or at a or b. Thus, it sufficies to evaluate $xW(x)$ at those points. ■

Proof (Lemma 11): The proof is by showing that $xW(x)$ eventually becomes monotonic increasing under the hypotheses. By assumption, $xW(x)$ has a finite number of local maxima. Let x_1 be the rightmost point where a local maximum occurs. Since $[0, x_1]$ is closed, $xW(x)$ has a global maximum B_1 on this interval, by Lemma 10, with $B_1 < B$ by hypothesis. Since B is the least upper bound, for any $0 < \epsilon < B - B_1$, there exists $x_2 > x_1$ such that $x_2 W(x_2) > B - \epsilon$. If there is $x_3 \in (x_1, x_2)$ such that $x_3 W(x_3) > x_2 W(x_2)$, $xW(x)$ has the global maximum on $[x_1, x_2]$ by Lemma 10. Since $x_3 W(x_3) > x_2 W(x_2) > x_1 W(x_1)$, the global maximum does not occur at boundaries, which implies the existence of another local maximum. This is contradiction, so $x_2 W(x_2) \geq xW(x)$ for any $x < x_2$. Then for any $x > x_2$, $xW(x) \geq x_2 W x_2 > B - \epsilon$, since if there is some x_4 such that $x_2 < x_4$ and $x_2 W(x_2) > x_4 W(x_4)$, and Lemma 10 guarantees the existence of a local maximum on $[x_2, x_4]$. That contradicts the choice of x_1 as the rightmost local maximum. ■

References

[1] D. Acemoglu and A. Ozdaglar. Price Competition in Communication Networks. In *Proc. IEEE INFOCOM*, pages 1–12, 2006.

[2] P. Alimonti and V. Kann. Hardness of approximating problems on cubic graphs. In *Proc. 3rd Ital. Conf on Algorithms and Complexity*, pp. 288–298. Springer, 1997.

[3] E. Anshelevich, B. Shepherd and G. Wilfong. Strategic network formation through peering and service agreements. in *Proc. IEEE Foundations of Computer Science (FOCS)*, 2006, pages 77–86.

[4] S. Arora, D. Karger and M. Karpinski. Polynomial time approximation schemes for dense instances of NP-hard problems. In *Proc. ACM Symposium on the Theory of Computing* (Las Vegas, NV), pp. 284-293, May 29–June 1, 1995.

[5] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation: Combinatorial optimization problems and their approximability properties*, Springer-Verlag, 1999.

[6] N. Balkan and A. Blum. Approximation Algorithms and Online Mechanisms for Item Pricing. *Theory of Computing*, 3(9):179-195, 2007.

[7] M. R. Garey, D. S. Johnson and L. Stockmeyer. Some simplified NP-complete problems. In *Proceedings of the sixth annual ACM symposium on Theory of computing*, 1974, pp 47-63.

[8] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, Nov. 1995.

[9] V. Guruswami, J. D. Hartline, A. R. Karlin, D. Kempe, C. Keynon and F. McSherry. On profit-maximizing envy-free pricing. In *Proc. ACM-SIAM symposium on Discrete algorithms*, pp. 1164–1173, 2005

[10] R. Johari and J. Tsitsiklis. Routing and peering in a competitive Internet. In *Proc. IEEE Conf. Decision and Control*, 2004.

[11] N. Karmarkar. A new polynomial time algorithm for linear programming, *Combinatorica*, 4(4):373–395, 1984.

[12] E. W. Lee, D. Buchfuhrer, L. L. H. Andrew, A. Tang and S. H. Low. Pricing in the Presence of Peering, in *Proc. Allerton Conference on Computing, Communications and Control*, 2007.

[13] S. Sahni and T. Gonzalez, P-Complete approximation problems, *Journal of the ACM*, 23(3):555–565, July 1976.

[14] S. Shakkottai and R. Srikant, Economics of network pricing with multiple ISPs, *IEEE Trans. Networking*, 14(6):1233-1244, Dec. 2006.

[15] B. Zhang, R. Liu, D. Massey, and L. Zhang. Collecting the internet AS-level topology. *ACM Computer Communication Review*, 35(1):53–61, Jan. 2005.