

Inefficiency in Provisioning Interconnected Communication Networks

Pedro Ferreira
School of Information Management and Systems
University of California at Berkeley
pedro@sims.berkeley.edu

Marvin Sirbu
Engineering and Public Policy Department
Carnegie Mellon University
sirbu@cmu.edu

ABSTRACT

We model the Internet as a collection of Internet Service Providers (ISPs) that transit and peer to meet an exogenous demand for the transport of IP traffic among an arbitrary set of cities. The prices of bandwidth and of transit are also exogenous. ISPs are assumed to behave rationally and selfishly and therefore they provision the network that maximizes their own profit regardless of how expensive the overall network becomes. We call optimal to a network that ISPs would build if they cooperated to reduce overall provisioning costs; and we call inefficiency cost the additional cost to provision a network at Nash equilibrium relative to the provisioning cost of an optimal network. We show that inefficiency cost is primarily related to transit agreements. In fact, in a world where only peering agreements exist there is no inefficiency cost. However, networks with only peering agreements forgo the efficiencies of traffic aggregation. Transit agreements help reduce provisioning costs by realizing benefits from economies of scale. We show, by example, that there exist Nash networks that are strictly more expensive than optimal networks even when ISPs choose transit prices and therefore the market to provision communication networks with interconnection agreements designed the way they are today is inefficient. We show that inefficiency cost may be reduced, for example, by enforcing side payments between ISPs. We conclude with an analysis of the difficulties a regulator would face in endeavoring to move a market from an inefficient Nash equilibrium to optimality.

Categories and Subject Descriptors

C.2 [Computer Systems Organization]: Computer Communication Networks

General Terms

Design, Economics, Theory

Keywords

Cost of Anarchy, Peering, Transit, Provisioning Networks

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

EC'05, June 5–8, 2005, Vancouver, British Columbia, Canada.
Copyright 2005 ACM 1-59593-049-3/05/0006 ...\$5.00.

1. INTRODUCTION

The Internet is a collection of Internet Service Providers (ISPs), or more precisely Autonomous Systems (ASes) [1, 15], which must interconnect to allow users of different ISPs to exchange traffic. Usually, ISPs use peering and transit agreements to interconnect and to exchange information about how to reach destinations in the Internet according to their interconnection strategies [35, 12]. Each ISP is a separate entity that routes traffic and provisions its network selfishly. In this paper, we will also assume that each ISP does so rationally. In this context, a non-cooperative game among ISPs arises and we look at the topology of the resulting network as that obtained at Nash equilibrium [20, 33, 24, 39, 2]. However, it is well known that the Nash equilibrium carries an inefficiency [28, 3, 20, 31]. For example, the aggregate cost to provision such a network may be greater than the cost to provision the network that would emerge if ISPs cooperated to reduce overall provisioning costs.

To start, we assume that demand, transit prices and the price of bandwidth are exogenously given. The goal of this paper is three-fold. First, we want to know how much more expensive a network at Nash equilibrium can be relative to the cheapest possible network ISPs could build to meet the same demand, that is, how large can inefficiency loss become. Second, we want to know how inefficiency loss relates to the fact that ISPs use transit and peering agreements to interconnect. Finally, we want to know if the market for provisioning networks that interconnect using transit and peering agreements is efficient, that is, if Nash networks become as inexpensive as optimal networks when, for example, ISPs are allowed to choose transit prices.

This paper is written at a time of dramatic reduction in the price of bandwidth [8, 36], which leads smaller ISPs to peer directly with each other, thus relying less frequently on transit from Tier-1 ISPs [5, 27] and creating a "peering-donut" around them. But, the declining trend in the price of bandwidth put pressure on transit prices, which also started falling significantly. This trend in the price of transit together with the steady increase in IP traffic worldwide (about 110% between 2003 and 2004 [38] in most major international backbones) resulted in a substantial increase in transit revenues [37]. Peering and transit agreements remain key constructs of today's communication networks and ISPs rely heavily on both of them to interconnect and exchange traffic across the globe. Thus, it is fundamental to study how these two types of interconnection agreements impact provisioning costs when ISPs provision their networks selfishly.

2. RELATED WORK

A significant literature has been recently published that looks at aspects of inefficiency in communication networks. These works, however, all define inefficiency in terms of latency and do not address inefficiency in provisioning networks. The seminal work in this area is provided by [20] who coined the term price of anarchy as an upper-bound on the ratio between the aggregated latency over all users in a network at Nash equilibrium, in which users route their traffic selfishly, and that latency in the fastest possible network, in which users would cooperate to route their traffic. It has been shown that in a network where end-users route their own traffic without using ISPs, the price of anarchy is $\sup_{\{e \in E\}} [(\lambda_e \mu_e + (1 - \lambda_e))^{-1}]$, where λ_e satisfies $l_e^*(\lambda_e r) = l_e(r)$, μ_e is defined by $\mu_e = l_e(\lambda_e r)/l_e(r)$, $l_e(x)$ is the latency function for link e and $l_e^*(x)$ is the marginal latency function for this link [33]. E is the set of links in the network. For the case of affine latency functions, this expression reduces to $4/3$. For posynomial latency functions of degree at most p , the price of anarchy is asymptotic to $[1 - p(p+1)^{-(p+1)/p}]^{-1}$ as $p \rightarrow \infty$. It has also been shown that this bound is tight; in fact it occurs in simple networks of parallel links [33, 28, 3], and independent of the network topology [34]. For the cases of M/M/1 and M/G/1 queues the price of anarchy approaches $+\infty$ as the sum of traffic rates approaches the smallest link capacity [33]. All these results are obtained assuming standard latency functions, for which $x.l_e(x)$ is convex.

It has been shown that these results can be extended to the case of networks with capacities and with discontinuous latency functions that are non-decreasing in the amount of traffic carried [6]. Particularly, it has been shown that the price of anarchy is given by $1/(1 - \nu(L))$, where $\nu(L) = \sup_{v \geq 0, l \in L} 1/v l(v) \max_{x \geq 0} \{x(l(x) - l(v))\}$ and l is a latency function in the class of latency functions L . $\nu(L)$ reduces to ρ when L is the set of standard latency functions. More recently, these concepts have been applied to the case of a resource allocation game. It has been shown that in a game where N users, endowed with the same concave utility function for the transport of traffic, submit individual payments for a single link they want to use, whose cost is a convex function of the rate, and anticipate the effects of their actions on the price to pay, the loss of efficiency is no more than $1/(2N + 1)$ [18]. This result has also been used to show that if costs are affine, efficiency loss is no more than 33% even if users have general concave utility functions. Different versions of these results have been extended to competitive network contexts. For example, it has been shown that when $N > 2$ firms compete (each facing convex production costs), to meet an exogenous demand D for the transport of traffic, by submitting to a fair clearing house a supply function of the form $D - w_n/p$, where w_n is a parameter chosen by firm n and p is the price charged to the users, efficiency loss is no more than $1/(N - 2)$ [17].

The price of anarchy is an upper bound on the efficiency loss in the network at Nash equilibrium and therefore it may convey little information about the actual loss experienced. The relevance of the price of anarchy depends on how often are we at the worst-case efficiency loss. This issue has been addressed by simulation over realistic Internet-like topologies and traffic demands [30] considering only intra-domain traffic. These simulations have shown that selfish routing achieves close to optimal average latency in such environments, however such performance benefit comes at the expense of significant increased congestion on certain links. Additionally, it has been shown that most traffic in the Internet uses TCP [29], which implements flow control. Hence, the rates at which users send traffic to the network depend on the congestion experienced.

It has been shown that the cost function users minimize over each TCP connection is approximately linear [20]. Thus, in this case, the price of anarchy is less than 33% and, in practice, one can expect losses to be small. In line with these arguments, it has been shown that, in a general network, the measure of the set of rates that imply a price of anarchy greater than φ is $O(\log(C(r))/\varphi)$, which is a set with a small Lebesgue measure. In this expression $C(r)$, defined as the ratio between the aggregate running cost at Nash equilibria to meet demand r and that to meet demand $r/2$, represents the criticality factor of the network and serves as an upper bound for its price of anarchy [11].

All these results do not take into consideration the existence of ISPs and assume that end-users route their own traffic, a situation that is not true whatsoever in today's communication networks. As such, the models currently used in the literature do not address the issue of interconnection, which is key for the development of the Internet. This paper examines how the Nash configuration of interconnected networks differs from the cheapest possible network ISPs could build if they cooperated.

3. INTERCONNECTION AGREEMENTS

Interconnection between ISPs is accomplished through peering and transit agreements. Peering is a bilateral business and technical arrangement, whereby two providers agree to accept traffic from one another, and from one another's customers recursively. Peering does not include the obligation to carry traffic to third parties [21, 16, 23, 25, 26]. Direct peering has often been offered on a Bill-and-Keep basis. However, there is an element of barter when both networks do not perceive a roughly equal exchange of value. For this reason, many of the largest ISPs impose minimum peering requirements that smaller providers willing to peer with them must meet. These requirements usually include a minimum number of locations for peering and a minimum bandwidth for the peering connections.

Network Access Points (NAPs) are meet points where several providers exchange traffic. ISPs deploy routing equipment at a NAP where they establish multiple bilateral connections. Each provider is responsible for the link from its premises to the meet point. The meet point is usually owned and run by a separate entity that leases space to the providers willing to co-locate there. NAPs can also be distributed. NAPs are a convenient solution for peering because many providers come to the same physical place and with short links within the meet point they can reach several other providers. Additionally, an ISP can run a fat link from its premises to the meet point that takes all the traffic to be delivered to peers at this point, which is cheaper than deploying one link per peering connection. However, the attractiveness of NAPs has resulted in congestion, which has been one of their major problems [19].

ISPs can also act as transit providers [21, 16, 23, 25, 26]. A transit agreement is a bilateral business and technical arrangement, where one provider, the transit provider, agrees to carry traffic to third parties on behalf of another provider. In most cases, the transit provider carries traffic to and from its customers, and to and from every destination in the Internet [21]. The major difference between a peering agreement and a transit agreement is that in the latter case the provider requesting the connection is seen as a customer of the provider offering the service. The customer pays a fee for transit service and expects some quality of service. In this paper, we are interested in studying how both transit and peering agreements impact efficiency loss in terms of provisioning costs.

4. NETWORK MODEL

Consider a set of ISPs $I = \{1, \dots, n\}$ and a set of cities $C = \{1, \dots, m\}$. ISP i serves cities $C_i \subseteq C$. Let $V = \{(i, u) : i \in I, u \in C_i\} \subseteq (I \times C)$ and consider a directed network $g_0 = (V, E, D)$ with vertex set V and edge set E . Define $D : E \rightarrow \mathbb{R}_0^+$ as a distance metric over set E . $D(e)$ indicates the length of edge e . Vertex (i, u) represents ISP i at city u . Edge $e = (i, u, j, v)$ represents a link between ISP i at city u and ISP j at city v used to carry traffic from ISP i to ISP j . Let $igi(e)$ and $igc(e)$ represent the ingress ISP and the ingress city of edge e , respectively. Also, let $egi(e)$ and $egc(e)$ represent the egress ISP and the egress city of edge e , respectively. $D(e)$ indicates the length of the shortest bandwidth link one can deploy between city $igc(e)$ and city $egc(e)$. Let $erd(e)$ represent the edge in the reverse direction of edge e , which always exists in g_0 . Let $E_i = \{e = (i, u, j, v)\} \subseteq E$ represent the set of all edges originated in ISP i .

Let $B = \{e = (i, u, j, v) : i \neq j\} \subseteq E$ represent the subset of edges in E between different ISPs. We require edges between different ISPs to have an extra attribute $a \in A = \{-1, 0, +1\}$, which defines the type of interconnection agreement. Edge $e = (i, u, j, v)$ represents a peering link if $a = 0$ and a transit link otherwise. If $a = -1(+1)$ this link is used under a transit agreement in which ISP i buys (sells) transit to ISP j . Let $B_a \subseteq B$, with $a \in A$, indicate the subset of edges in B with attribute a . If edge $e \in B_a$, then edge $erd(e) \in B_{-a}$. Let $B_a(i, j) \subseteq B$ indicate the set of edges in B_a used to interconnect from ISP i to ISP j . Let $\zeta(e) = a$, defined for $e \in B$, indicate the attribute of edge e .

Let $P(i, u, j, v)$ represent the set of (directed) paths in g_0 from vertex (i, u) to vertex (j, v) and let $P = \bigcup_{i \in I} \bigcup_{u \in C_i} \bigcup_{j \in I} \bigcup_{v \in C_j} P(i, u, j, v)$ represent the set of all these paths in the network. A path p in this set is written as $p = [e_1, \dots, e_z]$, where $e_k \in E$, for $1 \leq k \leq z$ and $igi(e_1) = i, igc(e_1) = u, egi(e_z) = j, egc(e_z) = v$ and $egi(e_{k-1}) = igi(e_k), egc(e_{k-1}) = igc(e_k)$, for $2 \leq k \leq z$. Additionally, let $\bar{P}(i, u, j, v) \subseteq P(i, u, j, v)$ contain all the paths in G_0 of the form $\bar{p} = [\bar{e}_1, \dots, \bar{e}_z]$ such that if $\exists k: 2 \leq k \leq z : egi(\bar{e}_{k-1}) \neq egi(\bar{e}_k)$ then $k = z$. That is, if $i = j$ then $\bar{P}(i, u, j, v)$ is the set of all paths internal to ISP i that originate in vertex (i, u) and terminate in vertex (i, v) . If, on the other hand, $i \neq j$ then $\bar{P}(i, u, j, v)$ is the set of all paths that originate at vertex (i, u) internal to ISP i plus an edge to interconnect to ISP j at city v . Finally, let $\bar{P} = \bigcup_{i \in I} \bigcup_{u \in C_i} \bigcup_{j \in I} \bigcup_{v \in C_j} \bar{P}(i, u, j, v)$ represent the set of all these paths in the network.

Let $\theta = (i', u', j', v')$ represent a type of traffic in g_0 and let $F = \{(i', u', j', v') : i', j' \in I, u' \in C_{i'}, v' \in C_{j'}\}$ represent the set of all possible types of traffic flows. Traffic of type $\theta = (i', u', j', v')$ originates from end-users of ISP i' in city u' and is destined to end-users of ISP j' in city v' . Let $egif(\theta)$ indicate the ISP to whom the traffic is destined. Additionally, let $f_\theta(p)$ represent the amount of traffic of type θ in path p . The total flow of traffic in path p is therefore given $f(p) = \sum_{\theta \in F} f_\theta(p)$. The flow of traffic of type θ in edge e is given by $f_\theta(e) = \sum_{p \in P: e \in p} f_\theta(p)$ and the total traffic in edge e is given by $f(e) = \sum_{p \in P: e \in p} f(p)$. Let $d(\theta)$ indicate the fixed exogenous demand for the transport of traffic of type θ and define $d(\theta, (i, u))$ in the following way: $d((i', u', j', v'), (i, u)) = -d(i', u', j', v')$ if $i = i' \wedge u = u'$, $d((i', u', j', v'), (i, u)) = d(i', u', j', v')$ if $i = j' \wedge u = v'$ and $d((i', u', j', v'), (i, u)) = 0$ otherwise.

Define $d = [d(\theta), \theta \in F]$ as the exogenous demand matrix and define $G_0(d)$ as the set of all possible graphs that allow for realizing demand d . Throughout this work, we will assume that $g_0 \in G_0(d)$, that is, it is possible to derive a network over g_0 to meet the exogenous demand for the transport of traffic d . Finally, note that the model introduced in this section is very general and therefore very flexible. For example, one can easily note that this model allows for considering NAPs when we allow for several ISPs to meet at the same city by deploying internal links to this city and none of them serves end-users there.

5. N-ISP NON-COOPERATIVE GAME

Consider the non-cooperative n -ISP game that arises in the network setting defined above. Let $\Gamma = [(S_1, \dots, S_n); (\Pi_1, \dots, \Pi_n)]$ represent the normal form of this game, where S_i is the strategy set of ISP i and Π_i represents its profit function. Additionally, let s_i represent a generic strategy for ISP i and let $s = \Pi_{i \in I} s_i$ represent a generic joint strategy for all ISPs. Let $S = \Pi_{i \in I} S_i$ represent the space of joint strategies for all ISPs. Sometimes we will write $s = (s_i, s_{-i})$, where $s_{-i} = \Pi_{j \in I \setminus \{i\}} s_j$ is the joint strategy of all ISPs other than ISP i . Similarly, we also write $S = S_i \times S_{-i}$, where S_{-i} represents the joint strategy space of all ISPs except ISP i . The remainder of this section specifies S_i and Π_i in detail.

ISP i chooses all traffic flows within its network and the traffic flows to send to other ISPs. In other words, a strategy for ISP i is $s_i = \{f_\theta(e) : e \in E_i, \theta \in F\}$. An assignment of traffic flows must however satisfy a number of constraints, which implicitly define S_i as a compact convex set of strategies with we will assume to be non-empty. Let $In(i, u) = \{e \in E : egi(e) = i, egc(e) = u\}$ represent the set of edges in the network that terminate at vertex (i, u) . Similarly, let $Out(i, u) = \{e \in E : igi(e) = i, igc(e) = u\}$ represent the set of edges in the network that originate at vertex (i, u) . The first two constraints ensure that traffic flows are non-negative and that flow is conserved at every vertex in the network:

- (1) $f_\theta(e) \geq 0, \forall \theta \in F, \forall e \in E$
- (2) $\sum_{e \in In(i, u)} f_\theta(e) - \sum_{e \in Out(i, u)} f_\theta(e) = d(\theta, (i, u))$
 $\forall \theta \in F, \forall (i, u) \in V$

The third constraint indicates that traffic flows must conform to the type of interconnection agreements established. Let $TB = \{f_\theta(e) : e \in B, \theta \in F\}$ represent the set of all traffic between ISPs. Let $TRS(j, TB)$ indicate the Transit Reachability Set of ISP j given TB . This set, which we will define formally later in this section, includes ISP j , all its customers and all the customers of customers of ISP j recursively. If ISP i peers with or sells transit to ISP j , then ISP i cannot send, over that connection, traffic destined to an ISP that does not belong to $TRS(j, TB)$. This constraint translates to:

- (3) $egif(\theta) \notin TRS(egic(e), TB) \Rightarrow f_\theta(e) = 0$
 $\forall \theta \in F, \forall e \in (B_0 \cup B_{+1})$

The fourth constraint indicates that an ISP can only introduce traffic into a peering link to another ISP, if the latter ISP signals that it also wants to peer with the former ISP. We will assume that an ISP signals that it wants to peer by sending traffic over a peering link. Therefore, an ISP can only introduce traffic into a peering link if it receives traffic over the peering link in the reverse direction. The same sort of constraint applies when an ISP sells transit to another ISP and, hence, the following constraint:

$$(4) f_\theta(\text{erd}(e)) = 0 \Rightarrow f_\theta(e) = 0, \forall \theta \in F, \forall e \in (B_0 \cup B_{+1})$$

The last constraint prevents an ISP, say ISP i , from sending traffic destined to itself over a transit link used under an agreement in which he sells transit to another ISP, say ISP j . If that was possible, ISP i could increase its profit indefinitely by sending traffic to ISP j which would later return to ISP i , at the expense of ISP j . Formally,

$$(5) \text{egif}(\theta) = \text{igi}(e) \Rightarrow f_\theta(e) = 0, \forall \theta \in F, \forall e \in B_{+1}.$$

To define formally $TRS(j, TB)$ let $TRS_h(j, TB)$ indicate the h -level TRS of ISP j given traffic flows TB , for any integer $h : 0 \leq h \leq n-2$. For an ISP belonging to this set there are exactly h other ISPs between him and ISP j . We have $TRS_h(j, TB) = \{j\}$ for $h = 0$. The following definition holds for $1 \leq h \leq n-2$:

$$TRS_h(j, TB) = \{j' \in I \setminus TRS_{h-1}(j, TB) : \exists j'' \in TRS_{h-1}(j, TB), \exists e \in B_{-1}(j', j''), \exists \theta \in F : f_\theta(e) > 0\} \quad (1)$$

Using this recursive definition for $TRS_h(j, TB)$, we can say that $TRS(j, TB) = \{j' \in I : \exists h \geq 0 : j' \in TRS_h(j, TB)\}$. It remains to define Π_i to fully specify the normal form for game Γ . Π_i is a function of the strategy of ISP i and of the strategies of the other ISPs and therefore we can write $\Pi_i(s_i, s_{-i})$. Let $PT(M, e, a)$, defined for $e = (i, u, j, v) \in (B_{-1} \cup B_{+1}) : i < j$ and for $a \in \{-1, +1\}$, represent the observed price per unit of traffic for a transit agreement to carry at most M Mbps over edge e , excluding the cost of the transit link. Also, let $PB(M, L)$ represent the observed price per unit of traffic for a link to carry M Mbps over a distance of L miles. We will assume that this price includes termination costs and that links are full-duplex and priced according to the maximum flow of traffic in either direction. The functions PT and PB were estimated in [7]. According to the results there, these functions are continuous and exhibit strict economies of scale in their arguments.

We can define $\Pi_i(s_i, s_{-i}) = \sum_{e \in E_i} T\Pi(e, s_i, s_{-i})$ by using the additional definitions in Table 1. Let $\eta(e)$ indicate how much of the cost of a link each ISP pays. Assuming that two ISPs that peer split equally the cost of the peering link and that the ISP buying transit pays the link entirely, that is,

$C(e, s_i, s_{-i})$	Cost/unit of traffic to provision edge e
$R(e, s_i, s_{-i})$	Revenue/unit of traffic over edge e
$\Pi(e, s_i, s_{-i}) = R(e, s_i, s_{-i}) - C(e, s_i, s_{-i})$	Profit/unit of traffic over edge e
$T\Pi(e, s_i, s_{-i}) = \Pi(e, s_i, s_{-i}) \cdot f(e)$	Total profit over edge e
$TC(e, s_i, s_{-i}) = C(e, s_i, s_{-i}) \cdot f(e)$	Total cost over edge e
$T\Pi'(e, s_i, s_{-i})$	Marginal profit over edge e
$TC'(e, s_i, s_{-i})$	Marginal cost over edge e

Table 1: Definition of profit and cost functions over edges.

$$\eta(e) = \begin{cases} 1, & \text{if } e \in (E \setminus B \cup B_{-1}), \\ 1/2, & \text{if } e \in B_0, \\ 0, & \text{if } e \in B_{+1}. \end{cases} \quad (2)$$

and using $mf(e) = \max\{f(e), f(\text{erd}(e))\}$ to represent the largest flow between edges e and $\text{erd}(e)$, the following definitions hold: $C(e, s_i, s_{-i}) = PB(mf(e), D(e)) \cdot mf(e)/f(e) \cdot \eta(e)$ and $R(e, s_i, s_{-i}) = PT(mf(e), e, \zeta(e)) \cdot mf(e)/f(e) \cdot \zeta(e)$ when $e = (i, u, j, v) \in (B_{-1} \cup B_{+1}) : i < j$, otherwise $R(e, s_i, s_{-i}) = PT(mf(e), \text{erd}(e), \zeta(\text{erd}(e))) \cdot mf(e)/f(e) \cdot \zeta(\text{erd}(e))$ when $e = (i, u, j, v) \in (B_{-1} \cup B_{+1}) : i > j$; and 0 when $e \in (E \setminus B \cup B_0)$. Finally, we have $T\Pi'(e, s_i, s_{-i}) = \partial T\Pi(e, s_i, s_{-i})/\partial f(e)$ if $f(e) \geq f(\text{erd}(e))$ and $T\Pi'(e, s_i, s_{-i}) = \partial T\Pi(e, s_i, s_{-i})/\partial f(\text{erd}(e))$

otherwise. A similar definition holds for $TC'(e, s_i, s_{-i})$ with respect to $C(e, s_i, s_{-i})$. We can also write $\Pi_i(s_i, s_{-i})$ in terms of flows over paths: $\Pi_i(s_i, s_{-i}) = \sum_{u \in C_i} \sum_{j \in I} \sum_{v \in C_j} \sum_{\bar{p} \in \bar{P}(i, u, j, v)} T\Pi(\bar{p}, s_i, s_{-i})$ together with the definitions in Table 2.

$C(\bar{p}, s_i, s_{-i}) = \sum_{e \in \bar{p}} C(e, s_i, s_{-i})$	Cost/unit of traffic to provision path \bar{p}
$R(\bar{p}, s_i, s_{-i}) = \sum_{e \in \bar{p}} R(e, s_i, s_{-i})$	Revenue/unit of traffic over path \bar{p}
$\Pi(\bar{p}, s_i, s_{-i}) = R(\bar{p}, s_i, s_{-i}) - C(\bar{p}, s_i, s_{-i})$	Profit/unit of traffic over path \bar{p}
$T\Pi(\bar{p}, s_i, s_{-i}) = \Pi(\bar{p}, s_i, s_{-i}) \cdot f(\bar{p})$	Total profit over path \bar{p}
$TC(\bar{p}, s_i, s_{-i}) = C(\bar{p}, s_i, s_{-i}) \cdot f(\bar{p})$	Total cost over path \bar{p}
$T\Pi'(e, s_i, s_{-i}) = \sum_{e \in \bar{p}} T\Pi'(e, s_i, s_{-i})$	Marginal profit over path \bar{p}
$TC'(e, s_i, s_{-i}) = \sum_{e \in \bar{p}} TC'(e, s_i, s_{-i})$	Marginal cost over path \bar{p}

Table 2: Definition of profit and cost functions over paths.

It will be helpful to note several important properties of the functions introduced before: i) $\Pi(e, s_i, s_{-i})$ can be positive or negative because it accounts for the transfer payments from selling and buying transit; ii) $C(e, s_i, s_{-i})$ is always positive (and only zero if $f(e) = 0$); iii) $TC(e, s_i, s_{-i})$ is increasing in any $f_\theta(e)$ in its argument; iv) $TC'(e, s_i, s_{-i})$ is non-increasing in any $f_\theta(e)$ in its argument. All functions defined heretofore are continuous. These properties extend trivially to the counterpart functions defined in terms of paths.

Finally, note that the problem we want to solve is completely defined by the network graph g_0 , the exogenous demand d (and we will assume $g_0 \in G_0(d)$) and the profit per-unit of flow functions $\Pi(e, s_i, s_{-i})$, which use the exogenous prices functions PB and PT . We will call (g_0, d, PB, PT) an instance of our problem. We will call $G(g_0, d)$ the set of all feasible networks for this problem, that is, all networks that can meet the exogenous demand d over g_0 , which we will assume to be non-empty.

6. ADDITIONAL DEFINITIONS

We want to study flows at Nash Equilibrium for an instance (g_0, d, PB, PT) of our problem ($g_0 \in G_0(d)$). Such flows are characterized by $s^* = (s_i^*, s_{-i}^*)$ with $\Pi_i(s_i^*, s_{-i}^*) \geq \Pi_i(s_i, s_{-i}^*)$, for all $s_i \in S_i$ and for all $i \in I$ [24] and induce a set of network $G^*(g_0, d, PB, PT)$, which includes a subset of links in g_0 , namely those that carry traffic. Also, intuitively, and if we assume that flows in the network are infinitely divisible [13, 14], we may expect each unit of such a flow (no matter how small it is) to be routed along a path of maximum profit per unit of flow, otherwise such a flow would be rerouted in a path that yields a higher profit per-unit of flow. This notion of equilibrium is called Wardrop equilibrium [39].

We can then write for our problem the following definition: a set of flows s^* feasible for instance (g_0, d, PG, PT) is at Nash equilibrium if and only if for every pair $\{(i, u) \in V; (j, v) \in V\}$ and \bar{p}_1, \bar{p}_2 in $\bar{P}(i, u, j, v)$ we have: if $f_\theta^*(\bar{p}_1) > 0$ for some $\theta \in F$ then $\Pi(\bar{p}_1, s_i^*, s_{-i}^*) \geq \Pi(\bar{p}_2, s_i^*, s_{-i}^*)$. Therefore, we can say that $\Pi(\bar{p}, s_i^*, s_{-i}^*) \leq \Pi_{i, u, j, v}^*$ for every path \bar{p} in $\bar{P}(i, u, j, v)$, where $\Pi_{i, u, j, v}^*$ represents the profit per unit of flow along any path \bar{p} in $\bar{P}(i, u, j, v)$ used to deliver traffic from (i, u) to (j, v) in the network at Nash equilibrium.

We call optimal any network that would emerge if ISPs were to cooperate to reduce overall provisioning costs, still using transit and peering agreements to interconnect. We denote this set of networks by $G^{**}(g_0, d, PB)$. This overall cost minimization problem can be written in the following way:

$$\begin{aligned} \min_s \sum_{i,u,j,v,\bar{p}} \sum_{\theta \in F} TC(\bar{p}, s_i, s_{-i}) &\Leftrightarrow \\ \Leftrightarrow \max_s \sum_{i,u,j,v,\bar{p}} \sum_{\theta \in F} T\Pi(\bar{p}, s_i, s_{-i}) &\quad (3) \end{aligned}$$

subject to the constraints (1)-(5) introduced before and $\sum_{i,u,j,v,\bar{p}}$ indicates $\sum_{i \in I} \sum_{u \in C_i} \sum_{j \in I} \sum_{v \in C_j} \sum_{\bar{p} \in \bar{P}(i,u,j,v)}$. Note that the equivalence in the above expression comes from the fact that the sum of all transit payments cancel out and therefore the provisioning cost of an optimal network is independent of the prices of transit.

Therefore, we can also say that for a feasible set of flows s^{**} optimal for instance (g_0, d, PB, PT) , we have: if $f_{\theta}^{**}(\bar{p}_1) > 0$, for some $\theta \in F$, then $T\Pi'(\bar{p}_1, s_i^{**}, s_{-i}^{**}) \geq T\Pi'(\bar{p}_2, s_i^{**}, s_{-i}^{**})$, for every pair $\{(i, u) \in V; (j, v) \in V\}$, for any two paths $\bar{p}_1, \bar{p}_2 \in \bar{P}(i, u, j, v)$. Also, note that this definition for s^{**} means that ISPs route traffic over paths of maximal profit per unit of flow. Therefore, we can say that $T\Pi'(\bar{p}, s_i^{**}, s_{-i}^{**}) \leq T\Pi'_{i,u,j,v}^{**}$, for every path \bar{p} in $\bar{P}(i, u, j, v)$, where $T\Pi'_{i,u,j,v}^{**}$ represents the marginal profit along any path used to deliver traffic from $(i, u) \in V$ to $(j, v) \in V$ in the optimal network.

The striking similarity between the characterization of optimal flows and flows at Nash equilibrium [2] provides an interpretation for optimal flows as flows at Nash equilibrium with respect to different profit-per-unit-of-flow functions. Formally, we can say that a set of flows s^{**} feasible for exogenous parameters g_0, d and profit functions Π is optimal if and only if it is at Nash equilibrium for exogenous parameters g_0, d and profit functions $T\Pi'$. Note that $T\Pi'(e, s_i, s_{-i}) = \Pi(e, s_i, s_{-i}) + f(e) \cdot \partial \Pi(e, s_i, s_{-i}) / \partial f(e)$. While the definition of Nash equilibrium takes into account only the term $\Pi(e, s_i, s_{-i})$, which captures the profit per unit of flow over edge e , in this case there is an additional term that accounts for the change in profit experienced by the inframarginal flow on this edge. This second term induces the flow on each edge to behave unselfishly and it is usually called the altruistic term.

Let $IC(g, g_0, d, PB)$ represent the inefficiency cost of network $g \in G(g_0, d)$, which is defined as the ratio between the costs to provision this network and the costs to provision any optimal network in $G^{**}(g_0, d, PB)$, that is,

$$\begin{aligned} IC(g, g_0, d, PB) &= \sum_{e \in E} TC(e, s_i, s_{-i}) / \sum_{e \in E} TC(e, s_i^{**}, s_{-i}^{**}) \\ &= \sum_{i,u,j,v} \sum_{\bar{p} \in \bar{P}(i,u,j,v)} TC(\bar{p}, s_i, s_{-i}) / \\ &\quad \sum_{i,u,j,v} \sum_{\bar{p} \in \bar{P}(i,u,j,v)} TC(\bar{p}, s_i^{**}, s_{-i}^{**}) \end{aligned} \quad (4)$$

where s and s^{**} are, respectively, the set of flows in network g and in any of the optimal networks. Clearly, $IC(g, g_0, d, PB) \geq 1$ for all possible networks $g \in G(g_0, d)$ and there is no inefficiency cost when this ratio equals 1.

We use the term cost of anarchy to refer to an upper bound on the inefficiency cost over all possible Nash networks. That is, the cost of anarchy is the worst possible inefficiency cost experienced at Nash equilibrium taken over all the possible assignments of (g_0, d, PB, PT) , as long as $g_0 \in G_0(d)$. Therefore, the cost of anarchy depends exclusively on the price of bandwidth, PB , because any other changes in the other exogenous parameters of our model will simply result in different topologies. Let $\kappa(PB)$ denote the cost of anarchy. Formally,

$$\begin{aligned} \kappa(PB) &= \\ \sup_{(g_0, d, PB, PT): g_0 \in G_0(d), g^* \in G^*(g_0, d, PB, PT)} \{IC(g^*, g_0, d, PB)\} &\quad (5) \end{aligned}$$

The cost of anarchy is also always greater than 1. A cost of anarchy of 1 means that there is never inefficiency cost no matter the network graph, the exogenous demand, the exogenous price of bandwidth or the exogenous prices of transit considered. Finally, consider the following definitions for $\lambda(\bar{p})$ and for $\mu(\bar{p})$ for all $\bar{p} \in \bar{P}(i, u, j, v)$ (similar to those introduced in [32]):

$$\begin{aligned} TC'(\bar{p}, s_i(\bar{p}, \lambda(\bar{p})f(\bar{p})), s_{-i}(prd(\bar{p}), \lambda(\bar{p})f(prd(\bar{p}))) &= \\ C(\bar{p}, s_i, s_{-i}) &\quad (6) \end{aligned}$$

$$\begin{aligned} \mu(\bar{p}) &= \\ C(\bar{p}, s_i(\bar{p}, \lambda(\bar{p})f(\bar{p})), s_{-i}(prd(\bar{p}), \lambda(\bar{p})f(prd(\bar{p})))) / C(\bar{p}, s_i, s_{-i}) &\quad (7) \end{aligned}$$

Existence of $\lambda(\bar{p}) \in (0, 1)$ follows from the Intermediate Value Theorem and the facts that $C(\bar{p}, s_i, s_{-i}) \geq TC'(\bar{p}, s_i, s_{-i})$ and that $TC'(\bar{p}, s_i(\bar{p}, 0), s_{-i}(prd(\bar{p}), 0)) \geq C(\bar{p}, s_i(\bar{p}, x), s_{-i}(\bar{p}, x))$, $\forall x \geq 0$, where $s(\bar{p}, x)$ indicates the same strategy as s expect that $f(\bar{p})$ is substituted by x . $\lambda(\bar{p})$ represents the shrinking factor that must be applied to flow $f(\bar{p})$ in order for its marginal cost to be equal to its cost per unit of traffic. $\mu(\bar{p})$ represents the ratio between the cost per unit of traffic of the reduced flow and that of the original flow.

7. THE COST OF ANARCHY

Existence of optimal network and Nash networks is established in [7], which also shows that these networks are not unique. We now state the main result of this section (refer to [10] for a detailed proof):

Result 1

$$\kappa(PB) \geq \sup_{\bar{p} \in \bar{P}} (1 - 1/\mu(\bar{p}) + \lambda(\bar{p})/\mu(\bar{p})) / (\lambda(\bar{p})\mu(\bar{p})) \quad (8)$$

Assume that price of bandwidth is given by an exponential function of the form $PB(M, L) \cdot M = \alpha M^\beta P(L)$ where $\alpha \geq 0$, $\beta \in (0, 1)$ represents the level of economies of scale and $P(L)$ represents the dependency of this price function on the length of the link. Empirical evidence to support this sort of price function was provided in [7]. In this case, we have that $\lambda(\bar{p}) = \beta^{-1}/(\beta-1)$ and $\mu(\bar{p}) = \beta^{-1}$ for every $\bar{p} \in \bar{P}(i, u, j, v)$, which imply $\kappa(\beta) \geq \beta^{1/(\beta-1)}(\beta + \beta^{2(\beta-1)} - \beta^2)$. Figure 1 shows how this lower bound for the cost of anarchy behaves as a function of the level of economies of scale in the price of bandwidth.

A little algebra shows that this lower bound for the cost of anarchy converges to 1 as β goes both to 0 and to 1. When the level of economies of scale is significant, which happens for β close to 0, ISPs observe the benefits from aggregating traffic, both within their networks and over interconnection agreements, and seize these benefits, which results in a low anarchy value. On the other hand, when economies of scale are not very significant, which happens for β close to 1, the cost per unit of bandwidth deployed is about the same everywhere in the network and therefore there is no immediate benefit from aggregating traffic. In other words, the optimal network is not much better than the network at Nash equilibrium and, again, anarchy is low. The maximum for this lower bound for the cost of anarchy is 26%. The results provided in [7] show that currently $\beta \simeq 0.5269$. According to Figure 1, the cost of anarchy for this level of β is at least 25%.




Figure 1: Lower bound for the cost of anarchy as a function of the level of economies of scale in the price of bandwidth.

Note that this lower bound for the cost of anarchy is tight. In fact, it occurs in networks where only peering agreements are used (in which $f^*(\bar{p}) = f^{**}(\bar{p})$ for every path \bar{p} in the network, and such networks exist at Nash equilibrium as discussed in [9]). Finally, note that the cost of anarchy converges to 1 as the level of economies of scale vanishes and therefore our results blend perfectly with those encountered in the literature, since in [32] a price of anarchy is computed for $\beta \geq 1$, which converges to 1 as β gets close to 1.

8. LOSS AND INTERCONNECTION

We want to study the relationship between inefficiency cost and the type of interconnection agreement. For this matter, consider the following result.

Result 2

For an assignment of the exogenous parameters g_0, d, PB, PT , such that $\partial\Pi(\bar{p}, s_i, s_{-i})/\partial f(\bar{p}) = (\beta - 1) \cdot \Pi(\bar{p}, s_i, s_{-i})/f(\bar{p})$ for every $\bar{p} \in \bar{P}(i, u, j, v)$ and for some constant $\beta \in [0, 1]$, with $g_0 \in G_0(d)$, we have $G^*(g_0, d, PB, PT) = G^{**}(g_0, d, PB, PT)$.

Proof: We will show that the flows of a strategy s^* at Nash equilibrium are also the flows of an optimal network and vice-versa. Consider two paths \bar{p}_1 and \bar{p}_2 in $\bar{P}(i, u, j, v)$ used to deliver traffic at Nash equilibrium. Therefore, we must have $\Pi(\bar{p}_1, s_i^*, s_{-i}^*) = \Pi(\bar{p}_2, s_i^*, s_{-i}^*) > \Pi(\bar{p}_3, s_i^*, s_{-i}^*)$, where \bar{p}_3 is any unused path in $\bar{P}(i, u, j, v)$. Recall that flows in the optimal network for paths $\bar{p}_1, \bar{p}_2, \bar{p}_3$ in $\bar{P}(i, u, j, v)$ are characterized by:

$$\begin{aligned} T\Pi'(\bar{p}_1, s_i, s_{-i}) &= \Pi(\bar{p}_1, s_i, s_{-i}) + f(\bar{p}_1) \cdot \partial\Pi(\bar{p}_1, s_i, s_{-i})/\partial f(\bar{p}_1) \\ &= \Pi(\bar{p}_2, s_i, s_{-i}) + f(\bar{p}_2) \cdot \partial\Pi(\bar{p}_2, s_i, s_{-i})/\partial f(\bar{p}_2) \\ &= T\Pi'(\bar{p}_2, s_i, s_{-i}) > T\Pi'(\bar{p}_3, s_i, s_{-i}) \\ &= \Pi(\bar{p}_3, s_i, s_{-i}) + f(\bar{p}_3) \cdot \partial\Pi(\bar{p}_3, s_i, s_{-i})/\partial f(\bar{p}_3) \end{aligned} \quad (9)$$

It follows that:

$$\begin{aligned} &\Pi(\bar{p}_1, s_i^*, s_{-i}^*) + f(\bar{p}_1) \cdot \partial\Pi(\bar{p}_1, s_i^*, s_{-i}^*)/\partial f(\bar{p}_1) \\ &= \Pi(\bar{p}_1, s_i^*, s_{-i}^*) + (\beta - 1) \cdot f(\bar{p}_1) \cdot \Pi(\bar{p}_1, s_i^*, s_{-i}^*)/f(\bar{p}_1) \\ &= \Pi(\bar{p}_1, s_i^*, s_{-i}^*) + (\beta - 1) \cdot \Pi(\bar{p}_1, s_i^*, s_{-i}^*) \\ &= \Pi(\bar{p}_2, s_i^*, s_{-i}^*) + (\beta - 1) \cdot \Pi(\bar{p}_2, s_i^*, s_{-i}^*) \\ &= \Pi(\bar{p}_2, s_i^*, s_{-i}^*) + (\beta - 1) \cdot f(\bar{p}_2) \cdot \Pi(\bar{p}_2, s_i^*, s_{-i}^*)/f(\bar{p}_2) \\ &= \Pi(\bar{p}_2, s_i^*, s_{-i}^*) + f(\bar{p}_2) \cdot \partial\Pi(\bar{p}_2, s_i^*, s_{-i}^*)/\partial f(\bar{p}_2) \end{aligned} \quad (10)$$

and therefore $T\Pi'(\bar{p}_1, s_i^*, s_{-i}^*) = T\Pi'(\bar{p}_2, s_i^*, s_{-i}^*)$. To show that s^* is also the set of flows in an optimal network, it suffices to show, for example, that $T\Pi'(\bar{p}_3, s_i^*, s_{-i}^*) < T\Pi'(\bar{p}_2, s_i^*, s_{-i}^*)$ for any unused path \bar{p}_3 in s^* . Noting that $T\Pi'(\bar{p}_2, s_i^*, s_{-i}^*) = \beta\Pi(\bar{p}_2, s_i, s_{-i})$ it follows that

$$\begin{aligned} T\Pi'(\bar{p}_3, s_i^*, s_{-i}^*) &= \Pi(\bar{p}_3, s_i^*, s_{-i}^*) + f(\bar{p}_3) \cdot \partial\Pi(\bar{p}_3, s_i^*, s_{-i}^*)/\partial f(\bar{p}_3) \\ &= \Pi(\bar{p}_3, s_i^*, s_{-i}^*) + (\beta - 1) \cdot f(\bar{p}_3) \cdot \Pi_{i,u,j,v}(s_i^*, s_{-i}^*)/f(\bar{p}_3) \\ &= \beta\Pi(\bar{p}_3, s_i^*, s_{-i}^*) < \beta\Pi(\bar{p}_2, s_i^*, s_{-i}^*) = T\Pi'(\bar{p}_2, s_i^*, s_{-i}^*) \end{aligned} \quad (11)$$

The condition on the shape of the profit function in this result requires, essentially, that the profit per unit of traffic follows an exponential law, but with the same level of economies of scale β everywhere in the network. In such a case there is no inefficiency cost at Nash equilibrium. Thus, inefficiency cost might arise from the fact that the market for the transport of traffic is not homogenous and there are different levels of economies of scale across cities. Links where the level of economies of scale is higher tend to attract more traffic, because ISPs realize the benefits from transporting aggregate traffic over these links as opposed to splitting such traffic over separate links. While these advantages may be fully explored when providers cooperate to minimize overall provisioning costs, that will hardly be the case when they behave selfishly. Each ISP seizes the benefits of strong economies of scale within its own network (and over interconnection agreements to other ISPs whenever possible), but this optimization is done locally at each network. As a consequence, traffic might be shifted differently across paths and might be handed-off between ISPs at sub-optimal interconnection points, which, on aggregate, results in higher provisioning costs.

Consider the model introduced in the previous section but now assume that the cost of every link used for a transit agreement becomes infinity, or in other words, assume that only peering agreements are allowed between ISPs. Mathematically, $C(e, s_i, s_{-i}) = \infty$ for $e \in (B_{-1} \cup B_{+1})$. In this modified version of our model, let $G_p^*(g_0, d, PB, PT)$ represent a network at Nash equilibrium and let $G_p^{**}(g_0, d, PB)$ represent an optimal network (always with $g_0 \in G_0(d)$). In this case, the profit function of each ISP reduces to a summation of costs, since no revenues can be realized. But the cost of an internal link and the cost of a peering agreement are the same because the cost of a peering agreement is just the cost of the bandwidth link between the peering ISPs. Therefore, the cost of a peering agreement exhibits the same level of economies of scale as internal links do. Hence, in a world where only peering agreements are allowed, no inefficiency cost arises between Nash and optimal networks because these networks are the same, that is, $G_p^*(g_0, d, PB, PT) = G_p^{**}(g_0, d, PB)$, for all possible assignments of (g_0, d, PB) with $g_0 \in G_0(d)$ and for any PT .

Another way to look at this issue is to notice that as a peering agreement is no more than establishing a direct link between the premises of the two interconnecting ISPs, its price reflects directly the true cost of the infrastructure needed to provision it. In other words, peering agreements do not introduce any distortion between prices and costs and when prices reflect costs, we expect ineffi-

ciency to be low. Inefficiency costs arise in the presence of transit agreements. According to [7], the price of a transit agreement exhibits a different level of economies of scale than internal links and peering agreements do. This is because transit agreements include a service charge that does not reflect the true cost that the ISP selling transit incurs to transport the traffic to and from the customer ISP. Usually, this fee is only a function of the amount of traffic exchanged and does not take into account, by any measure, the distribution of sources and destinations for such traffic. In practice, this fee is mainly determined by market conditions and thus it is poorly related to the true managerial and routing costs that the ISP selling transit incurs. Therefore, a transit agreement includes a certain element of hidden information because the ISP selling transit commits to provide a service for a fee, established a-priori, without knowing exactly the true cost it will incur to provide such service. We can then expect this distortion between costs and prices to result in some inefficiency.

9. MANIPULATING TRANSIT PRICES

We showed that when only peering agreements are allowed no inefficiency cost arises. This only-peering situation can be modelled by adding a new constraint to the model before. To rule out transit agreements, we can simply say that traffic in any transit link in the network must equal zero. Therefore, the overall cost minimization problem that defines the optimal network is, in the case of only peering, a minimization problem over a feasible set that is a subset of the feasible set of the overall cost minimization problem when both peering and transit agreements are allowed. Therefore, it is trivial to show that the costs to provision the optimal network cannot increase when transit becomes available.

In reality, transit agreements contribute significantly to reduce provisioning costs because they allow for better aggregating traffic that is destined to the same places in the network, which might in turn result in networks with fewer links. For example, a network with n nodes and only peering requires on the order of $n(n-1)/2$ links, whereas a network with only transit agreements will most likely result in a tree structure that only needs on the order of $n \cdot \log(n)$ links. NAPs have also emerged as a way to reduce the number of links in the network. An ISP can run a fat single link to the NAP, which it uses to carry all the traffic to all the ISPs with whom it interconnects at the NAP, instead of running several separate links, one to each of these ISPs.

In the case of networks with both peering and transit agreements, the Nash equilibrium is not unique and different Nash networks can exhibit different provisioning costs. The optimal network is not unique as well. However, all optimal networks will have the same overall provisioning cost. More importantly, an optimal network does not depend on the prices of transit. This is because the optimal network is a cost minimization problem over the entire network and whatever payments ISPs make to each do not really matter (see expression 7). The fact that sometimes the optimal network uses transit links depends only on how using these links allows for aggregating more traffic thus reducing costs.

We are now interested in knowing if there are cases in which there exist Nash networks that are strictly more expensive than any optimal network. If this is the case it will be interesting to know if there are transit prices that make a network at Nash equilibrium look more like an optimal network. If so, a regulator could help reduce overall provisioning costs by manipulating such prices.

The following is a very simple example that shows that a Nash network can be more expensive than the optimal network. Consider ISP i serving city u and city v and ISP j serving city r . There is an exogenous demand of 10 Mbps from users of ISP j at city r to users of city ISP i at city v (we can assume that there is also a demand of traffic less than 10 Mbps from the latter users to the former users, but this changes nothing in the what come next). All other demand in the network is set to zero. Furthermore, assume that the length of shortest links between city u and city r , between city v and city r and between city u and city v are, respectively, 2105, 2125 and 332 miles, which is clearly a realizable network. Assume that the price of bandwidth links is given by the function $PB(M, L) = 248.29M^{0.5269}L^{0.3774}$ which we have estimated in previous work [7]. Figure 2 illustrates this network.




Figure 2: Example of two ISPs serving three cities in which there is a Nash network strictly more expensive than an optimal network.

In this case, the optimal network is to send 10 Mbps from ISP j at city r to ISP i at city v for a total provisioning cost of \$15050. Now, consider the network in which ISP i at city u and ISP j at city r peer. In this latter case, ISP i and ISP j pay a total cost of \$14970 and \$7500, respectively, for a total provisioning cost of \$22740. We can show that this network, which is more expensive than the optimal network, is at Nash equilibrium. For that, we will show that ISP i and ISP j cannot agree to change it.

An alternative network would be to have ISP i at city v and ISP j at city r peer. In this case, both ISP i and ISP j would have to pay \$7525. Therefore, ISP i is not willing to proceed with the change. Other alternative networks include transit agreements. Consider the network in which ISP i at city v buys transit from ISP j at city r . In this case, ISP i would have to pay $15050+T$ and ISP j would have receive T , where T represents the transit payment. Therefore, ISP i would always be worse off and would not agree to change. A similar reasoning applies when ISP i at city v sells transit to ISP j at city r . The remaining two network configurations involve transit between ISP i at city u and ISP j at city r . Consider the network in which ISP i at city u buys transit from ISP j at city r . In this case, ISP i would have to pay $22470+T$ and ISP j would have receive T , where again T represents the transit payment. Therefore, ISP i would always be worse off and would not agree to change. Finally, consider the case in which ISP i at city u sells transit to ISP j at city r . In this case, ISP i would incur a cost of $7470-T$ and ISP j would incur a cost of $15000+T$. Therefore, ISP j would always become worse off and would not agree to change. In sum, ISP i and ISP j cannot find an alternative network to which they would both agree to change and the network with the peering link between ISP i at city u and ISP j at city r (and the internal link in ISP i 's network between city u and city v) is a Nash equilibrium, which is more expensive to provision than the optimal network.

Finally, note that not all possible networks in the previous example are at Nash equilibrium. For example, the network in which ISP i at city u buys transit from ISP j at city r is not a Nash equilibrium. To see this, note that under that network ISP i pays $\$22470+T_1$ and ISP j receives T_1 , where T_1 is the transit payment. But, consider the network in which ISP i at city v buys transit from ISP j at city r . In this case, ISP i pays $\$15050+T_2$ and ISP j receives T_2 , where T_2 is the transit payment in this case. Therefore, any positive payment T_2 between T_1 and $T_1 - 7470$ will make both ISPs agree to change to the latter network and hence the former one is never a Nash equilibrium.

The example above proves that the market is not efficient to provision interconnected communication networks because even when ISPs choose transit prices there exist Nash networks that are (strictly) more expensive than the optimal networks. The following result shows that for each optimal network there exist side payments among ISPs that make this network also a Nash network.

Result 3

For each assignment of the exogenous parameters g_0, d, PB, PT , with $g_0 \in G_0(d)$, for each optimal network g^{**} in $G^{**}(g_0, d, PB)$ there exist a set of side payments among ISPs $T(i, u, v, j)$, defined for $i \in I, j \in I : i < j, u \in C_i, j \in C_j$, that makes g^{**} at Nash equilibrium.

Proof: Using the language introduced before, let g^{**} represent any optimal network in $G^{**}(g_0, d, PB)$ and let g represent any network in $G(g_0, d)$, therefore feasible, since $g_0 \in G_0(d)$. Additionally, let $\Pi(s_i, s_{-i})$ represent the profit of ISP i under g and let $\Pi(s_i^{**}, s_{-i}^{**})$ represent the profit of this ISP under g^{**} , for any $i \in I$. Also, write $\Pi(s_i^{**}, s_{-i}^{**}) = \Pi(s_i, s_{-i}) + \Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**})$. We do not know the sign of each $\Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**})$, but we know that

$$\sum_{i \in I} \Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) \geq 0 \quad (12)$$

We will show that there exist a set of side payments among ISPs $T(i, u, v, j)$ that make every ISP at least as good off in network g^{**} as in network g , which is sufficient to prove that network g^{**} is at Nash equilibrium once these side payments are realized. In this new setting, the profit of ISP i becomes $\Pi(s_i^{**}, s_{-i}^{**}) = \Pi(s_i, s_{-i}) + \Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) + \Delta T(i)$ where $\Delta T(i)$ is defined by

$$\Delta T(i) = - \sum_{j=1}^{i-1} \sum_{u \in C_i} \sum_{v \in C_j} T(j, v, i, u) + \sum_{j=i+1}^n \sum_{u \in C_i} \sum_{v \in C_j} T(i, u, j, v) \quad (13)$$

for $i = 1, \dots, n$. Now, proving our theorem reduces to showing that there exist these side payments $T(i, u, j, v)$ such that $\Delta T(i) + \Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) \geq 0$ for all $i = 1, \dots, n$. We will show this is true by induction on n . For $n = 1$ and $n = 2$ these statements are trivial. Assume these statements are true for some number n of ISPs, that is there exist $T(i, u, j, v)$ such that

$$\Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) - \sum_{j=1}^{i-1} \sum_{u \in C_i} \sum_{v \in C_j} T(j, v, i, u) + \sum_{j=i+1}^n \sum_{u \in C_i} \sum_{v \in C_j} T(i, u, j, v) \geq 0 \quad (14)$$

for every $i = 1, \dots, n$. We want to show that this implies the existence of $\bar{T}(i, u, j, v)$ such that

$$\Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) - \sum_{j=1}^{i-1} \sum_{u \in C_i} \sum_{v \in C_j} \bar{T}(j, v, i, u) + \sum_{j=i+1}^{n+1} \sum_{u \in C_i} \sum_{v \in C_j} \bar{T}(i, u, j, v) \geq 0 \quad (15)$$

for every $i = 1, \dots, n+1$. Let $\bar{T}(i, u, j, v) = T(i, u, j, v)$ when j is less than $n+1$, and our proof reduces to show that $\bar{T}(i, u, n+1, v)$ with $v \in C_{n+1}$ exist. Summing expression 15 over $i = 1, \dots, n$ and isolating the last term in the last summation, we obtain

$$\begin{aligned} & \sum_{i=1}^n \sum_{u \in C_i} \sum_{v \in C_{n+1}} \bar{T}(i, u, n+1, v) \geq \\ & \sum_{i=1}^n -\Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) + \\ & \sum_{i=1}^n [\sum_{j=1}^{i-1} \sum_{u \in C_i} \sum_{v \in C_j} T(j, v, i, u) - \\ & \sum_{j=i+1}^n \sum_{u \in C_i} \sum_{v \in C_j} T(i, u, j, v)] \Leftrightarrow \\ & \sum_{i=1}^n \sum_{u \in C_i} \sum_{v \in C_{n+1}} \bar{T}(i, u, n+1, v) \geq \\ & \sum_{i=1}^n -\Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) \end{aligned} \quad (16)$$

In addition, expression 15 for $i = n+1$ implies that

$$\Delta\Pi_{n+1}(s_{n+1}, s_{-(n+1)}, s_{n+1}^{**}, s_{-(n+1)}^{**}) \geq \sum_{j=1}^n \sum_{u \in C_{n+1}} \sum_{v \in C_j} \bar{T}(j, v, n+1, u) \quad (17)$$

Note that the summation on the right side of this inequality is the same as the summation on the left side of the last inequality in expression 17 and hence this summation is bounded from below and above by expressions 16 and 17. Therefore, for $\bar{T}(i, u, n+1, v)$ with $v \in C_{n+1}$ to exist we need to have

$$\Delta\Pi_{n+1}(s_{n+1}, s_{-(n+1)}, s_{n+1}^{**}, s_{-(n+1)}^{**}) \geq \sum_{i=1}^n -\Delta\Pi_i(s_i, s_{-i}, s_i^{**}, s_{-i}^{**}) \quad (18)$$

which is obviously true, since this is just expression 12 for the case of $n+1$ ISPs. Hence, side payments $T(i, u, j, v)$ exist that make optimal network g^{**} a Nash network. When these side payments are realized, this optimal network can be obtained as a result of the non-cooperative interaction among ISPs. This result can be seen as an instance of what is known as the Coase theorem [4]: if there are ways to find further gains from trade then agents will effect these trades over bargaining [22]. In this case, bargaining is achieved through the side payments.

10. DISCUSSION AND CONCLUSIONS

This paper addresses the issue of economic inefficiencies in terms of provisioning costs for interconnected communication networks. A number of ISPs, assumed to behave rationally, engage in a non-cooperative game to build their networks and to establish transit and peering agreements to interconnect and meet some exogenous demand for the transport of IP traffic between an arbitrary set of cities. A model of interconnected networks based on the traditional models of multi-commodity flows is provided to study how the configuration of Nash networks differs from that of a network that minimizes overall provisioning costs. We use the term cost of anarchy to refer to the worst-case efficiency loss in provisioning costs. A lower and tight bound for the cost of anarchy varies as an inverted U-shaped function of the level of economies of scale in the price of bandwidth. Anarchy is low when the latter are either low or high. For the current level of economies of scale in the price of bandwidth, the cost of anarchy is at least 25%.

The cost of anarchy can be seen as the worst-case social cost to allow several providers to provision the network independently and selfishly. However, a market with several providers promotes competition, which is likely to drive down the price of the services

offered over the network. If ISPs were to cooperate or to collude to provision the cheapest possible network, we would be in the presence of some sort of over-arching centralized authority that dictates where ISPs should interconnect and how they should route traffic within their networks. But one could also expect this entity to maximize profits with the services offered over the network and users would then be facing a monopolist. This monopolist would likely abuse its market power and extract monopoly rents from users. Some sort of regulation, most likely price regulation, would then have to be in place to prevent such behavior. The appropriate regulation to mitigate the possible harmful effects of the monopolistic situation described above would come as a cost to the regulator. If this cost is greater than the cost of anarchy, overall society will be worse off (where, in this case, society comprises the regulator, the ISPs and end-users). If this is the case, it is better to let the market run and bear the inefficiencies in terms of provisioning costs that will certainly arise. This will still be cheaper than regulating the market and will free up the regulator to look at other, possibly more pertinent, policy issues.

We have also shown that inefficiency costs are primarily related to transit agreements. In fact, in a world where only peering agreements are allowed no inefficiency arises. However, networks with only peering agreements become rather expensive. Transit agreements enable the benefits of economies of scale to be realized and reduce overall provisioning costs. Nevertheless, we showed, by example, that there are Nash networks that are strictly more expensive than optimal networks when ISPs choose transit prices and therefore the market for provisioning communication networks with interconnection agreements such as peering and transit designed they way they are today is inefficient. We have also shown that for every optimal network there is a set of side payments that makes this network a Nash network. However, it is not guaranteed that ISPs will realize such payments and overall inefficiency cost is likely to arise.

The market failure identified above suggests that there is scope for the regulator to intervene, for example, by enforcing side payments or regulating transit prices. However, to do so, the regulator would face a series of difficulties that lead us to believe that it is better to allow the market to run and to give us a good allocation of bandwidth and interconnection agreements yet possibly not the optimal one. First, we do not know how often inefficient situations occur and it would be clearly unwise to have the regulator focus on welfare loss cases that occur seldom. Second, the regulator is at a severe disadvantage in terms of accurate knowledge about the topology of the network and actual traffic flows. In the face of such lack of information, no regulator could create a model of the network good enough to allow him to identify the appropriate changes to reduce provisioning costs. Moreover, the changes in the network that the regulator could promote to reduce provisioning costs depend on the demand for the transport of traffic, which in our model was considered exogenous. Realistically, a model with endogenous demand is needed to incorporate the impact on demand of changes in network costs and, presumably, prices.

Third, and assuming that the regulator can identify these changes, he would have to devise ways to intervene in the market to change the behavior of ISPs, which could include designing the proper incentive mechanisms to induce ISPs to route traffic flows and/or to interconnect differently. However, the process of incentive design and implementation takes time and it is likely that the the state of the world will have changed by the time the regulator's policies are, which could render them useless, if not harmful. Forth, ISPs will

certainly behave differently once they anticipate that the regulator will act upon the market to mitigate welfare losses. Thus, the model that the regulator should use to predict the behavior of ISPs would have to be more complex than the model we presented in this paper, which does not account for the role of the regulator.

Finally, consider the case of international networks, for which no regulator holds overall jurisdiction. Networks that span different countries operate under the regulatory regime defined by different regulators who would then have to cooperate to reduce overall costs, which is a task that might become quite hard to implement. Moreover, different regulators might have different incentives and each of them might behave selfishly to meet his goals, which might not be aligned with the overarching goal of reducing overall provisioning costs. Taking into account all the issues listed above, it is reasonable to say that it is extremely difficult for the regulator to act upon the market inefficiency identified in this paper. It is therefore fair to acknowledge that it might be better to allow the market to work, which will result in a good yet possibly not optimal allocation of bandwidth and of interconnection agreements.

11. REFERENCES

- [1] T. Bates, E. Gerich, L. Joncheray, J. Jouanigot, D. Karrengerg, M. Terpstra, and J. Yu. Representation of IP Routing Policies in a Routing Registry. RFC 1786, 1995.
- [2] M. Beckmann, C. McGuire, and C. Winsten. *Studies in the Economics of Transportation*. Yale University Press, 1956.
- [3] D. Braess. Uber ein Paradoxon der Verkehrsplanung. *Unternehmensforschung*, 12:258–268, 1968. available from <http://homepage.ruhr-uni-bochum.de/Dietrich.Braess/>.
- [4] R. Coase. The Problem of Social Cost. *Journal of Law and Economics*, 1:1–44, 1960.
- [5] G. Cook. Peering, Transit and IXs: Economics and Policy. available at www.cookreport.com, November 2002.
- [6] J. Correa, A. Schulz, and N. Stier. Selfish Routing in Capacitated Networks. In *INFORMS Annual Meeting*, October 2003. Atlanta, Georgia.
- [7] P. Ferreira. Interconnected Communication Networks Provisioned Selfishly. Phd dissertation, Carnegie Mellon University, 2004.
- [8] P. Ferreira, J. Mindel, and L. McKnight. Why Bandwidth Trading Markets Have not Matured? Analysis of Technological and Market Issues. *International Journal of Technology, Management and Policy*, 2, 2004.
- [9] P. Ferreira and M. Sirbu. The Inefficiency of Markets for Provisioning Communication Networks. In *32nd Annual Telecommunication Policy Research Conference*, October 1-3, Arlington, VA 2004. School of Law, George Mason University.
- [10] P. Ferreira and M. Sirbu. The Price of Anarchy for Interconnected Communication Networks Provisioned Selfishly. In *Conference on Competition in Networks: Wired and Wireless*, May 2004. London Business School, London, UK.
- [11] E. Friedman. A Generic Analysis of Selfish Routing. Working Paper, available at <http://www.orie.cornell.edu/friedman/papers.htm>, 2001.
- [12] S. Halabi and D. McPherson. *Internet Routing Architectures*. Cisco Press, 2000.
- [13] A. Haurie and P. Marcotte. On the relationship between Nash-Cournot and Wardrop equilibria. *Networks*, 15:295–308, 1985.

- [14] A. Haurie and P. Marcotte. A Game-theoretic Approach to Network Equilibrium. *Mathematical Programming Study*, 26:252–255, 1986.
- [15] J. Hawkison and T. Bates. Guidelines for Creation, Selection and Registration of an Autonomous System (AS). RFC 1930, 1996.
- [16] G. Huston. *ISP Survival Guide: Strategies for Running a Competitive ISP*. John Wiley & Sons, 1999.
- [17] R. Johari. Efficiency loss in market mechanisms for resource allocation. Phd dissertation, Massachusetts Institute of Technology, 2004.
- [18] R. Johari and J. Tsitsiklis. Efficiency loss in a network resource allocation game. *Mathematics of Operations Research*, 3:407–435, 2004.
- [19] M. Kende. The Digital Handshake: Connecting Internet Backbones. Working Paper 32, Office of Plans and Policy, Federal Communications Commission, FCC, September 2000.
- [20] E. Koutsoupias and C. Papadimitriou. Worst-case Equilibria. In *16th Annual Symposium on Theoretical Aspects of Computer Science*, pages 404–413, 1999.
- [21] J. Marcus. Global Traffic Exchange Among Internet Service Providers (ISPs). In *OECD - Internet Traffic Exchange*, June 2001. Berlin, Germany.
- [22] A. Mas-Colell, M. Whinston, and J. Green. *Microeconomic Theory*. Oxford University Press, 1995.
- [23] T. McGarty. Peering, Transit and Interconnection: Internet Access in Central Europe. MIT Internet Telephony Consortium, Working Paper, January 2002.
- [24] J. Nash. Non-cooperative Games. *Annals of Mathematics*, 54(2):286–295, 1951.
- [25] W. Norton. Interconnection Strategies for ISPs. White Paper, Equinix, 1999.
- [26] W. Norton. Internet Service Providers and Peering. White Paper, Equinix, 1999.
- [27] M. O’Kelly and T. Grubescic. Backbone Topology, Access and the Commercial Internet: 1997-2000. *Environmental and Planning B*, 29:533–552, 2002.
- [28] C. Pigou. *The Economics of Welfare*. Macmillan, 1920.
- [29] J. Postel. The Transport Control Protocol(TCP). RFC 793, 1985.
- [30] L. Qiu, Y. Yang, Y. Zang, and S. Shenker. On Selfish Routing in Internet Like-Environments. In *SIGCOMM’03*, August 25-29 2003. Karlsruhe, Germany.
- [31] T. Roughgarden. Designing Networks for Selfish Users is Hard. *42nd Annual Symposium on Foundation of Computer Science*, pages 472–481, 2001.
- [32] T. Roughgarden. The Price of Anarchy in Networks with Polynomial Edge Latency. Technical report tr2001-1847, Cornell University, 2001.
- [33] T. Roughgarden. Selfish Routing. PhD Dissertation, 2002.
- [34] T. Roughgarden. The Price of Anarchy is Independent of the Network Topology. In *34th Annual ACM Symposium on the Theory of Computing*, 2002.
- [35] J. Stewart. *BGP-4: Inter-Domain Routing in the Internet*. Addison-Wesley Publishers Company, 1998.
- [36] Telegeography. . Light Reading, September 2004.
- [37] Telegeography. . Light Reading, August 2004.
- [38] Telegeography. Annual Traffic Report. , 2004.
- [39] J. Wardrop. Some Theoretical Aspects of Road Traffic Research. In *Institute of Civil Engineers*, volume 1, pages 325–378, 1952. Part II.