

Peering vs Transit: a Game Theoretical Model for Autonomous Systems Connectivity

Giovanni Accongiagioco*, Eitan Altman[†], Enrico Gregori[‡] and Luciano Lenzini[§]

*IMT Institute for Advanced Studies, Lucca, Italy giovanni.accongiagioco@imtlucca.it

[†]INRIA, Sophia Antipolis, France eitan.altman@inria.fr

[‡]Institute of Informatics and Telematics, Italian National Research Council, Pisa, Italy enrico.gregori@iit.cnr.it

[§]Information Engineering Department, University of Pisa, Italy l.lenzini@iet.unipi.it

Abstract—We propose a model to analyze the decisions taken by an Autonomous System (AS) when joining the Internet. We first define a realistic model for the interconnection costs incurred and then we use this cost model to perform a game theoretic analysis of the decisions related to the creation of new links in the Internet. The proposed model doesn't fall into the standard category of routing games, hence we devise new tools to solve it by exploiting peculiar properties of our game. We prove analytically the existence of multiple equilibria for specific cases, and provide an algorithm to compute the stable ones. The analysis of the model's outcome highlights the existence of a Price of Anarchy (PoA) and a Price of Stability (PoS), originated by the non-cooperative behavior of the ASes, which optimize their cost function in a selfish and decentralized manner. We further observe the presence of competition between the facilities providing either transit or peering connectivity, caused by the cost differences between these two interconnection strategies.

Index Terms—Internet Modeling, Complex Networks, Game Theory, AS-level Internet Topology, Supermodular Games

I. INTRODUCTION

The Internet ecosystem is made of tens of thousands Autonomous Systems, interconnected together in a complex and dynamic manner. Roughly speaking, an Autonomous System is a network under a single administrative control. ASes can be grouped in different tiers and categories, depending on the service they offer and the organization they belong to: content providers, access providers, transit providers and so on [1].

The late twentieth-century Internet ecosystem was largely dominated by transit links, where the relationship between the connecting ASes was of “customer to provider” type. This kind of relationship produces a hierarchical pricing scheme, where the customer AS pays its provider for the traffic flowing on the link, both incoming and outgoing; in return, the latter provides a default gateway to reach all Internet's routes. For example, an access provider AS wishing to grant Internet access to the eyeballs (i.e. its end users), needs to establish a link with a transit provider, and pay for the traffic flowing on this link. Transit providers are also known as Network Service Providers (NSP) [2] and their pricing strategy is typically volume-based, metered using the 95th percentile traffic sampling technique (this allows customer ASes to burst, for a limited time period, beyond their committed base rate) [3].

This work has been partially supported by the European Commission within the framework of the CONGAS project FP7-ICT-2011-8-317672.

ISBN 978-3-901882-58-6 © 2014 IFIP

The beginning of twenty-first century brought a new paradigm into the environment, since more and more ASes found it beneficial to establish peering links between them [4]. This kind of relationship is “settlement-free”, meaning that the two ASes mutually agree to exchange traffic for free between them, and the only cost they incur is that of laying out the physical interconnection. Peers must agree to each other's policy, which is used to avoid abuse of the peering relationship. Typical clauses include prohibition of using the peer as default gateway (therefore peers cannot be used to reach other Internet's routes) and traffic ratio balancing, meaning that the ratio between incoming and outgoing traffic over the link must not exceed some value (e.g. 2:1) [5].

The exponential growth of peering links was made possible mainly thanks to the deployment of Internet Exchange Points (IXPs) [6]. These interconnection points are facilities through which ASes can exchange traffic, typically by settlement-free (i.e. peering) relationships. The growth of IXPs, in number and in size, made it easy to establish more and more public peering relationships. In fact, by joining an IXP, an AS can potentially peer with all (usually a subset) of the other ASes connected to the same IXP. The pricing strategy of an IXP, with respect to its customers, is typically flat. Each one of them pays a monthly-based fee, depending on the size (speed) of the port they are using and the cost of maintaining the equipment. Thanks to this mechanism, the IXP can share maintenance costs among all its participants [7]. It is worth noting that this pricing strategy doesn't allow standard cost function modelization (like in [8]), since the addition of new participants potentially brings down the costs of an IXP customer.

When an Autonomous System joins the Internet, he needs to decide the connections to lay out with other ASes. While in the last century, as shown above, the decision space for ASes was substantially small, today they have many alternatives: transit or peering, joining one or more IXPs, dealing with distance and geographic issues. In principle, the best strategy for an AS is the one yielding the lowest cost. However, the outcome of its strategy also depends on that of other ASes dealing with the same problem, thus we find it straightforward to analyze the problem in a game-theoretical framework. We propose a model keeping into account the above factors, which can be used to compute the outcome of this problem and the strategy of the players. Realistic modeling of the whole decision space

of an AS is an extremely difficult task, therefore in our work we restrict our analysis to the problem of peering versus transit. Nevertheless, a proper understanding of this problem is fundamental to get insight on the behavior of ASes in the Internet environment, as shown by the results of the analysis.

This work brings contributions both from a game theoretic perspective and an engineering perspective. First of all this is, to the best of our knowledge, a novel model to analyze the strategic choices of ASes living in an Internet environment with both technological and economic constraints. The modelization takes into account many realistic elements, which do not fall into standard frameworks, yet tries to keep the problem mathematically manageable. From a game-theoretic perspective, we prove that our game falls in a specific category for which we both demonstrate the existence of equilibria and provide an algorithm for computing stable solutions. From an engineering perspective, the outcome of the analysis is highly insightful as it shows both the suboptimality of the decentralized solution and the emerging competition, first observed in [9], between the two facilities enabling either transit or peering connectivity: Network Service Providers and Internet Exchange Points.

The remainder of this paper is organized as follows: in Section II we describe the related work, Section III defines the general model, in Section IV we analyze a simplified model, derive the existence of equilibria and the algorithm to compute them, in Section V we show inefficiencies of the decentralized solution, Section VI shows simulative results for more general scenarios, then we conclude in Section VII.

II. RELATED WORK

This work relates to the characterization and modeling of the Internet AS-level topology. In this field the majority of graph theoretic models try to reproduce observed Internet topological properties, such as its power-law degree distribution [10], the small-world property [11] and other structural properties (communities, cliques, etc.) [12]. The graph can be reconstructed [13] by either defining some attachment criteria, as in [11], [14], [15], [16] or solving constrained optimization problems for the different nodes, like in [17], [18].

In our paper we rather try to understand network formation as the result of a game between ASes. In this context, it relates to game-theoretic network formation models, which populate both computer science and economics literature (see books [19], [20]). This research branch focuses on proving the existence of equilibria in networks with a fixed number of agents, where links are formed taking into account their preferences in the form of utility functions. The need for mathematical tractability requires simplifications that make these models unrealistic and unable to study real life networks such as the Internet. Other models simulating the dynamics of network formation are agent-based computational models such as GENESIS [21]. In this case authors can include more realistic considerations by skipping the analytical part and simulating the behavior of each agent, hoping to find one of the possibly many equilibria.

In order to keep the problem analytically tractable we do not aim at modeling the whole network formation process. We rather study the interaction between ASes which connect to an existing network in order to serve some demands. A possible modelization of a network where access providers need to select a subset of content providers and fetch traffic from them in a cost-efficient manner is given in [22]. However, the aim of this work is quite different from ours, as it concentrates on the economic analysis of neutral/non-neutral network features, without taking into account the difference between traffic and peering agreements. In [23], authors perform an interesting analysis on network pricing and analyse the economics of private internet exchanges. This kind of peering, known as private peering, has different rules and costs compared with public peering. As explained in the introduction, nowadays Internet is largely dominated by public peering, occurring at IXPs, therefore in our work we concentrate on this last phenomenon, which allows us to give different insights on the present difference between transit and peering.

III. GENERAL SCENARIO

In the following we describe the general scenario under investigation and derive the cost function.

A. Description

As said earlier, ASes can be grouped in different categories. Throughout the paper we will use the following categorization, closely resembling the one provided by PeeringDB [2]:

*Internet Service Provider (ISP)*¹: this node gives eyeballs access to the Internet and its contents. Each service provider has a traffic demand, hereafter demand, which represents the amount of traffic (uplink+downlink) that it handles.

*Content Provider (CP)*¹: this node has physical access to the contents users are looking after, therefore an ISP with a demand for his specific content, has to connect to it in order to serve this demand.

*Internet Exchange Point (IXP)*¹: this is a facility that provides peering connection to all its participants. This means that all the nodes connected to a given IXP can potentially communicate with each other.

*Network Service Provider (NSP)*¹: this node is located at the highest hierarchical level of the network, meaning that each CP can be reached through it. ISPs can reach CPs by establishing a transit connection with an NSP.

We consider a network with $i \in \{1, \dots, I\}$ ISPs, $n \in \{1, \dots, N\}$ CPs and $l \in \{1, \dots, L\}$ transmit facilities (TF), that can be either the NSPs or the IXPs. Without loss of generality, we impose that TFs $j_1 \in \{1, \dots, l_1\}$ are NSPs, while TFs $j_2 \in \{l_1 + 1, \dots, L\}$ are IXPs.

The main difference between transmit facilities is that while links to NSPs are established through transit connections, links to IXPs are established through peering connections.

In a *transit connection*, or customer-to-provider (C2P) connection, the cost to the customer is a function of the amount of traffic that crosses the link (typically expressed as \$/Mbps).

¹ISPs, NSPs and CPs are typically ASes. IXPs are not ASes, even if their infrastructure is under a single administrative control.




Figure 1: General Model

In a *peering connection*, the price is generally flat and depends on the size of the port the customer buys. Moreover, when peering connections are maintained by an IXP, the costs are shared among all the participants.

Each ISP i has a demand for a CP n , which we indicate as ϕ_i^n . The players of our game are the ISPs, which need to decide how to split their demand among all possible transmit facilities. Figure 1 depicts our network scenario.

We indicate with $x_{i,l}^n$ the flow from ISP i to CP n via TF l . The strategy of ISP i is given by the vector $\mathbf{x}_i = (x_{i,1}^1, \dots, x_{i,l}^n, \dots, x_{i,L}^N) \in \mathbb{R}^{L \times N}$, while the strategy of all the other players is expressed as $\mathbf{x}_{-i} \in \mathbb{R}^{(L-1) \times L \times N}$. The goal of each player is to serve, at the minimum possible cost, his demand $\phi_i = (\phi_i^1, \dots, \phi_i^N) \in \mathbb{R}^N$ by splitting it into several flows $x_{i,l}^n$. We also indicate as:

- $x_{i,l} = \sum_{n=1}^N x_{i,l}^n$ the total flow from ISP i to TF l
- $x_l = \sum_{i=1}^I \sum_{n=1}^N x_{i,l}^n$ the total flow at TF l

Each player, say i , for each transmit facility (NSP or IXP) it connects to, incurs some costs:

TF usage cost : this cost depends on the transmit facility used. If it is an NSP, then it is a function of $x_{i,l}$, the flow from the player to the NSP. Otherwise the TF is an IXP and the cost is shared among all the participants, therefore it also depends on the other players, in the form of x_l , the total flow at the IXP. Consequently, this cost can be written as a function:

$$t^l(x_{i,l}, x_l) \quad (1)$$

TF capacity cost : each link between an ISP and a TF has a fixed capacity $c_{i,l}$; this means that we have a constraint of the form $x_{i,l} \leq c_{i,l}$. While we may introduce it in the problem “as is”, this would make the model less manageable. Moreover, due to performance and congestion issues, network operators typically avoid reaching the capacity limit and keep a margin for traffic fluctuations. We can think of this performance degradation as a “virtual cost” for the ISP, and therefore model the constraint as a cost, that increases as the flow over the link approaches the capacity limit (as is typically done in the literature for M/G/1 Processor Sharing queues [8]):

$$\frac{1}{c_{i,l} - x_{i,l}} x_{i,l} \quad (2)$$

We are aware that, in reality, network operators adjust this capacity when there is more demand for it, and the inter-connection cost grows accordingly. However, this situation

can be avoided as long as our working region is sufficiently far away from the saturation point. We will always assume that capacities are symmetric w.r.t. the players, therefore $c_{i,l} = c_l \forall i$. Typically the capacity of the NSP can be assumed to be much larger than that of IXPs: $c_{NSP} \gg c_{IXP}$ (see [3] and [7]).

CP reachability cost : let's indicate with b_l^n the cost of transporting one unit of flow from TF l to CP n . This cost is not relevant from the player's perspective (it is paid by the CP), however it can be used to express the reachability of a given CP. In fact, while all the CPs are connected to the NSPs, an IXP can be connected only to a subset of CPs. This phenomenon can be expressed by putting:

$$b_l^n = \begin{cases} 0 & \text{if } l \leq l_1 \text{ or } IXP_l \rightarrow CP_n \\ \infty & \text{otherwise} \end{cases} \quad (3)$$

Thanks to all these considerations, the cost function for player i can be expressed as the sum of (1), (2) and (3):

$$C_i(\mathbf{x}_i, \mathbf{x}_{-i}) = \sum_{l=1}^L \left(t^l(x_{i,l}, x_l) + \frac{1}{c_l - x_{i,l}} x_{i,l} \right) + \sum_{l=1}^L \sum_{n=1}^N x_{i,l}^n b_l^n \quad (4)$$

In order to serve all the demands, each player i has to satisfy the flow constraint: for every CP, the total flow has to be equal to the demand ϕ_i^n . Therefore player i 's best response $BR_i(\mathbf{x}_{-i})$ is obtained by minimizing cost function (4), subject to the flow constraints (5):

$$\begin{cases} BR_i(\mathbf{x}_{-i}) = \arg \min_{\mathbf{x}_i} C_i(\mathbf{x}_i, \mathbf{x}_{-i}) \\ \text{s.t.} \quad \sum_l x_{i,l}^n = \phi_i^n \quad \forall n \end{cases} \quad (5)$$

The vector $\mathbf{x}^* = (\mathbf{x}_1^*, \dots, \mathbf{x}_I^*) \in \mathbb{R}^{I \times L \times N}$ is an equilibrium of the game if and only if $\mathbf{x}_i^* \in BR_i(\mathbf{x}_{-i}^*) \forall i$, that is, if the strategy of each player is a best response to the strategies of other players.

Throughout the paper we will always refer to the description of Fig. 1, however, mutatis mutandis, the results are still valid for scenarios where players are CPs or a mix of CPs and ISPs, as long as the demand is changed accordingly.

B. Transmit Facility Usage Cost

The TF usage cost is different between the NSPs and the IXPs. More specifically the NSP usage cost is linear in the amount of flow that each player sends to it [3]. Therefore we can write:

$$t^l(x_{i,l}, x_l) = a_l x_{i,l} \quad l \leq l_1 \quad (6)$$

where a_l , $l \leq l_1$ represents the transit price of NSP l per unit of flow. We are aware that, due to economies of scale in the traffic delivery, transit costs are subadditive in reality. However, introducing this aspect would overcomplicate the model, hiding the truly interesting differences between transit and peering. Nevertheless, we are able to show that some of our results still hold for more generic transit cost functions (see below, theorem 5).




Figure 2: IXP Port Costs for MIX (log-log scale)

For the IXP usage [1], each player has to pay a share of the total cost of IXP maintenance. This share can be expressed as the ratio between the flow sent by player i on IXP l and the total flow crossing that IXP: $\frac{x_{i,l}}{x_l}$. Assume we can write the cost of maintenance of IXP l as a function h_l of the total flow through the IXP, therefore the usage cost is:

$$t^l(x_{i,l}, x_l) = \frac{x_{i,l}}{x_l} h_l(x_l) \quad l > l_1 \quad (7)$$

The cost of maintaining the equipment of an IXP is, in general, a non-linear function of several parameters. In order to keep the problem manageable, we will approximate this cost with that of a single port which handles x_l , the entire flow over the IXP. The cost of a port is a step-wise increasing function, as shown in Figure 2 for the MIX², an Italian IXP. This type of cost functions can be modeled (see [24]) by using a function like x^α with $\alpha \in [0.4; 0.7]$. For simplicity, we take $\alpha = 0.5$ as this value provides a fairly accurate fit (shown in Fig. 2). Therefore, we express the maintenance cost as:

$$h_l(x_l) = a_l \sqrt{x_l} \quad (8)$$

where $a_l, l > l_1$ is a constant relating the total flow through IXP l with its maintenance cost. By putting together definitions (6), (7) and (8), the cost function (4) can be rewritten as:

$$C^i(\mathbf{x}_i, \mathbf{x}_{-i}) = \sum_{l=1}^{l_1} a_l x_{i,l} + \sum_{l=l_1+1}^L \left(\frac{a_l}{\sqrt{x_l}} x_{i,l} \right) + \sum_{l=1}^L \left(\frac{1}{c_l - x_{i,l}} x_{i,l} \right) + \sum_{l=1}^L \sum_{n=1}^N x_{i,l}^n b_l^n \quad (9)$$

Now, we define these new functions:

$$\begin{cases} f^l \left(\sum_i \sum_n x_{i,l}^n \right) = \begin{cases} a_l & l \leq l_1 \\ \frac{a_l}{\sqrt{\sum_i \sum_n x_{i,l}^n}} & l > l_1 \end{cases} \\ g^l \left(\sum_n x_{i,l}^n \right) = \frac{1}{c_l - \sum_n x_{i,l}^n} \end{cases} \quad (10)$$

By using (10) and recalling that $x_{i,l} = \sum_n x_{i,l}^n$ and $x_l = \sum_i \sum_n x_{i,l}^n$, we can rewrite (9) as:

$$C^i(\mathbf{x}_i, \mathbf{x}_{-i}) = \sum_l \sum_n x_{i,l}^n \left[f^l \left(\sum_i \sum_n x_{i,l}^n \right) + g^l \left(\sum_n x_{i,l}^n \right) + b_l^n \right] \quad (11)$$

²Milan IXP - public peering costs available online: <http://www.mix-it.net>




Figure 3: Minimal Complexity Model

Equation (11) is the most general expression of the cost function for each player. Please note that (11) is in general a non-convex function of $x_{i,l}$, and therefore we cannot directly establish existence of pure equilibria. In particular it does not comply with the general assumptions used for link cost functions in the framework described in [8]. Nevertheless, we cannot avoid dealing with functions of this shape if we want to properly grasp the difference between transit and peering strategies offered, respectively, by NSPs and IXPs.

IV. MINIMAL COMPLEXITY MODEL (MCM)

In order to gain insights on the problem solution, in this section we analyze a simpler model, which we call Minimal Complexity Model (MCM), where $I, L, N = 2$. The MCM has two ISP players, two CPs (therefore each player only has two demands, namely ϕ_i^1 and ϕ_i^2) and two transmit facilities available, either the NSP ($l = 1$) or the IXP ($l = 2$), as depicted in Figure 3. With some algebraic manipulations, explicitly shown in the full version [25] of the paper, we can rewrite cost function (11) for player 1 of the MCM as:

$$\begin{aligned} C^1(x_1, x_2, y_1, y_2) &= \quad (12) \\ &= (\phi_1^1 + \phi_1^2 - x_1 - x_2) \left(a_1 + \frac{1}{c_1 - (\phi_1^1 + \phi_1^2 - x_1 - x_2)} \right) + \\ &+ (x_1 + x_2) \left(\frac{a_2}{\sqrt{x_1 + x_2 + y_1 + y_2}} + \frac{1}{c_2 - (x_1 + x_2)} \right) + \\ &+ (\phi_1^1 - x_1) b_1^1 + (\phi_1^2 - x_2) b_1^2 + x_1 b_2^1 + x_2 b_2^2 \end{aligned}$$

where x_n is the flow sent from player 1 to CP n through the IXP and $\phi_1^n - x_n$ is, by constraint, the flow sent through the NSP. The same applies to y_n for player 2.

Assume now that the topology is fully connected, meaning that $b_l^n = 0 \forall l, n$. In this case, from the player's perspective, the cost does not depend on the facility used for a specific CP, but rather on the total amount of flow going through a specific TF, independently from the destination. Once again, with some algebraic manipulations (shown in [25]) we rewrite cost function (12) for both players as:

$$\begin{cases} C^1(x, y) = (\phi_1 - x) \left(a_1 + \frac{1}{c_1 - (\phi_1 - x)} \right) + \\ \quad + x \left(\frac{a_2}{\sqrt{x+y}} + \frac{1}{c_2 - x} \right) \\ C^2(x, y) = (\phi_2 - y) \left(a_1 + \frac{1}{c_1 - (\phi_2 - y)} \right) + \\ \quad + y \left(\frac{a_2}{\sqrt{x+y}} + \frac{1}{c_2 - y} \right) \end{cases} \quad (13)$$

where x is the cumulative flow of ISP_1 through the IXP and ϕ_1 is its cumulative demand. The same applies to y and ϕ_2 for ISP_2 . The best response of player i is thus obtained by minimizing $C^i(x, y)$ defined in (13).

While simple, the MCM is interesting on its own as it provides a clear way to study the fundamental difference between transit and peering agreements, shedding light on the emerging competition between NSPs and large IXPs (see [9]).

A. Theoretical Results

Definition 1. Supermodular games [26]

Consider a generic game G , where players' payoffs are given by an utility function $u : \mathbb{R}^k \rightarrow \mathbb{R}$. The game is said supermodular if the utility function is supermodular, that is:

$$u(x \vee y) + u(x \wedge y) \geq u(x) + u(y) \quad \forall x, y \in \mathbb{R}^k$$

where $x \vee y$ denotes the componentwise maximum and $x \wedge y$ the componentwise minimum of x and y . If u is twice continuously differentiable, this property is given by the following condition:

$$\frac{\partial^2 u}{\partial z_i \partial z_j} \geq 0 \quad \forall i \neq j$$

In our case we consider costs rather than utilities and minimization instead of maximization, therefore a game like ours is supermodular if:

$$\frac{\partial^2 C(x, y)}{\partial x \partial y} \leq 0 \quad (14)$$

Theorem 1 of [27] proves the existence of equilibria for supermodular games, moreover it provides a way of computing them. The proof is based on showing that **best response sequences are monotone** and therefore converge to a limit which is then shown to be a Nash Equilibrium Point (NEP). The monotonicity is a consequence of the "strategic complementarity" of the players: if one of them chooses a strategy x that decreases its own cost, this decision is beneficial for the cost of the other players too.

Here we relax the results on the existence of equilibria and convergence of best response sequences in supermodular games.

Definition 2. Symmetric supermodularity

We define as symmetric supermodular games, those for which (14) holds for all strategies $x = y$, meaning that the property holds along the symmetric axis.

Definition 3. Symmetric best response sequence

We call symmetric best response sequence a path $(x_0, y_0), (x_1, y_1), \dots$, where $x_0 = y_0$ and $\forall i, (x_i, y_i)$ satisfies $x_i = y_i$.

Theorem 4. *In symmetric supermodular games, pure equilibria exist and are given as the limit of symmetric best response sequences.*

Proof: Consider a sequence of best responses $(x_0, y_0), (x_1, y_1), \dots$. Due to symmetry we can choose this path to be a symmetric best response sequence. From definition 2 and by applying the same reasoning as in the original proof [27], we shall get monotone sequences whose limits are equilibria. \square

Theorem 4 not only proves the existence of equilibria for symmetric supermodular games, but also gives an algorithm for computing them. Please note that, for this theorem to

hold, the game does not need to satisfy (14) for all possible strategies, but just along the symmetric path. This result can be applied to our game thanks to Theorem 5 and Corollary 6.

Theorem 5. *The game defined in (13) is symmetric supermodular.*

Proof: Consider the cost function of our game. We compute the mixed second derivatives:

$$\begin{cases} \frac{\partial^2 C^1(x, y)}{\partial x \partial y} = -\frac{a_2}{2(x+y)^{\frac{3}{2}}} + \frac{3a_2x}{4(x+y)^{\frac{5}{2}}} \\ \frac{\partial^2 C^2(x, y)}{\partial x \partial y} = -\frac{a_2}{2(x+y)^{\frac{3}{2}}} + \frac{3a_2y}{4(x+y)^{\frac{5}{2}}} \end{cases} \quad (15)$$

Since we are interested only in their sign, we can multiply both derivatives in (15) by $(x+y)^{\frac{3}{2}}$, which is always positive as long as the flows are positive. Therefore we have:

$$\begin{cases} \text{sgn} \left(\frac{\partial^2 C^1(x, y)}{\partial x \partial y} \right) = \text{sgn} \left(\frac{a_2}{4} \cdot \frac{x-2y}{x+y} \right) \\ \text{sgn} \left(\frac{\partial^2 C^2(x, y)}{\partial x \partial y} \right) = \text{sgn} \left(\frac{a_2}{4} \cdot \frac{y-2x}{x+y} \right) \end{cases} \quad (16)$$

Consider now the symmetric axis, where $x = y$. With this equality both mixed second derivatives in (16) become negative, therefore due to condition (14) the game is symmetric supermodular. \square

It is interesting to note that, as long as the transit cost function t^l of one ISP does not depend on the other ISP, the mixed second derivative (15) does not change. Therefore, symmetric supermodularity can be applied to game (13) even for more general transit cost functions (as outlined in section III-B). Please note that without the symmetric assumption, the game is neither supermodular, nor submodular, because we cannot say anything about the sign of the mixed derivatives.

Corollary 6. *The game defined in (13) has at least one pure equilibrium for symmetric demands, given as the limit of a symmetric best response sequence.*

Proof: By hypothesis the demands satisfy $\phi_1 = \phi_2$. If we put this condition in system (13), the two players become symmetric, therefore we get this result combining Theorems 4 and 5. \square

B. Cost Function Analysis

In order to gain insights on the outcome of the best response sequence algorithm, here we analyze the cost function. Consider the cost function of player 1 and suppose that the strategy y of player 2 is fixed, so:

$$\begin{aligned} C^1(x) &= (\phi - x) \left(a_1 + \frac{1}{c_1 - (\phi - x)} \right) + \\ &+ x \left(\frac{a_2}{\sqrt{x+y}} + \frac{1}{c_2 - x} \right) \end{aligned} \quad (17)$$

Lemma 7 and Theorem 8, whose proofs can be found in the full version [25], tell us the shape of the cost function.

Lemma 7. *The second derivative of the cost function (17) is a monotonically increasing function.*

Theorem 8. *The cost function (17) can be either: always concave, always convex, or first concave and then convex.*




Figure 4: Cost Function

Figure 4 shows the possible cases. Please note that the shape depends both on the parameters and the strategy y of the other player: while for specific values of y the function might be convex as in Fig. 4b, it can also be concave (Fig. 4a) and in general is neither convex nor concave, as shown in Fig. 4c.

C. Simulation Results

The analysis of the cost function, performed in section IV-B, suggests that in our game (13), even if the best response procedure converges to a NEP, there might be multiple equilibria, because of the presence of multiple local minima. This assumption can be verified via simulation, where we show a specific case in which the NEP reached can change, depending on the starting point of the algorithm.

We implemented in MATLAB [28] the general model (11) described in section III. Iteratively, each player performs its best response to the set of other players' strategies as shown in algorithm 1. If the simulation converges (this has not been proven for the general game (11)), the output `newx` is the NEP for the given input parameters, which are:

- The number of ISPs, TFs and CPs, respectively I, L, N .
- The cost function parameters a_l, c_l, b_l^n and demands ϕ_l^n .
- The *tolerance* and the *startingpoint*.

We can use the implemented algorithm on the MCM, which has fully connected topology and symmetric demands, by putting $I = 2, L = 2, N = 2, b_l^n = 0 \forall l, n$ and $\phi_i^n = \phi^n$. Given the selected scenario and the best response sequence algorithm, Theorem 6 ensures the convergence of the simulation for whatever cost function coefficients. Moreover, thanks to the symmetric property, we can just investigate the strategy of player 1 (x), because player 2 will show exactly the same behavior. We simulate the following parameters: $a_1 = 4, a_2 = 5, \phi^1 + \phi^2 = 4, c_1 = 10, c_2 = 5$. In this case, as previously shown in Figure 4c, the cost function could present multiple local minima, depending on the players' strategies. The simulation has multiple outcomes: if we start from the mean point ($x = 2, \phi - x = 2$) we end up in an equilibrium where traffic is split between the IXP and the NSP: $x^* = 3.64; \phi - x^* = 0.36$. The IXP is preferred because the usage cost is shared among the two players, however it is not

used exclusively due to its smaller capacity not being able to serve all the traffic. With a bigger capacity, all the traffic would have been routed through the IXP. Otherwise, if we start from a strategy where the majority of traffic is routed through the NSP ($x = 0.4, \phi - x = 3.6$), we end up in an equilibrium where all the traffic flows through the NSP: $x^* = 0; \phi - x^* = 4$. This happens because when the IXP is routing a small amount of traffic, the flat port cost is too high to justify its use, therefore the players prefer the NSP. Once the NSP is serving all the traffic no player has an incentive to deviate, because he would pay the whole IXP cost by himself.

As we see, the outcome of the game is highly dependent on the starting point: the IXP is preferred only if it already has, at the beginning, a good amount of flow passing through him, otherwise all players will stick to the NSP. This result is consistent with reality, in fact, the necessary condition for an IXP to emerge is that it has a critical mass (represented by a fraction of the traffic/users in the Internet) which makes the value perceived by a potential participant greater than the cost he would incur in by joining the facility [7].

D. The Best Response Behavior

With the purpose of understanding the number and position of NEPs, we draw the Best Response Intersection (BRI) picture. In this graph, shown in Figure 5, the line with tick marks represents the best response x of player 1 as a function of player 2's strategy y , while the line with cross marks does the exact opposite. The intersection points on the graph mean that both players are playing their best responses, therefore they are Nash Equilibrium Points. As we can see, there are three NEPs and, as expected due to the symmetric property, they are all on the symmetric axis [29]:

Left Equilibrium: is in $x = x_L^* = 0$ and corresponds to the scenario where all the traffic is routed through the NSP.

Right Equilibrium: is for $x = x_M^* = 1.43$ and is the one where traffic is split between the IXP and the NSP.

Middle Equilibrium: happens for $x = x_M^* = 0.31$. This is however a repulsive equilibrium, in fact, as soon as one of the two players deviate, they will never return to this point and reach instead one of the two others equilibria.

Algorithm 1 Best Response Sequence

```

1: startingpoint = ...           ▷ Initial strategies
2: tolerance = ...             ▷ NEP stationarity tolerance
3: newx = startingpoint
4: repeat
5:   oldx = newx             ▷ current step game strategies
6:   for  $i = 1; i < I; i++$  do
7:      $\mathbf{x}_{-i} = \mathbf{oldx}_{-i}$    ▷ other players strategy
8:      $\mathbf{x}_i = \arg \min_{\mathbf{x}_i} C^i(\mathbf{x}_i, \mathbf{x}_{-i})$    ▷  $i$  strategy
9:       s.t.  $\sum_l x_{i,l}^n = \phi_i^n \quad \forall n$ 
10:    newx $_i = \mathbf{x}_i$          ▷ next step game strategies
11:  end for
12: until  $\|\mathbf{newx} - \mathbf{oldx}\| < \mathit{tolerance}$ 

```

These three equilibria can be understood by observing Figure 4c: x_L^* and x_R^* are attractive, and correspond to the minima of the cost function, while x_M^* corresponds to the maximum of the cost function, and is thus repulsive. Of course, the last picture corresponds to the cost function for a specific strategy, therefore it cannot assert the position or the existence of equilibria, however it gives an insight on their meaning.

As we change the game parameters we observe that the shape of the best response is always the same, while the position of x_M^* and x_R^* changes. In particular, as shown in Figure 6, if the ratio a_2/a_1 increases (meaning that IXP cost w.r.t. NSP cost increases) then x_M^* gets nearer to x_R^* , making the left equilibrium is easier to reach. Moreover, we observe that if the capacity c_2 of the IXP is large enough, than in the right equilibrium all the traffic will flow through him.

V. PRICE OF ANARCHY, STABILITY AND FAIRNESS

A. Social Optimum

We now compare the performance of the distributed system, where each Service Provider acts on its own, with that of an ideal centralized system where decisions are taken by some external entity. In this case the objective is to minimize the total cost of the two players, given by the summation of the two costs in (13):

$$\begin{aligned}
 C(x, y) &= \sum_i C^i(x, y) = \\
 &= (\phi_1 + \phi_2 - x - y) a_1 + \frac{\phi_1 - x}{c_1 - (\phi_1 - x)} + \\
 &+ \frac{\phi_2 - y}{c_1 - (\phi_2 - y)} + \frac{x + y}{\sqrt{x + y}} a_2 + \frac{x}{c_2 - x} + \frac{y}{c_2 - y}
 \end{aligned} \tag{18}$$

Theorem 9 and Corollary 10, whose proofs can be found in the full version of the paper [25], explain how to optimize this cost function.

Theorem 9. *The cost function (18) has a global minimum point. For symmetric demands this minimum point is attained at symmetric strategies, and it is either the left endpoint of the strategy space or the unique local minimum point of its convex part.*

Corollary 10. *The global minimum point of (18) is, for symmetric demands, either the left endpoint of the strategy space or the output of a standard algorithm for convex function optimization that starts from the right endpoint.*




Figure 5: BRI - Case 1: $a_1 = 2, a_2 = 2, \phi = 2, c_1 = 10, c_2 = 3$

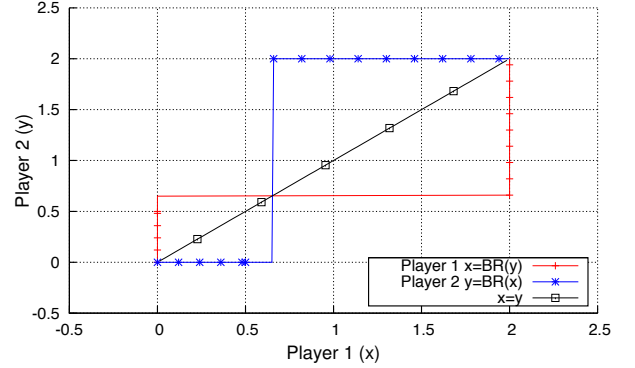


Figure 6: BRI - Case 2: $a_1 = 2, a_2 = 3, \phi = 2, c_1 = 10, c_2 = 4$

The globally optimal solution to problem (18) can thus be computed by comparing the two candidate points highlighted in Corollary 10.

B. Alpha-Fair Solution

Another metric for comparison comes from the theory of fairness. A unifying mathematical formulation, known as α -fairness [30], says that given a set of players and utility functions $U_i(x)$, the α -fair solution to the problem of maximizing their utilities is given by:

$$\max_x \left(\sum_i \frac{U^i(x)^{(1-\alpha)} - 1}{1-\alpha} \right)$$

For $\alpha = 0$, this is the same as maximizing the sum of the utilities, thus it gives the social optimum for the problem. The case $\alpha \rightarrow 1$ yields the proportional fair share assignment, however this solution is not feasible when we have to deal with cost function rather than utilities, and for $\alpha \rightarrow \infty$ it is equivalent to the max-min fairness. For $\alpha = 2$, the formula gives us the ‘‘harmonic mean fair’’ solution, which we investigate numerically in the next subsection.

C. PoA/PoS/PoF Comparison

As is usually done in the literature ([22], [30]), we define the Price of Anarchy (PoA) as the ratio between the worst decentralized solution (equilibrium) and the social optimum. Similarly, the Price of Stability (PoS) is defined as the ratio between the best equilibrium and the social optimum:

$$PoA = \frac{C(x_{worst}^*)}{C(x_{opt})} \geq 1 \quad PoS = \frac{C(x_{best}^*)}{C(x_{opt})} \geq 1$$

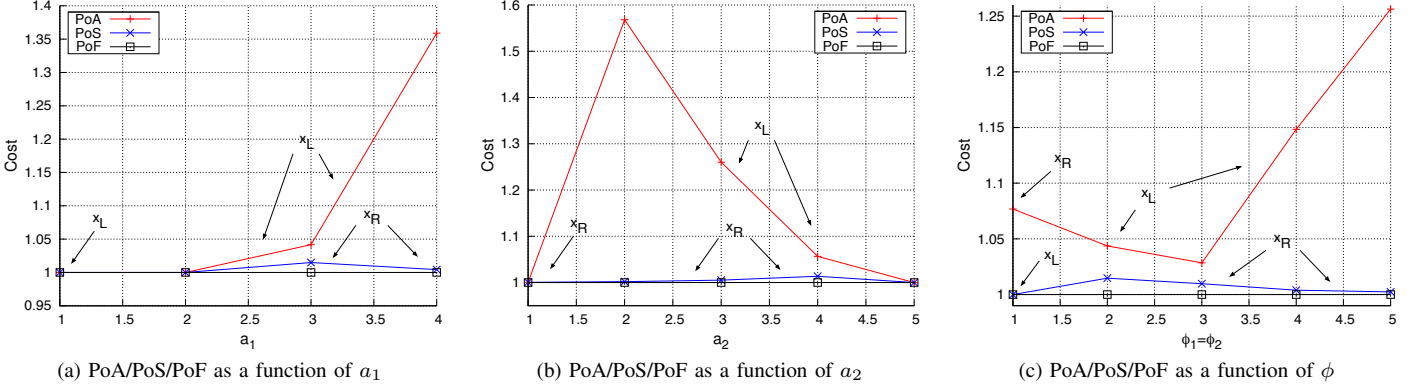


Figure 7: Price of Anarchy, Stability and Fairness

In our case we have just two attractive equilibria, therefore the best and worst equilibria are either x_L^* or x_R^* . Following the same path, we define the price of fairness as the ratio between the fair and the optimal solution:

$$PoF = \frac{C(x_{fair})}{C(x_{opt})} \geq 1$$

Algorithm 1 has been extended to include numerical computation of the above defined Prices of Anarchy, Stability and Fairness. We use as general configuration: $a_1 = 3$, $a_2 = 4$, $\phi = 2$, $c_1 = 10$, $c_2 = 3$ (except for Fig. 7c, where c_2 increases as ϕ increases), and show PoA, PoS and PoF as parameters a_1 , a_2 and ϕ change. Results are reported in Figure 7. As we see, it is always the case that $PoF = 1$, meaning that the harmonic mean fair solution is equal to the social optimum.

The PoA almost always corresponds to the left equilibrium. An exception to this is the case where there is a small amount of total traffic, shown in Fig. 7c for $\phi = 1$: in this case the left equilibrium outperforms the right one, meaning that for small amounts of flow it is not convenient to share costs at the IXP. As ϕ increases, the advantages of sharing become obvious. Figures 7a and 7b show that the PoA increases as a_1 increases and decreases as a_2 increases. An exception to this is the case $a_2 = 1$ of Fig. 7b: with these parameters the cost function resembles that of Fig. 4b, therefore we have only one equilibrium. The PoS is almost always very low, and it is always caused by the fact that the competition between ISPs reduces the amount of traffic through the IXP, thus reducing their opportunities to share costs.

VI. GENERALIZATIONS

While the results obtained by the MCM are interesting on their own, as they shed light on the competition between an Internet Exchange Point and a Network Service Provider, driven by the clear differences between transit and peering, one might argue that this topology is quite small and simplified to represent the Internet. In this section, we use our MATLAB implementation to check what happens in more general topologies. While at this stage we do not have convergence guarantees for these topologies, we know that if the best response sequence algorithm converges, then it converges to an equilibrium [27]. Therefore we can use our simulator both

to study our system in more general cases, and to assess convergence for specific cases.

The base configuration used is $I = 2$, $L = 2$, $N = 2$, $b_l^n = 0 \forall l, n$ and $\phi_i^n = \phi^n = 2$, and all tests have been performed with fully connected topology and symmetric demands. The cost coefficients used are: $a_1 = 1$, $a_2 = 1.5$, $c_1 = 10$, $c_2 = 6$. As long as flows and capacities are properly balanced, the existence of multiple CPs does not seem to affect the results of the simulations, therefore here we check what happens when we have either more ISPs or IXPs.

a) *Generic Number of ISPs*: When the number I of players increase, the benefits of joining an IXP increase as well, due to the fact that costs are shared among multiple participants: in fact, as shown in Figure 8 on the y1 axis, the fraction of traffic flowing through the IXP at the equilibrium increases with I . We recall from section IV-C that IXPs need a critical mass to be used, which in our case corresponds to a fraction of the total traffic in the system. Very interestingly, the y2 axis of Figure 8 shows that this fraction decreases as the number of player grows.

b) *Generic Number of IXPs*: In order to have an interesting case study as L grows, we test a scenario where the global IXP capacity does not change, therefore $c_l = c_2 / (L - 1) \forall l \neq 1$. This means that instead of having one “large” IXP with a high capacity, we have multiple IXPs with less capacity. In order to have meaningful capacities for the small IXPs, we increased global flows and capacities to: $c_1 = 50$, $c_2 = 25$, $\phi_i^n = \phi^n = 10$. As shown in Figure 9 the fact that IXPs only offer small ports is detrimental for the players, and after a certain point they will all stick to the NSP.

While simulation allowed us to only check a limited amount of cases, in all these cases we converged to an equilibrium, therefore we conjecture that the properties shown here for the MCM can be generalized to more complex scenarios.

VII. CONCLUSIONS AND FUTURE WORK

The proposed model gives insight into the economy of different types of Autonomous Systems and the driving forces behind the decisions they make when joining the Internet. The peculiar pricing strategies of players doesn’t allow standard modelization, however, by exploiting peculiar properties of

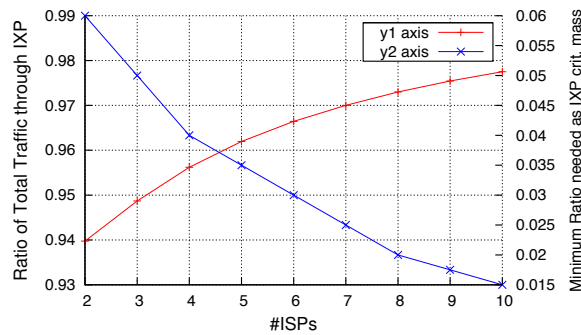


Figure 8: Traffic ratios and equilibrium breakpoint as I grows




Figure 9: Traffic ratio as L grows

the game, we are able to prove analytically the existence of multiple equilibria, and provide an algorithm to compute the stable ones. From a game theoretic perspective, while the theory on supermodularity is well-developed, we relaxed this concept and introduced the new category of Symmetric Supermodular games. Thanks to this we were able to prove existence of equilibria and convergence of best response sequences in our game. This is the first case, to the best of our knowledge, where results on supermodularity are applied even if the property does not hold for the game in general, by showing that it holds along the symmetric path. From an engineering perspective, the outcome of the analysis is highly insightful as it shows different interesting aspects. First of all, we observe the sub-optimality of the decentralized solution, originated by the non-cooperative behavior of the ASes, by showing the existence of a Price of Anarchy and Stability. Even more interestingly, we highlight the growing competition between IXPs, providing customers the ability to lay out peering connections, and NSPs, high level providers selling transit connections.

The work calls for extensions in at least two directions, actively pursued at the time of writing. The first is extending the results of this paper by checking if analytic results on the MCM can be applied to more general scenarios, relaxing assumptions such as the complete connectivity, and improving the simulative part to those cases which weren't covered here. The second follows directly from the dynamic nature of the Internet. In fact, while here we examine a static situation, both transit and peering costs can be renegotiated over time. In order to keep into account this phenomenon, we could add as players the NSPs and IXPs, define their utilities accordingly, and then analyze a dynamic game where their strategies are intertwined with that of ISPs.

REFERENCES

- [1] W. Norton, *The Internet Peering Playbook: Connecting to the Core of the Internet*. 2011.
- [2] "Peeringdb - <http://www.peeringdb.com>."
- [3] W. B. Norton, "Internet transit prices - historical and projected," tech. rep., Dr. Peering White Paper, 2010.
- [4] A. Dhamdhere and C. Dovrolis, "The internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proceedings of the 6th International Conference, Co-NEXT*, ACM, 2010.
- [5] W. B. Norton, "A study of 28 peering policies," tech. rep., Dr. Peering White Paper, 2010.
- [6] E. Gregori, A. Improta, L. Lenzi, and C. Orsini, "The impact of IXPs on the AS-level topology structure of the Internet," *Computer Communications*, vol. 34, pp. 68–82, Jan. 2011.
- [7] W. B. Norton, "The art of peering - the ix playbook," tech. rep., Dr. Peering White Paper, 2010.
- [8] A. Orda, R. Rom, and N. Shimkin, "Competitive routing in multiuser communication networks," *IEEE/ACM Transactions on Networking*, vol. 1, pp. 510–521, Oct. 1993.
- [9] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a large european ixp," in *Proc. of the 2012 ACM SIGCOMM Conference*, pp. 163–174, 2012.
- [10] M. Faloutsos, P. Faloutsos, and C. Faloutsos, "On Power-Law Relationships of the Internet Topology," *ACM SIGCOMM Computer Communication Review*, 1999.
- [11] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks.," *Nature*, vol. 393, pp. 440–2, June 1998.
- [12] E. Gregori, L. Lenzi, and C. Orsini, "k-clique Communities in the Internet AS-level Topology Graph," *31st International Conference on Distributed Computing Systems Workshops*, pp. 134–139, June 2011.
- [13] H. Haddadi, M. Rio, G. Iannaccone, A. Moore, and R. Mortier, "Network Topologies: Inference, Modeling and Generation," *IEEE Communications Surveys*, vol. 10, no. 2, pp. 48–69, 2008.
- [14] S. Yook, H. Jeong, and A. Barabási, "Modeling the Internet's large-scale topology," in *PNAS*, pp. 1–15, 2002.
- [15] X. Wang and D. Loguinov, "Understanding and Modeling the Internet Topology : Economics and Evolution Perspective," *IEEE Transactions on Networking*, vol. 18, no. 1, pp. 257–270, 2010.
- [16] G. Accongiagioco, E. Gregori, and L. Lenzi, "A Structure-Based Topology Generator for the Internet's Core," in *4th Workshop of Italian group on Quantitative Methods in Informatics (InfQ)*, pp. 1–8, 2013.
- [17] A. Fabrikant, E. Koutsoupias, and C. H. Papadimitriou, "Heuristically optimized trade-offs: A new paradigm for power laws in the internet," in *Proceedings of the 29th ICALP '02*, Springer-Verlag, 2002.
- [18] H. Chang, A. Arbor, P. Ave, and F. Park, "To Peer or not to Peer: Modeling the Evolution of the Internet's AS-level Topology," in *Proceedings of 25th IEEE Conference on Comput. Commun., INFOCOM '06*.
- [19] M. Jackson, *Social and Economic Networks*. Princeton Univ.Press, 2010.
- [20] S. Goyal, *Connections: an introduction to the economics of networks*. Princeton University Press, July 2007.
- [21] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "GENESIS: An agent-based model of interdomain network formation, traffic flow and economics," *2012 Proceedings IEEE INFOCOM*, pp. 1197–1205, Mar. 2012.
- [22] T. Jiménez, Y. Hayel, and E. Altman, "Competition in access to content," in *Proceedings of IFIP Networking*, pp. 211–222, 2012.
- [23] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple ISPs," *IEEE/ACM Trans. Netw.*, vol. 14, Dec. 2006.
- [24] M. Motiwala, A. Dhamdhere, N. Feamster, and A. Lakhina, "Towards a cost model for network traffic," *ACM SIGCOMM Comput. Commun. Rev.*, vol. 42, pp. 54–60, Jan. 2012.
- [25] G. Accongiagioco, E. Altman, E. Gregori, and L. Lenzi, "Peering vs transit: a game theoretical model for autonomous systems connectivity," tech. rep., [Online] Available: <http://tinyurl.com/naj7o4t>, 2013.
- [26] D. D. Yao, "S-modular games, with queueing applications," *Queueing Systems*, vol. 21, pp. 449–475, Sept. 1995.
- [27] E. Altman and Z. Altman, "S-modular games and power control in wireless networks," *IEEE Trans. on Automatic Control*, vol. 48, 2003.
- [28] Mathworks, "Matlab - <http://www.mathworks.it/products/matlab/>."
- [29] S. fen Cheng, D. M. Reeves, Y. Vorobeychik, and M. P. Wellman, "Notes on equilibria in symmetric games," in *Proceedings of the 6th International Workshop GTDT*, pp. 71–78, 2004.
- [30] T. Lan, D. Kao, M. Chiang, and A. Sabharwal, "An axiomatic theory of fairness in network resource allocation," *INFOCOM '10*.