Peering Strategy Adoption by Transit Providers in the Internet: A Game Theoretic Approach*

Aemen Lodhi
School of Computer Science
Georgia Institute of Technology
aemen.lodhi@gatech.edu

Amogh Dhamdhere
CAIDA
University of California San Diego
amogh@caida.org

Constantine Dovrolis
School of Computer Science
Georgia Institute of Technology
constantine@gatech.edu

1. INTRODUCTION

The Internet is composed of tens of thousands of Autonomous Systems (ASes) which interconnect with one another through customer-provider (transit) and settlement-free peering links. Their interconnection objectives are a function of their business type e.g., transit providers, content providers, enterprise customers etc. In order to achieve these objectives the ASes adopt a set of criteria which are used to assess potential and existing settlement-free peering relationships. These relationships are bilateral in nature, i.e., for two ASes x and y to establish a peering relationship, they must satisfy each other's peering criteria. The set of criteria used by an AS for assessing potential peering relationships is referred to as its *peering policy*. We commonly observe three peering policies publicized by ASes in PeeringDB [1], an open portal where ASes voluntarily share information about their peering policies, – *Restrictive*, *Selective* and *Open*. These peering policies are explained in section 2.

The peering strategy adoption by ASes of different categories taken from a recent snapshot of PeeringDB showed that *Open* is the dominant peering strategy among all AS categories, with more than 60% of ASes in each category using *Open* peering. The fact that 64% of NSPs (transit providers) use *Open* peering is especially surprising since transit providers prefer other ASes as their customers rather than peers. Why do transit providers tend to peer openly?

In our work [2] we use an agent-based computational model, GENESIS, to study peering strategy adoption by transit providers. Our computational model incorporates most of the real world constraints e.g., geographic co-location, skewed distribution of traffic, economies of scale, multiple transit prices per AS etc. We employed computational modeling as incorporation of all these constraints in an analytical model quickly renders the model intractable. In that work we find that peering decisions are interdependent and myopic decisions and lack of coordination among ASes results in *Open* peering as an attractor among peering strategies for transit providers. Interestingly, we observe that this adoption of *Open* peering results in loss of economic fitness for a majority of transit providers. Further, large scale adoption of $Open$ peering results in stable equilibria.

In this paper we use game theoretic analysis to gain further insight into peering strategy adoption by transit providers in the Internet. We employ a much simplified variant of GENESIS in our current work to keep the analytical approach tractable. Our analytical results corroborate our previous simulation based results. Our results show that when transit providers have complete information about co-located providers they can optimize their economic fitness by adopting *Selective* peering strategy. Further, any uncertainty in the system causes the providers to gravitate towards *Open* peering, a suboptimal equilibrium.

*This research was financially supported by Google and the National Science Foundation (grant CNS-1017139). Its contents are solely the responsibility of the authors and do not necessarily represent the official views of Google or NSF.

2. MODEL DESCRIPTION

In this section, we briefly describe our network formation model. A detailed description of GENESIS can be found in our previous work [3, 4]. The nodes in the network can broadly be divided into two classes – transit providers and stubs. Transit providers have at least one transit customer whereas stubs have none. Stubs have a passive role in the model. They do not change their assigned transit provider and *openly* peer with any node offering to peer with them.Thus, they do not engage in strategic decision making. Transit providers engage in strategic decision making by choosing the peering strategy that maximizes their economic fitness. We next describe each component of the model in brief.

Co-location: Two nodes are "co-located" if they are present in at least one common geographic location. Co-location is necessary to establish any type of link between two nodes. The set G_m denotes the set of ASes co-located with m.

Traffic components: The traffic sent from node m to n is given by V_{mn}. The traffic exchanged between nodes m and n in both directions is given by V'_{mn} Thus,

$$V'_{mn} = V_{mn} + V_{nm}$$

Note that $V'_{mn} = V'_{nm}$. Traffic is routed using the real world *Valley-free customer-prefer-then-peer* routing policy.

Economic attributes: Each transit provider m charges a transit price P_m \$/Mbps from its customers. If node n exchanges traffic V'_{Tn} with the entire network through provider m then n makes a transit payment TC_n to its transit provider m which is given by:

$$TC_n = P_m \times V'_{Tn} \quad (1)$$

The sum of all transit payments from the customers of m constitutes its transit revenue TR_m. Our model captures public peering relationships such as those at Internet Exchange Points. Public peering incurs a cost on its participants as well. If V'_{Pm} is the total peering traffic for node m then its peering cost is given by:

$$PC_m = P_P \times V'_{Pm} \quad (2)$$

where P_P \$/Mpbs is the universal peering cost.

The economic fitness of a node m represents its net profit. In our simplified model it is a function of the peering policy of m and the peering policies of all co-located nodes. We denote the set of possible strategies of node m by S_m, its peering policy with PS_m and its set of peers by PP_m.

$$\pi_m(PS_m, PS_n, ...) = TR_m - TC_m - PC_m \quad \forall \quad n \in G_m \tag{3}$$

The objective of each node is to maximize its fitness by choosing the best peering strategy.

Settlement-free peering: Nodes enter into settlement-free peering relationships with one another based on their peering policies. In this paper we consider the following three peering policies, based on the dominant strategies published and widely discussed at PeeringDB [1] and NANOG [5]:

1. *Restrictive (R)*: A node that uses this policy does not peer with any other node. However, the Tier-1 nodes while using this policy peer among themselves to prevent loss of connectivity in the Internet.

2. *Selective (S)*: A node x that uses this strategy agrees to peer with nodes of similar "size". We use total traffic volume as a measure of similarity of the nodes.

3. *Open (O)*: A node that uses this strategy agrees to peer with any other colocated node except direct customers. All stubs follow this policy only.

3. REFERENCE NETWORK AND PARAMETERS

We analyze a simple network shown in figure 1. The network consists of three transit providers: x, y and T. We assume that T is a tier-1 transit provider i.e., it does not require a transit provider. Nodes x and y are its customers. We also assume that T neither generates nor consumes any traffic. a and b are stubs and customers of x and y respectively. We assume the following for the co-location of nodes in the network: $G_x = \{T, y, a, b\}$; $G_y = \{T, x, a, b\}$; $G_a = \{x, y\}$, $G_b = \{x, y\}$. Thus, there are two players x and y in the model whose peering policies affect the network. In the given network, the traffic exchanged between providers and their direct customers i.e. V'_{xa} and V'_{yb} is not affected by the peering policies of the providers. Therefore, we can ignore these traffic quantities. We also assume that the total traffic volume of x and y is such that a and b do not qualify to become their peers. The total traffic exchanged between x and T i.e., V'_{Tx} and y and T i.e., V'_{Ty} for the given network is given as follows:

$$V'_{Tx} = V'_{xy} + V'_{xb} + V'_{ab} + V'_{ya} \tag{4}$$

$$V'_{Ty} = V'_{xy} + V'_{xb} + V'_{ab} + V'_{ya} \tag{5}$$

In the Internet, transit costs generally exceed peering costs for the same volume of traffic. Thus, we assume that $P_i > P_P \ \forall i \in (x, y)$.

4. RESULTS

4.1 Static Game of Complete Information

In this section we assume that both players, x and y, play simultaneously and both are completely aware of each other's strategies and payoffs. Given that there are 2 players, each with 3 strategies, there are 9 possible outcomes.

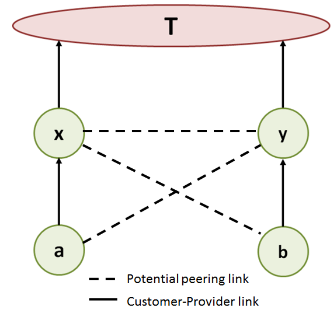

Figure 1: Reference Network

Since the peering policies of stubs a and b are fixed ($Open$) we express the fitness of the two providers x and y as a function of their peering policies only. Let $\pi_x(PS_x, PS_y)$ and $\pi_y(PS_x, PS_y)$ denote the fitness of x and y respectively. Table 1 gives the peers of each provider and Table 2 gives the payoffs for all strategy combinations.

Analysis: An analysis of the payoffs reveals that $\pi_i(R) \leq \pi_i(A) \ \forall i \in (x, y), \forall A \in (S, O, R)$ i.e., $Restrictive$ peering strategy is dominated by both $Selective$ and $Open$ strategies. Thus, there is no incentive for the providers not to engage in settlement free peering with other nodes.

Additionally, the game has two Nash equilibria: (S, S) and (O, O). Of the two equilibria, (S, S) stands out as the compelling solution to the game since $\pi_i(S, S) > \pi_i(O, O) \forall i \in (x, y)$. Thus, (S, S) is the payoff dominant equilibrium while (O, O) is the risk dominant equilibrium. Why is (S, S) payoff dominant? In order to optimize their economic fitness both providers would like to retain their transit revenues while reducing their transit and peering costs. For the strategy pair (S, S) there is no loss in transit revenue of either provider and both providers are able to put all their transit traffic on peering links, thus significantly reducing their costs. For the strategy pair (O, O), both providers are able to put all their transit traffic on peering links similar to the (S, S) configuration. However, under the (O, O) configuration some fraction of customer traffic bypasses the providers and is exchanged directly between their customers and peers. For example, under (O, O) traffic $a \to b$ bypasses x resulting in loss of revenue. Similarly, traffic $b \to a$ bypasses y. While stubs benefit from this configuration, the providers have to suffer a loss in revenue. Why is (O, O) an equilibrium even though an optimal strategy pair exists? Suppose x uses $Open$ strategy while y uses $Selective$. x establishes a peering link with b which causes traffic $a \to b$ and $b \to a$ to flow over the peering link $x - b$. Thus, y's transit traffic goes to zero resulting in a complete loss of revenue. If however, y also adopts $Open$ strategy it establishes a peering link with a. This partially alleviates y's loss, as traffic $a \to b$ now passes through it. Thus, any unilateral deviation from $Open$ strategy under (O, O) will result in complete loss of revenue for the deviating provider. No rational provider will attempt to do so.

The (S, S) strategy pair is payoff dominant and hence the rational choice for both players. However, an underlying assumption for a provider x to choose S over O is the condition that it has sufficient traffic volume to qualify to be y's peer if $PS(y) = S$ (we refer the reader to our previous work [3] for detailed discussion of *Selective* peering strategy). If x does not qualify to be the peer of y under $PS(y) = S$ then x would choose O. If $PS(x) = O$ then the rational choice for y is also to choose O. Hence the Nash equilibrium (O, O).

		Player *x*		
		R	**S**	**O**
Player *y*	**R**	$PP_x = \emptyset$ $PP_y = \emptyset$	$PP_x = \emptyset$ $PP_y = \emptyset$	$PP_x = \{b\}$ $PP_y = \emptyset$
	S	$PP_x = \emptyset$ $PP_y = \emptyset$	$PP_x = \{y\}$ $PP_y = \{x\}$	$PP_x = \{y, b\}$ $PP_y = \{x\}$
	O	$PP_x = \emptyset$ $PP_y = \{a\}$	$PP_x = \{y\}$ $PP_y = \{x, a\}$	$PP_x = \{y, b\}$ $PP_y = \{x, a\}$

Table 1: Peers of providers under Static Game of Complete Information

		Player *x*		
		R	**S**	**O**
Player *y*	**R**	$\pi_x = P_x \times (V'_{ya} + V'_{ab}) - P_T \times V'_{Tx}$ $\pi_y = P_y \times (V'_{xb} + V'_{ab}) - P_T \times V'_{Ty}$	$\pi_x = P_x \times (V'_{ya} + V'_{ab}) - P_T \times V'_{Tx}$ $\pi_y = P_y \times (V'_{xb} + V'_{ab}) - P_T \times V'_{Ty}$	$\pi_x = P_x \times (V'_{ya} + V'_{ab}) - P_T \times (V'_{xy} + V'_{ya}) - P_P \times (V'_{xb} + V'_{ab})$ $\pi_y = -P_T \times (V'_{xy} + V'_{ya})$
	S	$\pi_x = P_x \times (V'_{ya} + V'_{ab}) - P_T \times V'_{Tx}$ $\pi_y = P_y \times (V'_{xb} + V'_{ab}) - P_T \times V'_{Ty}$	$\pi_x = P_x \times (V'_{ya} + V'_{ab}) - P_P \times V'_{Tx}$ $\pi_y = P_y \times (V'_{xb} + V'_{ab}) - P_P \times V'_{Ty}$	$\pi_x = P_x \times (V'_{ya} + V'_{ab}) - P_P \times V'_{Tx}$ $\pi_y = -P_P \times (V'_{xy} + V'_{ya})$
	O	$\pi_x = -P_T \times (V'_{xy} + V'_{xb})$ $\pi_y = P_y \times (V'_{xb} + V'_{ab}) - P_T \times (V'_{xy} + V'_{xb}) - P_P \times (V'_{ya} + V'_{ab})$	$\pi_x = -P_P \times (V'_{xy} + V'_{xb})$ $\pi_y = P_y \times (V'_{xb} + V'_{ab}) - P_P \times V'_{Ty}$	$\pi_x = P_x \times V_{ba} - P_P \times (V'_{yx} + V'_{xb} + V_{ba})$ $\pi_y = P_y \times V_{ab} - P_P \times (V'_{xy} + V'_{ya} + V_{ab})$

Table 2: Payoffs Static Game of Complete Information

4.2 Dynamic Game of Complete Information

Our model of the static game can be extended to a dynamic game as well. Using the same network as shown in figure 1, we analyze a two stage game of complete information. In the sequential game, x plays first followed by y. Both players are rational and have complete information about each other's payoffs for each possible strategy. The extensive form representation of this game is shown in figure 2. The payoffs for each pair of moves are the same as in section 4.1. Since both players are rational and have complete information about each other, x calculates that y will respond by choosing $Selective$ strategy if it chooses $Selective$ strategy. Thus, both providers optimize their fitness by choosing $Selective$ strategy when they play sequentially. Hence, in this model only (S, S) constitutes the Nash equilibrium. (O, O) is not an equilibrium in this model since x (the first player) knows that if it uses O then y will respond with O resulting in a suboptimal situtation of (O, O). Thus, in the sequential game of complete information the players will reach the optimal equilibrium only.

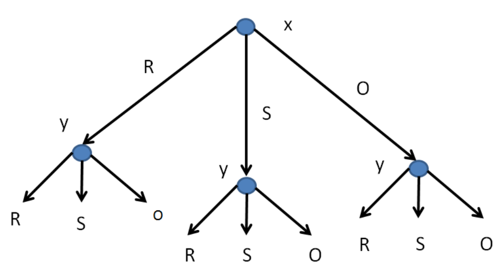

Figure 2: Dynamic Game: Extensive Form Representation

4.3 Static Game of Incomplete Information

In this section we again assume that transit providers play simultaneously. However, the players are uncertain about each other's payoffs. Each player has three possible *types* corresponding to the three peering policies. Thus, the type space of x is: $T_x = \{t_{xR}, t_{xS}, t_{xO}\}$ where t_{xA} denotes that x uses peering policy A.

Analysis of a single stage static game with incomplete information can be done as follows. Using the same set of players and actions we define the following probability distributions on the types of each player:

$P[t_{xR}] = \theta_{xR}$ (probability that x uses strategy R)
$P[t_{xS}] = \theta_{xS}$
$P[t_{xO}] = 1 - \theta_{xR} - \theta_{xS}$
$P[t_{yR}] = \theta_{yR}$ (probability that y uses strategy R)
$P[t_{yS}] = \theta_{yS}$
$P[t_{yO}] = 1 - \theta_{yR} - \theta_{yS}$

These probability distributions reflect the beliefs of players about each other's strategy choices. Given the topology of the network, each provider x and y can predict its payoff for each strategy. For example, if x uses $Selective$ strategy then it can compute its expected payoff given θ_{yR}, θ_{yS} and $1 - \theta_{yR} - \theta_{yS}$. The objective of each provider is to choose the peering strategy that maximizes its expected payoff.

The expected payoff of x if it uses $Restrictive$ strategy is given by:

$$\begin{aligned} E[\pi_x(R)] &= \theta_{yR} \times \pi_x(R, R) + \theta_{yS} \times \pi_x(R, S) \\ &+ (1 - \theta_{yR} - \theta_{yS}) \times \pi_x(R, O) \end{aligned}$$

The objective of provider x is to choose the peering strategy such that:

$$PS(x) = \arg \max_{s \in S(x)} \{E[\pi_x(s)]\}$$

Since $Restrictive$ strategy is always dominated by $Selective$ and $Open$ strategies for both providers we can assume that $\theta_{xR} = \theta_{yR} = 0$ i.e. neither provider will use $Restrictive$ strategy. We focus on the remaining cases. Table 3 gives the expected payoffs for both providers.

Analysis: We carry out the analysis for x. x uses $Open$ strategy if $E[\pi_x(S)] < E[\pi_x(O)]$. In order to check if this condition is true we assume that $E[\pi_x(O)] > E[\pi_x(S)]$. If $E[\pi_x(O)] > E[\pi_x(S)]$ then:

$$P_x \times (V_{ba}) - P_P \times (V'_{xy} + V'_{xb} + V_{ba}) > -P_P \times (V'_{xy} + V'_{xb})$$

$$\implies P_x \times (V_{ba}) - P_P \times (V_{ba}) > 0$$

	Player x	Player y
S	$E[\pi_x(S)] = \theta_{yS} \times \{P_x \times (V'_{ya} + V'_{ab}) - P_P \times (V'_{Tx})\} + (1 - \theta_{yS}) \times \{-P_P \times (V'_{xy} + V'_{xb})\}$	$E[\pi_y(S)] = \theta_{xS} \times \{P_y \times (V'_{yb} + V'_{Ty}) - P_P \times (V'_{Ty})\} + (1 - \theta_{xS}) \times \{P_y \times (V'_{yb}) - P_P \times (V'_{xy} + V'_{ya})\}$
O	$E[\pi_x(O)] = \theta_{yS} \times \{P_x \times (V'_{ya} + V'_{ab}) - P_P \times (V'_{Tx})\} + (1 - \theta_{yS}) \times \{P_x \times (V_{ba}) - P_P \times (V'{}_{yx} + V'_{xb} + V_{ba})\}$	$E[\pi_y(O)] = \theta_{xS} \times \{P_y \times (V'_{yb} + V'_{Ty}) - P_P \times (V'_{Ty})\} + (1 - \theta_{yS}) \times \{P_y \times (V'_{yb} + V_{ab}) - P_P \times (V'_{xy} + V'_{ya} + V_{ab})\}$

Table 3: Expected Payoffs Static Game of Incomplete Information

$$\implies P_x > P_P$$

which is the fundamental condition for settlement-free peering in the first place. Thus, x will adopt $Open$ peering strategy. The same can be said of y leading to the Bayesian equilibrium (O, O). The above evaluation implies that as long as peering cost is less than transit cost and the likelihood of the other player adopting $Open$ peering is non-zero, the result would be gravitation towards $Open$ peering. It is interesting to note x adopted $Open$ peering regardless of the magnitude of θ_{yO}. This happens because even the remote likelihood that the other player will adopt $Open$ peering implies a complete loss in its transit revenue. Given that transit costs (and revenue) are much higher than peering costs, the uncertain player will always choose $Open$ peering so as to avoid a complete loss in revenue. Therefore, the uncertainty about each other's payoffs leads to both providers adopting $Open$ peering strategy.

5. CONCLUSIONS

In this paper, we used game theoretic analysis to explore the apparently counterintuitive adoption of *Open* peering by transit providers. Our analysis showed that peering strategy adoption is not an isolated decision. We showed that in the presence of complete information, transit providers optimize their economic fitness by adopting *Selective* peering strategy. We showed that in the absence of complete information about other providers, transit providers adopt *Open* peering which leads to a suboptimal equilibrium.

6. REFERENCES

[1] "PeeringDB," http://www.peeringdb.com, October 2011.

[2] A. Lodhi, A. Dhamdhere, and C. Dovrolis, "Gravitation of Internet Providers towards Open peering: Dynamics and Economic Consequences," http://www.cc.gatech.edu/~lodhi/peering/conext.pdf.

[3] ——, "GENESIS: An agent-based model of interdomain network formation, traffic flow and economics." in *Proceedings of the IEEE INFOCOM*, 2012.

[4] "GENESIS," http://www.cc.gatech.edu/~dovrolis/genesis.

[5] "North American Network Operators' Group," http://www.nanog.org, 2011.