Route Distribution Incentives

Joud Khoury[1], Chaouki T. Abdallah[1], Kate Krause[2], and Jorge Crichigno[1]

[1] ECE Department, MSC01 1100,
1 University of New Mexico, Albuquerque NM 87131
{jkhoury, chaouki, jcrichigno}@ece.unm.edu
[2] Economics Department, University of New Mexico
1915 Roma NE/Economics Bldg., Albuquerque NM 87131
kkrause@unm.edu

Abstract. We present an incentive model for route distribution in the context of path vector routing protocols and focus on the Border Gateway Protocol (BGP). We model BGP route distribution and computation using a game in which a BGP speaker advertises its prefix to its direct neighbors promising them a reward for further distributing the route deeper into the network. The result of this cascaded route distribution is an advertised prefix and hence reachability of the BGP speaker. We first study the convergence of BGP protocol dynamics to a unique outcome tree in the defined game. We then study the existence of equilibria in the *full information* game considering competition dynamics focusing on the simplest two classes of graphs: 1) the line (and the tree) graphs which involve no competition, and 2) the ring graph which involves competition.

1 Introduction

The Border Gateway Protocol (BGP) [13] is a policy-based path vector protocol and is the de-facto protocol for Internet interdomain routing. BGP is intrinsically about distributing route information about destinations (which are IP prefixes) to establish paths in the network. Path discovery, or simply discovery hereafter, starting with some destination prefix is the outcome of route distribution and route computation. Accounting for and sharing the cost of discovery is an interesting problem and its absence from current path discovery schemes has led to critical economic and scalability concerns. As an example, the BGP control plane functionality is oblivious to cost. A node (BGP speaker) that advertises a provider-independent prefix (identifier) does not pay for the cost of being discoverable. Such a cost, which may be large given that the prefix is maintained at every node in the Default Free Zone (DFZ), is paid by the rest of the network. For example, Herrin [6] has preliminarily analyzed the non-trivial cost of maintaining a BGP route. Such incentive mismatch in the current BGP workings is further exacerbated by provider-independent addressing, multi-homing, and traffic engineering practices [12]. The fact that the number of BGP prefixes in the global routing table (or RIB) is constantly increasing at a rate of roughly $100,000$ entries every 2 years and is expected to reach a total of $388,000$ entries in 2011 [7], has motivated us to devise a model that accounts for distribution incentives in BGP.

A large body of work has focused on choosing the right incentives given that Autonomous Systems (AS) are self-interested, utility-maximizing agents. Most previous work has ignored the control plane incentives [1] (route advertisement) and has instead focused on the forwarding plane incentives (e.g. transit costs). One possible explanation is based on the fact that a node has an incentive to distribute routes to destinations since the node will get paid for transiting traffic to these destinations, and hence route distribution is ignored as it becomes an artifact of the transit process. We argue that this assumption is not economically viable by considering the arrival of a new customer (BGP speaker). While the servicing edge provider makes money from transiting the new customer's traffic to the customer, the middle providers do not necessarily make money while still incurring the cost to maintain and distribute the customer's route information. In this work, we separate the control plane incentives (incentives to distribute route information) from the forwarding plane incentives (incentives to forward packets) and use game theory to model a BGP distribution game. The main problem we are interested in is how to allow BGP prefix information to be distributed globally while aligning the incentives of all the participating agents.

Model and Results We synthesize many of the ideas and results from [3,5,9,10] into a coherent model for studying BGP route distribution incentives. A destination d is willing to invest some initial amount of money r_d to get its route information to be globally distributed. Since d may only advertise its prefix to its direct neighbors, d must incentivize them to further distribute the route. The neighbors then incentivize their neighbors, and so on. While this work takes BGP as the motivating application, we are interested in the general setting of distributing a good to a set of agents. In this paper, we define a *BGP distribution game* by building upon the general model for studying BGP devised by Griffin et. al in [5]. We assume *full information* since our main goal is to study the existence of equilibria rather than how to reach the equilibrium. Studying the equilibria for arbitrary graph structures is difficult given the complexity of the strategic dependencies and the competition dynamics. Since are not aware of general existence results that apply to our game, we initially focus on two simple graphs: 1) the line (and the tree) graphs which involve no competition, and 2) the ring graph which involves competition. Our results are detailed in section 3, and more fully in [8].

Related work The Simple Path Vector Protocol (SPVP) formalism [5] develops sufficient conditions for the outcome of a path vector protocol to be stable. A respective game-theoretic model was developed by Levin [10] to capture these conditions and incentives in a game theoretic setting. Feigenbaum et. al study incentive issues in BGP by considering least cost path (LCP) policies [2] and more general policies [3]. Our model is fundamentally different from [2] (and other works based in mechanism design) in that the prices are strategic, the incentive structure is different, and we do not assume the existence of a central "designer" (or bank) that allocates payments to the players but is rather completely distributed as in real markets. The bank assumption is limiting in a distributed setting, and an important question posed in [3] is whether the bank can be replaced by direct payments by the nodes. Li et. al [11] study an incentive model for query relaying in peer-to-peer (p2p) networks based on rewards, upon which Kleinberg

[1] In this paper, we use the term"control plan" to refer only to route prefix advertisements (not route updates) as we assume that the network structure is static.

et. al [9] build to model a more general class of trees. In [9], Kleinberg and Raghavan allude to a similar version of our distribution game in the context of query incentive networks. They pose the general question of whether an equilibrium exists for general Directed Acyclic Graphs (DAGs) in the query propagation game. Both of these probabilistic models do not account for competition. While we borrow the basic idea, we address the different problem of route distribution rather than information seeking.

2 The General Game

Borrowing notation from [3,10], we consider a graph $G = (V,E)$ where V is a set of n nodes (alternatively termed players, or agents) each identified by a unique index $i = \{1,\ldots,n\}$, and a destination d, and E is the set of edges or links. Without loss of generality (WLOG), we study the BGP discovery/route distribution problem for some fixed destination AS with prefix d (as in [5,3,10]). The model is extendable to all possible destinations (BGP speakers) by noticing that route distribution and computation are performed independently per prefix. The destination d is referred to as the *advertiser* and the set of players in the network are termed *seekers*. Seekers may be distributors who participate in distributing d's route information to other seeker nodes or consumers who simply consume the route. For each seeker node j, Let $P(j)$ be the set of all routes to d that are known to j through advertisements, $P(j) \subseteq \mathscr{P}(j)$, the latter being the set of all simple routes from j. The empty route $\phi \in \mathscr{P}(j)$. Denote by $R_j \in P(j)$ a simple route from j to the destination d with $R_j = \phi$ when no route exists at j, and let $(k, j)R_j$ be the route formed by concatenating link (k, j) with R_j, where $(k, j) \in E$. Denote by $B(i)$ the set of direct neighbors of node i and let $next(R_i)$ be the next hop node on the route R_i from i to d. Finally, define node j to be an *upstream* node relative to node i when $j \in R_i$. The opposite holds for a *downstream* node. The general distribution game is as follows: destination d first exports its prefix (identifier) information to its neighbors promising them a reward $r_d \in \mathbb{Z}^+$ which directly depends on d's utility of being discoverable. A node i, a player, in turn receives offers from its neighbors where each neighbor j's offer takes the form of a reward r_{ji}. We use $r_{next(R_i)}$ to refer to the reward that the upstream parent from i on R_i offers to i.

Strategy Space: Given a set of advertised routes $P(i)$ where each route $R_i \in P(i)$ is associated with a promised reward $r_{next(R_i)} \in \mathbb{Z}^+$, a *pure strategy* $s_i \in S_i$ of an autonomous node i comprises two decisions:

First, after receiving offers from neighboring nodes, pick a single "best" route $R_i \in P(i)$ (where "best" is defined shortly in Theorem 1);

Second, pick a reward vector $r_i = [r_{ij}]_j$ promising a reward r_{ij} to each candidate neighbor $j \in B(i)$ that it has not received a competing offer from (i.e., such that $r_{ji} < r_{ij}$ where $r_{ji} = 0$ means that i did not receive an offer from j). Then export the route and reward to the respective candidate neighbors. The distribution process repeats up to some depth that is directly dependent on the initial investment r_d as well as on the strategies of the players.

Cost: The cost of participation is local to the node and includes for example the cost associated with the effort spent in maintaining the route information. We assume that every player i incurs a cost of participation c_i and for simplicity we take $c_i = c = 1$.

Utility: A strategy profile $\mathbf{s} = (s_1, \ldots, s_n)$ and a reward r_d define an outcome of the game. [2] Every outcome determines a set of paths to destination d given by $O_d = (R_1, \ldots, R_n)$. A utility function $u_i(\mathbf{s})$ for player i associates every outcome with a real value in $\mathbb{R}$. We use the notation s_{-i} to refer to the strategy profile of all players excluding i. A simple class of utility functions we experiment with rewards a node linearly based on the number of sales that the node makes. This model incentivizes distribution and potentially requires a large initial investment from d. More clearly, define $N_i(\mathbf{s}) = \{j \in V \backslash \{i\} | i \in R_j\}$ to be the set of nodes that pick their best route to d going through i (nodes downstream of i) and let $\delta_i(\mathbf{s}) = |N_i(\mathbf{s})|$. Let the utility of a node i from an outcome or strategy profile $\mathbf{s}$ be:

$$u_i(\mathbf{s}) = (r_{next(R_i)} - c_i) + \sum_{\{j|i=next(R_j)\}} (r_{next(R_i)} - r_{ij})(\delta_j(\mathbf{s}) + 1) \tag{1}$$

The first term $(r_{next(R_i)} - c_i)$ of (1) is incurred by every participating node and is the one unit of reward from the upstream parent on the chosen best path minus the local cost. Based on the fixed cost assumption, we often drop this first term when comparing player payoffs from different strategies since the term is always positive when $c = 1$. The second term of (1) (the summation) is incurred only by distributors and is the total profit made by i where $(r_{next(R_i)} - r_{ij})(\delta_j(\mathbf{s}) + 1)$ is i's profit from the sale to neighbor j (which depends on δ_j). A rational selfish node will always try to maximize its utility by picking $s_i = (R_i, [r_{ij}]_j)$. There is an inherent tradeoff between $(r_{next(R_i)} - r_{ij})$ and $(\delta_j(\mathbf{s}))$ s.t. $i = next(R_j)$ when trying to maximize the utility in Equation (1) in the face of competition as shall become clear later. A higher promised reward r_{ij} allows the node to compete (and possibly increase δ_j) but cuts the profit margin. Finally, we implicitly assume that the destination node d gets a constant marginal utility of r_d for each distinct player that maintains a route to d - the marginal utility of being discoverable by any seeker - and declares r_d truthfully to its neighbors i.e., r_d is not strategic.

Assumptions: We take the following simplifying assumptions to keep our model tractable:

1. the advertiser d does not differentiate among the different players (ASes).
2. the advertised rewards are integers and are strictly decreasing with depth i.e. $r_{ij} \in \mathbb{Z}^+$ and $r_{ij} < r_{next(R_i)}, \forall\, i, j$ and let 1 unit be the cost of distribution.
3. finally, our choice of the utility function isolates a class of policies which we refer to as the Highest Reward Path (HRP). We assume for the scope of this work that transit costs are extraneous to the model.

Convergence under HRP. Before proceeding with the game model, we first prove the following theorem which results in the Highest Reward Path (HRP) policy. All proofs may be found in the full version of this paper [8].

Theorem 1 *In order to maximize its utility, node i must always pick the route R_i with the highest promised reward i.e. such that $r_{next(R_i)} \geq r_{next(R_l)}, \forall\, R_l \in P(i)$.*

---

[2] We abuse notation hereafter and we refer to the outcome with simply the strategy profile $\mathbf{s}$ where it should be clear from context that an outcome is defined by the tuple $< \mathbf{s}, r_d >$.

Theorem (1) implies that a player could perform her two actions sequentially, by first choosing the highest reward route R_i, then deciding on the reward vector r_{ij} to export to its neighbors. Thus, we shall represent player i's strategy hereafter simply with the rewards vector $[r_{ij}]$ and it should be clear that player i will always pick the "best" route to be the route with the highest promised reward. When the rewards are equal however, we assume that a node breaks ties consistently. Given the asynchronous nature of BGP, we ask the question of whether the BGP protocol dynamics converge to a unique outcome tree T_d under some strategy profile $\mathbf{s}$ [5]. From Theorem (1), it may be shown that the BGP outcome converges under any strategy profile $\mathbf{s}$, including the equilibrium (see [8] for proof). This result allows us to focus on the existence of equilibria.

2.1 The Static Multi-Stage Game with fixed schedule

We restrict the analysis of equilibria to the simple line and ring graphs. In order to apply the correct solution concept, we fix the *schedule* of play (i.e. who plays when?) based on the inherent order of play in the model. We resort to the *multi-stage game with observed actions* [4] where stages in our game have no temporal semantics. Rather, stages identify the network positions which have strategic significance due to the strictly decreasing rewards assumption. Formally, and using notation from [4], each player i plays only once at stage $k > 0$ where k is the distance from i to d in number of hops. At every other stage, the player plays the "do nothing" action. The game starts at stage 1 after d declares r_d. Players at the same stage play simultaneously, and we denote by $a^k = (a_1^k, , \dots, a_n^k)$ the set of player actions at stage k, the stage-k action profile. Further, denote by $h^{k+1} = (r_d, a^1, \dots, a^k)$, the *history* at the end of stage k which is simply the initial reward r_d concatenated with the sequence of actions at all previous stages. We let $h^1 = (r_d)$. Finally, $h^{k+1} \subset H^{k+1}$ the latter being the set of all possible stage-k histories. When the game has a finite number of stages, say $K+1$, then a terminal history h^{K+1} is equivalent to an outcome of the game (which is a tree T_d) and the set of all outcomes is H^{K+1}.The pure-strategy of player i who plays at stage $k > 0$ is a function of the history and is given by $s_i : H^k \to \mathbb{R}^{m_i}$ where m_i is the number of direct neighbors of player i that are at stage $k+1$ (implicitly, a player at stage k observes the full history h^k before playing). We resort to the multi-stage model (the fixed schedule) on our simple graphs to eliminate the synchronization problems inherent in the BGP protocol and to focus instead on the existence of equilibria. By restricting the analysis to the fixed schedule, we do not miss any equilibria (see [8]). The key concept here is that it is the *information sets* [4] that matter rather than the time of play i.e. since all the nodes at distance 1 from d observe r_d before playing, all these nodes belong to the same information set whether they play at the same time or at different time instants.

Starting with r_d (which is h^1), it is clear how the game produces actions at every later stage based on the player strategies resulting in a terminal action profile or outcome. Hence, given r_d, an outcome in H^{K+1} may be associated with every strategy profile $\mathbf{s}$ and so the definition of Nash equilibrium remains unchanged (see [4] for definitions of *Nash equilibrium*, *proper subgame*, and *subgame perfection*). In our game, each stage begins a new subgame which restricts the full game to a particular history. For example, a history h^k begins a subgame $G(h^k)$ such that the histories in the subgame are restricted to $h^{k+1} = (h^k, a^k)$, $h^{k+2} = (h^k, a^k, a^{k+1})$, and so on. Hereafter, the general notion of equilibrium we use is the Nash equilibrium and we shall make it clear when we generalize to subgame perfect equilibria. We are only interested in pure-strategy equilibria [4] and in studying the existence question as the incentive r_d varies.

3 Equilibria on the Line Graph, the Tree, and the Ring Graph

In the general game model defined thus far, the tie-breaking preferences of the players is a defining property of the game, and every outcome (including the equilibrium) depends on the initial reward/utility r_d of the advertiser. In the same spirit as [9] we inductively

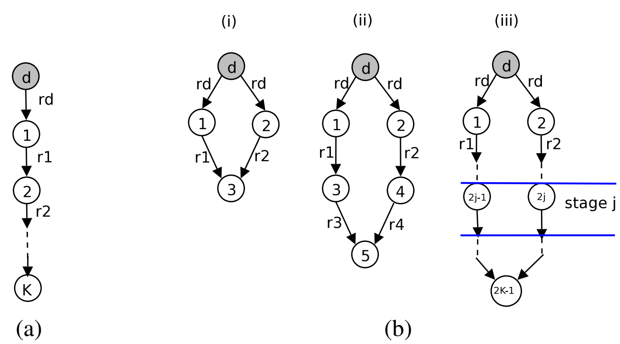

Fig. 1. (a) Line graph: a player's index is the stage at which the player plays; d advertises at stage 0; $K = n$; (b) Ring graph with even number of players: (i) 2-stage game, (ii) 3-stage game, and general (iii) K-stage game.

construct the equilibrium for the line graph of Figure 1(a) given the utility function of Equation (1). We present the result for the line which may be directly extended to trees. Before proceeding with the construction, notice that for the line, $m_i = 1$ for all players except the leaf player since each of those players has a single downstream neighbor. In addition, $\delta_i(\mathbf{s}) = \delta_j(\mathbf{s}) + 1, \forall i, j$ where j is i's child ($\delta_i = 0$ when i is a leaf). We shall refer to both the player and the stage using the same index since our intention should be clear from the context. For example, the child of player i is $i+1$ and its parent is $i-1$ where player i is the player at stage i. Additionally, we simply represent the history $h^{k+1} = (r_k)$ for $k > 0$ where r_k is the reward promised by player k (player k's action). The strategy of player k is therefore $s_k(h^k) = s_k(r_{k-1})$ which is a singleton (instead of a vector) since $m_i = 1$ (for completeness, let $r_0 = r_d$). This is a *perfect information* game [4] since a single player moves at each stage and has complete information about the actions of all players at previous stages. Backward induction may be used to construct the subgame-perfect equilibrium. We construct the equilibrium strategy s^* inductively as follows: first, for all players i, let $s_i^*(x) = 0$ when $x \leq c$ (where c is assumed to be 1). Then assume that $s_i^*(x)$ is defined for all $x < r$ and for all i. Obviously, with this information, every player i may compute $\delta_i(x, s_{-i}^*)$ for all $x < r$. This is simply due to the fact that δ_i depends on the downstream players from i who must play an action or reward strictly less than r. Finally, for all players i we let $s_i^*(r) = \arg\max_x (r-x)\delta_i(x, s_{-i}^*)$ where $x < r$.

Theorem 2 *The strategy profile s^* is a subgame-perfect equilibrium.*

The proof may be directly extended to the tree since each player in the tree has a single upstream parent as well and backward induction follows in the same way. On the tree, the strategies of the players that play simultaneously at each stage are also independent.

Competition: the ring. We present next a negative result for the ring graph. In a ring, each player has a degree $= 2$ and $m_i = 1$ for all players except the leaf player. We consider rings with an even number of nodes due to the direct competition dynamics. Figure 1(b) shows the 2-, the 3-, and general K-stage versions of the game. In the multi-stage game, after observing r_d, players 1 and 2 play simultaneously at stage 1 promising rewards r_1 and r_2 respectively to their downstream children, and so on. We refer to the players at stage j using ids $2j-1$ and $2j$ where the stage of a player i, denoted as $l(i)$, may be computed from the id as $l(i) = \lceil \frac{i}{2} \rceil$. For the rest of the discussion, we assume WLOG that the player at stage K (with id $2K-1$) breaks ties by picking the route through the left parent $2K-3$. For the 2-stage game in Figure 1(b)(i), it is easy to show that an equilibrium always exists in which $s_1^*(r_d) = s_2^*(r_d) = (r_d - 1)$ when $r_d > 1$ and 0 otherwise. This means that player 3 enjoys the benefits of *perfect competition* due to the Bertrand-style competition [4] between players 1 and 2. The equilibrium in this game is independent of player 3's preference for breaking ties. We now present the following negative result,

Claim 1 *The 3-stage game induced on the ring (of Figure 1(b)(ii)) does not have a subgame-perfect equilibrium. Particularly, there exists a class of subgames for $h^1 = rd > 5$ for which there is no Nash equilibrium.*

The value $r_d > 5$ signifies the breaking point of equilibrium or the reward at which player 2, when maximizing her utility $(r_d - r_2)\delta_2$, will always oscillate between competing for 5 (by playing large r_2) or not (by playing small r_2). This negative result for the game induced on the 3-stage ring may be directly extended to the general game for the K-stage ring by observing that a class of subgames $G(h^{K-2})$ of the general K-stage game are identical to the 3-stage game. While the full game does not always have an equilibrium when $K > 2$ stages, we shall show next that there always exists an equilibrium for a special subgame.

Growth of Incentives, and a Special Subgame. We next answer the following question: Find the minimum incentive r_d^*, as a function of the depth of the network K (equivalently the number of stages in the multi-stage game), such that there exists an equilibrium outcome for the subgame $G(r_d^*)$ that is a spanning tree. We seek to compute the function f such that $r_d^* = f(K)$. First, we present a result for the line, before extending it to the ring. On the line, K is simply the number of players i.e. $K = n$, and $f(K)$ grows exponentially with the depth K as follows:

Lemma 1. *On the line graph, we have $f(0) = 0$, $f(1) = 1$, $f(2) = 2$, and $\forall\ k > 2$, $f(k) = (k-1)f(k-1) - (k-2)f(k-2)$*

We now revisit the the K-stage game of Figure 1(b)(iii) on the ring and we focus on a specific subgame which is the restriction of the full game to $h_1 = r_d^* = f(K)$, and we denote this subgame by $G(r_d^*)$. Consider the following strategy profile $\mathbf{s}^*$ for the subgame: players at stage j play $s_{2j-1}^*(h^j) = f(K-j)$, and $s_{2j}^*(h^j) = f(K-j-1), \forall$ $1 \le j \le K-1$, and let $s_{2K-1}^*(h^K) = 0$.

Theorem 3 *The profile* $\mathbf{s}^*$ *is a Nash equilibrium for the subgame* $G(r_d^*)$ *on the K-stage ring,* $\forall\, K > 2$.

This result may be interpreted as follows: if the advertiser were to play strategically assuming she has a marginal utility of at least r_d^* and is aiming for a spanning tree (global discoverability), then $r_d^* = f(K)$ will be her Nash strategy in the game induced on the K-stage ring, $\forall\, K > 2$ (given $\mathbf{s}^*$). We can now extend the growth result of Lemma (1) to the ring denoting by $f_r(K)$ the growth function for the ring.

Corollary 1. *On the ring graph, we have* $f_r(k) = f(k)$ *as given by Lemma (1).*

In this paper, we have studied the equilibria existence question for a simple class of graphs. Many questions remain to be answered including extending the results to general network structures (and to the Internet *small-world* connectivity graph), relaxing the fixed cost assumption, quantifying how hard is it to find the equilibria, and devising mechanisms to get to them. All these questions are part of our ongoing work [8].

References

1. Blume, L., Easley, D., Kleinberg, J., Tardos, E.: Trading networks with price-setting agents. In: EC ’07, pp. 143–151. ACM, New York, NY, USA (2007)
2. Feigenbaum, J., Papadimitriou, C., Sami, R., Shenker, S.: A bgp-based mechanism for lowest-cost routing. Distrib. Comput. **18**(1), 61–72 (2005)
3. Feigenbaum, J., Ramachandran, V., Schapira, M.: Incentive-compatible interdomain routing. In: EC ’06, pp. 130–139. ACM, New York, NY, USA (2006)
4. Fudenberg, D., Tirole, J.: Game Theory. MIT Press (1991)
5. Griffin, T.G., Shepherd, F.B., Wilfong, G.: Policy disputes in path-vector protocols. In: ICNP ’99, p. 21. IEEE Computer Society, Washington, DC, USA (1999)
6. Herrin, W.: What does a bgp route cost? http://bill.herrin.us/network/bgpcost.html (2008)
7. Huston, G.: Bgp in 2008. http://www.potaroo.net/ispcol/2009-03/bgp2008.html (2008)
8. Khoury, J., Abdallah, C.T., Krause, K., Crichigno, J.: Route distribution incentives. arXiv:0909.3558v1 [cs.GT]
9. Kleinberg, J., Raghavan, P.: Query incentive networks. In: FOCS ’05, pp. 132–141. IEEE Computer Society, Washington, DC, USA (2005)
10. Levin, H., Schapira, M., Zohar, A.: Interdomain routing and games. In: STOC ’08, pp. 57–66. ACM, New York, NY, USA (2008)
11. Li, C., Yu, B., Sycara, K.: An incentive mechanism for message relaying in unstructured peer-to-peer systems. In: AAMAS ’07, pp. 1–8. ACM (2007)
12. Meyer, D., Zhang, L., Fall, K.: Report from the iab workshop on routing and addressing. Internet RFC 4984, Sep 2007
13. Rekhter, Y., Li, T., Hares, S.: RFC 4271: A border gateway protocol 4 (bgp-4) (2006)