

Notes on the Topological Consequences of BGP Policy Routing on the Internet AS Topology

Dávid Szabó, András Gulyás

► **To cite this version:**

Dávid Szabó, András Gulyás. Notes on the Topological Consequences of BGP Policy Routing on the Internet AS Topology. 19th Open European Summer School (EUNICE), Aug 2013, Chemnitz, Germany. pp.274-281, 10.1007/978-3-642-40552-5_25 . hal-01497022

HAL Id: hal-01497022

<https://hal.inria.fr/hal-01497022>

Submitted on 28 Mar 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Notes on the Topological Consequences of BGP Policy Routing on the Internet AS Topology

Dávid Szabó, András Gulyás[†]

Dept. of Telecommunications and Media Informatics,

[†] MTA (Hungarian Academy of Science) - BME Information Systems Research Group

Budapest University of Technology and Economics,

Magyar Tudósok krt. 2, 1117 Budapest, Hungary

{david.szabo,gulyas}@tmit.bme.hu

Abstract. On the Internet AS level topology, BGP policy routing is in charge of dictating the characteristics of routes which can be used for packet transmission. Furthermore the peculiarities of the BGP policies clearly affect the peering strategies of the ISP-s, hereby influencing the emerging topology. This paper takes a first step towards identifying the topological footprint issued by these policies. For this purpose we modify the framework of network formation games to support the most fundamental BGP policy called valley-free routing. We show the topological properties of the equilibrium topologies of this game in an analytical manner.

Keywords: Network formation games, BGP policy routing, AS level topology

1 Introduction

Resembling many of real world networks (e.g. social networks, biological networks etc.), the Internet AS level topology is a complex network exhibiting similar macroscopic features such as power-law degree distribution, high clustering coefficient and logarithmic diameter [1]. Exploiting this amazing structural resemblance, general complex network models are widely used for generating Internet-like topologies as an input for evaluating Internet related architectures, business models, protocols etc. For obtaining more realistic AS topology models one should incorporate also some microscopic features, which cannot be captured by high level statistics but contribute fundamentally to the final shape of the network.

In this paper we try to grab some of the microscopic features of the AS level topology. First we notice that there is a mutual correspondence between the network topology and the used routing policy. Topology trivially affects routing since one cannot choose routes having edges not contained in the topology. On the other side of the coin, the routing policy can fundamentally influence the topology of the network. For the illustration of this idea consider the simple example on Fig. 1.

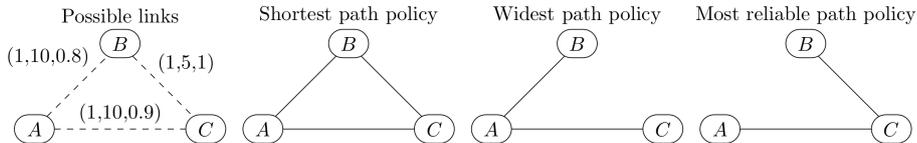


Fig. 1. Policy drives topology: We have three nodes (A, B, C) among which we decide to build a network. The characteristics of the possible edges (A, B), (A, C), (B, C) are defined by a triplet: latency, bandwidth and reliability (l, b, r) (rightmost). Now we have to decide what topology we should build. Clearly if the routing policy is the shortest path policy, then we have to use all the links. If the policy is the "widest path" meaning that we are shooting for the highest bandwidth between the nodes, then the third topology from the left is the appropriate choice. Finally if the policy is the "most reliable path" policy then the corresponding network will be the leftmost topology.

According to our best knowledge, the topological footprint¹ of the Internet's routing policy, called the BGP [5] policy routing, are not fully understood and hence not incorporated in the current structural models of the AS level topology². Our goal in this paper is therefore to take a first step towards characterizing the topological artifacts imposed by the BGP policies. We show the footprint of the most fundamental BGP policy called *valley-free routing* on the AS topology in an analytical manner. Since the AS level the topology of the Internet is formed by the business interactions between the ISP-s, such an analysis should be highly incentive-oriented. For this purpose we generalize the concept of Network Formation Games (NFG) and define a network game accommodating the valley-free routing policy. We study the Nash equilibrium topology of this game and characterize its features.

The rest of the paper is organized as follows. In Section 2 we give a brief overview on network formation games and the valley-free routing policy. In Section 3 after modifying the NFG framework to incorporate the valley-free policy, we examine the topological consequences of valley-free routing in an analytical manner. Section 4 discusses our findings and describes future work.

2 Background

Network Formation Games – Network Formation Games constitute a game theoretic framework which provides considerable insight into the mechanisms that forms the topology of networks [13]. In the NFG the players are identified

¹ Several other aspects (e.g. stability [2], memory requirement [3], path diversity [4] etc.) of BGP policy routing have been analyzed so far.

² There are "computational" models [6, 7] which can generate very realistic AS level topologies based on fully emulating the business decisions of the AS-s fed by complex datasets. Although these models generate considerably realistic networks they can scale only for around a few hundred nodes.

as the nodes of the network, while the strategies of the players is to create some set of links towards an arbitrary subset of the other players. The goal of the players is to minimize their cost function which consists of two parts: link costs and communication costs. The cost function for some player u was originally defined by Fabrikant [13] as:

$$c_u = \underbrace{\sum_{\forall u \neq v} d_{G(s)}^{\text{sh}}(u, v)}_{\text{communication costs}} + \underbrace{\alpha |s_u|}_{\text{link costs}}, \quad u, v \in \mathcal{P}, \quad (1)$$

where $d_{G(s)}^{\text{sh}}(u, v)$ denotes the length of the shortest path between players u and v on the graph $G(s)$ characterized by the union the strategies of the players, s_u stands for the set of nodes towards which u creates links, while α represents the cost of building one edge. Thus the NFG effectively analyzes the trade-off between link costs and communication costs for rational (selfish) players, as the key incentives of building specific topologies. As defined in (1), distance is usually measured as the average length of shortest path from the given node to all other nodes. An important state of the game is called as *Nash equilibrium* in which no player (node) has anything to gain by unilaterally changing it's strategy.

The valley-free routing policy – BGP is the core inter-domain routing protocol of the Internet. It is responsible for distributing available forwarding paths between Autonomous Systems (ASes) and lets each one to choose the path it prefers according to its own special interests. The business relationships of different ASes can be diverse, based on exclusive contracts, service-level agreements and other policy issues. Still we can classify most AS-AS links into basically just two major groups [14]: in a *customer-provider* relationship one AS pays another for forwarding its traffic, while in a *peering* relationship neighboring ASes voluntarily exchange traffic with each other in a settlement-free manner³. In this paper we focus on the topological consequences of the most fundamental policy called valley-free routing [15] [16]. This policy encompasses the basic economic fact that the flow of traffic must obey the flow of cash. The policy dictates that AS A can use a link to a neighboring AS B to route traffic if and only if either the incoming traffic is from a customer or B is a customer of A . All paths that induce valley-free routing have the same preference; the only concern this policy addresses is not violating elemental business rules.

3 Topological consequences of valley-free routing

For the analysis of the topological footprint generated by valley-free routing we modify the framework of network formation games as follows:

Players and valley-free routing – Let \mathcal{P} be the a set of players (identified as network nodes) with cardinality N . Recall that the rules of the valley-free policy dictate that a player u can forward traffic originated from player w to

³ For the sake of simplicity, we omit sibling and backup relationships in this paper.

a neighboring player v if and only if: (i) the incoming traffic of u is from a customer (in this case the relationship between u and v is indifferent and can be both provider-customer or peering), or (ii) v is customer of u (in this case the type of relationship between w and u becomes indifferent). In other words: *after a provider-customer edge or a peer edge, the path can not traverse customer-provider edge or another peer edge.*[17]

Strategies and topology – The strategy space of player $u \in \mathcal{P}$ is to create some set of undirected edges to other players in the network. The created edges can be of types customer-to-provider (p) and peer (r) edges in accordance with the relationships of the valley-free routing. The r edges are paid at both sides, however p edges are paid by the customer. Note that in some of the following arguments, according to the direction of traversing customer-to-provider links, we may write provider-to-customer link but these terms are referring to the same edge. Thus the complete strategy space of player u is $S_u = 3^{\mathcal{P} \setminus \{u\}}$, where the number 3 covers the third choice of node u : creating no edge. Let s be a strategy vector containing the strategies of all players hereby representing the current state of the game: $s = (s_0, s_1, \dots, s_{N-1}) \in (S_0, S_1, \dots, S_{N-1})$. Then the graph $G(s) = \cup_{i=0}^{N-1} (i \times s_i)$ represents the topology between the players.

Payoff – The goal of the players is to minimize their costs. The cost of player u is defined as:

$$C_u = \sum_{\forall v \neq u} d_{G(s)}(u, v) + \varphi_p u_p + \varphi_r u_r, \quad u, v \in \mathcal{P} \quad (2)$$

where φ_x is the cost an edge of type x , u_x is the number edges of type x and $d_{G(s)}(u, v)$ is the communication cost between u and v over $G(s)$ given by:

$$d_{G(s)}(u, v) = \begin{cases} 0 & \text{if valley free path exists between } u \text{ and } v \\ \infty & \text{otherwise.} \end{cases} \quad (3)$$

In what follows we identify the Nash equilibrium of the game in different settings of the parameters. The Nash equilibrium is one of the most fundamental states of a game, because it shows a stable state in which none of the nodes would alter its strategy, though they are rational (selfish) players being interested in their own utility.

Definition 1 (Valley-free footprint (VFF)) *A graph is a valley-free footprint if consists of (i) a clique K_r comprising peer (r) links only, and (ii) trees rooted at some subset of $V(K_r)$ having customer-to-provider links (p) only, such that the provider in the relationship is always closer to the root than the customer (see Fig. 2).*

Theorem 1 *Every Nash equilibrium of the game is a VFF having $|V(K_r)| = \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor + 1$.*

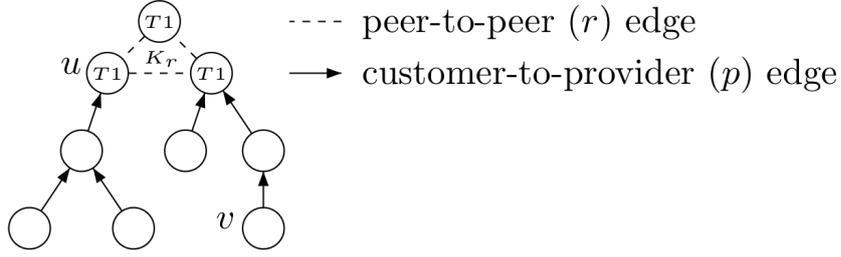


Fig. 2. Example for a VFF topology. In such a topology there could be two type of nodes, T1 and none T1. T1s are connected by r edges, which are counted on both sides in the cost function, however p edges are paid only by the customer, who requested it. The flow of cash is visualized by arcs. Conversely there are two possible cost functions: (i) $C_u = \varphi_r u_r = \varphi_r(|V(K_r)| - 1)$ and (ii) $C_u = \varphi_p$.

Proof. The proof consists of three parts: first we show that a VFF with $|V(K_r)| = \lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1$ is a Nash equilibrium, secondly that a VFF with $|V(K_r)| \neq \lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1$ is not and finally that a non-VFF topology is also not a Nash equilibrium.

1. The VFF topology with $|V(K_r)| = \lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1$ is a Nash equilibrium: We show that no player has anything to gain by changing his own strategy if others don't change theirs. Let C_u represent the cost of player u residing in the VFF. It is clear that in a VFF every player can reach others through valley-free paths, hence in the cost functions the communication cost is always zero ($\sum_{v \neq u} d_{G(s)}(u, v) = 0$). Moreover it is clear that any meaningful state of the game permits valley-free paths between arbitrary pairs of players (otherwise the cost of some player would be infinite), so we will omit writing this tag hereafter.

(a) if $u \in K_r$:

$$C_u = \varphi_r u_r = \varphi_r(|V(K_r)| - 1) = \varphi_r \lfloor \frac{\varphi_p}{\varphi_r} \rfloor \quad (4)$$

If u wants to deviate then it will use some other strategy. In this case the corresponding cost is given by $C_u' = \varphi_r u_r' + \varphi_p u_p'$. If $u_p' \geq 1$ then $C_u \leq C_u'$ since:

$$C_u = \varphi_r \lfloor \frac{\varphi_p}{\varphi_r} \rfloor \leq \varphi_r \frac{\varphi_p}{\varphi_r} = \varphi_p \leq C_u'. \quad (5)$$

If $u_p' = 0$ then again $C_u \leq C_u'$ since $u_r \leq u_r'$ must hold to ensure valley-free connectivity to all other nodes. This is because if $u_r > u_r'$ then there is at least one node $v \in K_r$ to whom u is not connected thus cannot be reached from u via valley-free path.

(b) if $u \notin K_r$:

$$C_u = \varphi_p \quad (6)$$

$u_p' \geq 1$ is trivially not an option for u since that case immediately implies $C_u \leq C_u'$. What remains is the case $u_p = 0$. In this case to ensure valley-free connectivity u has to create r edges to all other nodes residing in K_r . This would mean $C_u' = \varphi_r(\lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1)$, but even then $C_u < C_u'$:

$$C_u = \varphi_p = \varphi_r \frac{\varphi_p}{\varphi_r} \leq \varphi_r \left\lceil \frac{\varphi_p}{\varphi_r} \right\rceil \leq \varphi_r (\lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1) = C_u' \quad (7)$$

2. The VFF topology with $|V(K_r)| \neq \lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1$ is NOT a Nash equilibrium:

Easily if $|V(K_r)| < \lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1$, then any leaf node u from a rooted tree can lower its cost by joining K_r as:

$$C_u' = \varphi_r(|V(K_r)| - 1) < \varphi_r \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor \leq \varphi_r \frac{\varphi_p}{\varphi_r} = \varphi_p = C_u \quad (8)$$

Similarly if $|V(K_r)| > \lfloor \frac{\varphi_p}{\varphi_r} \rfloor + 1$ then some node $u \in K_r$ can lower its cost by leaving K_r and connect to some other $w \in K_r, w \neq u$ with a p edge.

$$C_u' = \varphi_p = \varphi_r \frac{\varphi_p}{\varphi_r} \leq \varphi_r \left\lfloor \frac{\varphi_p}{\varphi_r} \right\rfloor < \varphi_r(|V(K_r)| - 1) = C_u \quad (9)$$

3. Finally we show that graphs differing from VFF can't constitute Nash equilibria. Let G be an arbitrary graph on which the valley-free connectivity is satisfied, i.e., $\forall u \in G : C_u \neq \infty$. It's obvious that without this property G can't constitute a Nash equilibrium. In such graph let u be a node whose strategy differs from the nodes in a VFF. The possible cases are:

- (a) u has p edge: The strategy of the nodes in the VFF having p edges is to pay for only a single p edge and nothing more (They can have other p edges attached to them but these are paid by their customers, see Fig. 2). Since u differs from this, the corresponding cost function is characterized as:

$$C_u = \varphi_p u_p + \varphi_r u_r, \quad u_p \geq 1, u_r \geq 0 : (u_p, u_r) \neq (1, 0). \quad (10)$$

Now let the edge (u, w) be a p edge of u . Since w can reach every other player via valley-free path, u also has valley-free paths to all others through w . This means that u can delete its edges except the (u, w) edge, giving $C_u > C_u' = \varphi_p$.

- (b) u doesn't have a p edge: The cost function of u is:

$$C_u = \varphi_r u_r. \quad (11)$$

The valley-free connectivity implies that every path from u to others starts with one r edge which can be followed only by *provider-to-customer* edges. Let t be the number of nodes having no p edges. Such nodes must be the neighbors of u otherwise u cannot reach them. This imposes $u_r \geq t$. Furthermore for differing from VFF $u_r \neq t$ must also hold. In summary we get $u_r > t$. In this case u has edges (u, w) where

w has a provider. Now these edges can be deleted since u can reach w through its provider too. This gives:

$$C_u = \varphi_r u_r > \varphi_r t = C_u' \quad (12)$$

According to Theorem 1 aside from two trivial cases – (i) $\varphi_p \geq (n-1)\varphi_r$, when the Nash equilibria is a tree, and (ii) $\varphi_r > \varphi_p$, when the Nash equilibria is a complete graph – the Nash equilibria of our very simple game exhibits some level of structural resemblance with the Internet AS level topology. It is well known, that on the AS level topology there are only a few nodes in the whole network which can reach every other node without purchasing service (i.e. without having a customer-to-provider edge), we call such node a tier-1 (T1). T1s are in a clique having peering agreements with each other, this is the top of the hierarchy. The rest of the nodes are customers of T1s either in a direct or in an indirect way. These topological features are clearly reflected by our results and now we can understand them as a clear consequence of the valley-free policy.

4 Discussion and future work

In this paper we have taken a first step towards identifying which topological features of the AS level Internet follow from BGP routing policy. We modified the framework of network formation games for support the valley-free routing, which is the most fundamental BGP policy, and defined a new game. We've given a mathematical analysis for the Nash equilibria. For positioning our work we recall that our main motivation is to gain deeper insight into the driving forces forming the topology of the Internet. Since BGP is the core inter-domain routing protocol of the Internet our long-term goal is to identify what topological features this protocol implies. As the first step we have analyzed the topological footprint of its core building block, the valley-free policy, and identified the VFF as its consequence. Meanwhile the topologies produced by our simple model seems promising they are far from being realistic. In the future we plan to broaden the spectrum of BGP policies and investigate the topological consequences of the *local preference rule* and the effect of the *AS pathlength*. The in depth analysis of these policies may eventually lead us to more realistic AS level topology generators which may worth the comparison with the existing models as well.

Acknowledgement This work was partially supported by the European Union and the European Social Fund through project FuturICT.hu (grant no.: TAMOP-4.2.2.C-11/1/KONV-2012-0013).

References

1. Xiao Fan Wang and Guanrong Chen. Complex networks: small-world, scale-free and beyond. *Circuits and Systems Magazine, IEEE*, 3(1):6–20, 2003.

2. Lixin Gao, Timothy G Griffin, and Jennifer Rexford. Inherently safe backup routing with bgp. In *INFOCOM 2001. Twentieth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, volume 1, pages 547–556. IEEE, 2001.
3. Gábor Rétvári, András Gulyás, Zalán Heszberger, Márton Csernai, and József J Bíró. Compact policy routing. In *Proceedings of the 30th annual ACM SIGACT-SIGOPS symposium on Principles of distributed computing*, pages 149–158. ACM, 2011.
4. P Arjona-Villicana, C Constantinou, and Alexander Stepanenko. The internet’s unexploited path diversity. *Communications Letters, IEEE*, 14(5):474–476, 2010.
5. Y. Rekhter and T. Li. A border gateway protocol 4 (bgp-4). RFC 1771 (Draft Standard), mar 1995. Obsoleted by RFC 4271.
6. A. Lodhi, A. Dhamdhere, and C. Dovrolis. Genesis: An agent-based model of interdomain network formation, traffic flow and economics. In *INFOCOM, 2012 Proceedings IEEE*, pages 1197–1205, 2012.
7. A. Dhamdhere and C. Dovrolis. An agent-based model for the evolution of the internet ecosystem. In *Communication Systems and Networks and Workshops, 2009. COMSNETS 2009. First International*, pages 1–10, 2009.
8. F. Papadopoulos, D. Krioukov, M. Bogua, and A. Vahdat. Greedy forwarding in dynamic scale-free networks embedded in hyperbolic metric spaces. In *INFOCOM, 2010 Proceedings IEEE*, pages 1–9, 2010.
9. Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguñá. Hyperbolic geometry of complex networks. *Phys. Rev. E*, 82:036106, Sep 2010.
10. D. Krioukov, F. Papadopoulos, M. Kitsak, M. Serrano, and M. Boguna. Popularity versus similarity in growing networks. *Bulletin of the American Physical Society*, 2012.
11. Marián Boguná, Fragkiskos Papadopoulos, and Dmitri Krioukov. Sustaining the internet with hyperbolic mapping. *Nature Communications*, 1:62, 2010.
12. CAIDA AS Ranking webpage (<http://as-rank.caida.org/>).
13. Elitza Maneva Alex Fabrikant, Ankur Luthra. On a network creation game. *PODC*, pages 347–351, 2003.
14. G. Huston. Interconnection, peering and settlements: Part I. *Internet Protocol Journal*, 2(1), JUN 1999.
15. L. Gao and J. Rexford. Stable internet routing without global coordination. *ACM/IEEE Transactions on Networking*, 9(6):681–692, 2001.
16. Xiaowei Yang, David Clark, and Arthur Berger. NIRA: A new inter-domain routing architecture. *IEEE/ACM Transactions on Networking*, 15(4):775–788, 2007.
17. Lixin Gao. On inferring autonomous system relationships in the internet. *IEEE/ACM Trans. Netw.*, 9(6):733–745, dec 2001.