

Quantifying the BGP routes diversity inside a tier-1 network

Steve Uhlig*, Sébastien Tandel

Department of Computing Science and Engineering
Université catholique de Louvain, Louvain-la-neuve, B-1348, Belgium
E-mail: {suh,sta}@info.ucl.ac.be

Abstract. Many large ISP networks today rely on route-reflection [1] to allow their iBGP to scale. Route-reflection was officially introduced to limit the number of iBGP sessions, compared to the $\frac{n \times (n-1)}{2}$ sessions required by an iBGP full-mesh. Besides its impact on the number of iBGP sessions, route-reflection has consequences on the diversity of the routes known to the routers inside an AS. In this paper, we quantify the diversity of the BGP routes inside a tier-1 network. Our analysis shows that the use of route-reflection leads to a very poor route diversity compared to an iBGP full-mesh. Most routers inside a tier-1 network know only a single external route in eBGP origin. We identify two causes for this lack of diversity. First, some routes are never selected as best by any router inside the network, but are known only to some border routers. Second, among the routes that are selected as best by at least one other router, a few are selected as best by a majority of the routers, preventing the propagation of many routes inside the AS. We show that the main reason for this diversity loss is how BGP chooses the best routes among those available inside the AS.

Keywords: BGP, iBGP, route-reflection, route diversity.

1 Introduction

The Internet consists of a collection of more than 21,000 domains called Autonomous Systems (ASs). Each AS is composed of multiple networks operated under the same authority. Inside a single domain, an independent Interior Gateway Protocol (IGP) [2] such as IS-IS or OSPF is used to propagate routing information. Between ASs, an Exterior Gateway Protocol (EGP) is used to exchange reachability information. Today, BGP [2] is the de facto standard interdomain routing protocol used in the Internet. BGP routers exchange routing information over BGP sessions. External BGP (eBGP) sessions are established over inter-domain links, i.e., links between two different ASes (BGP peers), while internal BGP (iBGP) sessions are established between the routers within an AS.

Route-reflection [1] was initially introduced as an alternative to the iBGP full-mesh that requires $\frac{n \times (n-1)}{2}$ iBGP sessions to be established inside an AS. This number of sessions required for propagating the routes learned from the neighbors of the AS to all routers inside the AS does not scale for large networks containing hundreds or thousands of BGP routers. Route-reflection [1] was thus

* Steve Uhlig is Postdoctoral fellow of the Belgian National Fund for Scientific Research (F.N.R.S).

introduced to limit the number of iBGP sessions for large sized networks. An advantage of an iBGP full-mesh is that all routers know about all the best routes of the other routers inside the network. This means that when some route is withdrawn, routers can typically switch to another route immediately, without waiting for BGP to converge. Without the use of an iBGP full-mesh on the other hand, routers might know only a single route to reach an external destination. When this route is withdrawn, then the concerned prefix will not be reachable until BGP re-converges and advertises an alternative route. BGP is known to suffer from slow convergence [3]. BGP routes diversity is thus important to understand if high availability of the reachability service is to be provided, as is typically the case in tier-1 providers.

Route-reflection inside an AS defines two types of relationships among BGP routers: client and non-client. These relationships among BGP peers define a loose hierarchy among routers, going from the bottom level routers that have no clients up to the largest route-reflectors that are not client of any other router. Note that this implicit hierarchy is not practically enforced, as iBGP sessions can be established between any two routers inside the AS, even under route-reflection.

The redistribution of the routes in BGP works according to well-defined rules. First recall that a route is never re-advertized to the peer that announced it. Consider a given prefix p for which a router inside the iBGP receives several routes from its peers (iBGP or eBGP). The router chooses among the possible ones towards p its best route using the BGP decision process [4].

How the best route is propagated to the neighbors of a router depend on whether the router acts as a route-reflector. If a router does not act as a route-reflector, i.e. it has no "client" peer, then the router advertises this route to all its iBGP peers if it is learned from an eBGP session, or to none of them if the route was learned from an iBGP session. If a router acts as a route-reflector [1] on the other hand:

- If the route was learned from a client peer (or eBGP peer), the route-reflector redistributes the route to all its clients and non-client peers (except the one from which the route was received).
- If the route was learned from a non-client peer, the route-reflector redistributes the route to its client peers only.

These rules driving the redistribution of the routes inside the iBGP imply an implicit filtering of the routes over the internal BGP signaling graph. Besides the rules defined in [1] when connecting route-reflectors to ensure a proper working of the iBGP propagation, there is no clear design rules known today as to how to design a proper iBGP graph. Guidelines for checking that a correct iBGP configuration have been discussed in [5]. [6] provides a tool to detect potential problems due to the iBGP configuration based on static analysis.

Route-reflection was initially proposed as an alternative to the full-mesh, but in practice it caused many problems and it is unclear what it actually performs on which routes are propagated compared to a full-mesh. In this paper, we thus aim at quantifying the diversity inside a tier-1 network that relies on route-reflection. We see our work as a first step towards a better understanding of the impact of route-reflection on route diversity.

2 Methodology

Unless one has complete data concerning the full topology, the configuration of the routers, and the eBGP routes learned by an AS, it is not possible to correctly

reproduce its routing state [7]. This is the main reason why typically, simulations have to be used to reproduce the routing of a large AS. The aim of this section is to sketch our methodology to reproduce the routing of the studied network.

We relied on CBGP [7] to model our tier-1 network. For this, we used the physical topology of the network (links and IGP weights), as well as the configuration of the BGP routers. We obtained the Adj-RIB-In's from the main route-reflectors of the studied network. Because the BGP routes present in the Adj-RIB-In's of internal routers do not always contain the information about which eBGP peer actually originated a route, some reverse-engineering of the route origin was necessary. We could of course keep the routes learned directly from eBGP sessions, as large route-reflectors also have a significant number of eBGP peerings (see Section 5). Two cases are possible when trying to find the entry point of a route:

- The BGP next-hop of the route is the IP address of an external peer. In this case we must pay attention to advertise this route from the external peer found to the internal router with which the external peer has established the eBGP session.
- The BGP next-hop of the route is the IP address of an internal router because this router has been configured with *next-hop-self*. We have to find the originating external peer that advertised the route to the internal router. To find it, we rely on the AS path information. We search for eBGP peers belonging to the leftmost AS on the AS path that have an eBGP peering with the internal router.

To ensure that our model was correct we validated the conversion by injecting the routes in the model and then checked the routes computed by the model against the original best routes seen in the route-reflectors. Due to space limitations, we do not provide these results here.

As even in our simulation model, it is not always possible to identify the eBGP peer from which a route has been advertised by looking at the route, all external routes in the C-BGP simulation had to be tagged with a special community value identifying the external router from which the route was learned. This made the analysis of the results easier. Once all external routes were identified, C-BGP [8] performed the propagation of the routes according to the internal iBGP structure of the network, and we retrieved the content of the Adj-RIB-In's of all routers inside the C-BGP simulation. Our analysis is based on the outcome of this simulation.

Among all prefixes of our input data, we selected a subset of them (940). Those 940 prefixes were learned from several locations in the network. In the analysis of this paper, only those 940 multiply-advertized prefixes are considered as measuring diversity for singly-advertized prefixes is meaningless. We selected the largest of them in terms of the amount of traffic sent towards them. These prefixes captured 80% of the total traffic according to the Netflow [9] statistics. 80,000 destination prefixes were present in the Netflow statistics, most of them representing an insignificant fraction of the total traffic.

The iBGP structure of the studied network consists of 3 levels of route-reflection according to which router is a client of which other router. This graph contains 105 nodes (routers) partitioned into 36 geographically distinct POPs and 169 undirected edges. This iBGP "hierarchy" is a static one, by design of the iBGP graph. To find out the hierarchy inside the route-reflection graph, we rely on a topological sort of a directed acyclic graph (DAG) [10]. The reason why we have to rely on this concept of a DAG is that the route-reflection graph is not a strict

hierarchy (a forrest). Contrary to a general misbelief, route-reflection does not require a strict hierarchy to work. A strict hierarchy is even not desirable for route diversity. The vertices of the route-reflection graph (rr_graph) are all routers inside the iBGP graph. An arc (i,j) of the route-reflection graph rr_graph connects a reflector (i) to a client router (j). The *level* in the route-reflection hierarchy is computed by finding out which reflectors are not clients of any other router in the reflection graph. These are given a level of 0 in the hierarchy, they are the top-level route-reflectors of the graph (16 routers). Route-reflectors which are clients of the top-level (0) reflectors have a route-reflection level of 1 (57 routers). Finally, clients of reflectors at level 1 are given a level of 2 (32 routers).

3 Example of route diversity loss

When relying on an iBGP full-mesh, all the external routes selected as best by the border routers are known to all other routers inside the AS. An iBGP full-mesh is thus "ideal" in terms of the diversity of the routes known to all routers inside an AS, at the cost of a large number of iBGP sessions. Even this "ideal" situation might hide some eBGP routes when a border router has multiple eBGP sessions or when it does not choose as its best route one among its eBGP-learned ones. This would happen if one of its non eBGP-learned routes has a higher local-pref or smaller AS path length than its eBGP-learned routes. A loss in diversity will thus occur only because of this order of the rules of the BGP decision process.




Fig. 1. Example of route loss inside iBGP.

For instance, Figure 1 illustrates the two main causes for loss of diversity on an example. Prefix p is advertised to AS4 by 3 neighboring ASes (AS1, AS2 and AS3), two of them at border router $BR2$ and another at border router $BR1$. eBGP sessions are indicated by solid lines, while iBGP sessions by dashed lines. Arrows indicate the propagate of a route from one router to another. Only the best route chosen by $BR2$, let us call it p_{BR2}^{best} , will be propagated inside AS4.

The best route propagated by $BR1$, assuming it is the external one (p_{BR1}^{best}), will also be propagated within AS4. Route reflector RR is on the iBGP propagation path of both routes p_{BR1}^{best} and p_{BR2}^{best} , hence it will choose at most one of these as best route, which we call p_{RR}^{best} . As we have one route reflector in AS4, all other routers are clients, hence because of the iBGP propagation rules RR will redistribute its best route to all its clients except the one from which it learned the route.

To prevent this loss of diversity, several solutions can be envisioned. First, one can change the location of the eBGP peerings so as to minimize the loss of the routes at the border routers. Changing the location of eBGP peerings is typically not practically feasible because it depends on the slots available on the routers and the geographical constraints about where peers can connect to the routers of the AS. Another solution is to reconfigure the iBGP graph by adding and removing iBGP peerings between routers, but this operation is tricky as it is difficult to predict its impact on the BGP propagation [5, 6]. Finally, redistributing more than a single route [11] could be seen as a solution. This would however require changes to the protocol at the risk of creating divergence. A proper understanding of route diversity is thus necessary before thinking about changed in how routes are propagated inside an AS.




Fig. 2. External routes lost at edge routers.

To show to what extent external routes can be lost only due to multiple eBGP peerings at the same border router, Figure 2 compares for each prefix the total number of known external routes with the number of routes that will never be selected as best due to multiple peerings at the edge routers, in the studied tier-1 network. Each border router may receive several external routes from its eBGP peers for a given prefix. The points labeled "lost" sums for each prefix (over all border routers) the number of external routes that cannot be chosen as best because several are received by a border router. 534 over 5018 routes are lost because of multiple external routes received by border routers. Hence more than 10% of the external routes cannot be considered just because of the location of the eBGP peerings inside the network. These lost routes concern 365 over the 940 prefixes, 40% of the considered prefixes for which several external routes are known.

4 iBGP structure of the studied network

In this section, we want to highlight two points. First, we want to make clear that the hierarchy induced by route-reflection and the propagation of the routes inside the AS are two very different things. Second, we want to discuss how much the location of a router inside the iBGP propagation graph varies across prefixes.

The propagation inside iBGP depends on from which border routers the routes were learned. Each prefix can be learned from a different set of border routers, even though most of the prefixes are typically learned from a small subset of all possible border routers. In the studied network, eBGP peerings can be attached to any router, from level-0 reflectors to routers at the edge of the network (level-2). Centrality in the reflection hierarchy hence does not match the centrality of a router inside the iBGP propagation graph.

Directly comparing the *level* of a router with its location in the signaling graph is problematic for two reasons. First, the *level* of a router is a very discrete variable taking only 3 different values. Second, the variation of the location of a router from the eBGP peering wherefrom the route has actually been learned by the AS varies a lot. We define the depth $depth(r, p)$ of a router r in the iBGP signaling graph for a given prefix p as the number of iBGP hops it took for the best route chosen by r towards p from the eBGP peer who advertised this route. The $depth(r, p)$ varies between 1 and 6 in our studied network. Still, the typical values of the depth lies around 2 and 3 for most routers. The routers having many eBGP peerings or that are central (level-0 reflectors) inside the iBGP graph tend to have a smaller depth than less central routers (level-2 reflectors).

5 Best route choice and route origin

An important factor to understand the propagation of the routes inside the iBGP is from what kind of peering the best routes of a router were learned by any router. Figure 3 provides the breakdown of the best routes chosen by each router according to what type of BGP peer advertised the route. A route can be learned either from an eBGP peer, a client peer (for route reflectors) and a non-client peer (both for reflectors and other routers). Routers on the x-axis of Figure 3 are ordered by their increasing level inside the route-reflection hierarchy, so the first 16 routers are level-0 reflectors, the next 57 level-1, and the last 32 are level-2 reflectors. This ordering of the x-axis was chosen because one might expect that more central routers like level-0 reflectors would have a larger fraction of their best routes learned from the iBGP. The y-axis of Figure 3 gives, for each router, the percentage of best routes of each type. For each router, we computed among the best routes it selected, the fraction of them that have been learned from eBGP sessions, client and non-client sessions. On Figure 3 we plot the fraction of client-learned routes, then the sum of client-learned routes and eBGP-learned ones. Non-client-learned routes are not shown on Figure 3 but make the rest of the 100% of the best routes.

It is easy to see that excepted for level-0 reflectors (the first 16 routers), most routes are non-client routes, i.e. routes learned from either a reflector from which the local router is a client or a regular iBGP peer. Only large reflectors (mainly level-0) select routes learned by client peers, as these routers also have the largest number of client peers. Note that routers for which it might seem on Figure 3 to




Fig. 3. Breakdown of best route choice by origin.

have only selected as best non-client-learned ones actually have typically a few eBGP-learned or client-learned routes as best. This is not apparent from the use of the percentage over all considered prefixes.

Figure 3 told that most best routes are learned from iBGP peers. However, this choice of the best routes might be biased by a lack of eBGP peerings at some routers. This is however not the case in the studied network, as non-client peerings represent 53% of the total peerings, client peerings about 23%, and eBGP peerings about 26%. More than one fourth of all BGP sessions are thus eBGP sessions, hence a lack of eBGP peerings is not the reason why routers do not select their best route from a eBGP peer-learned one.

95% of the best routes are learned from non-client peers, about 2% from client peers, and 3% from eBGP peers. Most routes chosen as best by the routers come either from a regular iBGP peer or a route-reflector of which the considered router is a client. The choice of the best route of a router thus depends a lot on the choice performed by the route-reflectors higher in the hierarchy. This phenomenon is caused by the relatively small number of locations from which a prefix is learned by the AS, hence the iBGP propagation graph is very important to understand which route will be propagated inside iBGP.

6 Measuring iBGP route diversity

To measure the diversity of the routes, we define two metrics: the *real diversity* and the *RIB diversity*. The choice of these metrics mainly reflects our own interest of understanding what fraction of the external routes is actually known to the routers inside the iBGP compared to those know to the whole AS. Let us insist on the fact that as the route-reflection graph is not a forest, a given eBGP-learned route can be propagated through different iBGP propagation paths. It thus makes sense to measure the difference between the number of actually distinct routes a router learns from its neighbors and how this number relates to from how many distinct eBGP peers those routes come.

The *real diversity* measures the proportion of the external routes known by the AS any router has learned. The *real diversity* $div_{real}(r, p)$ counts for each router r and prefix p the number of unique external routes (learned from distinct eBGP

peers) r has in its Adj-RIB-In's divided by the total number of eBGP routes that have been learned by routers of the AS:

$$div_{real}(r, p) = \frac{routes_{unique}(r, p)}{routes(p)}. \quad (1)$$

$routes(p)$ denotes the total number of distinct eBGP routes (learned from different eBGP peers) known by all routers of the AS and $routes_{unique}(r, p)$ the number of distinct eBGP routes r has in its Adj-RIB-In's for prefix p . Even in an iBGP full-mesh, some routers will not forcibly have a div_{real} of 1 when they learn multiple eBGP routes since they can propagate only a single route inside the iBGP.

The *RIB diversity* $div_{RIB}(r, p)$ on the other hand counts for each router r and prefix p the number of unique external routes (learned from distinct eBGP peers) r has in its Adj-RIB-In's divided by the total number of entries in its Adj-RIB-In's:

$$div_{rib}(r, p) = \frac{routes_{unique}(r, p)}{rib(r, p)}. \quad (2)$$

$rib(r, p)$ denotes the number of Adj-RIB-In entries router r has for prefix p . $div_{rib}(r, p)$ takes values in the $]0, 1]$ range. If r has no route towards p then its *RIB diversity* will be undefined. The closer to 1 the value of div_{rib} , the less redundancy there is among the routes r knows towards p .

In practice, one would like as high a value of both metrics. If many external routes are known inside the AS, then the value of div_{real} will be low so that a low value of div_{real} is not an indication of a "bad" diversity. A value of div_{rib} smaller than 1 indicates that among the several routes a router learns, some of them are duplicates and will thus be withdrawn if the corresponding external route is withdrawn. Such redundant iBGP routes protect a router from the failure of one of the routers that advertise this route.

7 Real and RIB diversity of the studied network

On Figure 4, we show for each considered prefix p the average over all routers of the network of $div_{real}(\cdot, p)$, $div_{rib}(\cdot, p)$, and $\frac{1}{routes(p)}$. Prefixes on the x-axis of Figure 4 are ordered by increasing value of $routes(p)$. Recall that $routes(p)$ denotes the number of different eBGP peers from which a route towards p is learned. The reason for plotting $\frac{1}{routes(p)}$ on Figure 4 is that it provides a lower bound on $div_{real}(\cdot, p)$, i.e. it is the value of $div_{real}(\cdot, p)$ if routers only know no more than a single unique external route towards p .

The main message from Figure 4 is how closely the *real diversity* curve follows the inverse of the number of total eBGP routes known to the AS. On average, routers know not much more than a single unique route (in terms of its eBGP origin) for any given prefix. This observation implies that the current iBGP structure of the studied network does not provide diversity in terms of the external routes. Furthermore, the value of the *RIB diversity* is about 0.5 for a large fraction of the prefixes. About half the entries in the Adj-RIB-In's are duplicate routes in terms of the eBGP peer who advertised the route inside the AS. This reflects the design choice of the studied network, which connects routers to several iBGP peers but the latter advertise the same eBGP-originated route. Note that as we ordered

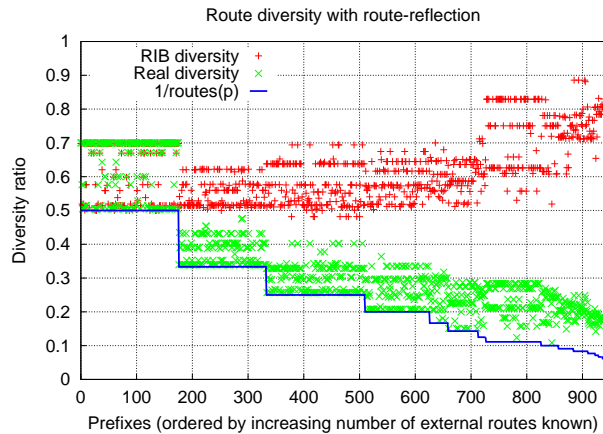


Fig. 4. External diversity for each prefix.

prefixes by increasing number of eBGP routes known, the large variations in this number of eBGP routes known for prefixes (up to 17) is pretty important, see the " $\frac{1}{routes(p)}$ " curve.

8 Route sampling performed by BGP route selection

Which routes are chosen as best by the routers inside an AS are another important factor that explain diversity inside the iBGP. Among all the routes advertised by eBGP peers, a subset of them are preferred by BGP routers because of the BGP decision process chooses the best route. From which kind of neighboring AS the route comes, its AS path length, and other attributes of the routes are a key factor for determining which routes will never be selected as best by any router inside the AS.

Figure 5(a) provides for each prefix the percentage of all external routes known that have been selected as best route by at least one router inside the network. The prefixes on both graphs of Figure 5 (x-axis) have been ordered by increasing number of external routes known ($route(p)$). There are three regions on Figure 5(a) that correspond to three different types of prefixes. The first type are those prefixes for which all external routes have been selected as best by at least one router inside the AS. For these prefixes, no external route is lost at the border routers. Note that most of the prefixes for which there is no loss of external routes at border routers are mainly those having only 2 external routes known. It is very unusual that prefixes having a higher number of external routes have all their external routes selected by at least one router.

The second type of prefixes are those for which only a single route is selected by all routers. 192 over the 940 prefixes having more than 2 external routes known inside the AS have only a single external route selected by all routers of the AS. The reason why these prefixes have only one route chosen as best by all routers is that routers prefer this single route over the others.

Finally, the third type of prefixes are in-between, with some routes lost at border routers, but more than one route is selected as best by at least one router. We can also observe on Figure 5(a) that prefixes for which a large number of external

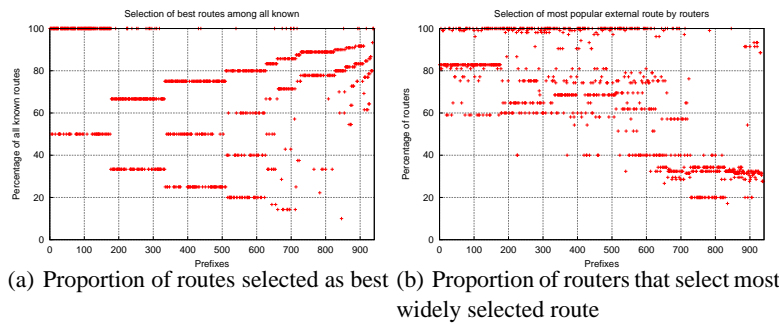


Fig. 5. Selection of best routes among all known.

routes are known tend to have a large fraction of these external routes selected as best by at least one router.

Figure 5(b) gives for each prefix, the fraction of all routers that selected the most popular among all the known routes. By *most popular route*, we mean the route which was selected as best by the largest number of routers inside the AS. Among the subset of the routes that are selected as best by at least one router, the one that is selected as best by the largest number of routers is chosen by a very large fraction of the routers compared to other routes. Obviously, all prefixes of the second type according to the previous paragraph will appear as points with 100% of the routers having selected the same route on Figure 5(b). We can see on Figure 5(b) that most points lie above 50%, except for prefixes having a very large number of external routes known. For the latter, the choice of the best route is less biased towards a single route.

9 Route diversity per router

Even though the previous section showed that the choice of the best routes inside the studied network favors a loss in route diversity across the iBGP graph, we would expect that diversity is still present somewhere in the AS. We might expect that the iBGP signaling graph under route-reflection limits the number of iBGP sessions compared to a full-mesh, but without removing all the route diversity known across the whole AS. In this section, we want to see whether there are differences among routers in terms of route diversity. A desirable goal would be that all routers know two unique routes for each prefix. In such a case, even if the current best route is withdrawn the router can switch immediately to the alternative route. Note that if the route is withdrawn due to a failure inside the AS or at the peering link over which the route was announced then local protection can be used.

Figure 6 show, for each router, how many unique external routes it knows towards any prefix. The y-axis of Figure 6 gives the 20 and 80 percentiles of this number of unique routes for each router over all considered prefixes. The ends of the bars show the 20 and 80 percentiles. Routers on the x-axis of Figure 6 are ordered by increasing median of their *real diversity*.

Figure 6 shows that some routers (the rightmost ones) have a large number of unique routes in their Adj-RIB-ins for most prefixes. These routers having diversity are both level-1 and level-0 reflectors. However, many routers have a value




Fig. 6. Unique routes known to routers.

of 1 both as their 20 and 80 percentile. This means that these routers only know 1 unique external route for most of the prefixes. These are mainly level-2 routers in the route-reflection hierarchy. This is something to be expected in a real network as most clients are topologically close to their route-reflectors, hence even though they might be connected to several higher level route-reflector, they will receive the same route (in eBGP origin) from the route-reflectors they are peering with. The iBGP structure of the studied network hence does not lack diversity, but diversity is very unevenly distributed among the routers. A few routers (top-level route-reflectors) have a very high diversity while most routers know only a single route.

10 Conclusion

In this paper, we quantified the diversity of the routes inside a tier-1 ISP. By building a model of the tier-1 ISP and reproducing its routing, we tried to better understand how its iBGP structure impacts its BGP route diversity.

We showed that the impact of the use of route-reflection on route diversity is significant. Most routers of our tier-1 network typically only know a single external route towards a destination prefix. Its iBGP graph propagated redundant routes that are not externally distinct from eBGP origin.

We identified two causes for this lack of diversity. First, some routes are never selected as best by any router inside the network, but are known only to one border router. Second, among the routes that are selected as best by at least one router, a few are selected as best by a majority of the routers, preventing diverse routes to propagate across the AS.

Our results point to the big distance in terms of route diversity between route-reflection and an iBGP full-mesh. Route-reflection thus reduces the number of iBGP sessions at a high cost in limiting the diversity of the routes inside the AS. Routes diversity inside an AS is important in case of failures, to ensure that all routers always have a route during the convergence of BGP after a failure. Our work hence calls for a deeper understanding of the possible trade-offs between iBGP route diversity, scalability and safety in the convergence of BGP.

Acknowledgments

The work of Sébastien Tandel was carried under France Télécom contract 46 126 699. We would like to thank Vincent Gillet from OpenTransit and Bruno Decraene from France Télécom R&D for their help to gather the data. We would also like to thank Olaf Maennel for many insightful suggestions about the presentation of this paper.

References

1. T. Bates, R. Chandra, and E. Chen, “BGP Route Reflection - An Alternative to Full Mesh IBGP,” Internet Engineering Task Force, RFC2796, April 2000.
2. B. Halabi and D. McPherson, *Internet Routing Architectures (2nd Edition)*, Cisco Press, January 2000.
3. C. Labovitz, A. Ahuja, A. Bose, and F. Jahanian, “An experimental study of Internet routing convergence,” in *Proc. of ACM SIGCOMM*, August 2000.
4. Cisco, “BGP best path selection algorithm,” <http://www.cisco.com/warp/public/459/25.shtml>.
5. Timothy G. Griffin and Gordon Wilfong, “On the correctness of iBGP configuration,” in *Proc. of ACM SIGCOMM*, August 2002.
6. N. Feamster and H. Balakrishnan, “Detecting BGP Configuration Faults with Static Analysis,” in *Proceedings of the 2nd Symposium on Networked Systems Design and Implementation (NSDI)*, May 2005.
7. B. Quoitin and S. Uhlig, “Modeling the routing of an Autonomous System with C-BGP,” *IEEE Network Magazine*, vol. 19, no. 6, November 2005.
8. B. Quoitin, “C-BGP, an efficient BGP simulator,” <http://cbgp.info.ucl.ac.be/>, September 2003.
9. Cisco, “NetFlow services and applications,” White paper, available from <http://www.cisco.com/warp/public/732/netflow>, 1999.
10. T. Cormen, C. Leiserson, R. Rivest, and C. Stein, *Introduction to Algorithms, Second Edition*, MIT Press and McGraw-Hill, 2001.
11. D. Walton, A. Retana, and E. Chen, “Advertisement of Multiple Paths in BGP,” Internet draft, draft-walton-bgp-add-paths-04.txt, work in progress, September 2005.