A Combinatorial Optimization Framework for the Design of Resilient iBGP Overlays

Cristina Mayr
Instituto de Computación
Universidad de la República
Montevideo, Uruguay
mayr@fing.edu.uy

Claudio Risso
Instituto de Computación
Universidad de la República
Montevideo, Uruguay
crisso@fing.edu.uy

Eduardo Grampín
Instituto de Computación
Universidad de la República
Montevideo, Uruguay
grampin@fing.edu.uy

***Abstract*—The Internet is an aggregation of Autonomous Systems (ASes) which exchange network prefixes reachability advertisements using the Border Gateway Protocol (BGP). ASes set up external BGP (eBGP) sessions between the AS border routers (ASBR) of neighboring ASes, while internal BGP speakers establish internal Border Gateway Protocol (iBGP) sessions to learn reachability for external prefixes. In order to avoid loops in the control and forwarding plane, and to ensure complete visibility and path diversity, routers within the same AS must deploy full-mesh BGP sessions, which causes scalability problems, both in the number of sessions and the resources (memory, CPU) consumed by BGP routers. Route Reflection is a widely accepted alternative to improve scalability, but requires careful design, as new issues may be introduced, such as: increased probability of loops, divergence and routing sub-optimality. In our previous work we presented Optimal Route Reflector Topology Design (ORRTD), a combinatorial optimization approach to tackle the problem of designing a consistent and yet optimal iBGP overlay, which minimizes the number of Route Reflectors (RRs), guaranteeing that no sub-optimal route is chosen, i.e., the routes selected with the designated RRs are those that would have been selected if instead of having RRs, the iBGP speakers were fully meshed. In this paper we propose a modification to ORRTD that addresses resilience, i.e., survivability to node or link failures.**

***Index Terms*—Internet Routing, BGP, Route Reflection, Network Design, Combinatorial Optimization, BGP resilience**

I. Introduction

The inter-domain routing is supported by the Border Gateway Protocol (BGP, [16]), which is used to exchange reachability information among Autonomous Systems (ASes). Intra-domain routing is fulfilled by the interaction of the Interior Gateway Protocol (IGP) and Internal BGP (iBGP): while the IGP builds connectivity for internal prefixes, iBGP is used to determine the exit gateway for those packets whose destination is external to the AS. A router running external BGP (eBGP) sessions is called Autonomous System Border Router (ASBR), while a router running only iBGP sessions is called Internal Router (IR).

As described in our previous works [4], [5] and related work op.cit., to avoid BGP loops and make sure that the complete routing information is disseminated, a full mesh of iBGP sessions between each pair of routers in the AS is required, resulting in $\frac{n\times(n-1)}{2}$ iBGP sessions for a domain of n routers, imposing large CPU and memory requirements to hold the Rib-In tables.

Route reflection [1] is an alternative in which one or more routers within the AS are designated as Route Reflectors (RRs) and they are allowed to re-advertise routes learned from an internal peer to other internal peers, while the rest plays the RR client role. With route reflection the number of iBGP sessions decreases to $n-1$ when using a sole RR, where n is the number of iBGP routers. Since a unique RR in an AS constitutes a single point of failure, at least two routers are to be selected as RRs.

BGP route selection process combines IGP and BGP routing information, and consequently RRs decisions are influenced by their locations within the AS. The problem of selecting which routers will have the RR role, following a consistent set of client-RR adjacencies (i.e., the RR topology), is known as *iBGP overlay design problem.* and has been extensively explored as in [3], [7].

In previous works [4], [5] we presented *Optimal Route Reflector Topology Design* (ORRTD), a novel combinatorial optimization approach to tackle the problem of designing a consistent and yet optimal iBGP overlay for an AS. The optimality criterion is to use the minimum number of route reflectors (RRs) and sessions, maintaining *correctness* and *full mesh optimality* [2], [9], [19], assuming that all prefixes matching a common gateway, or a set of equally preferred gateways are clustered into *classes* of prefixes (or labels).

This paper complements [4], introducing resilience, and is organized as follows: Section II explains what is an iBGP overlay, describes some design solutions, and explains the basis of ORRTD, our novel solution to design the overlay, section III proposes a mathematical approach for RRs selection and explains how to introduce resilience in the model, Section IV presents experimental results over some network topologies, Section V discusses the problem complexity and finally, Section VI summarizes our main conclusions and lines for further research.

II. iBGP overlay design based on route reflection

The selection of the best BGP route at each router depends on the IGP path cost to the BGP next hop announcing the route due to the *Hot potato* routing, where the preferred route is the one with shortest IGP path (the closest exit point).

Previous research works about RR selection, focus mainly in reliability, such as [22], [12], [19], or in reducing the number of sessions [24], or in modifying RR behavior or avoid them, as in [3], [14], [6]. Alternatives or variations to classical RR architecture have also been proposed to improve BGP reliability, including Multi-path [21], BGP Advertise-Best-External [10], Add-Path [21] and Diverse-Path [13].

In addition to the previously referred works, which mainly focus upon reliability issues associated with Route Reflection, this work also aims on those problems derived from the lack of optimality.

Control variables for designing a reliable iBGP route reflection topology should answer the following questions:

1) Which routers are to be chosen as route reflectors;
2) How clients are to be connected with route reflectors.

The objective function to be minimized counts the number of RRs, which also determines the number of BGP sessions. Constraints are introduced to avoid the problems described in [4], not only for steady/non-faulty state, but also to preserve such attributes after each possible single node or link failure. We will show in section V that the problem is NP-Complete.

The object of this work is to design a reliable iBGP overlay with minimum number of RRs and sessions, resilient to single node or link failure, and yet *route optimal* for a steady configuration of eBGP messages upon an internal given topology.

The technique introduced in this work is called *Optimal Route Reflector Topology Design* (ORRTD). It aims on being optimal in terms of routing, and in the number of RRs and sessions to keep. As we see in section III, the technique relies upon an integer programming problem formulation, whose constraints have been chosen to always select IGP optimal routes. In [4], [5], we demonstrate that this technique preserves correctness, and besides, *optimality principle* applies to all internal routers (by construction).

III. From the raw problem to its Integer Programing Formulation

This section introduces a mathematical formulation of the problem in two steps. The first (simpler) approach focuses upon optimization concerns of the problem. The second extends the basic formulation to integrate resilience to the design, which constitutes the main contribution of this article.

With ORRTD no additional functionality or BGP sessions are needed, and no changes to BGP process are suggested.

We consider an AS with a collection of BGP speaking routers, either IRs or ASBRs, connected by a pure IP network (i.e. hop-by-hop routing); and we want to determine which of the IRs will be designated as RRs. We assume that only an IR can be RR, every Client-IR (i.e. not RR) must be peer of a unique RR per class of prefixes, whenever optimal for some IR, every ASBR must be connected to at least one RR per class of prefixes and an ASBR cannot peer with a Client-IR. We also assume that all external prefixes are learned through BGP and they have been filtered according to the path selection algorithm to get to a set of prefix classes;

Suppose we have the graph associated with the network of some AS, like that represented in graph in Fig. 1.

This graph is undirected and the weight of each link is the IGP cost. ASBRs *A* and *AB* receive prefixes class *A*, while *B* and *AB* receive prefixes class *B*.

From that information, an optimal internal-to-border router graph can be build ([4], [5]) for each class of prefixes, which are respectively sketched on the left and right on Fig. 2. Complementary, IRs that share a common ASBR for a common prefix class, could serve as the reflector of each other for that class (Fig. 3).

A. Resilience considerations

Failures can occur at the links, or at the nodes (IRs or ASBRs). Even if the failure is in any IR or link, they could be in the shortest path calculated for deriving the graph in Fig. 2 and Fig. 3. To really ensure resilience, disjoint paths between each ASBR and the corresponding RR are needed.

In the present work we propose a solution when any element at a time fails, either a link, an IR, or an ASBR. We are interested in computing the smallest number of additional links (or nodes) that need to be added in order to increase the resilience of a network against random failures.

In order to make it possible to propose a resilient solution, we assume the original IGP graph is at least 2-node-connected, which translates into the existence of two node (and link) independent paths between every pair of nodes.

This guarantees in turn the existence of a detour against every possible single failure. More generally, k-edge (node) connectivity refers to the minimum number of edges (nodes) to be removed so that the graph becomes disconnected. Both problems are NP-complete.

If a graph is k-node-connected it can be proved that there are k node-disjoint paths between any pair of nodes.

B. Resilient ORRTD

Suppose we have a best p path from certain $u \in IR$ to $v \in ASBR$ for prefix class A. Let $p = u, x_1, ..., x_h, v$.

To add resilience to ORRTD we consider every type of failure:

1) *link failure* - an edge $e = (x_i, x_{i+1})$ fails. Suppose that, without this edge, the new closest ASBR from u is w. Then create a fictitious prefix class C_l advertised by w.

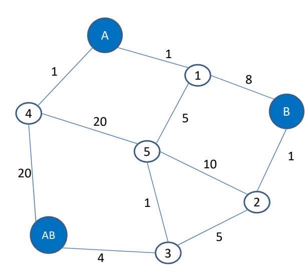

Fig. 1. Graph with 2 prefix classes (A and B)

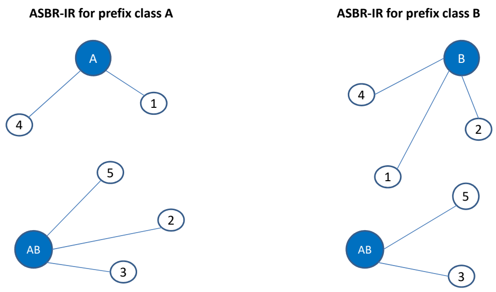

Fig. 2. IR-ASBR Adjacency graphs - 2 prefix classes (A and B)

2) *internal router failure* - if $x_i \in IR$ fails and it belongs to some best path p from another $u \in IR$ to $v \in ASBR$, then a new best path to some ASBR must be calculated. If the new best path ends in a different node $w \in ASBR$, proceed as in the previous case.
3) *border router failure* - if $v \in ASBR$ fails and it is the best exit for some $u \in IR$, proceed as in the first case.

We will analyze the case depicted in Fig. 1. The best ASBR for IR 5 and prefix class B is AB, by using the path $5-3-AB$. If the link $(5,3)$ from IR 5 to ASBR AB fails, then the best ASBR for IR 5 and prefix class B in $G' = (V, E \diagdown (5,2))$ is B, an the path is $5-2-B$ (Fig. 4). Then we add a prefix class B_j advertised by B, and an affinity set of nodes corresponding to the new best path, similar to those presented in Fig . 2. Observe that we add only one *fictitious prefixes class* for each combination of: IR $\times$ original prefix class $\times$ new ASBR in failure scenario. This guarantees that each IR gets optimal prefixes for all of those ASBRs for which it is necessary to keep optimality after each possible node or link failure. Although now we have to ensure that this does not introduce sub-optimal paths to other IRs. This can be achieved by ensuring that the fictitious prefix classes appear only in the routing tables of the nodes belonging to the alternative path considered. Note that in this path, an ASBR can appear as an intermediate node. This does not introduce any problem, as the fictitious prefix class is advertised by eBGP, so the next hop remains unaltered. When considering the next steps, the prefix class comes through iBGP, so, unless the receiving router is a RR, it cannot re-advertise that prefix class. The next step consists in assembling all these pieces into a single combinatorial optimization problem.

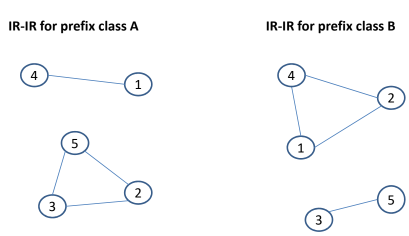

Fig. 3. Internal to Internal IR affinities graphs for prefixes classes A and B

$$
\left\{
\begin{array}{lll}
\min \displaystyle\sum_{i \in IR} x_i & & \\
\text{Subject to :} & & \\
\displaystyle\sum_{(ij) \in S'^k} y_{ij}^k \geq 1, & \forall i \in BR, k \in \mathcal{C}', \; S'^k \neq \emptyset & (i) \\
x_j - y_{ij}^k \geq 0, & \forall i \in BR, k \in \mathcal{C}', \; (ij) \in S'^k & (ii) \\
x_j + \displaystyle\sum_{(ij) \in T'^k} z_{ij}^k \geq 1, & \forall j \in IR, k \in \mathcal{C}' & (iii) \\
x_i + x_j - z_{ij}^k \geq 0, & \forall i \in IR, k \in \mathcal{C}' \; (ij) \in T'^k & (iv) \\
x_i + x_j + z_{ij}^k \leq 2, & \forall i \in IR, k \in \mathcal{C}' \; (ij) \in T'^k & (v) \\
\displaystyle\sum_{(jh) \in S'^k} y_{jh}^k - z_{ih}^k \geq 0, & \forall j \in IR, k \in \mathcal{C}' \; (ih) \in T'^k & (vi) \\
\displaystyle\sum_{i \in IR} x_i \geq 2, & \forall i \in IR & (vii) \\
w_{gh}^l \geq y_{ij}^k, & \forall i \in BR, j \in IR, k \in \mathcal{C}', \; g,h \in FC^l & (viii) \\
\displaystyle\sum_{(ij) \in P^l} y_{ij}^l \geq 1, & \forall i \in BR, l \in \mathcal{FC}, & (ix) \\
x_i, y_{ij}^k, z_{ij}^k, w_{gh}^l \in \{0,1\}, & \forall i,j \in V, k \in \mathcal{C}', l \in FC &
\end{array}
\right. \quad (1)
$$

Equations (1) have the following **input sets**:

$\mathcal{C}$	:	set of prefixes classes
$\{S^k\}$	:	set of border-to-internal BGP affinity matrices $S_{ij}^k = 1$ if and only if $j \in$ ASBR-to-IR for prefix class k, with $k \in \mathcal{C}$, $i \in$ BR, $j \in$ IR
$\{T^k\}$	:	set of internal-to-internal BGP affinity matrices
FC	:	set of fictitious prefix classes
$\{P^l\}$	:	set of new BGP best path nodes from internal-to-BR affinity matrices
$\{Q^l\}$	:	set of new BGP best path IR-to-IR affinity matrices

Equations (1) have the following **parameters** to support resilience:

BR	:	set of all Autonomous System Border Routers
IR	:	set of all Internal Routers
S'	:	$\{S^k\} \cup \{FC^l\}$
C'	:	$C \cup FC$
T'	:	$\{T^k\} \cup \{Q^l\}$

and the following boolean **variables**:

x_i	:	1 if router i is to be a RR and 0 otherwise;
y_{ij}^k	:	1 if ASBR i is to be iBGP adjacent to IR j for prefixes class k and 0 otherwise;
z_{ij}^k	:	1 if IR i is to be iBGP adjacent to IR j for prefix class k and 0 otherwise;
w_{gh}^l	:	1 if nodes $g,h \in P^l$, i.e., the alternative best path

The objective function in (1) pushes down to get the minimum number of RRs. But this objective has several

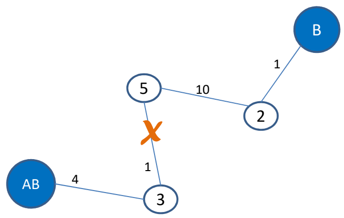

Fig. 4. New path in failure case, for prefix class B

constraints, stated in equation groups (i) to (ix). It is worth mentioning that RRs are globally selected, that is, they are common to all prefixes classes.

Equation groups (i) to (vi) are similar to those described in [4] and [5], but now considering the input sets and parameters described above; the rest of the equations are added to force resilience.

Equations (vii) ensure there is more than one RR, so the RR in not a single point of failure. Finally, equations $(viii)$ and (ix) ensure that only the nodes in the best alternative path learn the fictitious prefix classes.

For small networks, solutions can be easily found by brute-force or quasi-exhaustive methods. For more complex networks the problem can be solved with any popular solver like GLPK or CPLEX, even for hundreds of nodes, while the number of prefixes classes is limited. For even more complex problems, with hundreds of classes, a heuristic approach should be used.

In summary, the problem formulation has an augmented set of prefix classes. The new quantity of classes k' is k plus all the combinations of links and routers that can fail for each prefix class. In this new scenario, solution might not be found, given the increase on the number of restrictions. Anyway, in dense graphs, paths tend to repeat, so the problem can be preprocessed to eliminate redundant conditions. It is expected that there will not be too many additional prefix classes and so the number of constraints will not grow excessively.

IV. Experimental results

In this section we present early results obtained in the emulation environment proposed by [17] which is based on Quagga[1], MiniNExT[2] and ExaBGP[3] for injecting BGP messages. Some of the topologies were taken from "The Internet Topology Zoo" repository ([18]) and slightly adapted, for example, to ensure no vertex has degree one, as this makes finding a resilient topology design non-viable, and other topologies are theoretical cases.

For the purpose of this test we assume there are two prefix classes, and we know in advance which ASBRs advertise each prefix class. BGP updates from ASBRs will be set with the *next-hop self* option, as internal network does not know about external routes. The experimental results for the nominal case were presented in our previous work [4], [5]. In all cases ORRTD results in a reduced number of RRs.

In Table I we show, for two prefix classes, the resulting RRs applying ORRTD for the nominal case, and for the resilient case applied to different network topologies. We also show the number of equations needed for each topology. It can be easily seen that it quickly increases as the network becomes bigger. These early results show that in many cases the final quantity of RRs remains the same, and the change is about which IRs are chosen as RR, and an increased number of iBGP sessions established among the routers, as can be seen in Fig. 5. In other cases the number of RRs does increase. We observe that this strongly depends on the underlying topology. We remark that we assume RRs are in fact connected in a full-mesh, as it is the standard, so it is not introduced in the model as a constraint, and so it is not shown in Fig. 5. In Table II we show the reduction in the number of BGP sessions in the resilient version of ORRTD compared to full mesh.

V. Problem complexity

We show that finding a minimal solution in $ORRTD$ is at least as hard as finding a solution for Minimum Vertex Cover (MVC) problem, which is known to be NP-complete [8] and in fact APX-complete ([11]).

Formally, a vertex cover S of an undirected graph $G = (V, E)$ is a subset of V such that $uv \in E \Rightarrow u \in S \vee v \in S$. We consider ORRTD where there is just one prefix class, as it seems natural that if there are more prefix classes, the problem will be even more difficult. The decision version of both problems is as follows:

1) π' - Given an undirected graph $G' = (V', E')$ and a constant k, is there a subset S of V' such that $uv \in E' \Rightarrow u \in S \vee v \in S$ with size $\leq k$?
2) π - Given a weighted undirected graph $G = (V, E)$, where $V = IR \cup BR$ and a constant k, is there a subset $RR \subseteq IR$, with size $\leq k$ constructed with ORRTD?

A reduction from π' to π can be built as follows:

- for each vertex $v \in V'$ of π' there will be a vertex $v \in IR$
- for each edge $uv \in E'$: add an edge $uv \in E$ between a pair of vertices $u, v \in IR$ with weight 1, add a new vertex x_{uv} to V and a pair of edges from x_{uv} to u and from x_{uv} to v to E with weight 1.
- let $x_{uv} \in BR$. Then by construction, every IR is at distance 1 to some ASBR.

TABLE I
ORRTD - Comparison of Nominal and Resilient Case

Topology	# IRs	# ASBRs	RRs Nom.	RRs Resil.	Eqs. Nom.	Eqs. Resil.
Abilene	8	3	2	2	114	419
AB5	5	3	2	3	55	131
AB10	10	3	2	2	250	478
Airtel	3	6	1	2	32	131
Garr	47	7	4	4	3080	3852
UniC	24	3	2	2	902	944
Uran	18	5	3	5	599	687
Jgn2Plus	11	6	2	6	399	453

[1]Quagga Routing Suite. Available at: https://www.quagga.net/. Accessed: 2018-09-01

[2]MiniNExT (Mininet ExTended). Available at: shttps://www.quagga.net/. Accessed: 2018-09-01

[3]https://github.com/Exa-Networks/exabgp

TABLE II
COMPARISON OF BGP SESSIONS

Topology	resilient ORRTD	Full Mesh
Abilene	20	55
AB5	18	28
AB10	23	78
Airtel	26	36
Garr	331	1431
UniC	24	351
Uran	57	253
Jgn2Plus	64	136

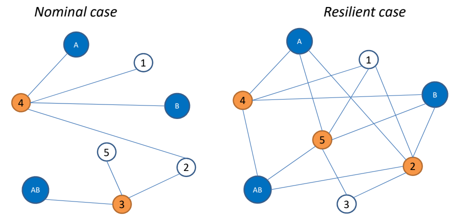

Fig. 5. More connections in resilient case

This graph in π has been designed to have a set of RRs if and only if a set cover exist in π', so $\pi' \preccurlyeq \pi$. Besides, this is a polynomial reduction.

VI. Conclusion

In this article we focus on the efficient usage of BGP, particularly in the intra-domain scope, though it suffers from serious scalability issues. With Route Reflection, a classic and simple approach, widely standardized over the Internet infrastructure, but requiring careful design, as it could lead to other kinds of issues, as described in section II. We based our proposal on overlay networks and present a novel mathematical approach to tackle several known problems of reflection, by means of a design that optimizes the scalability. The technique has been called Optimal Route Reflector Topology Design, or ORRTD for short. Among other advantages, with ORRTD there is no need to modify or augment existing BGP standards. Early experimental results in emulation environments demonstrate the theoretical consistency of ORRTD, even in the event of fails over single nodes or links. Besides, ORRTD outperforms other heuristic approaches, and according to our experimental results with known topologies, the number of RRs does not increase significantly, and even remains the same, while augmenting the BGP sessions needed.

It is also worth to mention that we assume that prefix classes categorization is a given input for the optimization process, and is done based on ISP policies, either static or dynamically. This classification may constitute a whole line of research, for example, using machine learning or other techniques to build the prefixes classes based on the dynamics of BGP updates. We also prove ORRTD is a NP-hard problem, which implies that when considering larger instances of the problem, some heuristic approaches should be considered to solve it, which introduces a new line for future research.

References

[1] T. Bates, E. Chen, and R. Chandra, "BGP route reflection: an alternative to full mesh internal bgp (iBGP)," RFC4456 (Draft Standard), 2006.
[2] M. O. Buob, S. Uhlig, and M. Meulle, "Designing Optimal iBGP Route-Reflection Topologies," LNCS4982, 2008.
[3] C. Pelsser, A. Masuda, and K. Shiomoto, "A novel internal BGP route distribution architecture," in IEICE General Conference, 2010.
[4] C. Mayr, and E. Grampín, and C. Risso,"Optimal route reflection topology design," in 10th Latinamerican Networking Conference (LANC '18). ACM, New York, NY, USA, pp. 65-72, 2018.
[5] C. Mayr, and E. Grampín, and C. Risso,"Optimal route reflection topology design," *Technical Report*, https://www.fing.edu.uy/~crisso/Optimal_Route_Reflection_Topology_Design.pdf, 2018.
[6] S. S. M Dakshayini,"Effect of route reflection on iBGP convergence and an approach to reduce convergence time," in International Journal of scientific research and management (IJSRM), 2016, vol. 4 (8).
[7] A. Flavel,"BGP, not as easy As 1-2-3," Ph.D. Dissertation, 2009, University of Adelaide, Australia.
[8] M. R. Garey and D. S. Johnson,"Computers and intractability: a guide to the theory of NP-completeness," W. H. Freeman, San Francisco, 1979.
[9] T. G. Griffin and G. Wilfong,"On the correctness of iBGP configuration," in SIGCOMM '02 Proceedings of the 2002 conference on Applications, technologies, architectures, and protocols for computer communications, pp. 17–29, New York, NY, USA, 2002. ACM.
[10] P. Marques, R. Fernando, E. Chen, P. Mohapatra, and H. Gredler,"Advertisement of the best external route in BGP, " draft-ietf-idr-best-external-05, 2017.
[11] C. H. Papadimitriou, and M. Yannakakis,"Optimization, approximation, and complexity classes," Journal of Computer System Sciences 43, pp. 425-440, 1991.
[12] J.H. Park. ,"Understanding the impact of internal BGP route reflection," PhD thesis, University of California, 2011.
[13] R. Raszuk, R. Fernando, K. Patel, D. McPherson, and K. Kumaki,"Distribution of diverse BGP paths," rfc6774, 2012.
[14] M. H. Raza, A. K. Kansara, and W. Robertson,"Effective iBGP operation without a full mesh topology," in International Refereed Journal of Engineering and Science (IRJES), vol. 5(8), pp. 16–23, 2016.
[15] E. Gutiérrez and D. Agriel and E. Saenz and E. Grampín,"RRLOC: A tool for iBGP RR topology planning and experimentation," in IEEE Network Operations and Management Symposium, NOMS 2014, Krakow, Poland, May 5-9, 2014, pp. 1–4. IEEE.
[16] Y. Rekhter and T. Li,"A Border Gateway Protocol 4 (BGP-4)," RFC1771 Obsoleted by RFC 4271, 1995.
[17] V. Solla, G. Jambrina and E. Grampín,"Route reflection topology planning in service provider networks," in 2017 IEEE URUCON, pp.1–4.
[18] S. Knight, H.X. Nguyen, N. Falkner, R. Bowden and M. Roughan,"The Internet Topology Zoo," Selected Areas in Communications, IEEE Journal, 2011, October, vol. 29(9), pp. 1765 -1775.
[19] S. Vissicchio, L. Cittadini, L. Vanbever, and O. Bonaventure,"iBGP deceptions: more sessions, fewer routes," in INFOCOM, IEEE, 2012.
[20] M. Vutukuru and P. Valiant and S. Kopparty and H. Balakrishnan,"How to construct a correct and scalable iBGP configuration," in Proceedings of the 25th IEEE International Conference on Computer Communications. INFOCOM, 2006, pp. 1-12.
[21] D. Walton, A. Retana, E. Chen, and J. Scudder,"Advertisement of multiple paths in BGP," RFC7911, July 2016.
[22] L. Xiao, J. Wang, and K. Nahrstedt,"Reliability-aware iBGP route reflection topology design," Proceedings - 11th IEEE International Conference on Network Protocols, ICNP, January 2003, pp.180–189.
[23] F. Zhao and X. Lu and P. Zhu and J. Zhao,"BGPSepD: an improved algorithm for constructing correct and scalable iBGP configurations based on vertexes degree," in Proceedings of the second International. Conference on High Performance Computing and Communications., Springer-Verlag, 2006, pp. 406-415.
[24] R. Zhang and M. Bartell,"BGP design and implementation," Cisco Press, Indianapolis, 2003, pp. 264–266.