Automatic Inference of BGP Community Semantics

Brivaldo Alves da Silva Junior[1], **Italo Scotá Cunha**[2], **Ronaldo Alves Ferreira**[1]

[1]FACOM– Federal University of Mato Grosso do Sul (UFMS)

[2]DCC – Federal University of Minas Gerais (UFMG)

{brivaldo.junior,ronaldo.ferreira}@.ufms.br, cunha@dcc.ufmg.br

Abstract. *The Border Gateway Protocol (BGP) enables communication between Autonomous Systems (ASes) on the Internet. BGP offers significant flexibility for traffic engineering through BGP communities, which are operator-defined tags that convey information or requests in route announcements. Unfortunately, the absence of standardized semantics or centralized repositories for BGP communities complicates and limits their use, hindering the effective management of interdomain routing. This thesis develops techniques to infer BGP community semantics using public BGP data from routing collectors, overcoming the lack of documentation and providing datasets that can be automatically updated. We first propose a set of techniques to infer location communities, which are communities related to entities or locations traversed by a route. We apply our techniques to billions of routing records from public BGP collectors and show that they produce high precision (ranging from 86% to 93%) and recall (ranging from 72% to 81%). We also design and evaluate algorithms to automatically uncover BGP* action communities *and ASes that violate standard practices, revealing undocumented relationships between them (e.g., sibling relationships). Our experimental evaluation uncovers previously unknown AS relationships and shows that our algorithm to identify action communities achieves average precision and recall of 92.5% and 86.5%, respectively.*

1. Introduction

The Internet consists of Autonomous Systems (ASes) that exchange reachability information using BGP, the *de facto* interdomain routing protocol [Rekhter 2006]. In BGP, the construction of a route begins with an *origin AS*, which announces an IP prefix to its neighbors. This announcement propagates via BGP *updates*, which contain mandatory and optional attributes. Mandatory attributes include the destination IP prefix, the next-hop router's IP, and the *AS-path*, which records the sequence of ASes a route traverses. Optional attributes can be transitive (forwarded by ASes) or non-transitive and include BGP communities and multi-exit discriminators (MEDs).

The BGP best-path selection algorithm allows operators to rank routes based on policies and economic agreements using parameters like LocalPref for route preference, MED for preferred interconnections, and intradomain cost minimization. However, these mechanisms are coarse-grained and only influence routes received from neighboring ASes. As networks demand greater reliability and performance, routing policies have become more dynamic and complex, exposing the limitations of BGP [Giotsas et al. 2014, Streibelt et al. 2018].

To overcome the limitations in BGP expressiveness, network operators increasingly use BGP communities, which are 32-bit tags whose meaning (*i.e.,* semantics) are

defined independently by each network. Network operators generally group BGP communities into two types: *informational* and *action*. *Informational communities*[1] convey details such as the country, city, PoP, or router where a route was learned, or the business relationship with the neighboring network that announced it [Giotsas et al. 2014]. *Action communities* request specific routing actions from upstream networks [Streibelt et al. 2018], such as AS path prepending to make a route less attractive or preventing a prefix from being advertised to certain peers to steer traffic away from low-performance ASes.

Unfortunately, BGP communities are opaque identifiers with no standardized semantics, allowing each network to define their own values and meanings. For example, network A may use A:X to trigger AS path prepending, while network B may use B:X to indicate that it received a route in New York. Some networks document their communities in Internet Routing Registries (IRR) databases [Tools 2024] or webpages, but most observed communities remain undocumented.

The lack of standardization and public databases that map community values to their semantics hinder traffic engineering and the development of tools to automate network management. Operators must rely on incomplete or outdated IRR records, webpages, or direct communication with AS operators. This manual and error-prone process makes it harder to integrate communities into routing decisions, often leading to suboptimal paths and limiting researchers' ability to analyze interdomain routing.

2. Problem Statement and Contributions

A recent study [Giotsas et al. 2017] uses natural language processing to infer the meaning of documented communities in IRR records and support webpages. This approach has two fundamental limitations: *(i)* it infers a small number of communities, as it depends on the descriptions provided by network operators; and *(ii)* it relies on incomplete and outdated data, leading to reduced precision and limited coverage of existing communities. Also, an AS can use the communities of a sibling AS, which complicates the understanding of BGP community usage, as an informational community may appear in a route announcement without the AS that defined it. Therefore, network operators still rely on documentation provided by each AS about their BGP communities and relationships with other ASes, which is often incomplete and insufficient for effective troubleshooting and understanding of Internet routing. Thus, our problem statement can be summarized as follows:

Problem Statement: *Networks do not publicly provide necessary and sufficient information about their BGP communities and relationships with other networks for effective troubleshooting and understanding of Internet routing.*

This thesis [da Silva Jr 2024] contributes to partially closing this gap by automatically building reliable databases to document a subset of communities that are actively being used on the Internet, *i.e.,* communities that appear in public BGP route collectors. It also presents mechanisms to uncover an undocumented type of confounding use of BGP communities in the wild in which an AS consistently uses the informational communities of another AS, which might help operators understand BGP community uses or uncover undocumented AS relationships. This behavior can impact previous research that infers AS relationships or the semantics of BGP communities [Giotsas et al. 2014, Krenc et al. 2023]. Specifically, this thesis focuses on the following two research questions (RQ):

[1] In this article, we use the term *informational community* interchangeably with *information community*.

RQ 1: *Can we build reliable databases of BGP community semantics using public routing data?*

We address this research question by developing techniques to automatically infer the semantics of BGP communities directly from publicly available route announcements collected by route collectors. We initially target *location communities* (§3), defined as communities that carry metadata about the location (*e.g.,* city, country, continent, router, PoP or link) where a route was learned. Location communities allow richer manipulation inside the tagging AS, but they would also be helpful to neighboring and remote ASes if their semantics were publicly available. For example, operators could use a tool that correlates BGP location communities and performance (*e.g.,* latency, jitter, etc.) to tune their route selection preferences at a finer granularity than possible with just AS paths.

We also present algorithms for identifying action communities from public routing data (§4). Our algorithms provide automatically updated databases of action communities that can benefit novel tools and models. For example, this information can help operators troubleshoot routing anomalies, *e.g.,* when routes with action communities follow unexpected paths, and identify opportunities for traffic engineering, *e.g.,* when an operator observes preferable routes induced by action communities not publicly documented. We show through longitudinal studies that our algorithms perform well over the years, even when ASes add new communities or decommission old ones, attesting to the reliability of the generated datasets [da Silva Jr et al. 2022, da Silva Jr et al. 2025][2].

RQ 2: *Can we use BGP communities to identify AS relationships?*

While sibling ASes may share communities, our analysis of routing announcements revealed cases where nonsibling ASes also use each other's communities. This behavior complicates the inference of location communities, as these tags can appear in route announcements without their defining AS in the AS path, deviating from expected patterns and making automated inference more challenging.

For inferring location communities, we build a heuristic based on the hitting set algorithm [Garey and Johson 1979] to detect the presence of these ASes that use the communities of others and prevent their presence from excluding location communities from the inference. We discovered that these relationships were not limited to sibling ASes, as some ASes use the communities of other ASes even when they are not siblings. We call this behavior *community squatting*[3] and identify the ASes involved as *AS squatters*.

We identify the squatting relationships and reduce the noise they introduce into the inference of action communities. We compare our inferred AS relationships with techniques that use public databases, such as PeeringDB [Arturi et al. 2023], and show that our algorithms capture relationships that the existing techniques missed, thus addressing Research Question 2. Our algorithm to uncover squatting relationships can complement techniques for validating AS-relationship inferences and tracking route changes.

[2]The papers presented at ACM SIGMETRICS are also published in the Journal Proceedings of the ACM on Measurement and Analysis of Computing Systems (POMACS).

[3]We borrow the term *squat* and its derived forms from "IP address squatting" [Salamatian et al. 2023], where a network uses another's IP address space internally for its own purposes.

2.1. Main Contributions

Our thesis contributes to the automatic identification of location (a subset of information communities) and action communities on the Internet. We treat these sets separately because they require different techniques and use BGP dumps from different time periods. We also design algorithms to identify ASes that engage in unconventional practices by consistently *squatting* on the BGP communities of other ASes, which we refer to as *squatting relationship*. This behavior can affect the validity of previous research that relies on BGP communities [Krenc et al. 2023, Giotsas et al. 2014, Streibelt et al. 2018].

2.1.1. Location Communities

Our approach to infer location communities (§3) [da Silva Jr et al. 2022] is fundamentally different from previous efforts, as our algorithms automatically infer communities from public route announcements observed by BGP route collectors [Meyer 1997, RIPE 2024] and generate databases of communities that can be regenerated any time to reflect additions of new communities or assignment changes. *Our work is the first to show that we can use public route announcements to infer, even at a coarse level, the semantics of BGP communities*.

We process over two billion route announcements from route collector projects and infer 15,505 location communities across 1,120 ASes, wich represents 19.67% of the communities that appeared in the BGP dumps in 2020. Our experimental evaluation shows that our methodology yields high precision (from 87% to 93%) and recall (from 72% to 81%). We compare our results with CAIDA's manually-built public database of BGP communities [CAIDA 2021] and show that our technique has higher recall and similar precision, with the advantage that it can be automatically updated as new BGP communities are defined or as definitions change over time.

2.1.2. Action Communities and AS Squatters

Similarly to our approach to inferring location communities, we infer action communities using only public BGP route announcements. The key difference between informational and action communities is their placement in routing paths [da Silva Jr et al. 2025]. *Informational communities* are added by an AS to convey details such as where it learned a route or its business relationship with the previous AS. Since they originate from the defining AS, they should always appear on routes that traverse it. In contrast, *action communities* request a routing action from another network, the *controlling AS*, which applies the request and typically removes the community, as prescribed by RFC 7454 [Durand et al. 2015]. This means action communities should rarely appear on routes that include their defining AS. Our algorithms leverage this distinction to classify action communities and detect *potential squatters*.

Our evaluation, based on billions of route announcements from 2018 to 2023, shows that our algorithm for identifying action communities achieves average precision and recall of 92.5% and 86.5%, respectively, across all communities covered by our manually-built ground truth. Analyzing 739 million announcements from December 2023, we inferred 19,564 action communities from 2,099 ASes, excluding 15,800 communities (14.86%) linked to private ASNs. Our squatter detection algorithm identified 54 squatting relationships involving 105 ASes that systematically used other ASes' communities. These cases may reveal undocumented AS relationships, including five sibling relationships missed by the state-of-the-art method in [Chen et al. 2023].

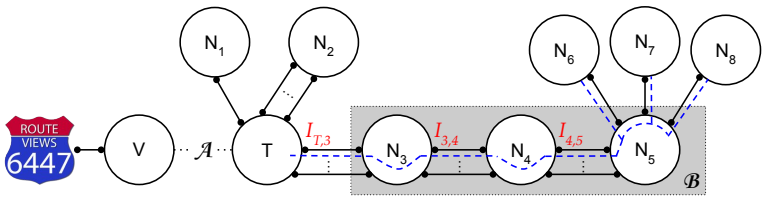

Figure 1. Example illustrating how long AS-paths between origin ASes and a target AS T constrain the locations where T receives and selects routes. We define $\mathcal{A}$ as the (possibly empty) AS sequence between the BGP collector peer V and T, and $\mathcal{B}$ (highlighted in gray) as the non-empty AS sequence that restricts where T receives BGP announcements. Solid black lines represent AS interconnections.

Our algorithms provide automatically updated metadata that can benefit novel tools and models. For example, action community information can help operators troubleshoot routing anomalies, *e.g.,* when routes that follow an unexpected or undesired path carry specific action communities, and identify opportunities for traffic engineering, *e.g.,* when an operator observes preferable routes induced by action communities not publicly documented. Our results show that operators use action communities much more extensively than publicly available documentation would indicate.

3. Inferring Location Communities

We infer location communities based on the fact that ASes peer at a finite set of locations and enforce dynamic but deterministic routing policies [Gao and Rexford 2001, Giotsas et al. 2014]. Consider a target AS T that tags received routes with location communities (see Figure 1). If AS T and AS N_1 interconnect at a single location, then T will tag *all* routes received from N_1 with the location community corresponding to their single interconnection. The idea that all routes received at a specific location will have the corresponding location community is the core of our algorithm. Unfortunately, we cannot simply infer communities that appear on all routes received from a neighbor N_1 as location communities. First, neighbor N_1 may tag all of its announcements with AS T traffic engineering communities, which would be incorrectly inferred as location communities. Second, when AS T and AS N_2 interconnect at multiple different locations (indicated by the multiple links between T and N_2 in Figure 1), then T may choose routes received from N_2 at any of these locations. Each chosen route will have a different location community corresponding to the interconnection over which it was received. Therefore, no community will appear in all routes, and no location community will be inferred.

We relax the requirement of a single interconnection and avoid the need to quantify the number of interconnections between the target AS T and neighboring ASes by analyzing paths that traverse multiple interconnections. Suppose that AS T and AS N_3 interconnect at multiple locations and that AS T receives a route with AS path $\langle N_3, N_4, N_5 \rangle$ (blue dashed line in Figure 1). Let $I_{T,3}$, $I_{3,4}$, and $I_{4,5}$ be the interconnections traversed by the route. Interconnection $I_{T,3}$ is constrained by the set of interconnections between ASes T and N_3 *and* their routing policies. Here is a non-exhaustive list of such constraints:

1. AS T might use multi-exit discriminators (MEDs) as a tie-breaker [Rekhter 2006] and choose routes from N_3 received at a particular interconnection. For example, if N_3 prefers to receive traffic from AS T towards $I_{3,4}$ at $I_{T,3}$, it may set lower MED values on routes exported at $I_{T,3}$, leading AS T to choose routes received at $I_{T,3}$ over routes received at other interconnections.

2. Routers systematically choose routes from the closest (lowest IGP cost [Rekhter 2006]) interconnection. For example, if $I_{T,3}$ is the closest interconnection to AS T's egress router towards the vantage point at V, then the egress router will choose and export routes from N_3 received at $I_{T,3}$.
3. Routes may not be accepted by AS T or exported by AS N_3 at some interconnections, especially in complex peering [Giotsas et al. 2014]. For example, if T and N_3 peer in Europe, but T buys transit from N_3 in the US, T will receive routes from N_3's peers and providers only in the US.

The constraints imposed by the set of interconnections and routing policies between each pair of ASes in a route compound over consecutive AS hops. In other words, interconnection $I_{3,4}$ is *also* constrained by the interconnections between ASes N_3 and N_4 and their routing policies. The same constraints apply to $I_{4,5}$. The implication is that chosen routes traversing a sequence of ASes (like $\langle N_3, N_4, N_5 \rangle$) will only be received by AS T at a small set of locations, possibly a single one. Looking at the problem another way, for AS T to receive routes traversing $\langle N_3, N_4, N_5 \rangle$ at different interconnections, then N_3 needs to receive and choose routes through $\langle N_4, N_5 \rangle$ at different interconnections, which implies N_4 receives and chooses routes from N_5 at different interconnections.

We sidestep incorrect inferences for origins that tag all their announcements with traffic engineering communities by combining observations on multiple routes from different origins. The chance that *all* these origins tag their announcements with AS T traffic engineering communities is low, which allows us to correctly remove traffic engineering communities from the set of inferred location communities. In our algorithm, we require routes from a configurable number of different origin ASes to infer location communities. The algorithm also requires other configurable parameters that dictate the minimum and maximum number of announcements satisfying specific properties to filter out possible noises [da Silva Jr et al. 2022].

3.1. Unknown Siblings

Another issue is that there are ASes that seem to tag routes with location communities of other ASes, with no apparent sibling relationship. For example, we observed announcements traversing AS20473 (Constant) tagged with location communities from AS1299 (Telia). We relax the heuristic to account for missing sibling ASes and cases where ASes reuse or incorrectly tag announcements with another AS's location communities. If a small set of ASes is responsible for tagging a target AS T's communities on routes that do not traverse T or its known siblings, we retain the inferred location communities. More precisely, let $\mathcal{R}_c$ be the set of routes tagged with community c from AS T, and let $\mathcal{R}_T$ be the set of routes whose AS paths traverse AS T or any of T's known siblings. We ignore routes that traverse T or any of T's siblings, and consider the route announcements $\mathcal{F}_c = \mathcal{R}_c \setminus \mathcal{R}_T$ when deciding whether to discard an inferred location community. We compute the *minimum hitting set* of $\mathcal{F}_c$ and discard c as a location community if the set contains more than K_{filter} ASes, where K_{filter} is a parameter of the algorithm.

In other words, we keep location community inferences only when few ASes are to blame for AS T's communities showing up on routes that do not contain T or any of T's siblings. The minimum hitting set is the smallest set of ASes $\mathcal{W}$ such that the intersection of $\mathcal{W}$ and each route $r \in \mathcal{F}_c$ is nonempty. The minimum hitting set problem is equivalent to the NP-complete minimum set cover problem [Garey and Johson 1979],

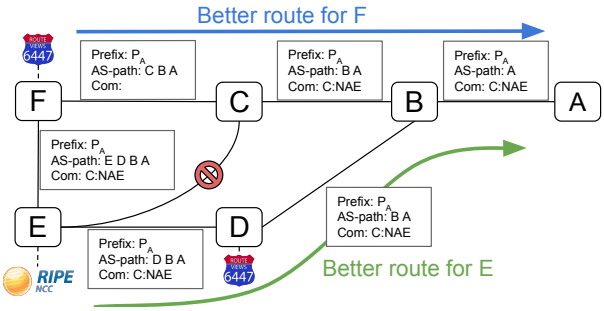

Figure 2. Example illustrating how an action community is more likely to appear in routes that do not include its controlling AS. The community C:NAE instructs AS C not to advertise routes to AS E. We can observe the community C:NAE on routes without AS C exported by ASes D, E, and F.

and we solve it using a greedy heuristic, which provides a tight approximation of the optimal solution [da Silva Jr 2024].

For the experiments and numerical results of our algorithms for inferring location communities, please see Chapter 4 of our thesis [da Silva Jr 2024].

4. Identifying Action Communities and AS Squatters

A basic premise of our work is that an action community should only appear if the route does not traverse its controlling AS (see Figure 2 for an example). In this section, we describe practical uses of BGP communities that violate this premise (§4.1). We then describe how we identify communities that rarely appear with their controlling ASes as action communities (§4.2) and how we use them to uncover other action communities that do not satisfy our premise (§4.3).

4.1. Identifying BGP Community Squatting

We observe that ASes may use BGP information communities defined by or belonging to other ASes. As an AS X is not supposed to tag routes with AS Y's information communities, we refer to this type of use as *squatting*. A common case is ASes using communities defined by one of their *siblings*, *i.e.,* another ASN under the control of the same organization [Gao 2001, Chen et al. 2023]. This behavior is particularly common after network mergers and could result from the homogenization of routing policies defined using BGP communities across the merged ASes (See [da Silva Jr et al. 2025] for examples).

We propose an algorithm to detect ASes that squat another AS's communities by identifying an AS X that systematically tags routes with *information* communities where the first 16 bits belong to another AS Y. The challenge is distinguishing between AS X improperly using AS Y's information communities and legitimate use of AS Y's action communities. We address this challenge by assuming that action communities are applied selectively for short-term traffic engineering while information communities are consistently added when announcements traverse a router. Thus, we identify an AS X that consistently appears with AS Y's communities as a potential squatter. The algorithm uses configurable parameters to select routes and relationships that satisfy the squatting criteria [da Silva Jr et al. 2025, da Silva Jr 2024].

4.2. Inferring BGP Action Communities

Our inference algorithm centers on checking how often a community is tagged on a route that does not traverse the controlling AS or any of its squatters, from now on collectively

referred to as *controlling ASes*. We design and evaluate different approaches to account for the lack of visibility and noise in observed community usage.

Handling squatting ASes To prevent misidentifying squatted communities as action communities, we first identify squatting ASes, and then we rewrite ASNs in the AS path and communities based on the squatting relationships, replacing each ASN with the smallest ASN in its squatting set. This ensures that if a route traverses a squatting AS X and carries a community from a squatted AS Y, both ASNs are consistently rewritten, effectively preventing squatted communities from being classified as action communities.

Filtering Low-Visibility Communities We exclude communities with limited visibility in public BGP dumps from our inferences. A community c must appear in at least two collector peers, with each peer observing it in at least four routes. These empirically chosen thresholds effectively filter out rarely seen communities without significantly impacting inference accuracy [da Silva Jr et al. 2025]. This filter removed 11,836 communities, accounting for less than 11% of those in BGP dumps. If these communities gain broader usage and visibility, our algorithm could classify them.

Inferring Action Communities Our algorithm analyzes each community independently by computing the fraction of routes tagged with a community from AS Y that do not traverse AS Y. It counts the total occurrences of each community c and how often those routes bypass c's controlling ASes. This method accounts for errors or cases where an action community remains tagged despite passing through its controlling AS (*e.g.,* if a customer fails to set it). Based on these counts, we classify communities as action communities when they are largely absent from routes traversing their controlling ASes.

4.3. Uncovering Missing Action Communities

Our algorithm requires a minimum number of observed announcements to classify a community as an action community confidently. However, limited route collector coverage and community filtering by some ASes reduce visibility. To address this problem, we use high-confidence inferences to construct a prefix tree based on community digits, enabling more accurate classification of low-visibility communities. ASes typically assign communities sequentially within a type, reserving contiguous blocks that share a common prefix. These prefixes vary in length depending on numbering schemes and may use fixed- or variable-size blocks per type. We validate this pattern using a manually built ground-truth dataset. Figure 3 illustrates a prefix tree for AS3257, where leaves denote community types: A for action and I for information. For instance, 3257:02XXX and 3257:1XXXX are action communities, while 3257:08XXX and 3257:3XXXX are informational.

For the experiments and numerical results of our algorithms for identifying action communities and community squatters, see Chapter 5 of our thesis [da Silva Jr 2024].

5. Related Work

BGP Communities. The use of BGP communities has expanded significantly [Streibelt et al. 2018], enabling tasks like network optimization, DDoS mitigation, and failure detection [Giotsas et al. 2017, Feldmann et al. 2004]. Yet, the lack of standardization and documentation hampers efforts to analyze routing dynamics and deploy advanced policies. Recent work has explored inferring community semantics [Giotsas et al. 2017, Krenc et al. 2023]. Giotsas *et al.* [Giotsas et al. 2017] apply NLP to extract meanings from IRRs,

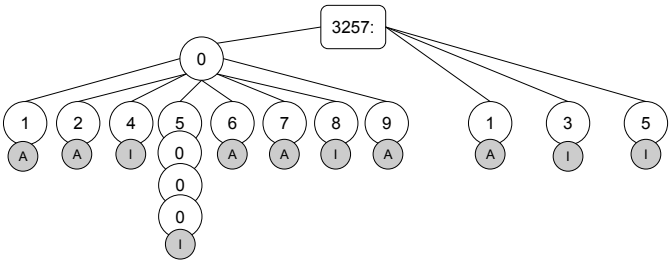

Figure 3. A prefix tree for the documented BGP communities from AS 3257. The branch 05000 is unusually long because it contains only one community, with no other communities sharing the 05* prefix.

websites, and AS documents. Krenc *et al.* [Krenc et al. 2023] propose a clustering approach to classify informational and action communities, but it depends on a ground-truth dataset to set separation parameters and evaluates on non-disjoint data, risking overfitting. Both rely heavily on AS documentation, which is often incomplete or outdated. In contrast, our approach uses documentation solely for ground-truth validation, allowing for independent inference of BGP community semantics.

AS Relationships. Characterizing AS relationships is challenging due to constant Internet changes and the lack of reliable public data, which often omits backup or regional connections not visible in route collectors [Gao 2001]. Additionally, AS relationships can be hybrid [Giotsas et al. 2014], varying by peering location. Identifying these relationships has practical applications, such as detecting route leaks, where a customer AS improperly exports routes from one provider to another, disrupting Internet traffic [Streibelt et al. 2018]. Recent studies infer sibling relationships using data from network operators [Chen et al. 2023, Arturi et al. 2023]. Our approach also leverages route collector data but extends beyond detection to identify ASes squatting on other ASes' communities, which may indicate sibling relationships or other agreements. By relying solely on public data, our method uncovers undocumented relationships that existing approaches, which depend on public documentation, may miss [Arturi et al. 2023].

6. Conclusion

This thesis addresses the challenges of inferring the semantics of BGP communities using only public route announcements. We propose automated classification techniques that work well in the wild for a subset of community types. Our algorithms perform well in identifying location communities, achieving a precision of 93% and a recall of 81% for major Internet providers (Tier-1 and Tier-2 ASes). Our method provides similar accuracy but identifies a far greater number of communities than CAIDA's manually built database. Also, our work automatically identifies action communities and community squatters. Analyzing data from December 2018 to 2023, our algorithm for identifying action communities achieves an average precision of 92.5% and an average recall of 86.5%, demonstrating the robustness of our approach over multiple periods.

Acknowledgments

This work was partially funded by CNPq procs. 420934/2023-5, 308101/2022-7, and 307061/2021-3; FAPESP procs. 2023/00812-7 and 2023/00811-0; FAPEMIG procs. APQ-02793-23 and APQ-02856-18; and CAPES Finance Code 001.

References

[Arturi et al. 2023] Arturi, A. et al. (2023). as2org+: Enriching AS-to-Organization Mappings with PeeringDB. In *Proc. of the PAM Conference*, pages 400–428.

[CAIDA 2021] CAIDA (2021). CAIDA's Geolocation Dataset. `https://www.caida.org/catalog/datasets/bgp-communities/`.

[Chen et al. 2023] Chen, Z. et al. (2023). Improving the Inference of Sibling Autonomous Systems. In *Proc. of the PAM Conference*, pages 345–372.

[da Silva Jr 2024] da Silva Jr, B. A. (2024). *Automatic Inference of BGP Community Semantics*. PhD thesis, Federal University of Mato Grosso do Sul.

[da Silva Jr et al. 2022] da Silva Jr, B. A. et al. (2022). Automatic Inference of BGP Location Communities. In *Proc. of ACM SIGMETRICS / IFIP Performance*.

[da Silva Jr et al. 2025] da Silva Jr, B. A. et al. (2025). Uncovering BGP Action Communities and Community Squatters in the Wild. In *Proc. of ACM SIGMETRICS*.

[Durand et al. 2015] Durand, J., Pepelnjak, I., and Döring, G. (2015). BGP7454: BGP Operations and Security. Technical report, RFC 7454, February.

[Feldmann et al. 2004] Feldmann, A. et al. (2004). Locating Internet Routing Instabilities. In *Proc. of ACM SIGCOMM*, page 205–218.

[Gao 2001] Gao, L. (2001). On Inferring Autonomous System Relationships in the Internet. *IEEE/ACM Transactions on Networking*, 9(6):733–745.

[Gao and Rexford 2001] Gao, L. and Rexford, J. (2001). Stable Internet Routing Without Global Coordination. *IEEE/ACM Transactions on Networking*, 9(6):12.

[Garey and Johson 1979] Garey, M. and Johson, D. (1979). *Computers and Intractability – A Guide to the Theory of NP-Completeness*. W. H. Freeman and Company.

[Giotsas et al. 2014] Giotsas, V. et al. (2014). Inferring Complex AS Relationships. In *Proc. of ACM IMC*, page 23–30.

[Giotsas et al. 2017] Giotsas, V. et al. (2017). Detecting Peering Infrastructure Outages in the Wild. In *Proc. of ACM SIGCOMM*, pages 446–459.

[Krenc et al. 2023] Krenc, T. et al. (2023). Coarse-grained Inference of BGP Community Intent. In *Proc. of ACM IMC*, pages 66–72.

[Meyer 1997] Meyer, D. (1997). University of Oregon Route Views Archive Project. `http://www.routeviews.org/`.

[Rekhter 2006] Rekhter, Y. (2006). RFC 4271: A Border Gateway Protocol 4 (BGP-4).

[RIPE 2024] RIPE (2024). RIS Project. `https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris/`.

[Salamatian et al. 2023] Salamatian, L. et al. (2023). Who Squats IPv4 Addresses? *ACM SIGCOMM CCR*, 53(1):48–72.

[Streibelt et al. 2018] Streibelt, F. et al. (2018). BGP Communities: Even More Worms in the Routing Can. In *Proc. of ACM IMC*, page 279–292.

[Tools 2024] Tools, I. (2024). WHOIS Servers List. `https://www.mobilefish.com/tutorials/whois_servers_list/whois_servers_list.html`.