

Real-Time Anomaly Detection in BGP: Challenges, IPv6 Considerations, and Machine Learning Opportunities

Shadi Motaali, Jorge E. López de Vergara, Luis de Pedro
Escuela Politécnica Superior, Universidad Autónoma de Madrid,
Madrid, Spain

shadi.motaali@estudiante.uam.es, jorge.lopez_vergara@uam.es, luis.depedro@uam.es

Abstract—Border Gateway Protocol (BGP) anomalies, such as route hijacks, misconfigurations, and worm-induced disruptions, significantly threaten global Internet stability. While machine learning (ML) methods have improved anomaly detection, critical challenges persist: limited availability of comprehensive IPv6/IPv4 labeled datasets and significant preprocessing delays that prevent real-time anomaly classification. This research deals with extending established dataset-generation methods to create robust, parallel datasets for IPv4 and IPv6 anomalies. It further evaluates advanced ML models, including LSTM, Transformers, and Graph Neural Networks (GNNs), specifically focusing on reducing detection latency. Our approach aims to integrate optimized preprocessing workflows, diversified datasets, and streaming-based inference. We expect this will improve anomaly detection accuracy and speed, moving closer to practical real-time BGP anomaly detection.

Index Terms—BGP, Anomaly Detection, IPv6, Real-Time Monitoring, Machine Learning, Feature Extraction, Dataset Generation

I. INTRODUCTION

The Border Gateway Protocol (BGP) manages inter-domain routing for thousands of Autonomous Systems (ASes) around the globe, making it essential for the stability of the global Internet. However, BGP is susceptible to issues like hijacks, route leaks, and misconfigurations, which can cause both partial and widespread disruptions [1], [2]. Machine learning (ML) and deep learning (DL) approaches—e.g., Random Forest, LSTM (Long Short-Term Memory), and Transformers—have shown promise for automating BGP anomaly detection [3], [4]. However, major challenges remain: (1) Publicly available BGP repositories (e.g., RouteViews, RIPE NCC) provide both IPv4 and IPv6 data; nevertheless, previous research has predominantly focused on IPv4 during preprocessing, leading to limited characterization of IPv6 anomalies [2]. (2) Existing ML-based detection pipelines often involve extensive preprocessing steps, which can take several minutes before feature extraction and classification, introducing delays that hinder real-time mitigation [4]. (3) The absence of standardized IPv4/IPv6-labeled datasets limits model generalization

and benchmarking [2], [4]. To address these issues, this paper proposes an approach to improve BGP anomaly detection by:

- Developing parallel IPv4 and IPv6 datasets, incorporating diverse BGP anomalies.
- Evaluating low-latency ML architectures (LSTM, Transformers, GNNs and Random Forest) to enhance real-time classification.
- Implementing feature-based anomaly localization, enabling more precise mitigation.

The remainder of this paper is organized as follows. Section II surveys the state-of-the-art in BGP anomaly detection and dataset generation. Section III clarifies existing gaps and their implications for IPv6 research. Section IV presents the principal research questions. Section V outlines our initial methodology, followed by a comprehensive approach. Section VI concludes the paper, highlighting how unified IPv4/IPv6 datasets and near real-time ML frameworks can strengthen BGP resilience, and showing future directions in this research.

II. STATE-OF-THE-ART REVIEW

Recent research from 2017 to 2024 on BGP anomaly detection has employed a range of methodologies, from traditional statistical methods to advanced machine learning and deep learning techniques. Statistical methods typically involve heuristic or threshold-based detection mechanisms and have demonstrated efficacy in identifying large-scale anomalies such as route leaks or hijacking events [1]. However, these approaches have difficulties detecting subtle or small-scale anomalies and often necessitate frequent and manual updates [2].

Machine Learning (ML) techniques such as Random Forests and Support Vector Machines (SVM) have significantly advanced anomaly detection accuracy and adaptability. For instance, in [2] they used feature selection combined with SVM models to achieve high accuracy in distinguishing anomaly types such as route table leaks and link failures. Furthermore, the work in [3] reviewed multiple ML techniques, highlighting the robustness of SVMs in detecting various BGP anomalies, including worms and ransomware attacks.

Deep Learning (DL) methods, notably LSTM networks and autoencoder architectures, have effectively modeled temporal

This Ph.D. work is partially funded by a grant from the Department of Electronics and Communication Technologies at Universidad Autónoma de Madrid, as well as by the Community of Madrid, Spain, under the project RAMONES-CM (TEC-2024/COM-504).

dynamics in BGP anomalies. The work in [5] demonstrated the superiority of an LSTM-based autoencoder in identifying anomalies such as route hijacks and misconfigurations, significantly outperforming traditional ML methods.

Recent advancements have also emphasized Graph Neural Networks (GNNs) and Graph Attention Networks (GATs), utilizing Autonomous System (AS) relationship graphs to enhance anomaly detection and localization precision. The work in [4] showed graph-based methods significantly improved detection accuracy for small-scale events. Similarly, in [6] the use of spatio-temporal graph attention models is proposed, further improving anomaly detection accuracy and providing better anomaly localization.

Despite these advancements, a critical gap remains in the application of ML techniques to IPv6 anomalies. Publicly available repositories, such as RIPE NCC [7] and RouteViews [8], contain both IPv4 and IPv6 routing data. However, research efforts have overwhelmingly focused on IPv4 datasets, leaving IPv6-specific anomaly detection largely unexplored [2]. The structural differences in IPv6 routing, including longer AS paths and different prefix allocation policies, necessitate specialized anomaly detection approaches [4].

Moreover, real-time anomaly detection remains a pressing challenge [9]. While current detection models achieve accuracy improvements, much of the delay originates from data preprocessing steps rather than model inference times. The overall pipeline, including feature extraction and classification, often takes several minutes, which remains inadequate for real-time mitigation [4]. Additionally, the computational complexity of deep learning models makes real-time implementations challenging, indicating a need for optimized and scalable solutions suitable for operational deployment [10].

III. GAPS AND LIMITATIONS

Despite recent advancements, several significant gaps persist in the domain of BGP anomaly detection, as stated above:

Limited IPv6 Research: Existing datasets and anomaly detection models predominantly focus on IPv4. The increasing adoption of IPv6 necessitates research explicitly targeting IPv6 anomalies, which remain underrepresented in current studies and datasets.

Real-time Detection Challenges: The primary bottleneck in real-time detection is not only the model inference itself, but also the preprocessing of raw BGP data, including feature extraction and transformation. This process can take up to seven minutes, significantly delaying anomaly detection. Optimizing preprocessing workflows and leveraging real-time streaming techniques are essential for achieving sub-minute detection latency [4].

Dataset Shortcomings: As highlighted by [2] and [4], there is a notable absence of standardized and adequately labeled datasets covering both IPv4 and IPv6 anomalies. This limitation significantly restricts the ability to develop universally applicable and robust anomaly detection methods.

Anomaly Localization and Identification: Most existing methods detect anomalies without adequately identifying their

precise origin or cause. Enhanced localization and root-cause analysis capabilities are critical for practical anomaly mitigation and response [6].

These gaps are uniquely tied to BGP's inter-domain routing complexity and IPv6-specific behaviors (e.g., Neighbor Discovery Protocol, NDP), which generic anomaly detection methods cannot address due to their lack of tailored features and real-time capabilities. Addressing these gaps is essential for improving the effectiveness and deployment of BGP anomaly detection systems, ensuring better network security and resilience.

IV. RESEARCH QUESTIONS

Considering the gaps identified and the findings of recent research [2]–[4], [6], [10], this study aims to address the following research questions:

- 1) How can BGP anomaly detection methods be extended and optimized to specifically address IPv6 traffic anomalies, considering the absence of standardized IPv6 datasets and benchmarks?
- 2) What machine learning or deep learning methods and architectures (e.g., Random Forest, SVM, LSTM, Transformers, GANs) are best suited to achieve anomaly detection speeds suitable for real-time response, improving current benchmarks?
- 3) What combination of statistical, temporal, and graph-based features extracted from repositories like RouteViews [8], RIPE NCC [7], and BGPStream [11] maximizes the accuracy and enhances the precision of anomaly localization and root cause analysis with faster response time?
- 4) How can standardized datasets with comprehensive labeling for both IPv4 and IPv6 anomalies be developed and validated to enable consistent and rigorous benchmarking across diverse anomaly detection techniques?
- 5) Can hybrid models combining traditional machine learning, deep learning, and graph-based methods effectively overcome current limitations related to dataset heterogeneity, feature complexity, and computational efficiency?

V. METHODOLOGY

This section outlines our integrated approach to address the BGP anomaly detection challenge, drawing on the dataset-generation methods introduced by [12] and [4]. Our goal is twofold: first, to produce parallel IPv4 and IPv6 datasets that accurately capture diverse anomalies, and second, to systematically evaluate machine learning (ML) methods for reduced detection latency. Figure 1 illustrates the structured workflow of our proposed methodology, from data acquisition to performance evaluation. Below, we describe each phase of our methodology in detail.

A. Dataset Replication and Extension

The dataset replication and extension phase establishes the foundational datasets required for our experiments. Initially,

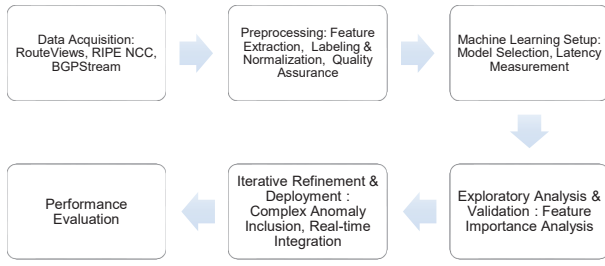


Fig. 1. Flowchart of the proposed methodology for BGP anomaly detection.

we collect relevant historical BGP data, replicate existing feature extraction pipelines, and subsequently extend these processes to include IPv6-specific characteristics. The detailed workflow includes:

- **Data Acquisition:** We collect historical BGP updates from multiple repositories—such as RouteViews [8], RIPE NCC [7], and BGPStream [11]—ensuring coverage of both IPv4 and IPv6 vantage points. Known anomaly events (e.g., route leaks, hijacks) are identified via official incident logs or prior literature.
- **Reproducing Existing Pipelines:** Following the feature-extraction processes by [12] and [4], we parse the raw BGP messages into standardized records that include features like AS path length, announcements/withdrawals, and route flaps.
- **Extending for IPv6:** These pipelines are tailored to handle variations in the generation of the IPv6 dataset, including differences in prefix length distribution, addressing structure, and unique IPv6 routing behaviors such as Multi-Protocol BGP extensions for 128-bit addresses, link-local next-hop resolution via NDP, and enhanced session security through IPsec support. While maintaining a parallel structure, this approach allows for additional data or slightly modified feature extraction steps for IPv6, thereby ensuring accurate representation and sufficient coverage. This careful consideration addresses the dataset gap highlighted in Section III while preserving the capacity to conduct meaningful comparative analyses between IPv4 and IPv6 anomalies.

B. Feature Extraction and Dataset Structuring

To assess whether IPv4 and IPv6 anomalies share comparable characteristics, we generate parallel datasets by applying an identical set of standardized preprocessing steps and extracting the same features from IPv4 and IPv6 data. This approach of creating parallel datasets facilitates a direct and equitable comparison, enabling us to ascertain whether IPv6 anomalies necessitate additional or specialized detection methods. Specifically, the features extracted include:

- **Feature Categories:**
 - *Volume-based metrics:* Announcement/withdrawal counts, flaps, and duplication rates.
 - *AS-path features:* Path length, edit distance, unique ASes, and path changes.

- *Graph-based features:* Node centrality, clustering coefficients, and AS relationship graphs for capturing topological shifts [6].
- *Address structure metrics:* Prefix length, address distribution, and other address-related attributes consistently captured for both IPv4 and IPv6.

- **Labeling and Consistency:** Each dataset entry is labeled based on already-known BGP anomaly events, such as route leaks (e.g., the Turk Telecom leak in 2004) and sub-prefix hijacks (e.g., the Pakistan Telecom YouTube hijack in 2008). We apply uniform time-binning to allow side-by-side comparisons of IPv4 and IPv6 events, ensuring consistent normalization across vantage points.
- **Quality Assurance:** We run exploratory data checks (e.g., outlier detection, missing data analysis) to validate the integrity of both IPv4 and IPv6 data subsets.

In previous works [4], [12], researchers often utilized extensive feature sets, such as 48 statistical or 15 graph-based features, to characterize BGP anomalies. In this study, we aim to investigate whether comparable detection accuracy can be achieved with a reduced feature set, thereby improving computational efficiency and model interpretability while maintaining robustness across IPv4 and IPv6 datasets.

C. Machine Learning Setup for Reduced Delay

Our machine learning setup specifically targets the reduction of detection latency to enable near real-time anomaly classification. To achieve this objective, we have established a structured pipeline covering the entire lifecycle—from selecting appropriate machine learning models and accurately measuring latency to evaluating the trade-offs between detection speed and classification accuracy. The following key phases guide our ML implementation strategy:

- **Model Selection:** We implement a pipeline to evaluate a range of ML/DL techniques—Random Forest, SVM, LSTM, Transformers, and potential GAN-based frameworks—based on their detection accuracy and computational requirements.
- **Latency Measurement:** We measure detection delay as the time elapsed from receiving a BGP update until an anomaly label is assigned. This end-to-end latency is crucial for real-time responsiveness, given that prior works report detection times often exceeding several minutes in practical environments.
- **Trade-off Analysis:** Each model would assess on accuracy, recall, F1-score, and inference speed. We seek approaches that deliver near real-time detection without significantly compromising classification performance. Latency profiling includes per-update processing times under different load scenarios, ensuring we balance accuracy with system throughput.

D. Exploratory Data Analysis and Validation

This phase includes evaluating feature relevance, optimizing model parameters, and verifying the generalizability of

anomaly detection techniques across both IPv4 and IPv6 protocols. Specifically, this process involves:

- **Feature Importance:** Before large-scale experiments, we use metrics like mutual information or random forest ranking to highlight the most discriminative features for IPv4 and IPv6 anomalies. This helps to guide any iterative feature engineering.
- **Hyperparameter Tuning:** We apply grid search or Bayesian optimization to refine model parameters (e.g., learning rates, tree depth) for improved detection efficiency.
- **Comparative Benchmarks:** We run initial validation on a balanced subset of IPv4/IPv6 anomalies to confirm generalization across protocols, ensuring that IPv6-specific differences are not overlooked.

E. Iterative Refinement and Deployment Roadmap

The iterative refinement and deployment roadmap outlines subsequent steps to progressively enhance our BGP anomaly detection solution. This iterative process ensures that the proposed methodology continuously evolves, integrating more complex scenarios, validating performance under realistic network conditions, and eventually transitioning into operational deployment. Specifically, future iterations will focus on:

- **Complex Anomaly Inclusion:** Future dataset iterations incorporate additional event types (e.g., worm-based traffic surges, policy misconfigurations) to enrich the training data for IPv4 and IPv6 contexts.
- **Scalability Tests:** We stress-test each model under high-volume scenarios to measure throughput and stability, examining potential bottlenecks in data ingestion or inference.
- **Real-time System Integration:** Ultimately, we plan to embed the detection pipeline into streaming architectures (Kafka, RabbitMQ) or SDN frameworks [13], facilitating automated anomaly mitigation upon early detection.

F. Performance Evaluation

Finally, we will assess our BGP anomaly detection system, focusing on key performance aspects in diverse environments, including varying network scales and IPv4/IPv6 deployments. This phase will evaluate accuracy (precision, recall, F1-score) versus latency, targeting sub-minute detection. Scalability will be tested using high-volume BGP data from RouteViews and RIPE NCC. We will explore challenges of real-world deployment, such as preprocessing delays and IPv6 routing behaviors (e.g., NDP), alongside practical implementation issues like computational resource demands and integration with streaming platforms (e.g., Kafka). Our ML approach will be compared with heuristic methods using historical events (e.g., the Turk Telecom leak). An implementation roadmap will guide integration into BGP monitoring systems, enhancing mitigation.

Overall, this methodology aims to solve two primary challenges: the lack of standardized IPv6/IPv4 anomaly datasets and the delay hindering real-time detection. By merging robust

dataset-generation techniques with carefully measured ML performance under realistic loads, our approach aspires to deliver a protocol-agnostic (either IPv4 or IPv6), near-real-time solution for BGP anomaly detection.

VI. CONCLUSION AND FURTHER RESEARCH

In this study, critical gaps in BGP anomaly detection were addressed, with a primary focus on dataset generation and real-time performance. By developing parallel and standardized IPv4/IPv6 datasets and optimizing preprocessing workflows, the quality and consistency of available data can be enhanced. Furthermore, through a systematic evaluation of multiple machine learning algorithms, detection latency would be reduced, achieving practical real-time applicability.

Although these efforts improved two significant gaps, challenges such as detailed anomaly localization and comprehensive IPv6-specific anomaly characterization remain open for future research. Continued exploration in these areas will further enhance the robustness and resilience of global Internet routing infrastructures.

By systematically pursuing this study, it is expected that both the depth and breadth of BGP anomaly detection research, particularly for IPv6, will be improved, moving closer to real-time, scalable, and automated protection of inter-domain routing.

REFERENCES

- [1] B. Al-Musawi, P. Branch, and G. Armitage, "BGP anomaly detection techniques: A survey," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 1, pp. 377–396, 2017.
- [2] N. H. Hammood and B. Al-Musawi, "Using BGP features towards identifying type of BGP anomaly," in *Int. Congr. of Advanced Technology and Engineering*, 7 2021.
- [3] A. H. Muosa and A. H. Ali, "Detecting BGP routing anomalies using machine learning: A review," in *Lect. Notes Netw. Syst.*, 2024, pp. 145–164.
- [4] K. Hoarau, P. U. Tournoux, and T. Razafindralambo, "BML: An efficient and versatile tool for BGP dataset collection," in *IEEE Int. Conf. Commun. Wkshps. (ICC Wkshps.)*, 6 2021.
- [5] A. H. Muosa and A. H. Ali, "Internet routing anomaly detection using LSTM based autoencoder," in *Proc. 2nd Int. Conf. Comput. Sci. Softw. Eng. (CSASE)*, 2022, pp. 319–324.
- [6] Z. Liu, H. Qiu, R. Wang, J. Zhu, and Q. Wang, "Detecting BGP anomalies based on spatio-temporal feature representation model for autonomous systems," *IEEE 22nd Int. Conf. Trust, Security and Privacy in Computing and Communications (TrustCom)*, 2023.
- [7] RIPE Network Coordination Centre, "RIPE routing information service (RIS)," <https://www.ripe.net/analyse/internet-measurements/routing-information-service-ris>, accessed: 2024-03-21.
- [8] U. of Oregon, "Routeviews project," <http://www.routeviews.org/>, accessed: 2024-03-21.
- [9] H. K. Takhar, "Machine learning classification of internet worms and ransomware attacks and effect of BGP feature properties," Master of Applied Science Thesis, Simon Fraser University, 2023.
- [10] M. Wichtlhuber, E. Strehle, D. Kopp, L. Prepens, S. Stegmüller, A. Rubina, C. Dietzel, and O. Hohlfeld, "IXP scrubber: Learning from blackholing traffic for ML-driven ddos detection at scale," in *Proc. ACM SIGCOMM Conference*, 8 2022, pp. 707–722.
- [11] OpenINTEL and CAIDA, "BGStream," <https://bgstream.caida.org/>, accessed: 2024-03-21.
- [12] P. Fonseca, E. S. Mota, R. Bennessby, and A. Passito, "BGP dataset generation and feature extraction for anomaly detection," in *IEEE Symp. Comput. Commun. (ISCC)*, 2019.
- [13] C. Yang and W. Jia, "BGP anomaly detection - a path-based approach," in *Proc. 3rd Asia-Pacif. Conf. Commun. Technol. Comput. Sci. (AC-CTCS)*, 2023, pp. 408–414.