SIDROPS Working Group
Internet-Draft
Intended status: Informational
Expires: 3 September 2026

Y. Xu
Y. Chen
K. Xu
Q. Li
J. Wu
Tsinghua University

2 March 2026

Structural Vulnerabilities in ASRank under Adversarial Conditions
draft-xu-sidrops-asrank-vulnerabilities-00

Abstract

This document analyzes the structural vulnerabilities of ASRank, a widely used algorithm for inferring Autonomous System (AS) business relationships from BGP routing data. ASRank plays a key role in security research and BGP operation, yet its inference process is highly sensitive to small changes in input data. This sensitivity introduces risks in adversarial conditions, where inference results may be manipulated without detection. This document outlines the design of ASRank, identifies its structural vulnerabilities, analyzes a minimal manipulation example, and discusses the security implications and potential countermeasures.

Status of This Memo

This Internet-Draft is submitted in full conformance with the provisions of BCP 78 and BCP 79.

Internet-Drafts are working documents of the Internet Engineering Task Force (IETF). Note that other groups may also distribute working documents as Internet-Drafts. The list of current Internet-Drafts is at https://datatracker.ietf.org/drafts/current/.

Internet-Drafts are draft documents valid for a maximum of six months and may be updated, replaced, or obsoleted by other documents at any time. It is inappropriate to use Internet-Drafts as reference material or to cite them other than as "work in progress."

This Internet-Draft will expire on 3 September 2026.

Copyright Notice

Table of Contents

1. Introduction

The Internet comprises over 80,000 ASes interconnected via the Border Gateway Protocol (BGP). These interconnections are governed by business relationships, primarily provider-to-customer (P2C) and peer-to-peer (P2P), which directly influence routing policies and global path selection.

Since AS business relationships are not publicly disclosed, operators and researchers rely on inference algorithms to identify them from observed BGP routing data. Among these, CAIDA's ASRank algorithm [ASRank] is the most widely used methods that underpins CAIDA's public AS relationship dataset [CAIDA_Dataset] and has supported hundreds of scientific studies and operational tools.

Despite its critical role, the security of ASRank under adversarial conditions remains largely unexplored. ASRank's inference process relies on AS triplets, which are sequences of three consecutive Autonomous Systems extracted from BGP paths, and applies a greedy ordering based on transit degree. This design introduces sensitivity to small input perturbations, which can alter inference order and lead to cascading misclassifications.

The objective of this document is to highlight the security risks introduced by structural vulnerabilities in ASRank, which may allow adversaries to manipulate AS relationship inference results at scale. Given ASRank's widespread use in both operational networks and research, such manipulation can undermine the reliability of critical datasets and downstream analyses. This document also discusses the security implications of these vulnerabilities and outlines potential countermeasures.

The rest of this document is structured as follows: Section 2 provides background on AS relationships and inference techniques. Section 3 presents an overview of the ASRank algorithm. Section 4 analyzes structural vulnerabilities in ASRank's inference process. Section 5 introduces a minimal manipulation example to illustrate the impact of input perturbations. Section 6 discusses the security implications of these vulnerabilities. Section 7 outlines potential mitigation strategies. Finally, Section 8 concludes the document.

2. Background and AS Relationship Inference

ASes interconnect under business agreements that define how traffic is exchanged between them. The two primary types of AS relationships are P2C and P2P. In a P2C relationship, the provider offers transit services to the customer in exchange for payment. In a P2P relationship, two ASes exchange traffic without settlement, typically to reach each other's customers.

These relationships play a central role in shaping BGP routing policies, including route selection and export behavior. However, AS relationships are generally considered proprietary and are not publicly disclosed. As a result, researchers and operators rely on inference algorithms to estimate these relationships from observed BGP routing data.

Most inference methods operate on BGP paths collected from public vantage points and apply heuristic or statistical techniques to classify AS links. Among these, CAIDA's ASRank algorithm is the most widely used and influential approaches.

3. ASRank Overview

ASRank is a multi-stage algorithm developed by CAIDA to infer AS relationships from empirical BGP routing data. It is used to generate CAIDA's public AS relationship datasets and has become a foundational tool in BGP researches and operations.

The algorithm begins by sanitizing BGP paths, removing those that contain artifacts such as loops, reserved AS numbers, or Internet Exchange Points (IXPs). It then sorts ASes in descending order of transit degree, computed from counts of AS triplets observed in the data. Based on this ordering, ASRank first identifies a clique of top-level ASes and infers the corresponding P2P relationships within the clique. After filtering poisoned paths further, it infers P2C relationships in a greedy, top-down manner. During this process, newly inferred relationships serve as a basis for subsequent inferences, thereby progressively determining P2C relationships until no additional ones can be resolved. The algorithm then handles special cases, such as ASes with atypical transit patterns. Finally, any remaining unresolved links are inferred as P2P relationships.

While ASRank’s design enables scalable and automated inference, its reliance on triplet-based statistics and order-sensitive heuristics introduces structural vulnerabilities, which are discussed in the following sections.

4. Vulnerabilities in ASRank Inference

A central component of ASRank’s inference process is the P2C classification phase, which accounts for the majority of inferred relationships. In this phase, ASRank processes ASes in descending order of their transit degree, a metric defined as the number of distinct neighbors adjacent to an AS when it appears in the middle position of observed AS triplets (X, Y, Z). For each AS, ASRank examines relevant triplets (X, Y, Z) and applies a heuristic: if the relationship between AS X and AS Y is P2P (denoted as X - Y) or P2C (denoted as X > Y), the algorithm infers a P2C relationship AS Y > AS Z.

A key characteristic of this process is its order-dependence: the relationships inferred earlier influence which triplets can trigger further inferences, shaping the final output. In practice, many ASes have similar transit degrees. As a result, even small changes in the input data can perturb the processing order, leading to different inference sequences and potentially altering the resulting AS relationships.

This sensitivity introduces a structural vulnerability. Minor perturbations in the input data can trigger cascading inference changes, resulting in qualitatively different relationship graphs. In open environments where BGP data is collected from public vantage points, this behavior creates an opportunity for inference results to be influenced or manipulated without directly modifying the algorithm. These characteristics raise concerns about the robustness of ASRank in adversarial or noisy settings, especially given its widespread use in BGP researches and operations.

5. A Minimal Example of Manipulations

Consider the following observed BGP paths, where we highlight only the contiguous AS-path segments that are relevant for inference:

```
path 1:  ... M N A B C D
path 2:  ... M N A B E
path 3:  ... M N A F
path 4:  ... M N A G
```

These paths represent a typical scenario where multiple downstream ASes are accessible through a shared chain of transit providers. ASRank initially infers a top-level clique of major transit ASes. For simplicity, we assume that M and N are part of this clique, with their mutual relationship already inferred as P2P (M-N).

From these paths, ASRank extracts AS triplets and computes the transit degrees (denoted by TD). AS A appears as the middle AS in four triplets and has four distinct neighbors, yielding TD(A)=4. Similarly, TD(B)=3 and TD(C)=2, while D, E, F, and G all have a transit degree of zero. ASRank thus visits ASes in the following order:

```
A -> B -> C -> D -> E -> F -> G
```

Starting with A, ASRank infers N>A from (M,N,A) as M-N is already known. It then moves to B, inferring A>B from the triplet (N,A,B), based on the previously inferred N>A. Following this sequence, ASRank infers the following P2C relationships:

```
{N > A, A > B, B > C, C > D, B > E, A > F, A > G}
```

Now consider adding two additional BGP paths:

```
path 5:  ... X A B D C Y
path 6:  ... X F D G Y
```

Due to the new introduced AS triplets, the transit degree of D increases to TD(D)=4, surpassing that of C, whose transit degree becomes TD(C)=3. As a result, D is visited before C, resulting in the updated inference order:

A -> B -> D -> C -> ...

This reordered traversal leads to a different inference chain. When visiting D, the relationship A>B has already been inferred, and (A,B,D) in path 5 therefore satisfies the heuristic and implies B>D. Later, when visiting C, (B,D,C) in path 5 together with the newly inferred B>D again matches the heuristic and implies D>C. Consequently, the final set of inferred relationships includes:

{N > A, A > B, B > D, D > C, B > C, B > E, A > F, A > G}

Compared to the original outcome, C>D is replaced by D>C, and a new P2C relationship B>D is inferred.

6. Security Implications

Because ASRank-derived datasets are widely used in both academic studies and operational tools, inaccuracies in inferred AS relationships can propagate into downstream systems, undermining the reliability of Internet measurements and topology analysis, and leading to incorrect assumptions about routing behavior and connectivity.

Manipulated or unstable inferences can lead to misleading views of Internet structure and behavior. For instance, incorrect relationships may distort the AS hierarchy or misrepresent network reachability. Such errors can reduce the accuracy of systems that depend on AS relationship data for tasks like anomaly detection, route validation, and security monitoring.

Moreover, inaccurate ASRank inferences can unfairly elevate an AS’s perceived importance, giving it undue visibility, influence, or business advantage. This can also distort peering decisions and mislead customers.

7. Mitigation Considerations

To address ASRank’s vulnerabilities, two areas require attention: improving inference robustness and limiting the injection of forged AS paths.

At the inference level, robustness can be enhanced by using organizational or business-affiliation data to flag relationships that don't align with expected economic logic. Rather than making early, fixed decisions, ASRank could analyze the full relationship graph to reduce the impact of inference order. Since forged paths may coexist with valid ones, detecting internal inconsistencies, like unexpected reversals in relationships, can help identify and verify suspicious inferences.

At the routing layer, deploying path validation mechanisms can reduce the risk of forged triplets. BGPsec [RFC8205] provides cryptographic protection for AS paths but is limited by low adoption and reduced effectiveness under partial deployment. Newer approaches like ASPA [I-D.ietf-sidrops-aspa-verification] and ASRA [I-D.geng-sidrops-asra-profile] offer more practical alternatives by verifying legitimate AS connections, though they are still in early stages and face known limitations.

8. Conclusion

This document analyzes the structural vulnerabilities of ASRank, a widely used algorithm for inferring AS business relationships from BGP routing data. We describe how ASRank's order-sensitive inference process can lead to unstable or incorrect relationship classifications under small input changes. A minimal example demonstrates how minor perturbations in BGP paths can trigger cascading inference shifts. These vulnerabilities raise concerns about the reliability of ASRank-derived datasets, especially in adversarial or noisy conditions. Strengthening inference robustness and deploying path validation mechanisms are essential steps toward improving the security and trustworthiness of AS relationship inference.

9. IANA Considerations

This document includes no request to IANA.

10. Security Considerations

The structural weaknesses in ASRank's inference process may be exploited to manipulate AS relationship outputs without altering the algorithm itself. Such manipulation can distort Internet topology views, mislead routing decisions, and undermine downstream systems that rely on inferred AS relationships. While this document does not describe specific attack strategies, it highlights the risks posed by inference sensitivity and the potential for undetected influence. Mitigation strategies are discussed in Section 7. Operators and researchers are encouraged to interpret ASRank-derived data with

caution and consider complementary validation mechanisms.

11. References

11.1. Normative References

[RFC8205] Lepinski, M., Ed. and K. Sriram, Ed., "BGPsec Protocol Specification", RFC 8205, DOI 10.17487/RFC8205, September 2017, <https://www.rfc-editor.org/info/rfc8205>.

11.2. Informative References

[ASRank] Luckie, M., Huffaker, B., Dhamdhere, A., Giotsas, V., and K. Claffy, "AS relationships, customer cones, and validation", IMC 2013, Proceedings of the 2013 Conference on Internet Measurement Conference, pp. 243256 , 23 October 2013, <https://doi.org/10.1145/2504730.2504747>.

[CAIDA_Dataset] CAIDA, "AS Relationships Dataset", 2025, <https://www.caida.org/catalog/datasets/as-relationships/>.

[I-D.ietf-sidrops-aspa-verification] Azimov, A., Bogomazov, E., Bush, R., Patel, K., Snijders, J., and K. Sriram, "BGP AS_PATH Verification Based on Autonomous System Provider Authorization (ASPA) Objects", Work in Progress, Internet-Draft, draft-ietf-sidrops-aspa-verification-24, 19 October 2025, <https://datatracker.ietf.org/doc/html/draft-ietf-sidrops-aspa-verification-24>.

[I-D.geng-sidrops-asra-profile] Geng, N., Sriram, K., and M. Huang, "A Profile for Autonomous System Relationship Authorization (ASRA)", Work in Progress, Internet-Draft, draft-geng-sidrops-asra-profile-02, 17 October 2025, <https://datatracker.ietf.org/doc/html/draft-geng-sidrops-asra-profile-02>.

Authors' Addresses

Yi Xu
Tsinghua University
30 Shuangqing Road
Beijing
100084
China

Email: y-xu22@mails.tsinghua.edu.cn

Yihao Chen
Tsinghua University
30 Shuangqing Road
Beijing
100084
China
Email: yh-chen21@mails.tsinghua.edu.cn

Ke Xu
Tsinghua University
30 Shuangqing Road
Beijing
100084
China
Email: xuke@tsinghua.edu.cn

Qi Li
Tsinghua University
30 Shuangqing Road
Beijing
100084
China
Email: qli01@tsinghua.edu.cn

Jianping Wu
Tsinghua University
30 Shuangqing Road
Beijing
100084
China
Email: jianping@cernet.edu.cn