Project Report 2

Jasper Eumann

Continuation of Reproducing and Validating of Spoofing Detection at IXPs

Betreuung durch: Prof. Dr. Thomas C. Schmidt
Eingereicht am: 21. November 2019

Fakultät Technik und Informatik
Department Informatik

Faculty of Computer Science and Engineering
Department Computer Science

Contents

1 Introduction

This project report is based on our reproducibility study [7] and extends and emends our previous work [6]. In that we tried to reproduce the results of the paper *Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses* presented at the ACM IMC 2017 [13].

After an exchange with the authors of the paper, we finally managed *(i)* to reproduce the algorithmic part of their methodology with a different team and setup [1]. Unfortunately, *(ii)* we cannot reproduce their results an pure algorithmic grounds, even though we explore various ways of inferring the customer cones.

In particular, spoofed traffic classified in our experiments exceeds the values of IMC'17 by orders of magnitude with a traffic mix that strongly indicates a dominant portion of false positives. We identify plausible reasons for these discrepancies from further analyses. It is worth noting that our insights appear independent of the vantage point and time but highlight intrinsic drawbacks of the previous methodology.

In the context of this work we present our results and some further analysis for spoofed traffic which we used to evaluate our results. As a first attempt to widen the accuracy of the customer cones based on BGP and CAIDA data we implement a cone that takes p2p relationships into account. In addition to our own approach, we analyze our and the IMC'17 results using the spoofer-ix methodology presented by Müller et al. [16].

We give a short summary of the methodology in §2, present some false positive indicators for spoofed traffic in §3, some malformed packet indicators in §4, our results in §5, and end with a short conclusion in §6.

2 Methodology

Each packet passing through an IXP is forwarded by a peering AS. The knowledge of whether one of the peering ASes can legitimately forward a packet from a specific source allows filtering spoofed packets. Looking at the prefixes owned by direct peers is not enough as they often provide transit for others.

2.1 Cone Approaches

The customer cone for a peer represents the relationships for an AS in a tree-like structure that ideally includes all ASes that might send packets via the peer. Building a cone is not straight forward as the necessary information are not easily available. IMC'17 utilize three approaches for their methodology (names taken from the paper):

1. **Naive Approach:** Built from public information, this approach considers a packet valid if it originates from an AS that is part of an announced path for its source prefix. It aims to reflect the topology but falls short of representing business relationships between ASes accurately. Live data provides enough information to deploy it.

2. **CAIDA Customer Cone:** In contrast to the naive cone it focuses on representing the business relationships rather than the topology. Its creation takes information such as community strings, directly reported relations, and historic information into account. More details are available in *AS Relationships, Customer Cones, and Validation* by Luckie et al. [15].

3. **Full Cone:** The idea for this cone is that ASes that appear as neighbors in an announcement are peering. Built from public BGP announcements this approach introduces transitive relationships between peers. Although it might miss some information specific to business relationships it results in the largest cone. The full cone approach is the main method examined by IMC'17 and the basis for most of their evaluation.

Multi-AS Organization Extension: This add-on can be combined with the two cones presented above. It adds information about sibling ASes by building connections between ASes belonging to the same multi-AS organization [2], thus allowing a bidirectional data exchange between them.

Full Cone with P2P Relations: We derived this cone from the full cone approach to take peer-to-peer (p2p) relationships into account. The p2p relationship information are taken from the CAIDA AS relationship dataset [3]. Without the CAIDA AS relationship data our implementation creates the same cone as the full cone implementation from IMC'17.

2.2 Manual Intervention

Lichtblau et al. added missing links to the full cone by hand, which they determined with the help of the whois information. We do not do this because we want to show the properties of the cone approaches without manual intervention. In our opinion only a fully algorithmic method is usable in practice.

2.3 Classification Pipeline

With the help of these approaches packets can be classified either as **invalid** (spoofed) or as **regular**. Before sorting packets into theses two categories the traffic is cleaned by filtering **bogon** packets, i.e., packets with addresses from private networks and other ineligible routable prefixes [18, 5, 19], as well as **unrouted** packets, i.e., packets from sources that do not have any announcements.

2.4 Spoofer-IX Methodology

An improved analysis based on the CAIDA cone is presented by Müller et al. [16]. The methodology filters traffic that flows from provider to customer as unverifiable. The provider forms the gateway for its customers and therefore forwards the traffic from the entire Internet to them. In consequence of this observation, the IMC'17 full cone misclassifies traffic arriving from outside the customer cone of the provider. on the previously proposed methodology.

The idea to exclude transit traffic from classification has already been mentioned in the work of Kováčik et al. [12]. They examined the traffic from the point of view of a single AS and excluded the traffic that was forwarded by this.

2.5 Reproduction Procedure

This study performed reproducibility work in two phases. First we applied scripts [14] kindly provided by the IMC'17 authors to our data sets. These scripts supported (1) constructing the *full cone* from BGP data, and (2) detecting *bogon* and *unrouted* packets. We extended the implementation with missing functionality, including tools to build the cones for the naive approach and CAIDA cone. These steps are partially part of our previous work [6]. We use data from a different Internet eXchange Point (IXP) and from a differ time. This should not point to causal affect the validity of the methodology.

For the extended measurements and further analysis in this work, we re-implemented the cones with enhanced features for classifying payloads of spoofed traffic using libpcap[1]. While carefully confirming consistency with the original scripts, the extended toolset allowed for a more accurate analysis of the classification, as discussed later in more detail.

3 False Positive Indicators

In our results, we obtain large amount of traffic that is possibly misclassified as invalid. We want to rate the quality of the approaches to get some clear characteristics for valid or invalid traffic. To achieve this, we fix some false positive indicators.

3.1 Existing TCP Connections

Established TCP connections are an indicator of false positives because injecting packets into existing TCP connections is not easily deployed on a large scale.

- **TCP packets carrying ACKs**: These packets might be less likely to be spoofed, but could possibly be misused to create a TCP ack storm. However, this is only possible with a man in the middle attack [10] and in this attack scenario the attacker is located between the victim and the measurement point. As a result it is not possible to detect the spoofed packets with this methodology

- **Established end-to-end encryption channel (SSL, SSH, ...)**: The presence of an encrypted channel only strengthens the assumption that a packet is not spoofed. Currently, we count only SSL- and TLS-based protocols during channel establishment and not the following encrypted packets. To do this, we would have to follow the TCP flows and find them in our sampled traffic data.

[1] https://www.tcpdump.org/

3.2 IPsec

IP Security (IPsec) is a suite of protocols that provides security to Internet communications at the IP layer [8] and uses the Authentication Header (AH) or the Encapsulating Security Payload (ESP) to mitigate reply attacks and provide security association. IPsec uses the Internet Key Exchange (IKE) to established a encrypted tunnel. The presence of this tunnel is an indicator of non-spoofed traffic.

However, we do not include IPsec as an indicator because it is not easily to detect in IPv4 traffic.

3.3 Established QUIC Connections

For established QUIC connection the same is true as for existing TCP connections. They are stronger false-positive indicators due to the structure of QUIC since the protocol uses address validation to avoid amplification attacks. Address validation is performed both during connection establishment and during connection migration [11].

We do not include established QUIC connections as an indicator because libpcap[2] cannot handle QUIC at the moment.

3.4 Packets Without Any Notable Negative Effect

Packets that can easily be dropped by the receiver and neither provoke a reply nor require action are not attractive for spoofing.

- **Application layer protocol responses**: These responses are typically not used for spoofing attacks because a response usually does not generate an exploitable action at the receiver
- **ICMP replies**: Echo, timestamp and information replies (the last one is deprecated [9])

[2] https://www.tcpdump.org/

- **Packets to ephemeral ports**: Many TCP or UDP packets to ports within the ephemeral port range are typically not used for spoofing attacks because the attacker wants to use a particular service and ephemeral ports cannot be clearly assigned to a service [4]. If no service is running on an addressed port, the packet is discarded directly

3.5 Summary

These indicators are not complete nor rigorous, but strong. For our work, they serve worth as indicates enough because they indicate that a packet is very likely regular. At our measuring point we get the first 120 bytes of packets, so it is not always possible to detect the application layer protocol. We only count packets with an application layer protocol that could be classified by libpcap.

4 Malformed Packet Indicators

In contrast to the False Positive Indicators the malformed packet indicators identify traffic that would disrupt communication and are more likely to be spoofed than part of regular traffic.

- **TCP or UDP packets that use port zero**: These packets cannot be interpreted correctly by the recipient and will be discarded directly
- **Identical source and destination address**: Packets with these sources addresses should stay in a local network and newer cross an IXP
- **Malformed source address**: For example the last byte of the source address is 0

5 Results

In our analysis, we use sampled traffic from two different time periods: February 2018 (19th-25th) and June 2019 (1st-7th). All results shown in this work are based on the week in Februrary'18, while we used the June'19 data to verify the stability of our results (some full cone results for June 2019 can be found in the appendix). We use BGP data from all available BGPStream [17] collectors for the corresponding weeks as well as one day before and one day after.

5.1 Contribution to Each Traffic Category

Table 1 compares the classification results for the different cones. It breaks down traffic per contributing IXP members, in total bytes, and in number of packets. Note that bogon and unrouted traffic filtering works independent of the cone approach and is thus only listed once per table.

Table 1: Comparison of the different classification results for anomalous traffic

		IMC 2017		Reproduced Results	
		Bytes	Packets	Bytes	Packets
	Bogon	0.003%	0.02%	0.0009%	0.0022%
	Unrouted	0.004%	0.02%	0.00001%	0.0001%
Invalid	Naive	1.1%	1.29%	0.579	1.537%
	CAIDA	0.19%	0.3%	0.955%	1.563%
	Full	0.0099%	0.03%	0.2%	0.488%
	Full with P2P	-	-	0.002%	0.02%
	Spoofer-IX	-	-	0.62%	0.91%

Overall we classify less traffic as invalid than IMC'17. While the contributing fraction of IXP members is similar, the number of bytes and packets differs significantly in some cases. We see two orders of magnitude less unrouted traffic than IMC'17 but classify much more traffic as invalid with the CAIDA and full cone approaches.

The multi-AS extension has a negligible effect on the classification which is only noticeable with the full cone. Plotting the packet count as a time series over the course of the week gives the graph in Figure 1. It shows the packet count adjusted for the sampling rate on a log scale as a function of the time in 12 hour steps. A static scaling factor allows us to plot the result range of IMC'17 in the same graph when available.

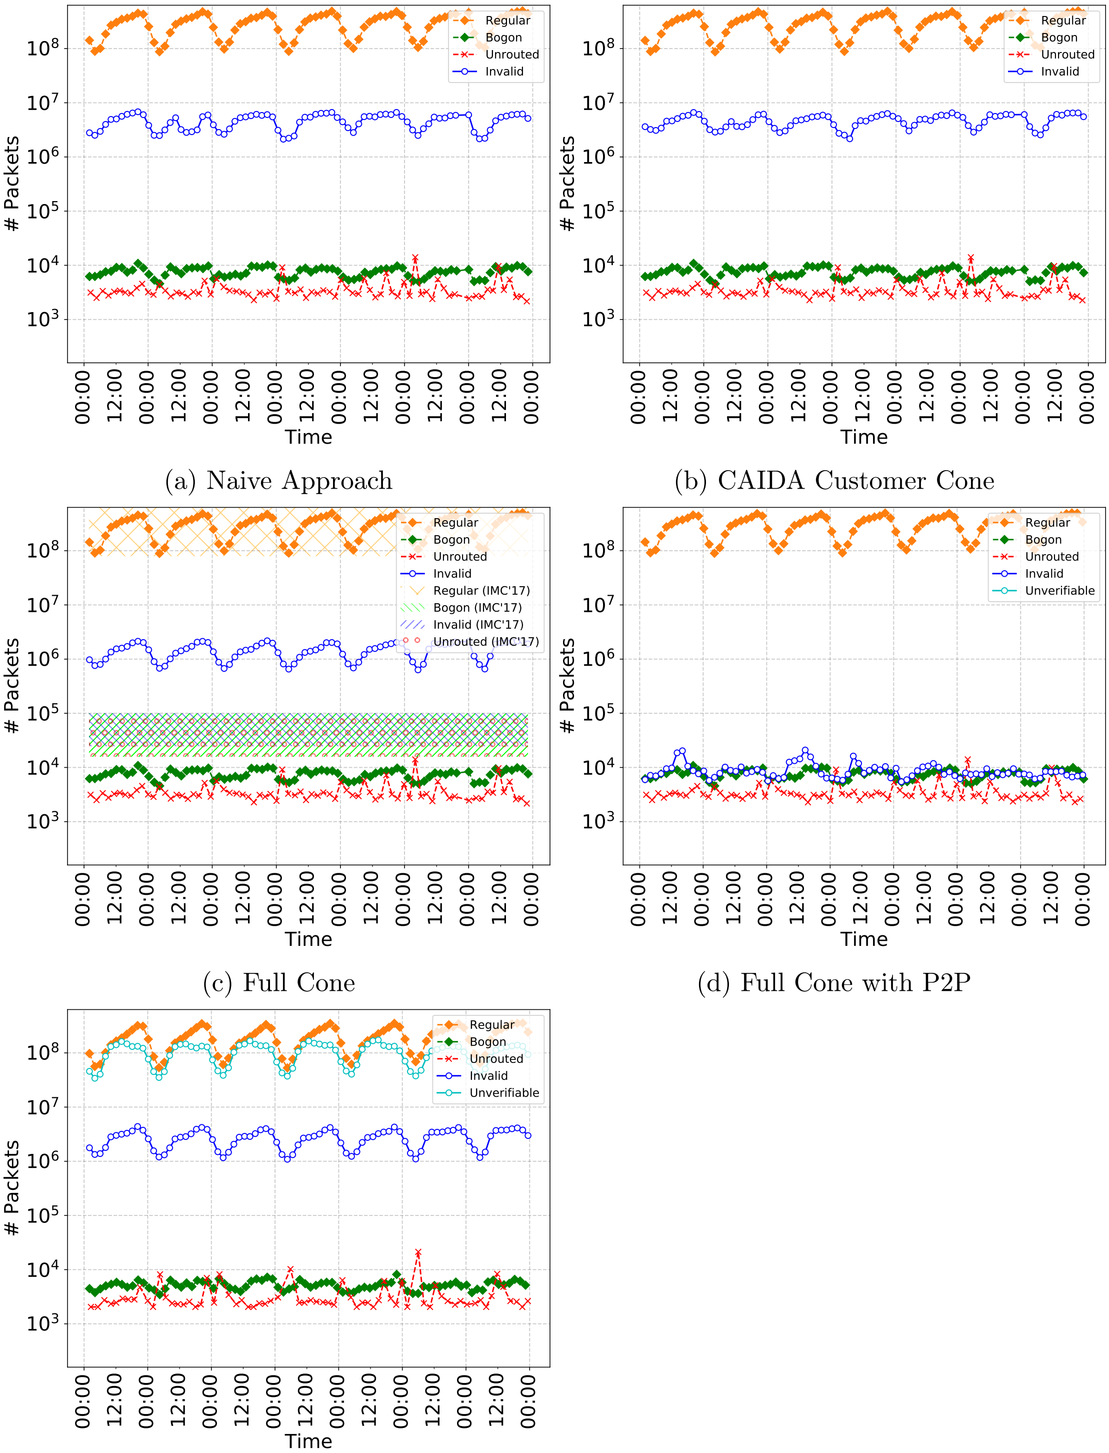

(a) Naive Approach

(b) CAIDA Customer Cone

(c) Full Cone

(d) Full Cone with P2P

(e) Spoofer-IX

Figure 1: Time series for the traffic distribution

All four graphs show the curves at similar values. Exceptions are the full cone based approaches which classified less packets as invalid. Figure 1d has three orders of magnitude less invalid traffic. Over each day the plots for regular and invalid traffic exhibit a repeating pattern where the traffic rises over the course of the day before if drops again at night. While a somewhat similar pattern between regular and unexpected traffic is not unexpected, the curves are similar in a way that suggest an overlap, i.e., regular traffic classified as invalid.

Figure 1e shows our results with without provider-to-customer traffic (Spoofer-IX). We still obtain many false positives that can be explained by the fact that traffic governed by (hidden) private contracts or bilateral agreements crosses the IXP platform while the related control information never enters public announcements. This traffic cannot be handled correctly with these approach.

5.2 Packet Sizes per Category

The next plot examines the observed packet sizes per category. Considering that spoofed packets are often used for amplification attacks, seeing a larger amount of small packets classified as invalid would support the classification results. Figure 2 shows the packet sizes in a cumulative distribution function (CDF) plot with the packet sizes in bytes on the x-axis and the percentage of fitting packets on the y-axis.

All three approaches exhibit a similar distribution of regular packet sizes with most packets larger than 1200 bytes. In contrast, bogon and unrouted traffic is overwhelmingly made up of small packets. Invalid packets tend to be smaller, but differ between approaches. While the naive approach and the full cone show around 70% of small packets, the invalid packets classified by the CAIDA customer cone approach tend to be a bit larger. Although the naive and CAIDA approach classify about the same number of packets as invalid, the specific packet sets diverge at least partially based on their sizes (see Table 1). The full cone approach with p2p relationships classifies many large packs as invalid.

IMC'17 observe a larger fraction of small packets in their invalid traffic than we do, see Figure 2c. In addition, the amount of big invalid packets is nearly negligible. This reflects the overall much smaller amount of bytes displayed in Table 1.

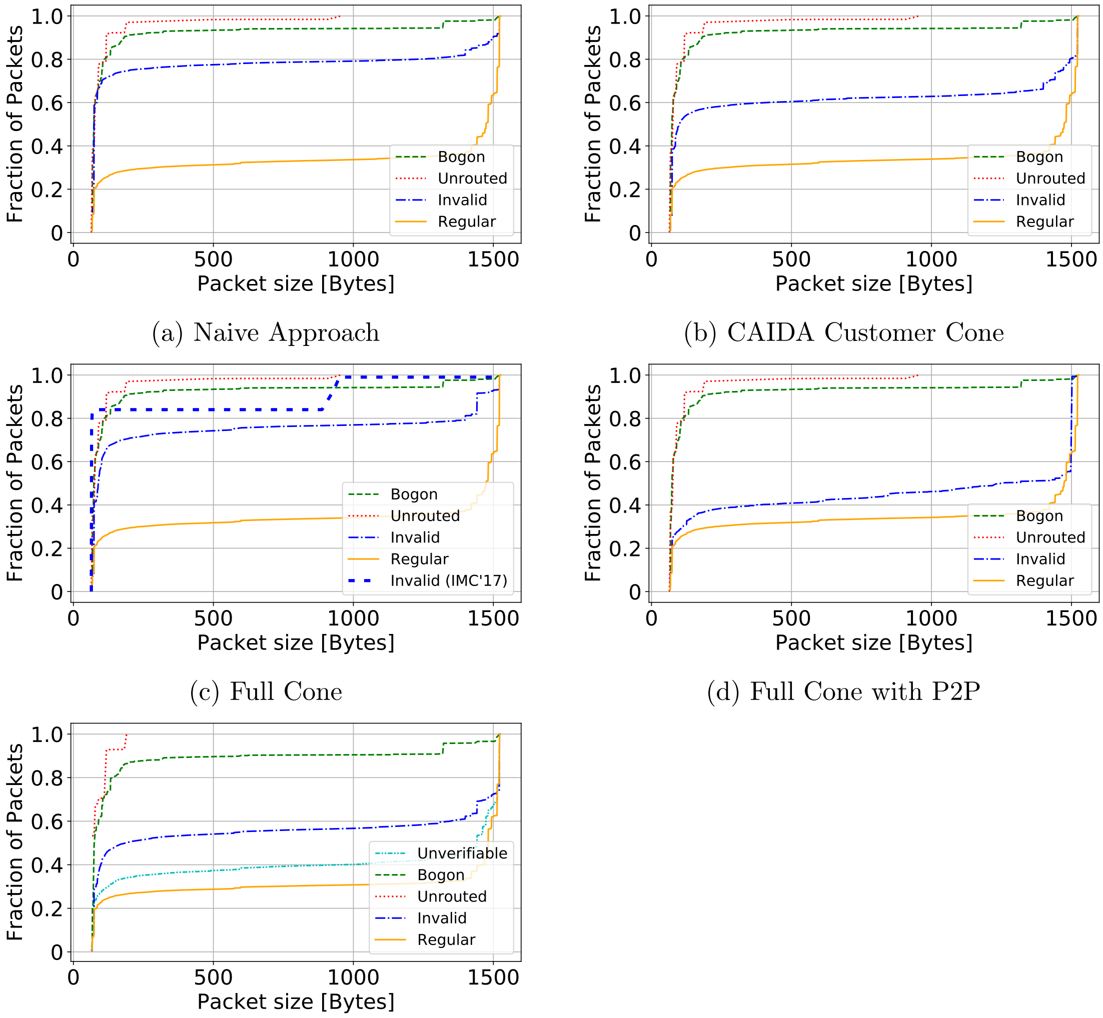

(a) Naive Approach

(b) CAIDA Customer Cone

(c) Full Cone

(d) Full Cone with P2P

(e) Spoofer-IX

Figure 2: CDF: Fraction of packets by size and category

5.3 AS Contribution to Traffic Classes

Figure 6 shows the percentage of IXP members that have a specific percentage of their traffic classified as unrouted, bogon, and invalid, represented as a complementary cumulative distribution function (CCDF).

No member has more than 1% of its traffic classified as either bogon, unrouted, or invalid. The graphs for the naive and CAIDA customer cone approach look similar although the later approach identifies slightly more members with a larger shares. Both approaches find invalid traffic from about 80% of the members.

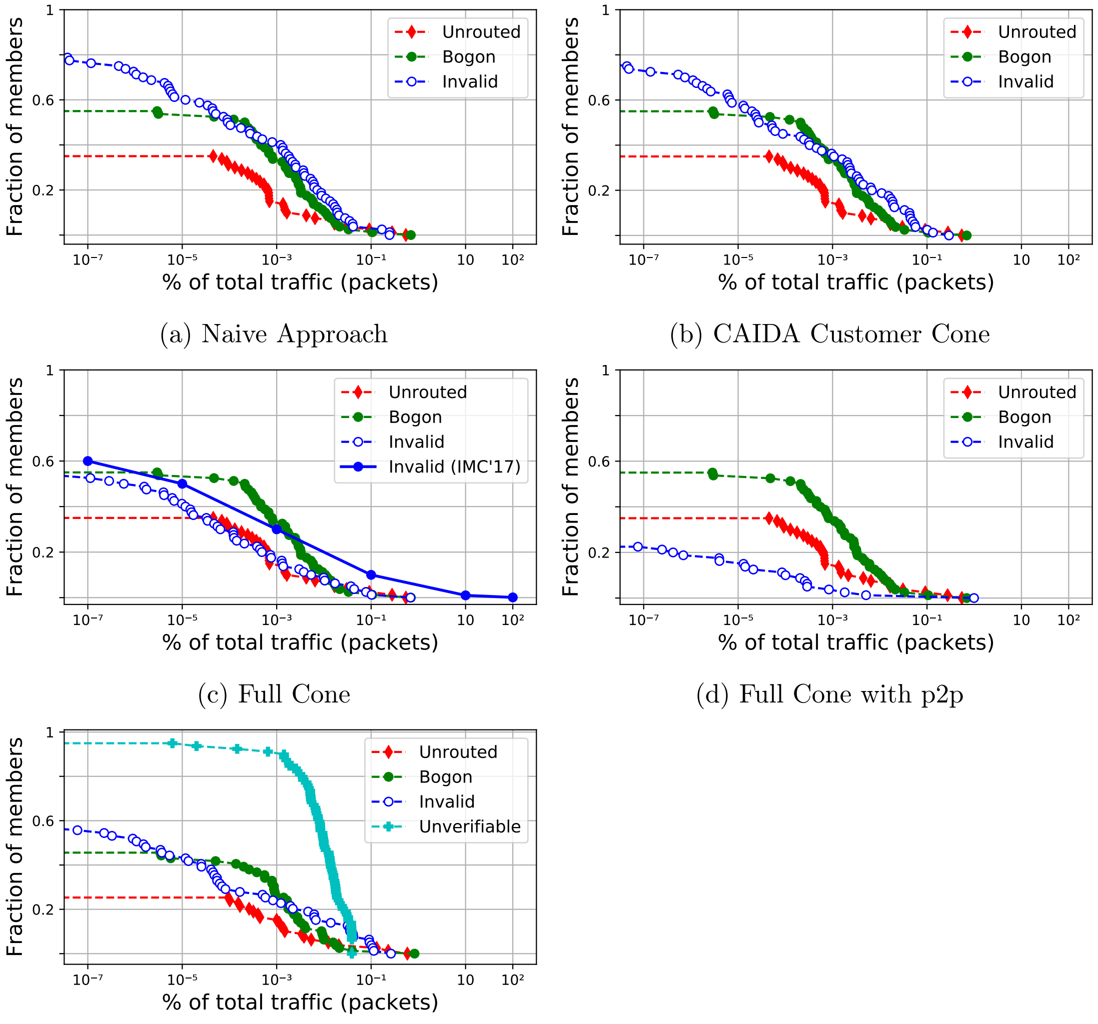

(a) Naive Approach

(b) CAIDA Customer Cone

(c) Full Cone

(d) Full Cone with p2p

(e) Spoofer-IX

Figure 3: CCDF: Fraction of bogon, unrouted and invalid of total traffic per IXP member AS

The full cone approach identifies about 60% of IXP members as contributing invalid traffic, similar to the measurement of IMC'17. In contrast to their measurement, we observe smaller shares of invalid traffic from members. With the full cone with p2p relations only 20% of all ASes produce invalid traffic (Graph 3d) . In contrast with the normal full cone, invalid traffic is visible in almost 60% of all ASes. (Figure 3c).

5.4 Port Distribution

We are now diving deeper into packet inspection of the traffic classified as invalid, and want to understand its characteristics. Table 2 explores the traffic mix with the full cone approach and lists the top destination port distributions of invalid UDP and TCP packets. We cannot equivalently compare to the IMC'17 results, as their traffic mix has not been fully disclosed.

Table 2: Traffic mix per protocol and destination port of invalid packets from the reproduced full cone

ICMP								total
								0.37 %
UDP	53	123	161	443	19302	ephem.	other	total
	1.18 %	< 0.1 %	0.35 %	19.73 %	0.18 %	0.94 %	0.81 %	20.36 %
TCP	80	443	27015	10100		ephem.	other	total
	3.50 %	62.29 %	0.00 %	0.00 %	–	6.75 %	13.67 %	79.45 %

Table 3 and Table 4 list the top port distribution of invalid UDP and TCP packets, separated by destination (DST) port. All values are rounded to one decimal place. Numbers smaller than 0.1, but greater than zero are represented as “<0.1”. The different classification approaches classify different sets of packets as invalid as can be seen by the port. As an example the Full and Full with multi-AS extensions approach label a larger proportion of UDP packets directed at 443 (Quic) and less traffic directed at 3075 (XBOX) as invalid.

We list the ephemeral port range for our data as an additional indicator to the classification performance. The reasoning is that spoofed packets are more likely targeted at well-known ports. Based on this indicator, the full cone shows better performance than the naive and CAIDA approach. The full cone with p2p relations classify lot of traffic as invalid that goes into the ephemeral port range. It produces the largest cone and much of the spoofed traffic is likely to be classified as false negative.

Table 3: Our top port UDP DST distribution of invalid packets

Naive	443	53	4500	3074	16759	1701	ephem.	other
	12.140%	4.040%	1.800%	1.218%	1.115%	1.011%	34.012%	44.664%
CAIDA	443	53	3074	1193	27005	2001	ephem.	other
	30.921%	3.637%	1.296%	0.951%	0.792%	0.714%	28.181%	33.507%
CAIDA	443	53	1193	27005	3074	2001	ephem.	other
(multi-AS)	30.941%	3.590%	1.315%	0.955%	0.782%	0.701%	28.195%	33.521%
Full	443	53	16759	161	3074	19302	ephem.	other
	77.174%	5.472%	1.645%	1.406%	1.016%	0.824%	5.129%	8.157%
Full	443	53	161	19302	19305	3074	ephem.	other
(multi-AS)	83.602%	5.805%	1.501%	0.896%	0.590%	0.190%	3.993%	3.422%
Full	23930	9306	53	40037	389	16393	ephem.	other
(with p2p)	8.139%	4.342%	3.057%	2.907%	1.944%	1.780%	32.682%	45.149%
Spoofer-IX	443	53	3074	16759	9307	123	ephem.	other
	48.580%	3.415%	2.003%	1.006%	0.819%	0.758%	21.268%	22.15%

Table 4: Our top port TCP DST distribution of invalid packets

Naive	80	443	25	440	22	1935	ephem.	other
	53.352%	20.883%	0.474%	0.386%	0.152%	0.119%	13.838%	10.796%
CAIDA	80	443	1935	5228	5900	22	ephem.	other
	29.536%	27.984%	0.357%	0.263%	0.225%	0.222%	25.298%	16.114%
CAIDA	80	443	1935	5228	5900	22	ephem.	other
(multi-AS)	29.536%	27.984%	0.357%	0.263%	0.225%	0.222%	28.195%	13.217%
Full	443	80	1935	5228	993	4070	ephem.	other
	78.848%	4.696%	1.107%	0.887%	0.608%	0.509%	8.304%	5.040%
Full	443	80	1935	5228	993	4070	ephem.	other
(multi-AS)	78.405%	4.401%	1.204%	0.886%	0.604%	0.509%	8.473%	5.518%
Full	443	80	3389	22	45148	25	ephem.	other
(with p2p)	38.446%	29.159%	4.080%	0.859%	0.326%	0.308%	11.042%	15.780%
Spoofer-IX	443	80	1935	5228	993	4070	ephem.	other
	44.613%	4.096%	0.622%	0.457%	0.263%	0.758%	33.442%	15.749%

Table 5 gives an overview of our UDP, TCP and ICMP traffic distribution of invalid packets. This distribution shows that our TCP-based indicators have more impact because we see much more invalid TCP than UDP traffic. The full cone with p2p relations classifies the least TCP packets as invalid.

Table 5: UDP,TCP and ICMP traffic distribution of invalid packets

	ICMP	UDP	TCP
Naive	0.118%	10.072%	87.832%
CAIDA	0.202%	17.090%	81.649%
CAIDA(multi-AS)	0.202%	17.088%	81.651%
Full	0.360%	21.498%	77.564%
Full(multi-AS)	0.369%	20.362%	79.718%
Full(with p2p)	0.056%	13.337%	40.635%
Spoofer-IX	0.199%	17.444%	82.017%

5.5 Checking Classification Results

Checking the classification results based on the collected flow samples is not easy. As part of the port distribution we looked at the ephemeral port range as one possible indicator. We further consider a few characteristics that could indicate false positive classification invalid traffic as described in Section 3.

Table 6: False positive indicators in traffic of the reproduced full cone

	SSL + TCP	HTTP resp.	ICMP echo reply	TCP ACK	malformed
Naive	3.985%	0.174%	0.056%	86.188%	0.000%
CAIDA	4.166%	0.134%	0.070%	69.197%	0.000%
CAIDA(multi-AS)	4.166%	0.134%	0.081%	80.148%	0.000%
Full	6.395%	0.117%	0.043%	76.079%	0.001%
Full(multi-AS)	6.512%	0.029%	0.044%	77.350%	0.001%
Full(with p2p)	2.599%	0.137%	0.317%	38.641%	0.000%
Spoofer-IX	5.776%	0.187%	0.052%	68.507%	0.000%

We analyze the invalid traffic for all approaches and count how many of these indicators exist as a fraction of all invalid packages. The results are shown in table 6.

A large amount for the false positive indicators suggests less spoofed traffic. While a high malformed indicator points to a higher amount of irregular traffic. Accordingly, the spoofed traffic classified with the full cone approach with p2p relations contains the least false positive and most malformed indicators and thus probably the most spoofed traffic.

6 Conclusion

In this work, we aimed at reproducing the results of methods for identifying spoofed traffic at IXPs presented at IMC'17 for different but equivalent data sets. Note that we rely on a purely algorithmic approach without manual intervention. Using both, scripts provided by the authors and an extended, independently written tool set, we could reproduce certain general properties of the paper, but failed to reproduce the core outcome of the method without manual intervention. Instead, on the basis of further analyses, we could derive strong indicators that the proposed method cannot be generalized beyond individually handcrafted examinations.

Accordingly, we find that the approach proposed at IMC'17 does not comply with common real-world deployment. Only effective manual intervention in the routing cones makes it possible to minimize false positives.

Our full cone approach with p2p relationships and the spoofer-ix methodology provide the least false positive indicators and the full cone approach with p2p relationships the most malformed indicators but both do not paint a clear picture. In our opinion, it is not possible to use cone-based approaches without further pre-filtering or other classification methods.

7 References

[1] Bajpai, Vaibhav ; Brunstrom, Anna ; Feldmann, Anja ; Kellerer, Wolfgang ; Pras, Aiko ; Schulzrinne, Henning ; Smaragdakis, Georgios ; Wählisch, Matthias ; Wehrle, Klaus: The Dagstuhl Beginners Guide to Reproducibility for

Experimental Networking Research. In: *ACM SIGCOMM Computer Communication Review* 49 (2019), January, Nr. 1, S. 24–30. – Editorial note

[2] Cai, Xue ; Heidemann, John ; Krishnamurthy, Balachander ; Willinger, Walter: Towards an AS-to-Organization Map. In: *Proc. of the 10th ACM IMC.* New York, NY, USA : ACM, 2010, S. 199–205

[3] : *The CAIDA AS Relationships Dataset, 11.09.2019.* http://www.caida.org/data/as-relationships/

[4] Cotton, M. ; Eggert, L. ; Touch, J. ; Westerlund, M. ; Cheshire, S.: Internet Assigned Numbers Authority (IANA) Procedures for the Management of the Service Name and Transport Protocol Port Number Registry / IETF. August 2011 (6335). – RFC

[5] Cotton, M. ; Vegoda, L.: Special Use IPv4 Addresses / IETF. January 2010 (5735). – RFC

[6] Eumann, Jasper: *Spoofing Detection at IXPs:A Reproducibility Study.* Jan 2019. – URL http://inet.haw-hamburg.de/teaching/ws-2018-19/project-class/Prj1_Jasper_Eumann.pdf

[7] Eumann, Jasper ; Hiesgen, Raphael ; Schmidt, Thomas C. ; Wählisch, Matthias: A Reproducibility Study of "IP Spoofing Detection in Inter-Domain Traffic" / Open Archive: arXiv.org. URL https://arxiv.org/abs/1911.05164, October 2019 (arXiv:1911.05164). – Technical Report

[8] Frankel, S. ; Krishnan, S.: IP Security (IPsec) and Internet Key Exchange (IKE) Document Roadmap / IETF. February 2011 (6071). – RFC

[9] Gont, F. ; Pignataro, C.: Formally Deprecating Some ICMPv4 Message Types / IETF. April 2013 (6918). – RFC

[10] Hubballi, N. ; Santini, J.: Detecting TCP ACK storm attack: a state transition modelling approach. In: *IET Networks* 7 (2018), Nr. 6, S. 429–434

[11] Iyengar, Jana ; Thomson, Martin: QUIC: A UDP-Based Multiplexed and Secure Transport / Internet Engineering Task Force. Internet Engineering Task Force, September 2019 (draft-ietf-quic-transport-23). – Internet-Draft. – URL https://datatracker.ietf.org/doc/html/draft-ietf-quic-transport-23. Work in Progress

[12] Kováčik, Michal ; Kajan, Michal ; Žádník, Martin: Detecting IP Spoofing by Modelling History of IP Address Entry Points. In: *7th International Conference on Autonomous Infrastructure, Management, and Security: Emerging Management Mechanisms for the Future Internet.* Berlin, Heidelberg : Springer-Verlag, 2013 (AIMS'13), S. 73–83

[13] Lichtblau, Franziska ; Streibelt, Florian ; Krüger, Thorben ; Richter, Philipp ; Feldmann, Anja: Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses. In: *Proceedings of the 2017 Internet Measurement Conference.* New York, NY, USA : ACM, 2017 (IMC '17), S. 86–99

[14] Lichtblau, Franziska ; Streibelt, Florian ; Krüger, Thorben ; Richter, Philipp ; Feldmann, Anja: *transitive closure cone.* 2018. – URL https://gitlab.inet.tu-berlin.de/thorben/transitive_closure_cone. – Accessed: 2019-08-28

[15] Luckie, Matthew ; Huffaker, Bradley ; Dhamdhere, Amogh ; Giotsas, Vasileios ; claffy, kc: AS Relationships, Customer Cones, and Validation. In: *Conference on Internet Measurement Conference.* New York, NY, USA : ACM, 2013 (IMC'13), S. 243–256

[16] Müller, Lucas ; Luckie, Matthew ; Huffaker, Bradley ; claffy kc ; Barcellos, Marinho: Challenges in Inferring Spoofed Traffic at IXPs. In: *Proc. of ACM CoNEXT.* New York, NY, USA : ACM, 2019, S. 96–109

[17] Orsini, Chiara ; King, Alistair ; Giordano, Danilo ; Giotsas, Vasileios ; Dainotti, Alberto: BGPStream: A Software Framework for Live and Historical BGP Data Analysis. In: *Proc. of the 2016 Internet Measurement Conference.* New York, NY, USA : ACM, 2016 (IMC '16), S. 429–444

[18] Rekhter, Y. ; Moskowitz, B. ; Karrenberg, D. ; Groot, G. J. de ; Lear, E.: Address Allocation for Private Internets / IETF. February 1996 (1918). – RFC

[19] Weil, J. ; Kuarsingh, V. ; Donley, C. ; Liljenstolpe, C. ; Azinger, M.: IANA-Reserved IPv4 Prefix for Shared Address Space / IETF. April 2012 (6598). – RFC

8 Appendix

Comparison of our full cone results for our two time periods

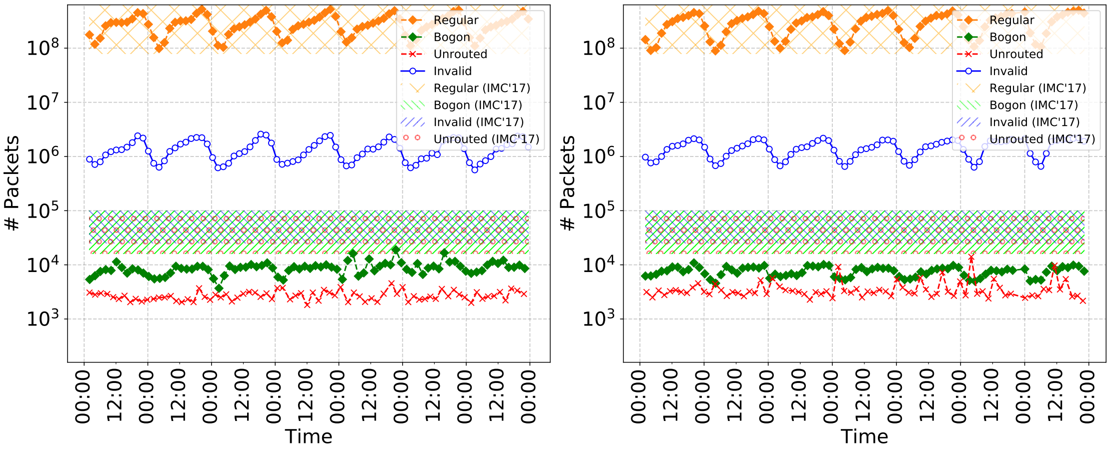

(a) 2019-06

(b) 2018-02

Figure 4: Time series for the traffic distribution

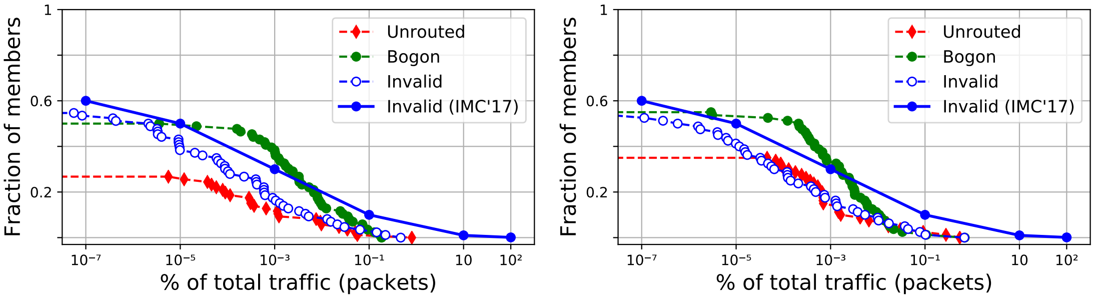

(a) 2019-06

(b) 2018-02

Figure 5: CCDF: Fraction of bogon, unrouted and invalid of total traffic per IXP member AS

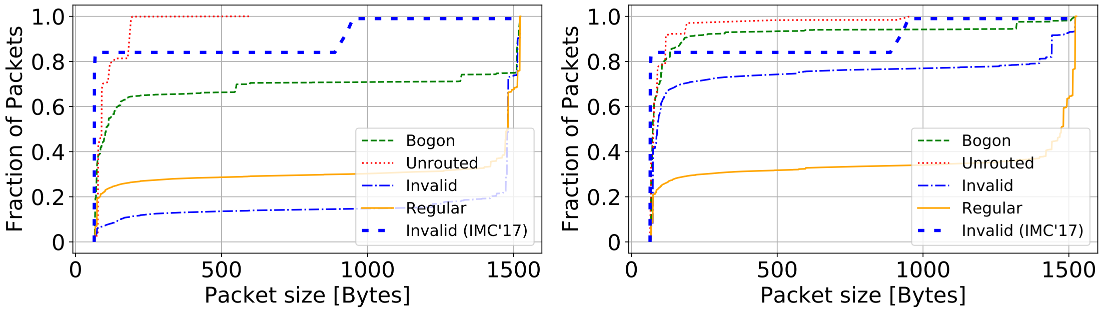

(a) 2019-06

(b) 2018-02

Figure 6: CDF: Fraction of packets by size and category