A Reproducibility Study of "IP Spoofing Detection in Inter-Domain Traffic"

Jasper Eumann, Raphael Hiesgen,
Thomas C. Schmidt, Matthias Wählisch

t.schmidt@haw-hamburg.de

Spoofing Detection in Interdomain Traffic

Starting Point:

- Lichtblau, Streibelt, Krüger, Richter, Feldmann: *Detection, Classification, and Analysis of Inter-Domain Traffic with Spoofed Source IP Addresses,* IMC 2017

Claim:

- Method to passively detect spoofed packets in traffic exchanged between networks in the inter-domain Internet that minimizes false positives

Application domain: IXP

- Measurements and Analyses preformed at a large European IXP

Our objective:

- Build a software infrastructure that can scrub spoofed traffic at IXPs in real-time
- First: Reproduce results with a different team, different setup, data and times

Our approach:

- Iterate methods and (provided) scripts at a large regional IXP
- Extend the analysis with additional BGP data sets and dig into classified traffic

The IMC‘17 Approach

Idea: If a valid packet leaves an AS, it must originate from the routable cone of the emitting AS, i.e., belongs to a prefix reachable through it

Three approaches to identify these cones:

- Naïve:
  A prefix P is in the cone of AS A, iff A appears on a BGP path for P
- CAIDA customer cone:
  All prefixes of customer ASes
- Full cone:
  Extends the naïve cone by assuming transitive relations between all neighboring ASes for all prefixes

Classification

Traffic types

- Regular
- Bogon: Private or multicast source addresses
- Unrouted: Source addresses from unannounced IP space
- Invalid: Classified as spoofed

		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%
Invalid	CAIDA	0.19%	0.3%	0.955%	1.563%
Invalid	Full	0.0099%	0.03%	0.2%	0.488%

Time Series for Classified Traffic

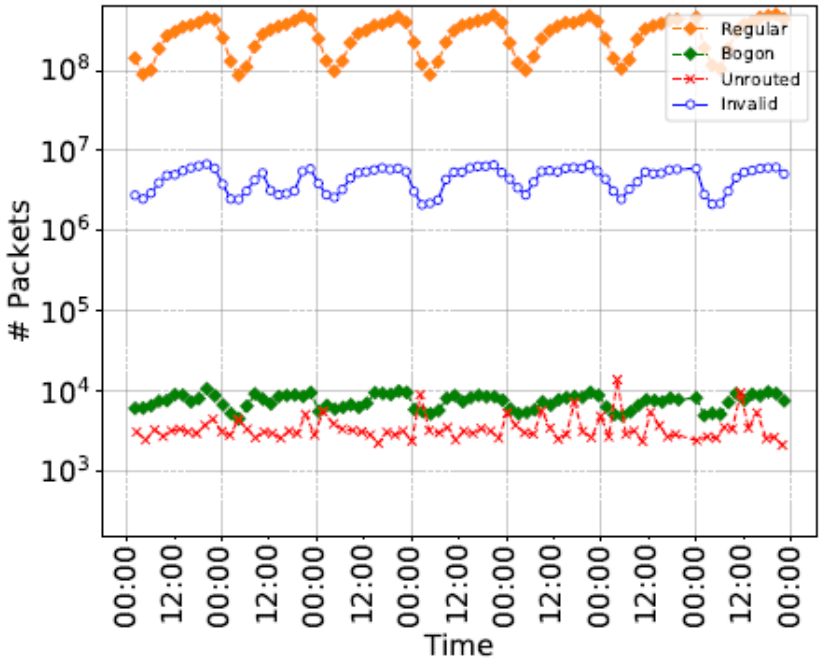

(a) Naive Approach

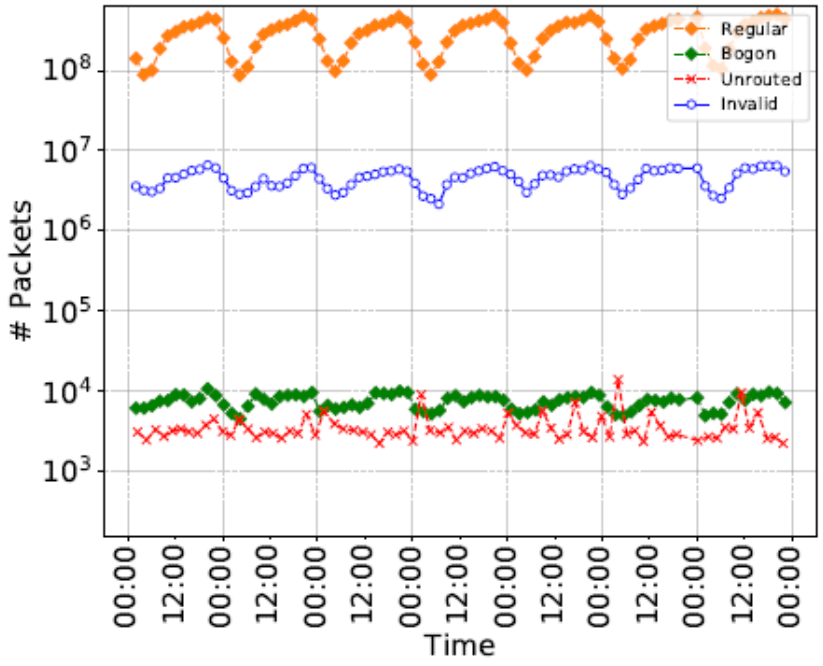

(b) CAIDA Customer Cone

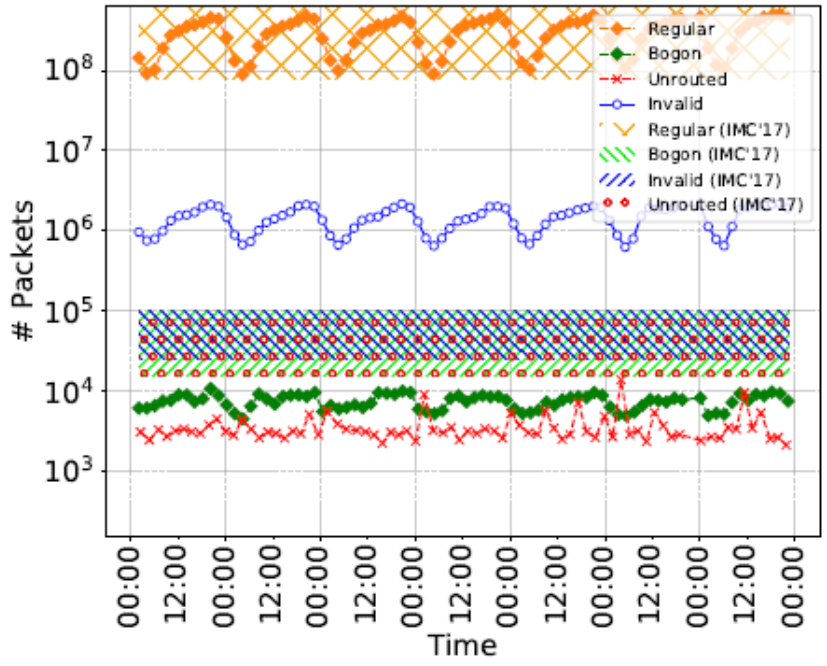

(c) Full Cone

Packet Properties

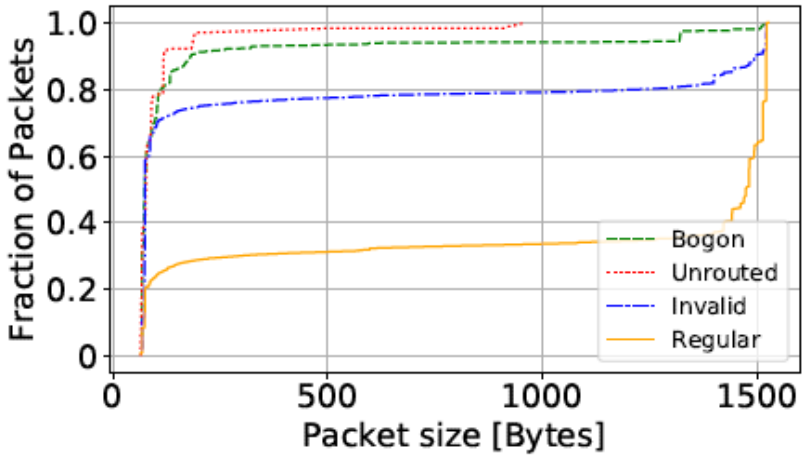

(a) Naive Approach

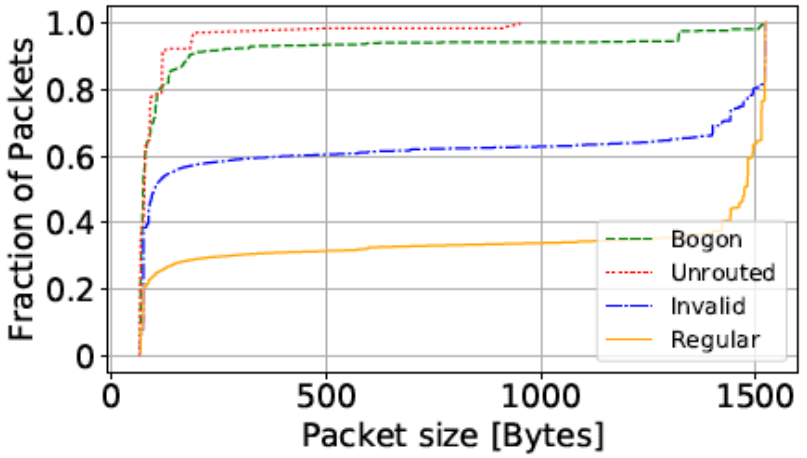

(b) CAIDA Customer Cone

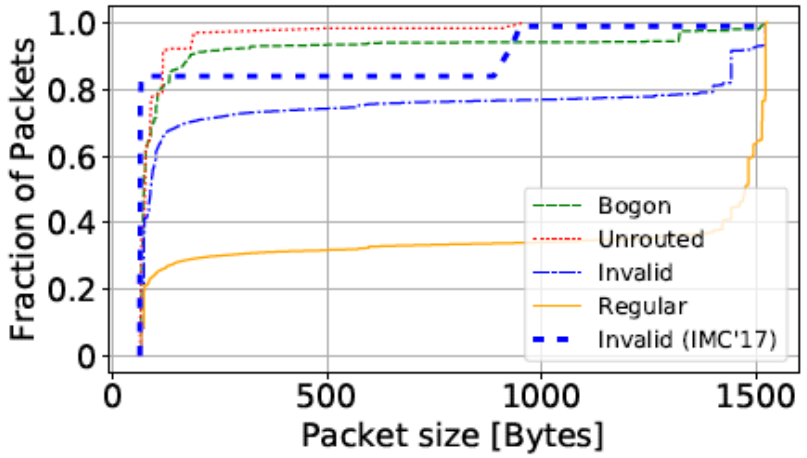

(c) Full Cone

IMC’17 sees 90 % of invalid UDP traffic to port 123 (NTP)

Looking Deeper in our Invalid Traffic

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

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

Table 3: False positive indicators in traffic of the reproduced full cone

	SSL over TCP	HTTP response	ICMP echo reply	TCP ACK	malformed
Naive Approach	3.985%	0.174%	0.056%	86.188%	0.000%
CAIDA Customer Cone	4.166%	0.134%	0.070%	69.197%	0.000%
CAIDA (multi-AS ext.)	4.166%	0.134%	0.081%	80.148%	0.000%
Full Cone	6.395%	0.117%	0.043%	76.079%	0.001%
Full (multi-AS ext.)	6.512%	0.029%	0.044%	77.350%	0.001%

Summary

- Results of IMC’17 could not be reproduced
  - Particular discrepancies for Full Cone approach
- Traffic classified as invalid appears mainly unspoofed
  - Majority of traffic seems HTTP(s) or Quick – not NTP or DNS
  - False positive indicators dominate
- Our impression: determination of cones not accurate enough
  - BGP visibility too low
  - Authors of IMC’17 manually added peerings after traffic inspection
- Approach seems unsuitable for operational deployment