How to Operate a Telescope Without Operating a Telescope

Daniel Wagner

Based on a IMC'23 paper

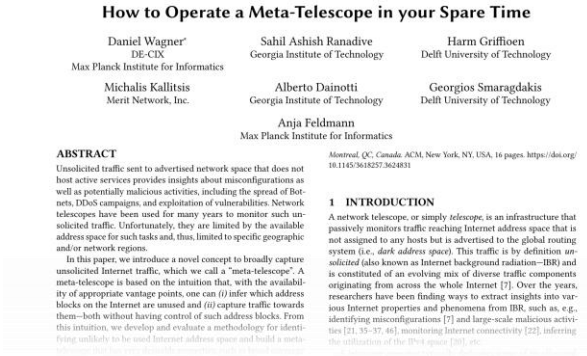

How to Operate a Meta-Telescope in your Spare Time

Daniel Wagner* DE-CIX Max Planck Institute for Informatics

Sahil Ashish Ranadive Georgia Institute of Technology

Harm Griffioen Delft University of Technology

Michalis Kallitsis Merit Network, Inc.

Alberto Dainotti Georgia Institute of Technology

Georgios Smaragdakis Delft University of Technology

Anja Feldmann Max Planck Institute for Informatics

ABSTRACT

Unsolicited traffic sent to advertised network space that does not host active services provides insights about misconfigurations as well as potentially malicious activities, including the spread of Botnets, DDoS campaigns, and exploitation of vulnerabilities. Network telescopes have been used for many years to monitor such unsolicited traffic. Unfortunately, they are limited by the available address space for such tasks and, thus, limited to specific geographic and/or network regions.

In this paper, we introduce a novel concept to broadly capture unsolicited Internet traffic, which we call a "meta-telescope". A meta-telescope is based on the intuition that, with the availability of appropriate vantage points, one can *(i)* infer which address blocks on the Internet are unused and *(ii)* capture traffic towards them—both without having control of such address blocks. From this intuition, we develop and evaluate a methodology for identifying unlikely to be used Internet address space and build a meta-

Montreal, QC, Canada. ACM, New York, NY, USA, 16 pages. https://doi.org/10.1145/3618257.3624831

1 INTRODUCTION

A network telescope, or simply *telescope*, is an infrastructure that passively monitors traffic reaching Internet address space that is not assigned to any hosts but is advertised to the global routing system (i.e., *dark address space*). This traffic is by definition *unsolicited* (also known as Internet background radiation—IBR) and is constituted of an evolving mix of diverse traffic components originating from across the whole Internet [7]. Over the years, researchers have been finding ways to extract insights into various Internet properties and phenomena from IBR, such as, e.g., identifying misconfigurations [7] and large-scale malicious activities [21, 35–37, 46], monitoring Internet connectivity [22], inferring the utilization of the IPv4 space [20], etc.

RIPE87, Rome

2023-12-01

Internet Telescopes

- Announced IP space

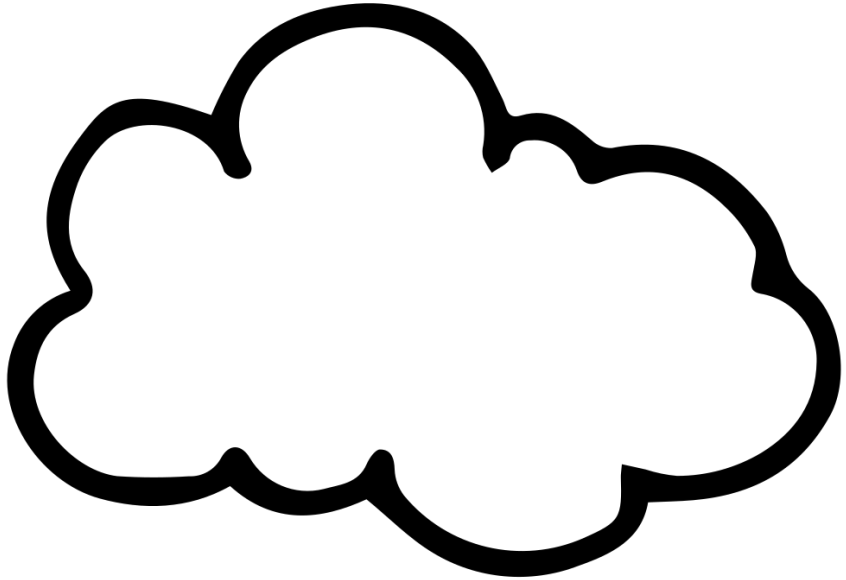

29.172.0.0/16

Internet Telescopes

- Announced IP space
  - Unused
  - Do not expect to see any traffic

29.172.0.0/16

Internet Telescopes

- Announced IP space
  - Unused
  - Do not expect to see any traffic
- Receives traffic, so called Internet Background Radiation (IBR)

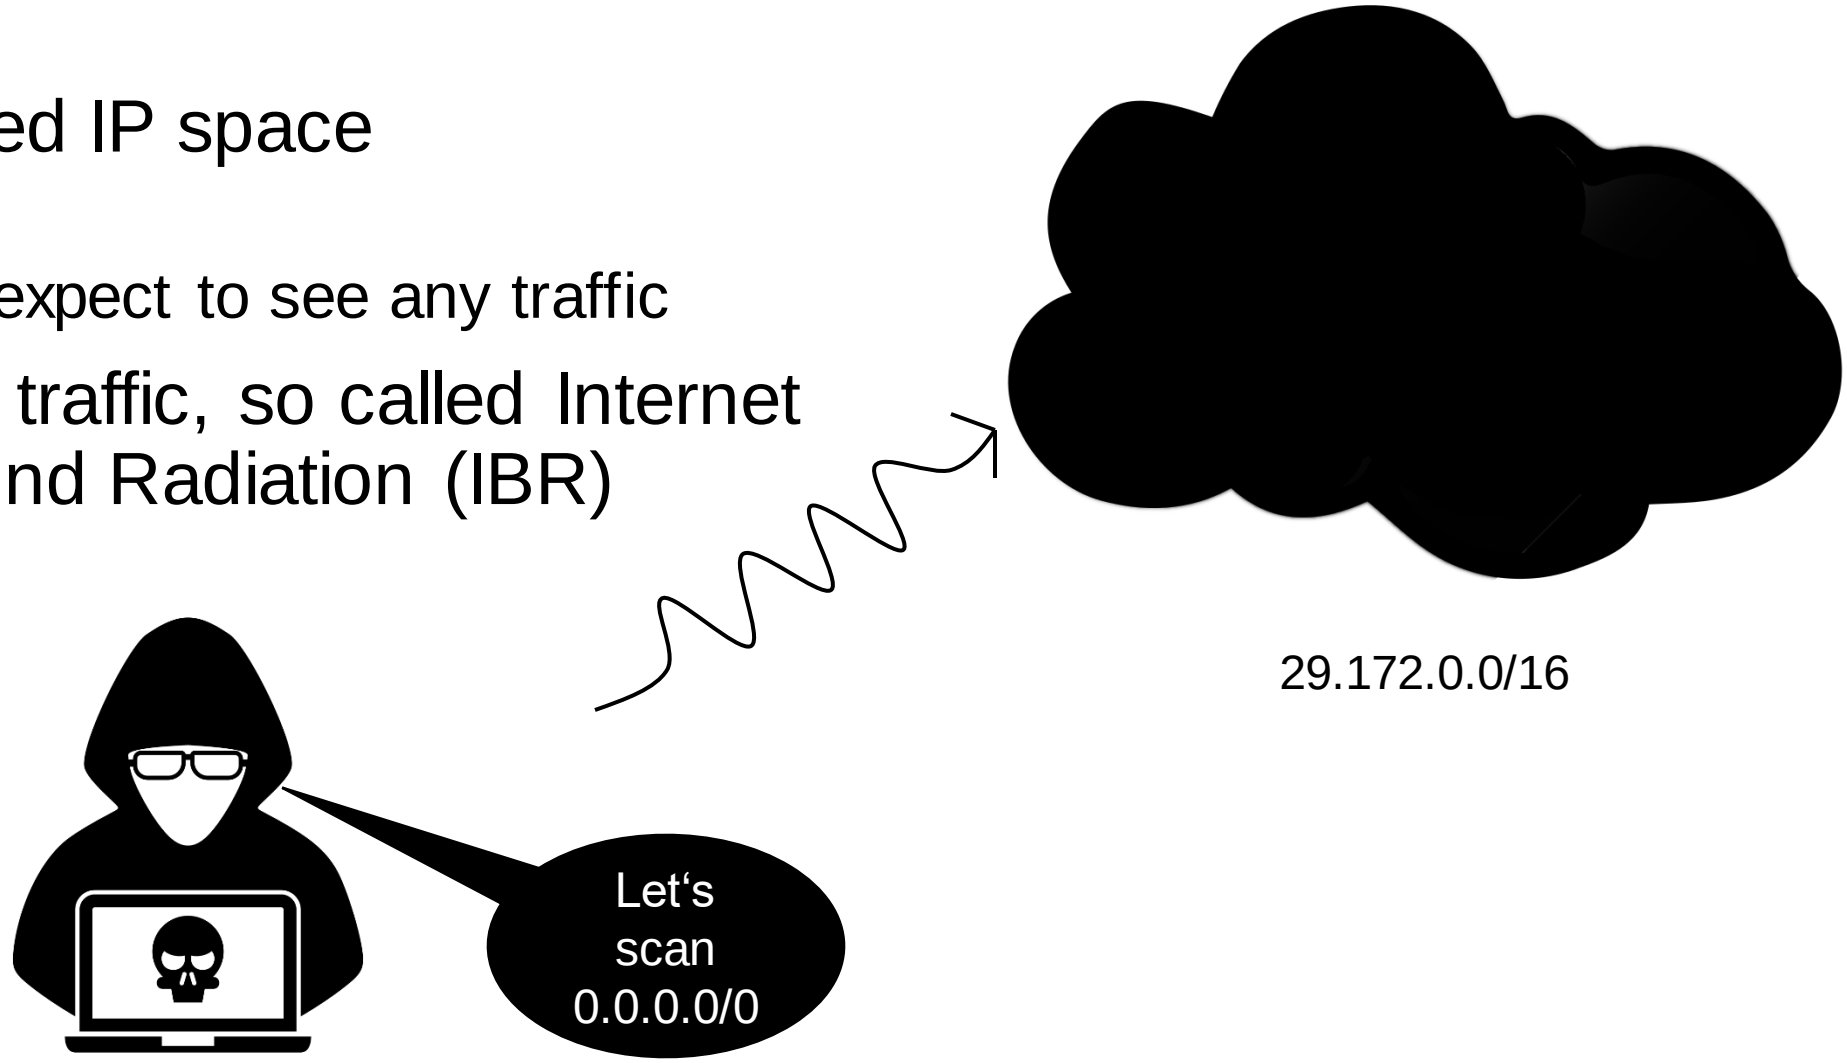

Internet Telescopes

- Announced IP space
  - Unused
  - Do not expect to see any traffic
- Receives traffic, so called Internet Background Radiation (IBR)
- Analogy to real “telescopes“

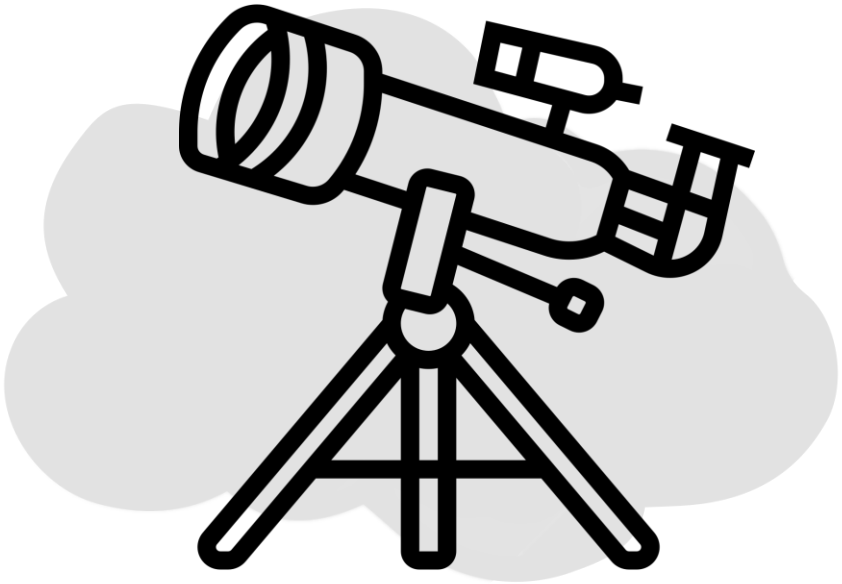

29.172.0.0/16

Internet Telescopes: Security Use Cases

- Obtain insights about recent scans
  - Who is scanning?
  - What ports are being scanned?
  - How many scanners are there?
  - Where are they coming from?
- Attack vector insights and prevention

[1] David Moore et al., IEEE Security and Privacy, "Inside the Slammer Worm", 2003
[2] David Moore et al., IEEE Security and Privacy, "The Spread of the Witty Worm", 2005
[3] David Moore et al., ACM IMC Workshop, "Code Red: A Case Study of the Spread and the Victims of an Internet Worm", 2002
[4] Stuart Staniford et al., ACM WORM, "The Speed of Flash Worms", 2004

Can You Run a Telescope without Owning a Prefix?

- Scans run through the Internet
  - Also through, e.g., IXPs
  - Advantage: visibility not limited to any announcement

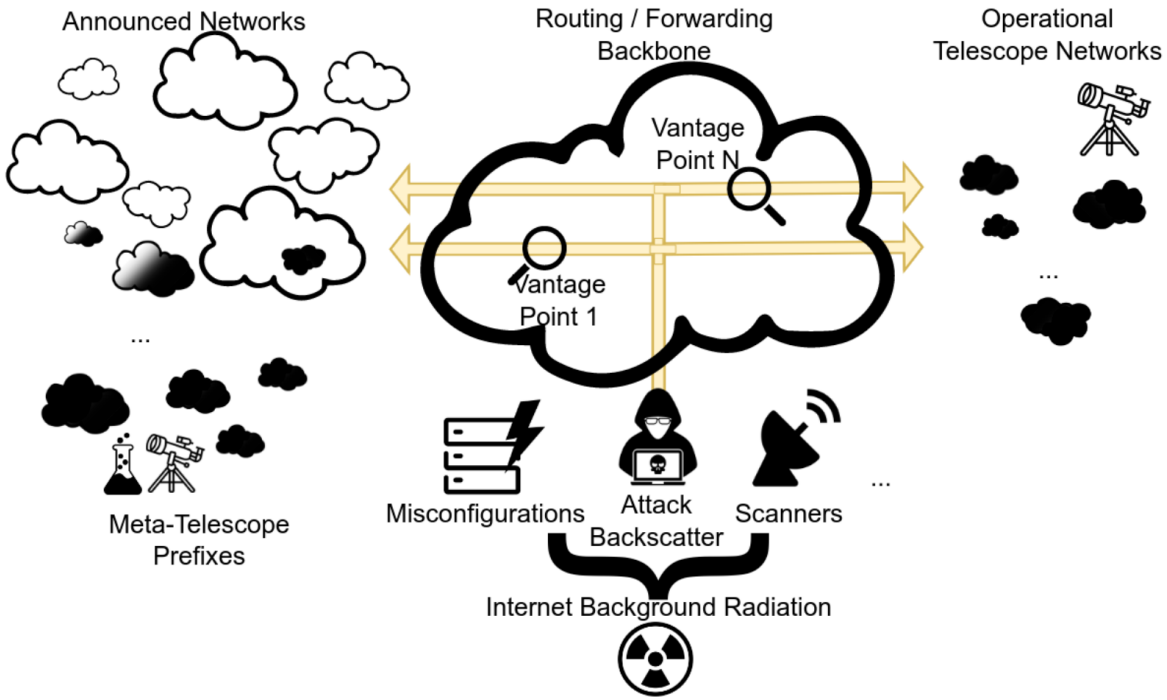

Main idea

Use telescopes to infer characteristics 

Develop methodology to detect scanned unused IP space ("meta-telescopes prefixes") at IXPs 

Overcome limitations of typical telescopes

1. Multiple prefix: evade blocklisting
2. Multiple ASes: evade network type bias
3. Multiple countries: evade locality bias

Telescope Characteristics Inference

- Telescope Data set
  - Full packets (.pcap)
  - 3 telescopes
    - 2x Europe (TEU1, TEU2)
    - 1x USA (TUS1)
  - Various prefix sizes
- Observation period:
  - 2023-04-24 (24 hours)

Telescope Characteristics Inference

- Outbound:
  - Nothing

Telescope Characteristics Inference

- Outbound:
  - Nothing
- Inbound:
  - ~90% TCP SYN
    - 20B IP header + 20B TCP header (+ 8B for one option)
    - Sensitivity analysis on packet size

Telescope Characteristics Inference

- Outbound:
  - Nothing
- Inbound:
  - ~90% TCP SYN
    - 20B IP header + 20B TCP header (+ 8B for one option)
    - Sensitivity analysis on packet size -> Avg. of 44B

Telescope Characteristics Inference

- Outbound:
  - Nothing
- Inbound:
  - ~90% TCP SYN
    - 20B IP header + 20B TCP header (+ 8B for one option)
    - Sensitivity analysis on packet size -> Avg. of 44B
  - Receiving no more than 1.7M packets per day and /24

Unused IP Space Inference

- Data set
  - 14 IXPs
    - Spread across Europe, North America, and Asia
    - Diverse member counts & peak traffic volumes
  - Sampled flow data

- Observation period: 2023-04-24 (24 hours)

Unused IP Space Inference

- Filter 1: TCP*
- Filter 2: Average <=44 Bytes
- Filter 3: No outbound traffic
- Filter 4: Reserved / private space
- Filter 5: Globally routed**
- Filter 6: Packet count < 1.7M

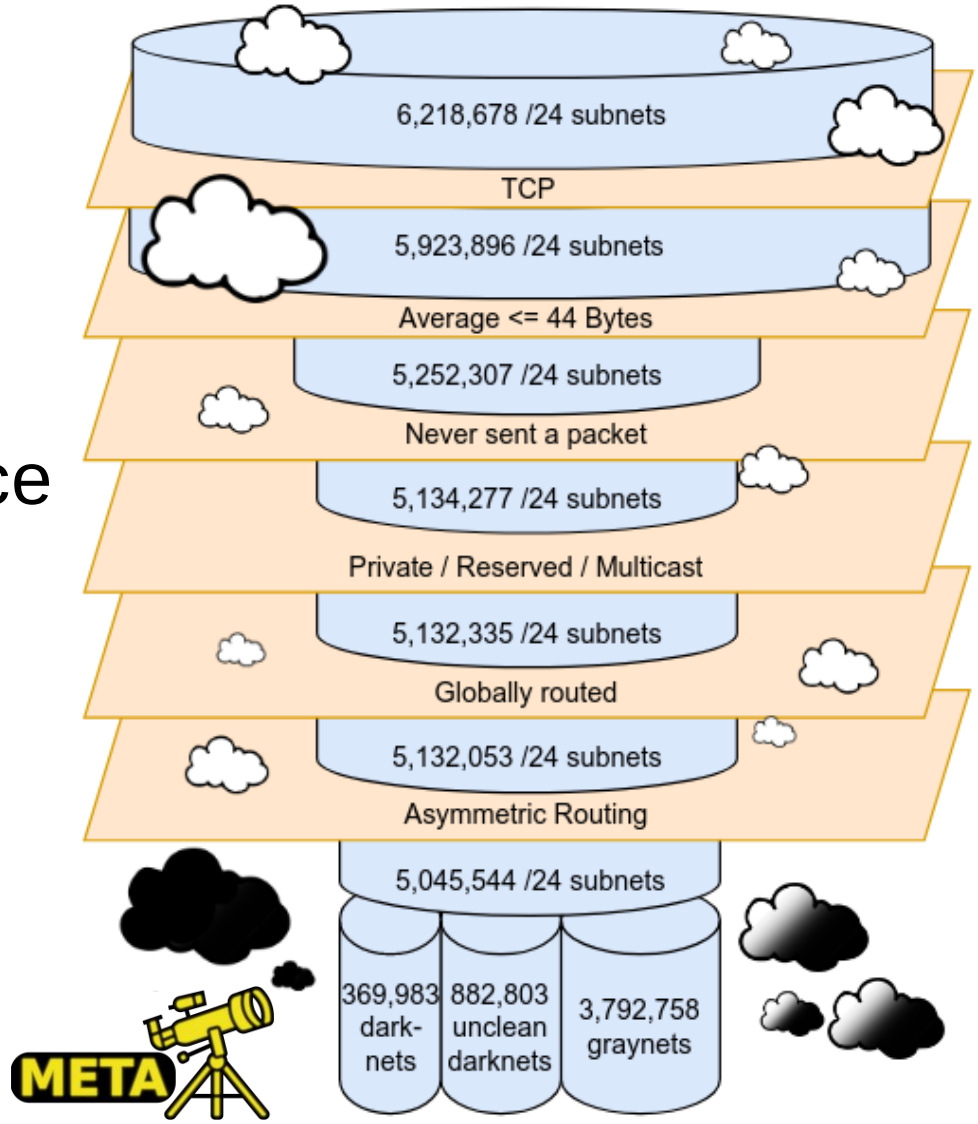

* We can‘t check for TCP flags

** According to Routeviews

Unused IP Space Inference

- Vantage point diversity:

IXP	#Inferred meta-telescope prefixes	#ASes	#Countries
CE1	397,000	8,529	201
CE2	21,340	1,597	124
CE3	61,607	3,982	173
CE4	2,178	455	84
NA1	395,585	8,960	198
NA2	12,489	919	102
NA3	262	128	17
NA4	1,054	299	74
SE1	34,222	2,269	152
	56,638	2,078	132
	3,782	729	97
	43,573	2,431	152
	1,949	667	104
	270	104	33
All	318,646	7,195	194

Combining IXP data increases chance to contain violating packets

Unused IP Space Inference

- Vantage point diversity:

IXP	#Inferred meta-telescope prefixes	#ASes	#Countries
CE1	397,000	8,529	201
CE2	21,340	1,597	124
CE3	61,607	3,982	173
CE4	2,178	455	84
NA1	395,585	8,960	198
NA2	12,489	919	102
NA3	262	128	17
NA4	1,054	299	74
SE1	34,222	2,269	152
SE2	56,638	2,078	132
SE3	3,782	729	97
SE4	43,573	2,431	152
SE5	1,949	667	104
SE6	270	104	33
All	318,646	7,195	194

Largest IXP finds unused space in the most different countries

Unused IP Space Inference

- Vantage point diversity:

IXP	#Inferred meta-telescope prefixes	#ASes	#Countries
CE1	397,000	8,529	
CE2	21,340	1,597	
CE3	61,607	3,982	
CE4	2,178	455	
NA1	395,585	8,960	198
NA2	12,489	919	102
NA3	262	128	17
NA4	1,054	299	74
SE1	34,222	2,269	152
SE2	56,638	2,078	132
SE3	3,782	729	97
SE4	43,573	2,431	152
SE5	1,949	667	104
SE6	270	104	33
All	318,646	7,195	194

2nd largest in the most different ASes

Unused IP Space Inference

- Validation:
  - Inferred IP address space of collaborating telescope

Unused IP Space Inference

- Validation:
  - Inferred IP address space of collaborating telescope

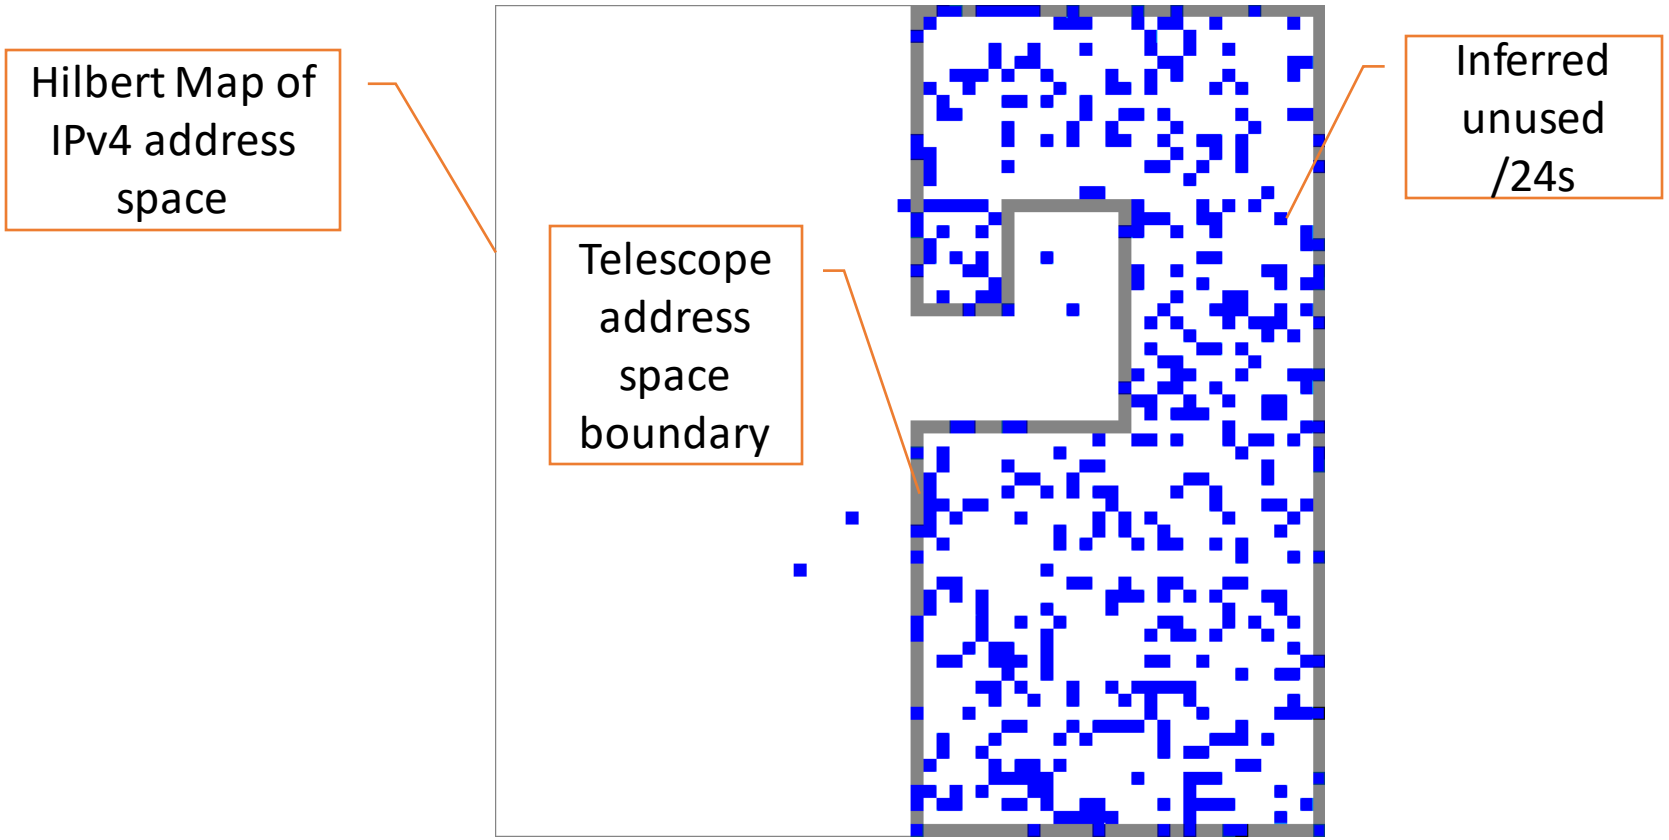

Unused IP Space Inference

- Found known telescopes

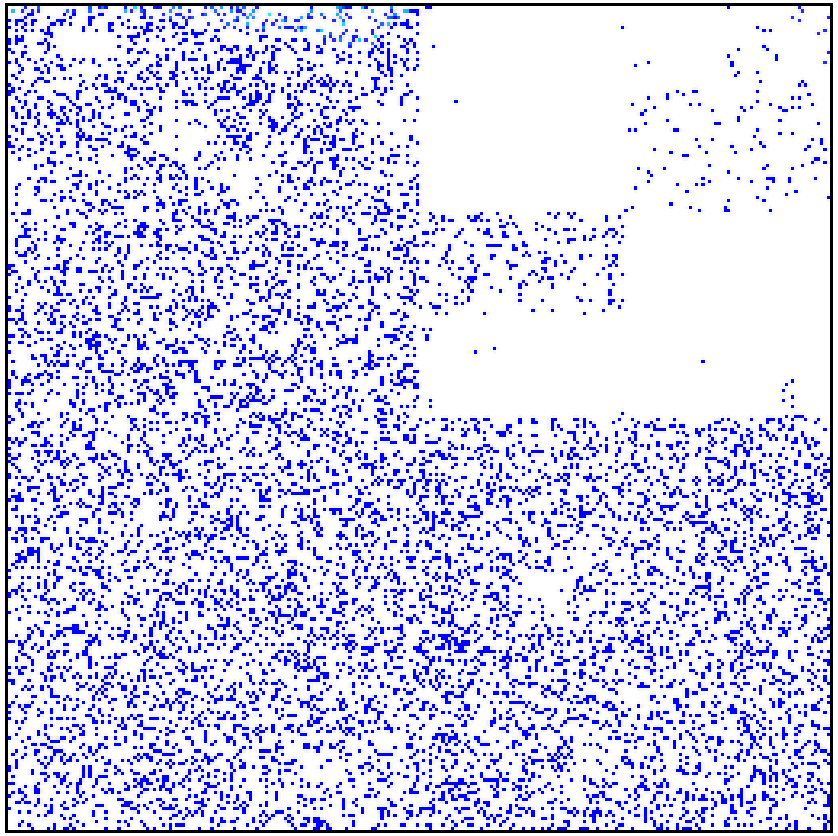

- Found unused space

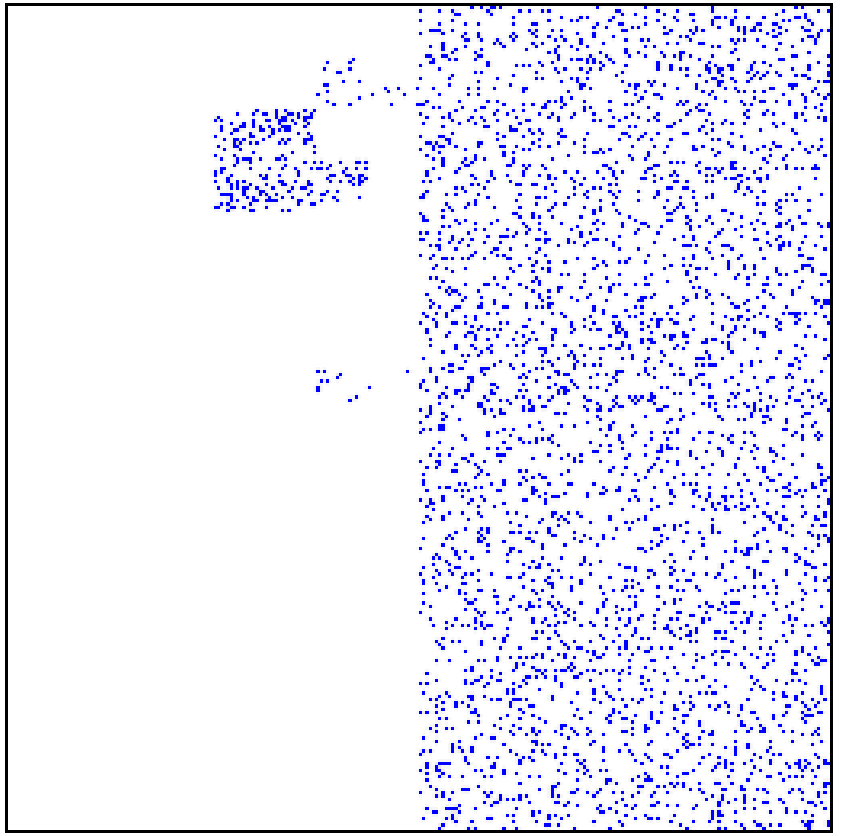

Insights

- Where are the most meta-telescope prefixes?

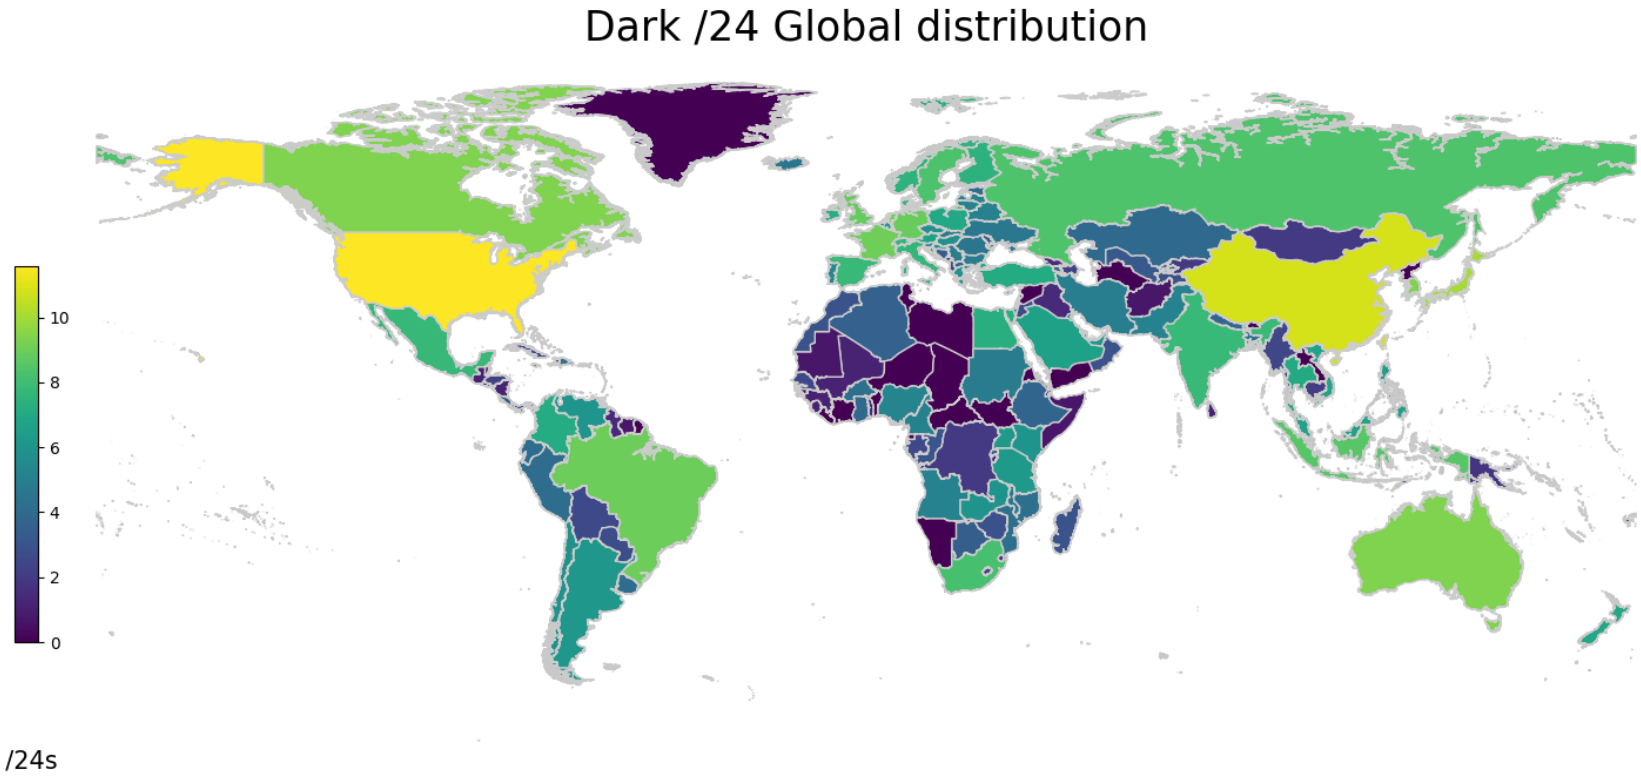

- Even in countries that no telescopes have ever been reported about

Prefix to country mapping with MaxMind GeoLite2

Insights

- Top scanned ports overall: **23 (telnet), 22 (SSH), 80 / 8080 (HTTP)**

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	23	22	23
#2	6379	80	22
#3	22	443	80
#4	80	8080	6379
#5	443	3389	445
#6	8080	5555	25565
#7	25565	60023	443
#8	5555	81	8080
#9	3389	8443	8090
#10	60023	2375	3389

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	8080	23	23
#2	23	0	22
#3	2083	22	80
#4	9001	53920	445
#5	1604	445	6379
#6	9480	80	0
#7	443	3389	5060
#8	143	5555	8088
#9	5900	49680	8090
#10	6000	8080	8080

Insights

- Top scanned ports overall: **Overlap**

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	23	22	23
#2	6379	80	22
#3	22	443	80
#4	80	8080	6379
#5	443	3389	445
#6	8080	5555	25565
#7	25565	60023	443
#8	5555	81	8080
#9	3389	8443	8090
#10	60023	2375	3389

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	8080	23	23
#2	23	0	22
#3	2083	22	80
#4	9001	53920	445
#5	1604	445	6379
#6	9480	80	0
#7	443	3389	5060
#8	143	5555	8088
#9	5900	49680	8090
#10	6000	8080	8080

Insights

- Top scanned ports overall: **Filtered ports visible**

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	23	22	23
#2	6379	80	22
#3	22	443	80
#4	80	8080	6379
#5	443	3389	445
#6	8080	5555	25565
#7	25565	60023	443
#8	5555	81	8080
#9	3389	8443	8090
#10	60023	2375	3389

Port Rank	Telescopes		
	TUS1	TEU1	TEU2
#1	8080	23	23
#2	23	0	22
#3	2083	22	80
#4	9001	53920	445
#5	1604	445	6379
#6	9480	80	0
#7	443	3389	5060
#8	143	5555	8088
#9	5900	49680	8090
#10	6000	8080	8080

Insights

- Top scanned ports by network type

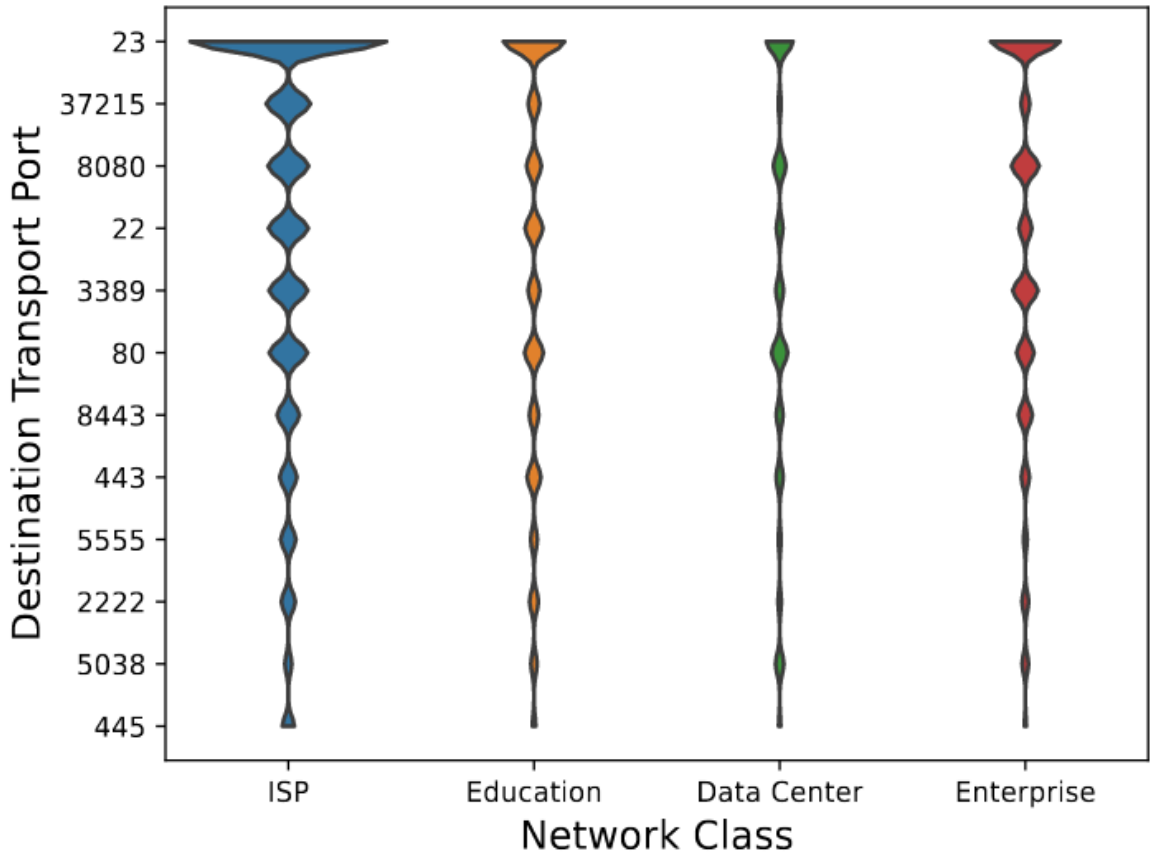

Prefix to country mapping with "IP to Company" by IPInfo

Insights

- Top scanned ports by continent

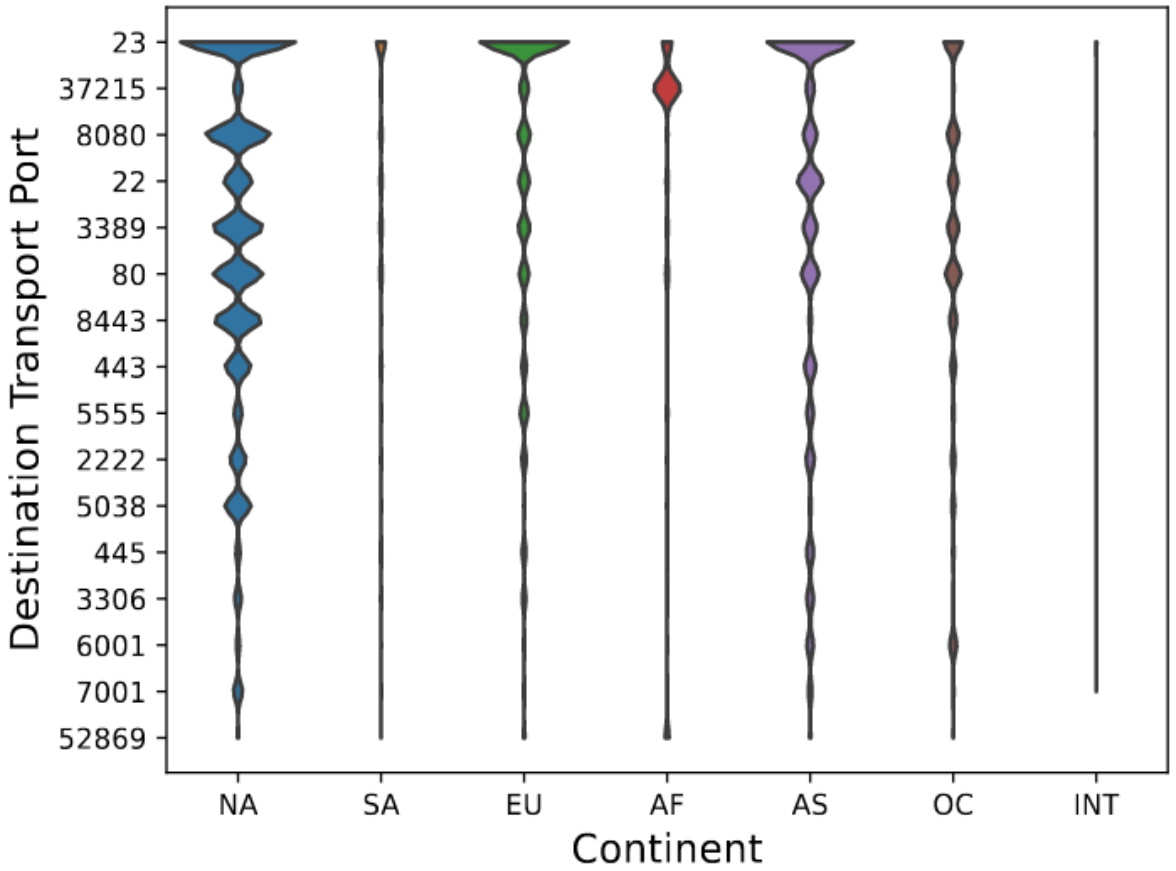

Conclusion

- /24 meta-telescope prefixes detectable with inferred filters
  - Around the globe
  - All network types
- Any network carrying IBR is theoretically suitable for our inference
- Helps to improve Internet security research

Challenges of Telescope Inference

- Effect of time: more unused blocks towards weekend

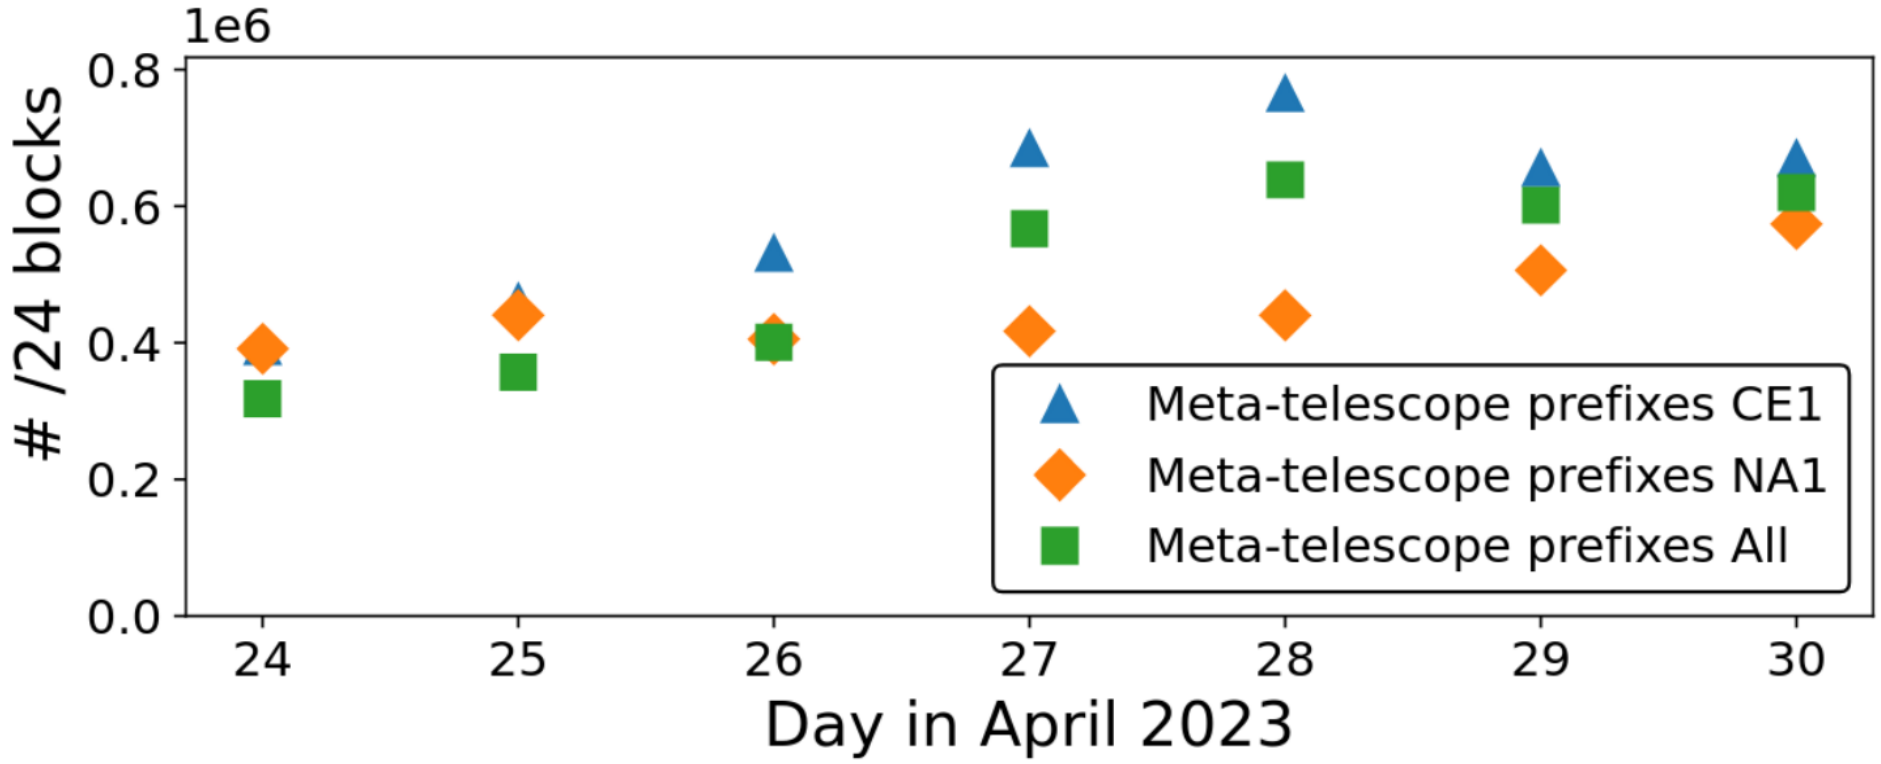

Challenges of Telescope Inference

- Effect of time
- Effect of sampling: Trade-off between FPR & #inferred blocks

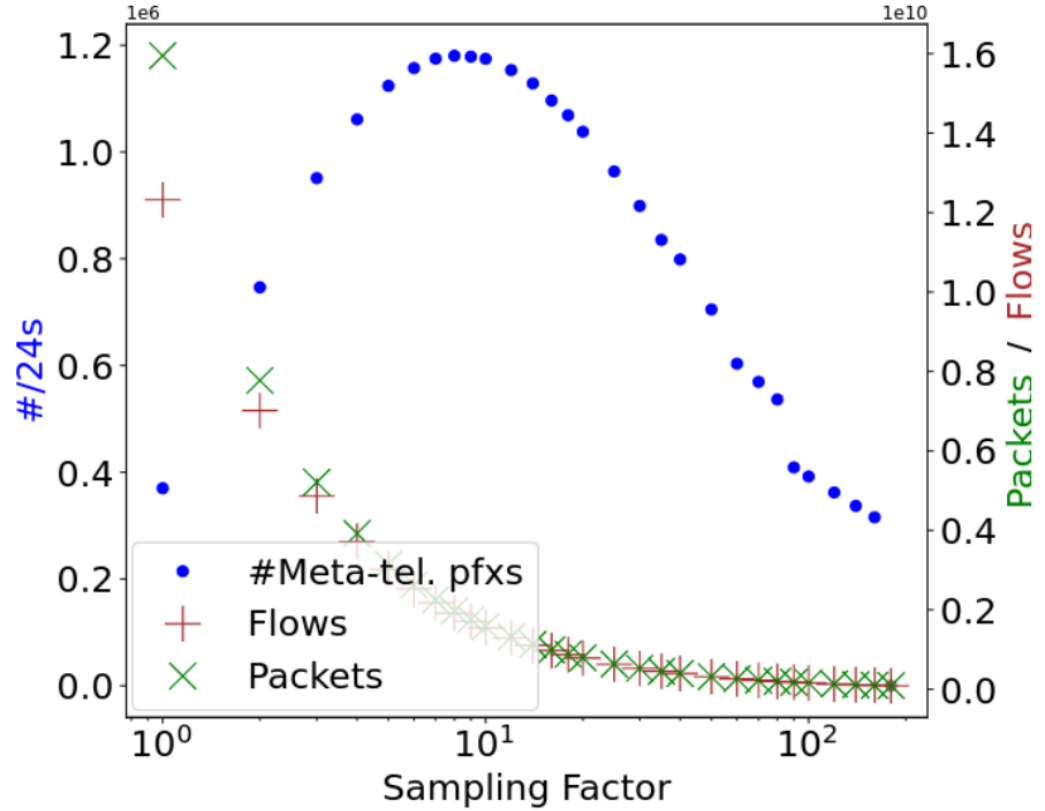

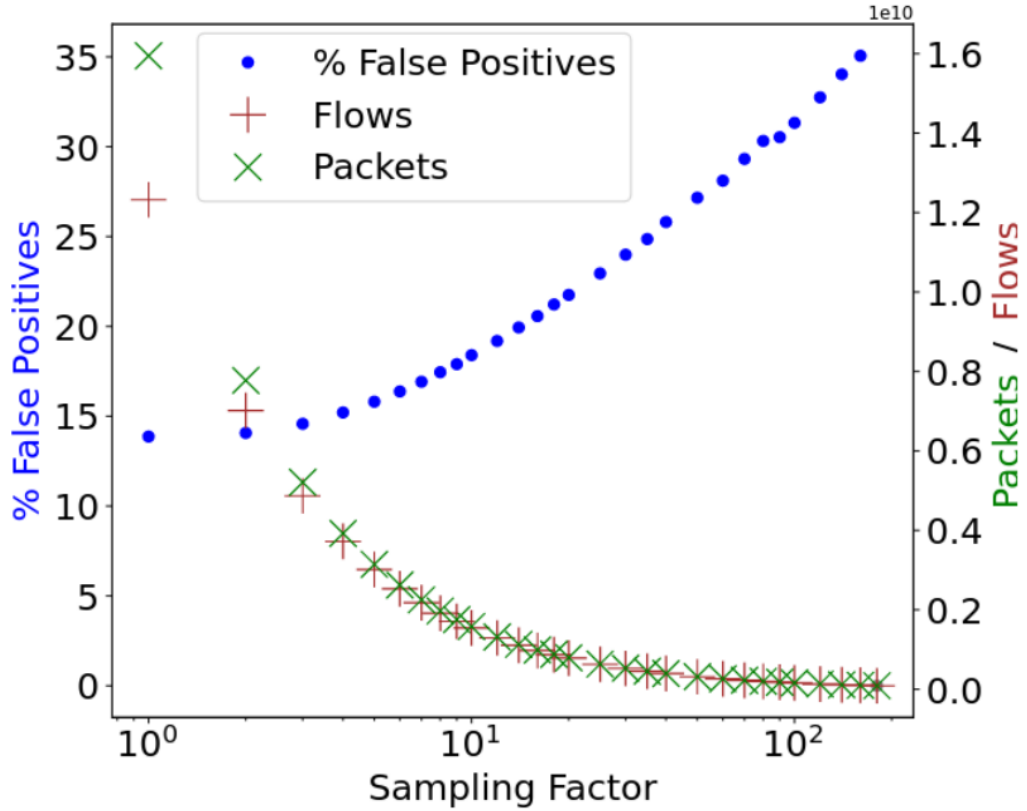

False positives as reported by Censys / NDT / ISI

Challenges of Telescope Inference

- Effect of time
- Effect of sampling
- Effect of spoofing: eliminates inferred blocks over time