An Active Telescope for Spoofing Detection

Raphael Hiesgen

INET, Hamburg University of Applied Sciences

Motivation

- Spoofing is a problem throughout the Internet
- Our focus: impact on measurements
  - Research and operations depend on reliable data
  - Source address often used for geolocation
- Application domain: UCSD Network Telescope

Goal

- Identify spoofed traffic in the IBR
- Challenges
  - One-way communication
  - Real-time processing

Probing to the Rescue

- Introduce active measurements to probe IBR sources
- Collect responses for a given source address
- Check if initial packet and replies have the same sender

Pseudo Source Address Validation

- *Idea:* Correlate initial IP ID with the IDs of probe replies
- Somewhat inaccurate (e.g., not all hosts reply to probes)
- Traditionally a system-wide counter
  - Can be used to attribute packets to the same host
  - Changed due to privacy concerns
  - Now often a counter per specific addresses + protocol tuple

IP ID Correlation

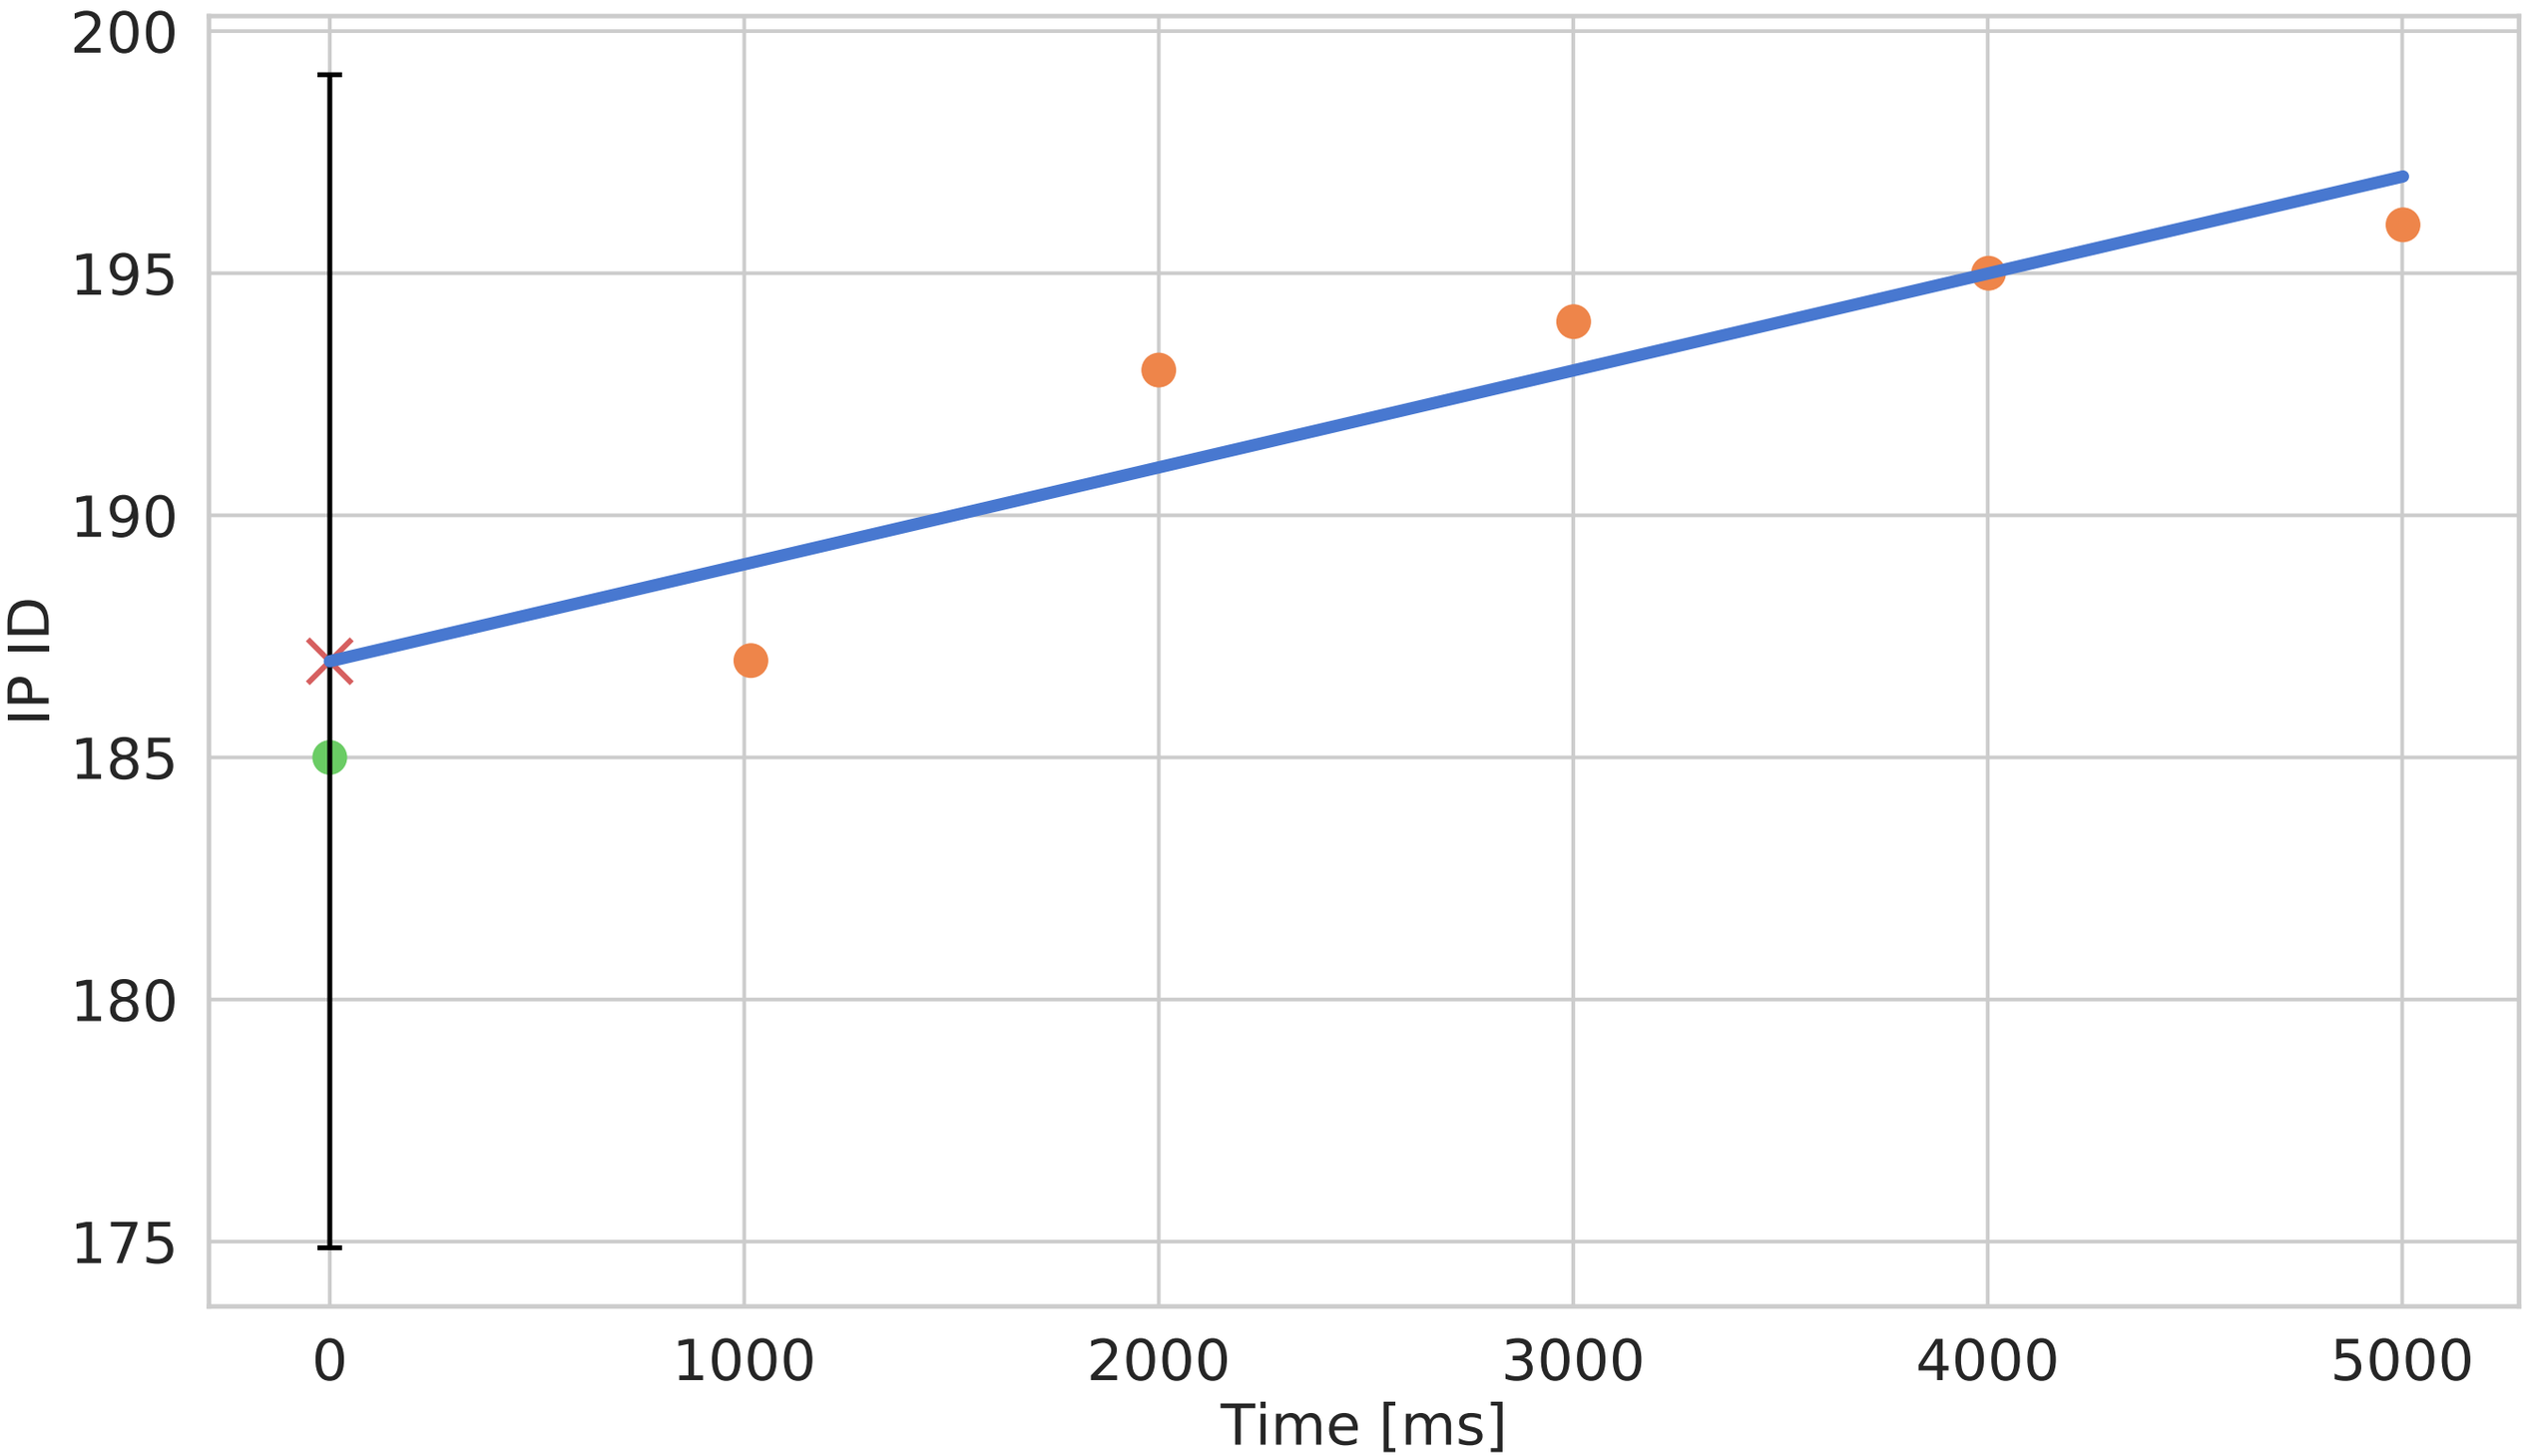

Handshake Continuation

- *Idea:* Accept TCP connections (SYN-ACK probing)
- High accuracy (only works if the target has state)
- Scanner behavior unclear
  - Some reply with RST, others establish the connection

Spoofing vs. Spoofing

- Both methods require probes from telescope addresses
- Replies mixed in with telescope traffic
- Impact on telescope traffic patterns unknown (so far)

Implementation: Spoki

- Native impl. based on the C++ Actor Framework (CAF)
- Parallel packet ingestion via libtrace
- Probing handled by scamper
- Deployed for two IP blocks:
  - 44.0.1.0/24 @UCSD
  - 91.216.216.0/24 @BCIX

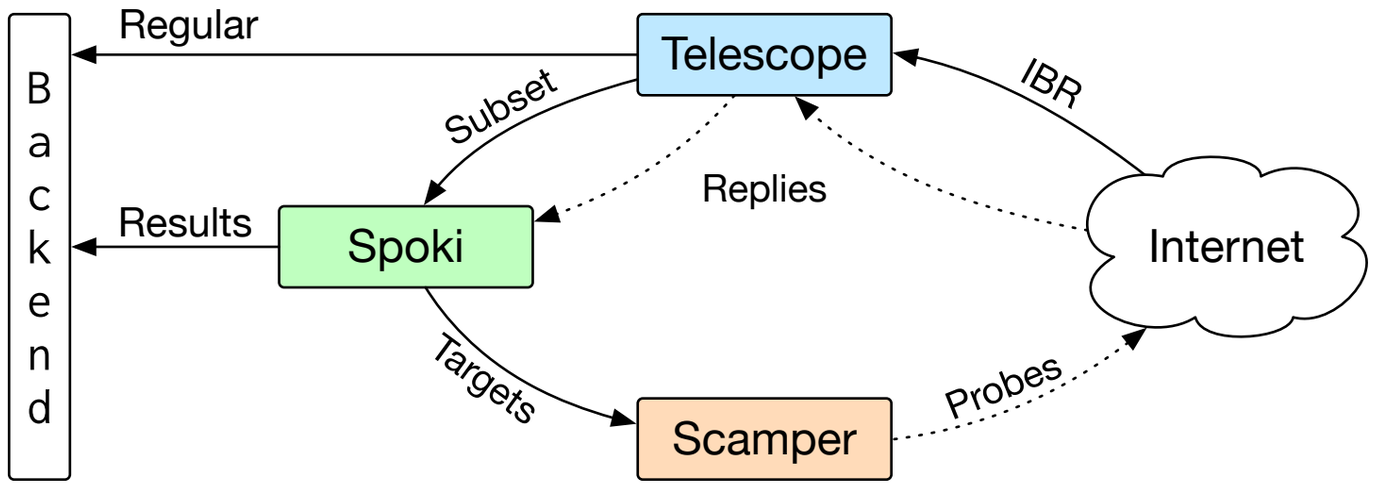

Challenges

- Reliably provoke replies
- Handle the data amount in real-time
- Identifies valid packets instead of spoofed ones

ICMP

- Probe with ICMP echo requests, analyze IP IDs of replies

	Events/Hour	Total Events	Got Reply	Validated
UCSD	40	573	346 (60%)	90 (16%)
BCIX	30	464	349 (75%)	85 (15%)

TCP

- Send SYN-ACK probe to complete the handshake

	Per Hour	Total Events	Got Reply	Validated
UCSD	5.439	78.705	65,651 (83%)	7,323 (9%)
BCIX	5.780	93.682	78,954 (84%)	10,146 (11%)

RST Replies

- 15 most targeted ports for events that replied with RST to probes

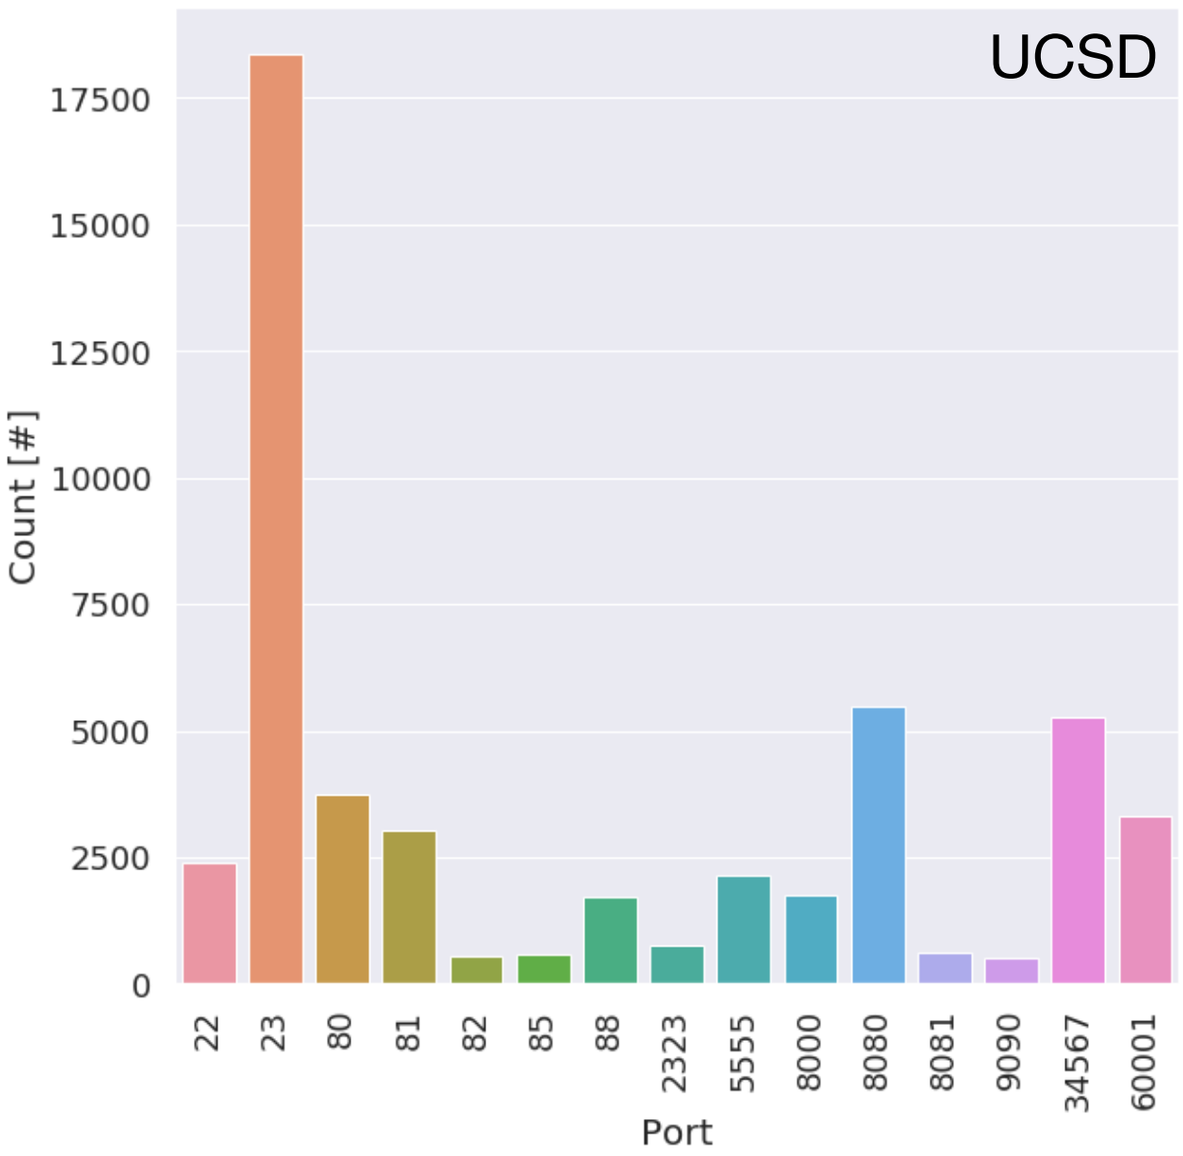

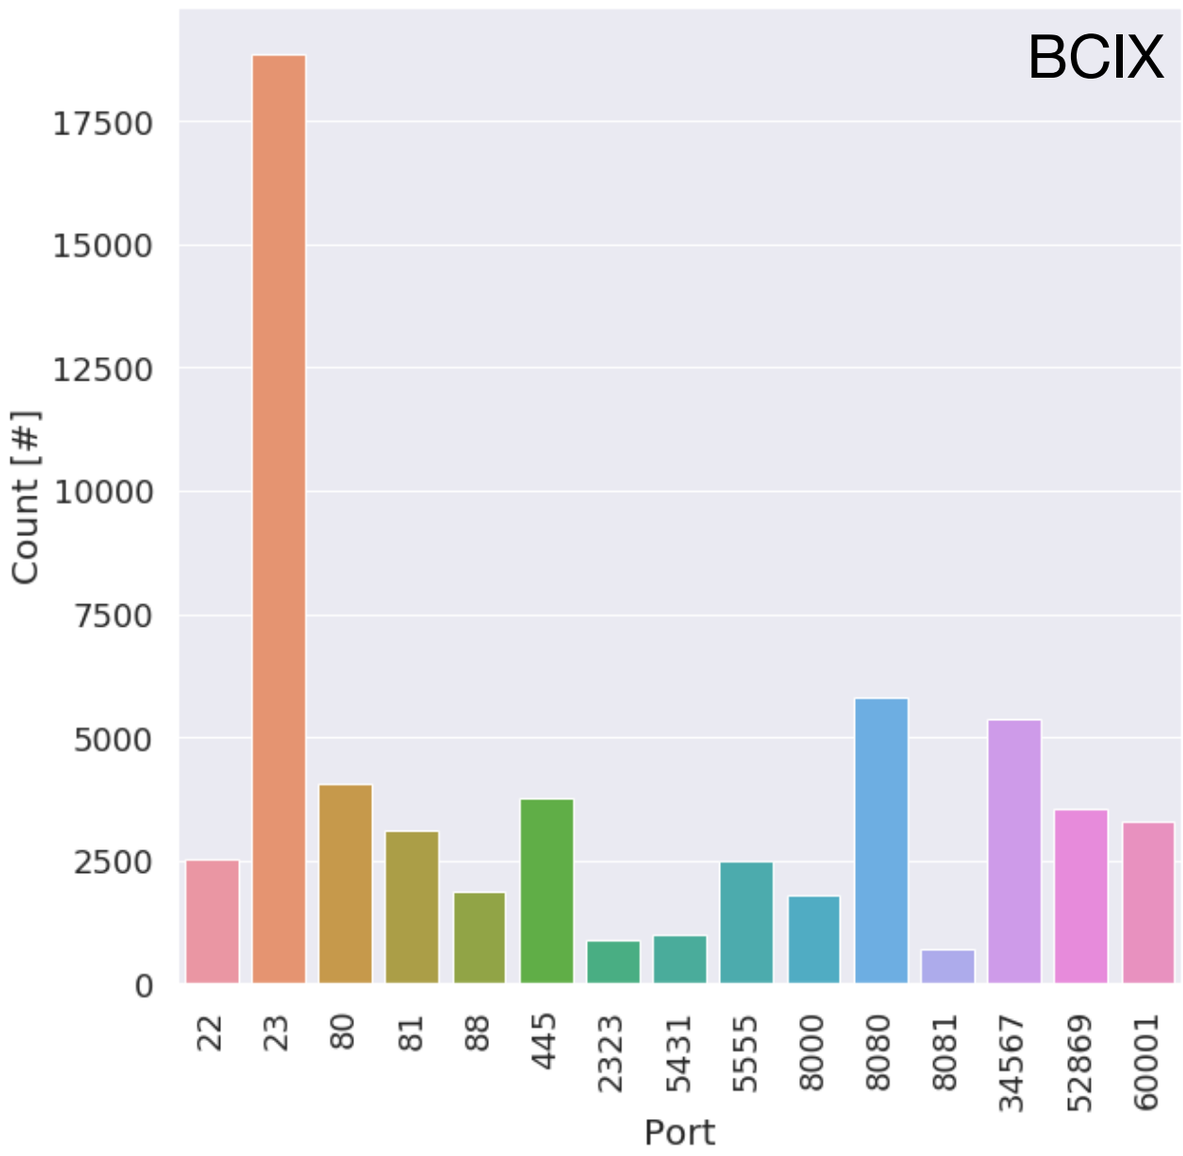

No Replies

- 15 most targeted ports for events did not get a reply

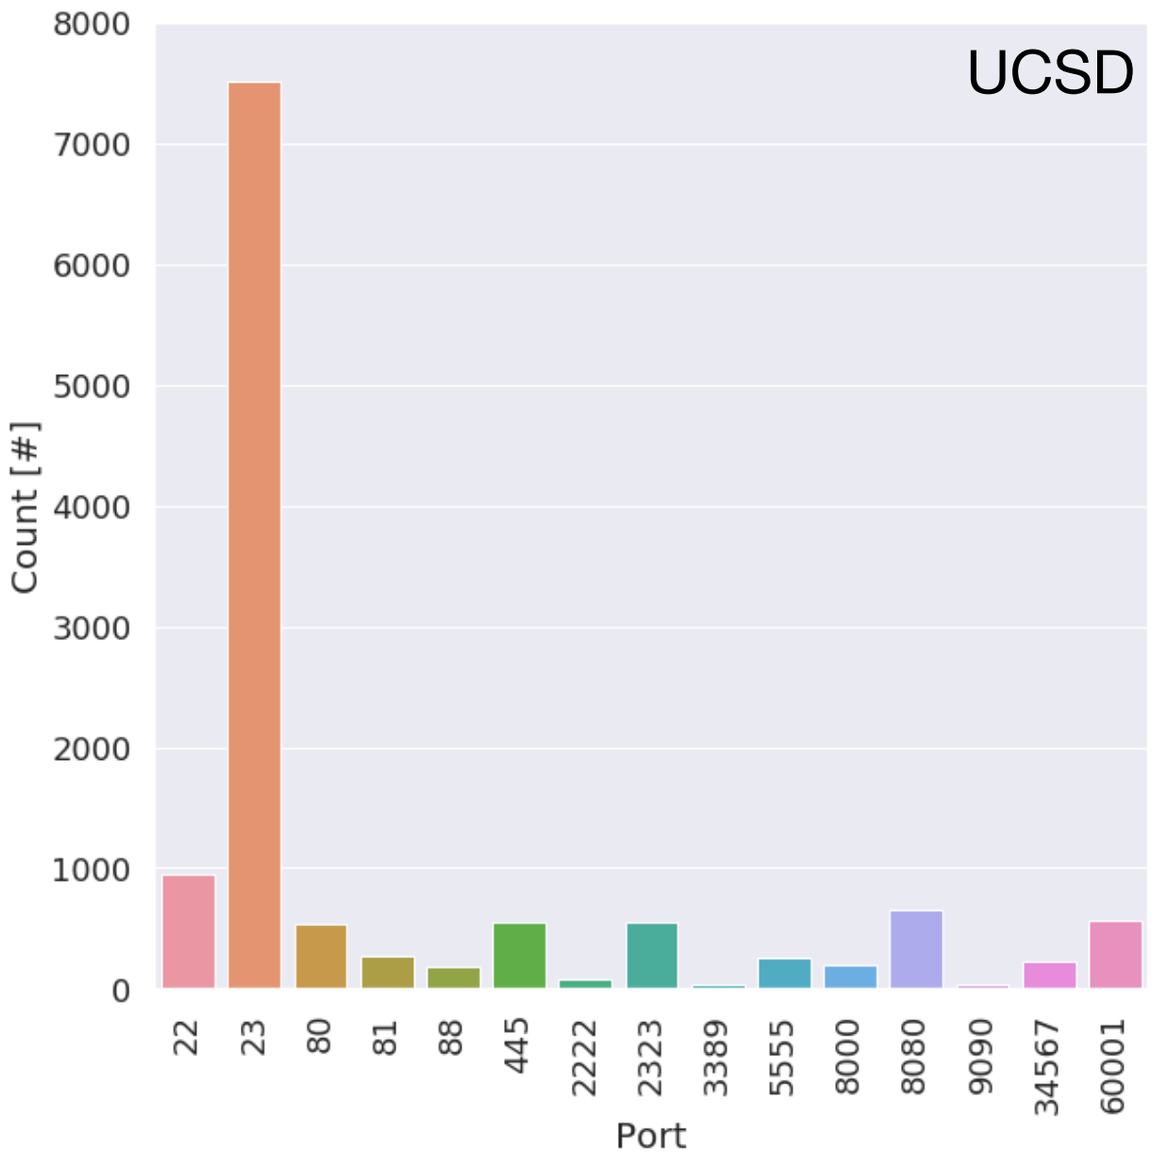

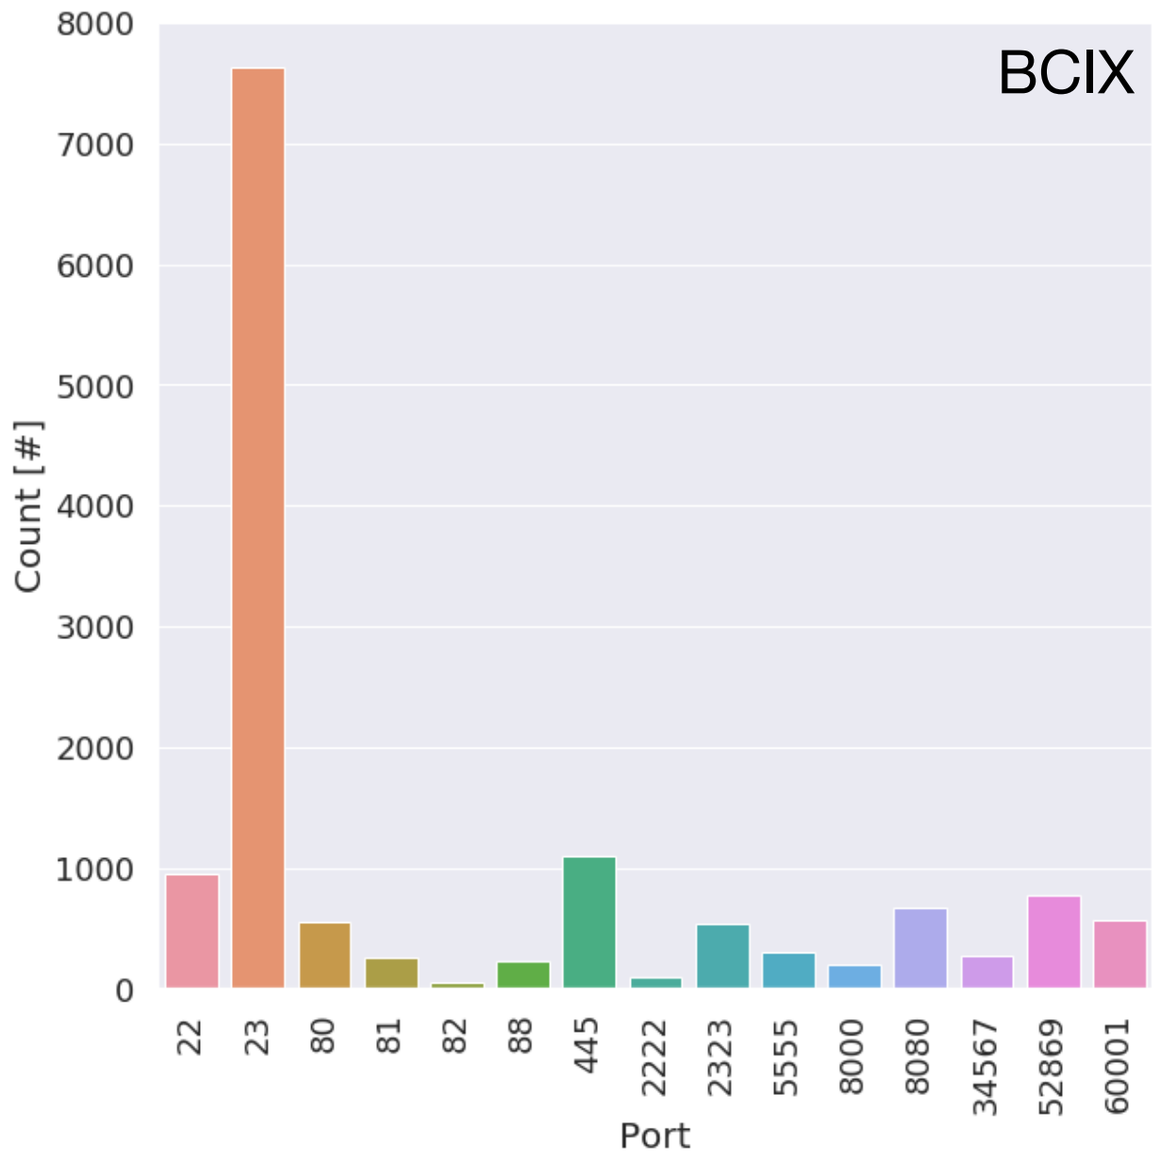

Regular Replies

- 15 most targeted ports for events that replied with non-RST to probes

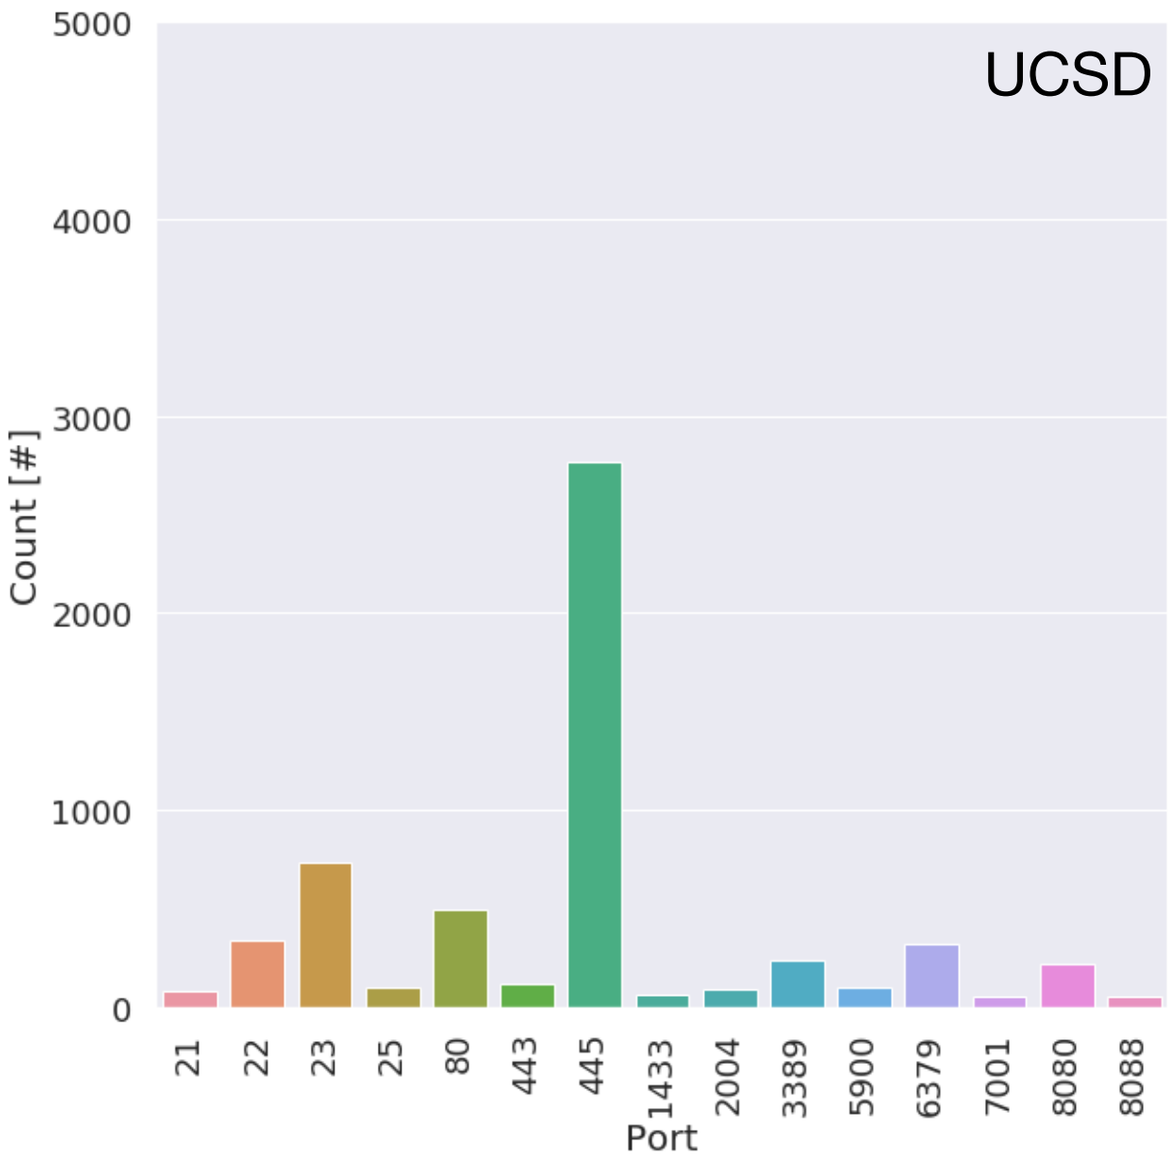

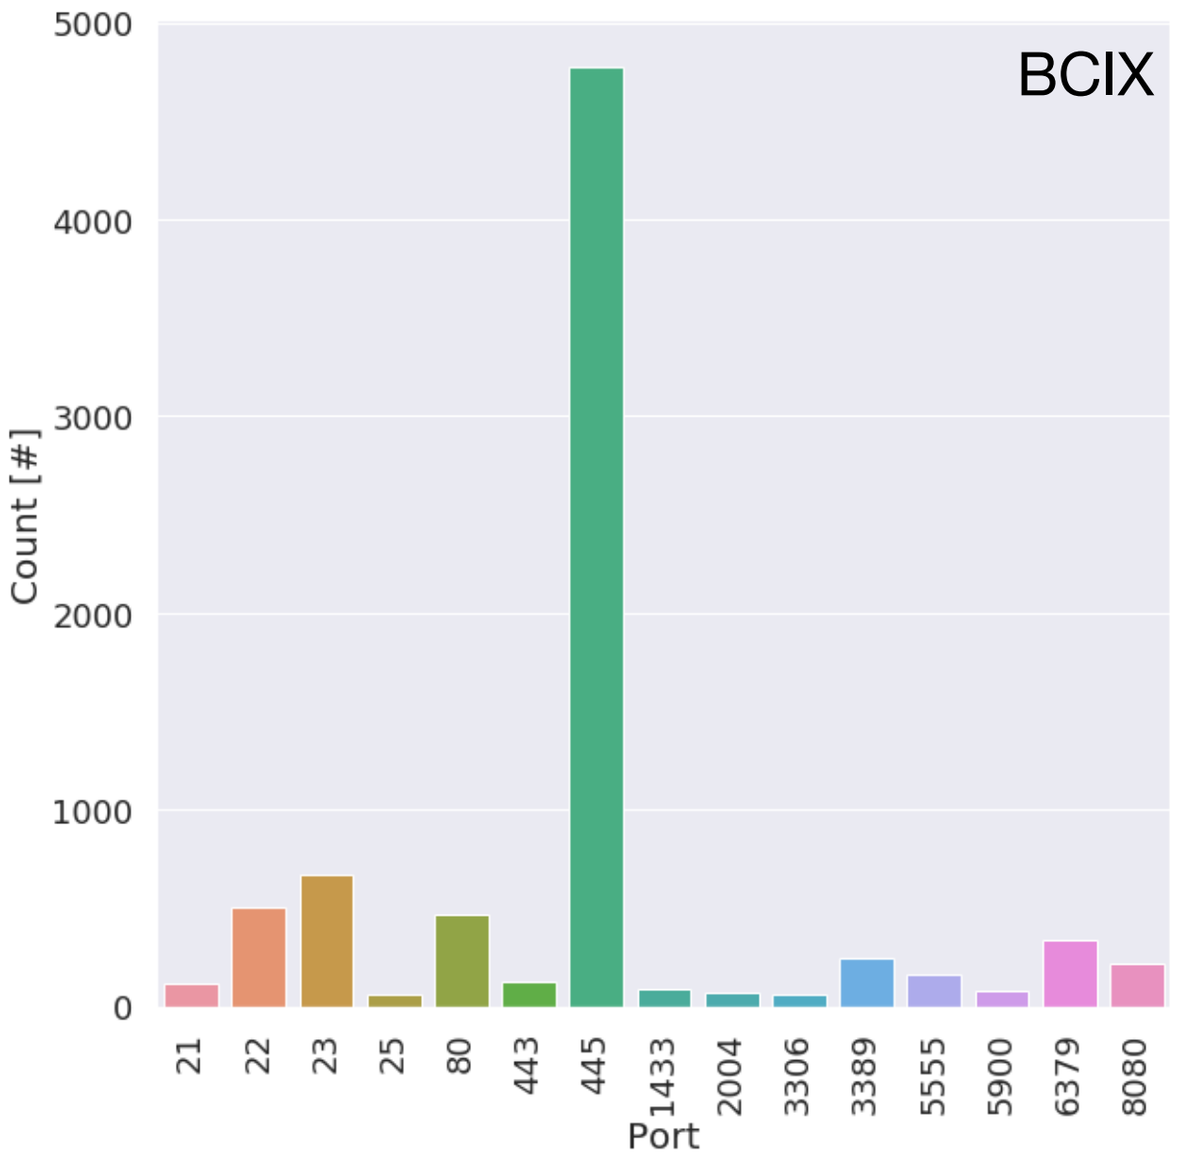

UDP

- Reflect payload, analyze IP IDs of replies

	Per Hour	Total Events	Got Reply	Validated
BCIX	215	3.241	175 (5%)	23 (1%)

Services

- 30 most targeted ports

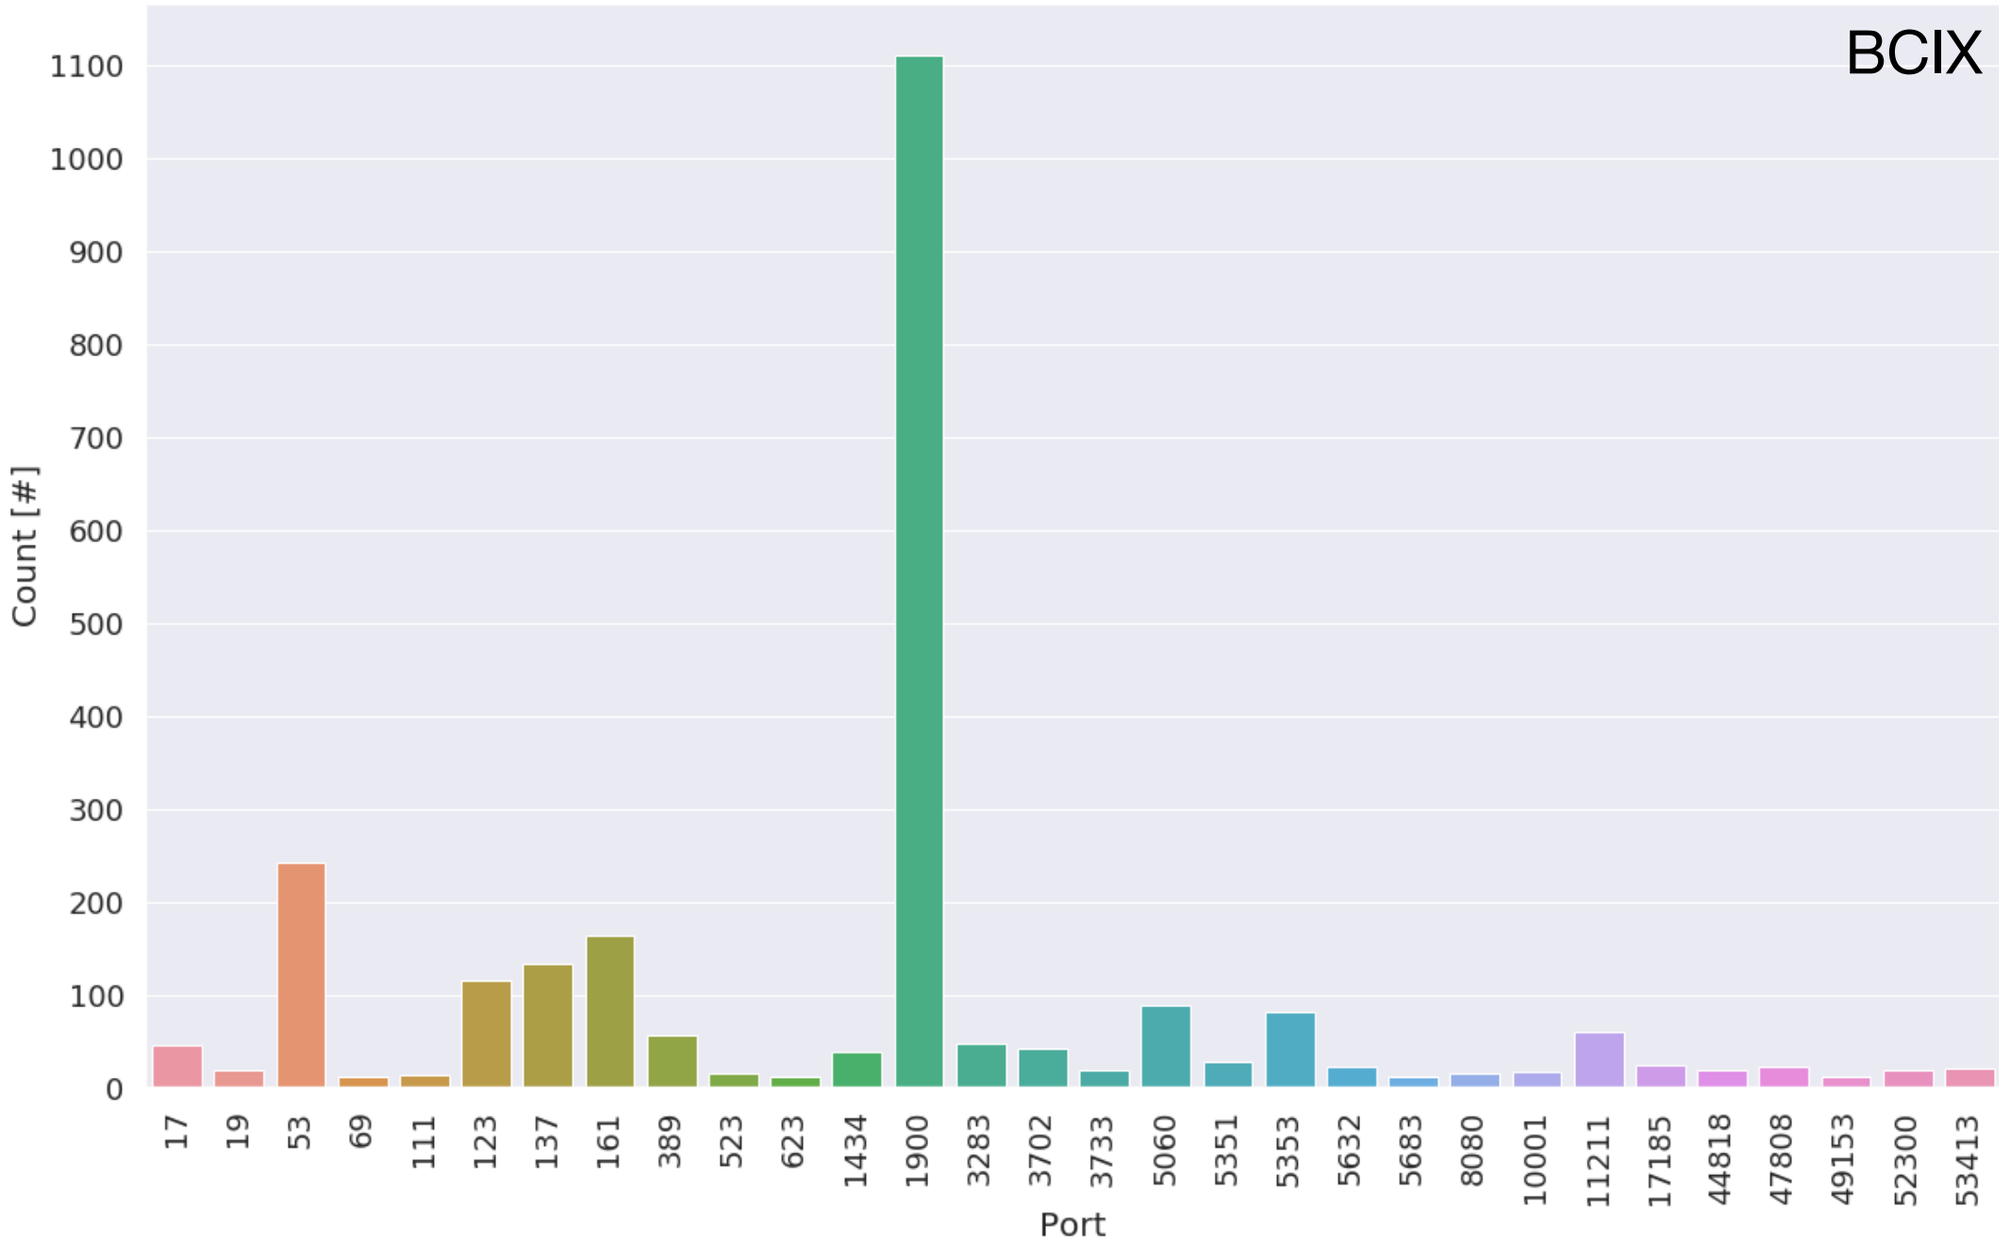

Provoking UDP Replies

- Problem: no standardized communication protocol
- Attempts so far:
  - Send service-specific probes
  - Send newline characters
  - Reflect payloads
  - Reply with ICMP destination unreachable — MTU exceeded

Next Steps

- Methodology
  - Validate the TCP results or find out how to improve them
  - UDP is very unstable and requires work
- How to extend the inferences to the entire /8?
- Can we transfer the technique into other contexts?