Advances in IPv6 Aliased Prefix Detection & A new Aliased Prefix Detection approach for IPv6 Scanning

Wei Zhang[1], **Gang Ren[1]**, Xia Yin[1], Lin He[1], Haoxiang Yang,[1] Haisheng Yu[2]

rengang@cernet.edu.cn

[1]Tsinghua University, [2]CNNIC

Feb, 2026

Contents

Introduction

IPv6 Scanning

- Why Scan IPv6?
  - IPv6 scanning is the cornerstone of next-generation cyberspace mapping and Cybersecurity Situational Awareness technologies.
- How to scan IPv6?
  - Target generation algorithm (TGA)
  - Aliased prefix detection

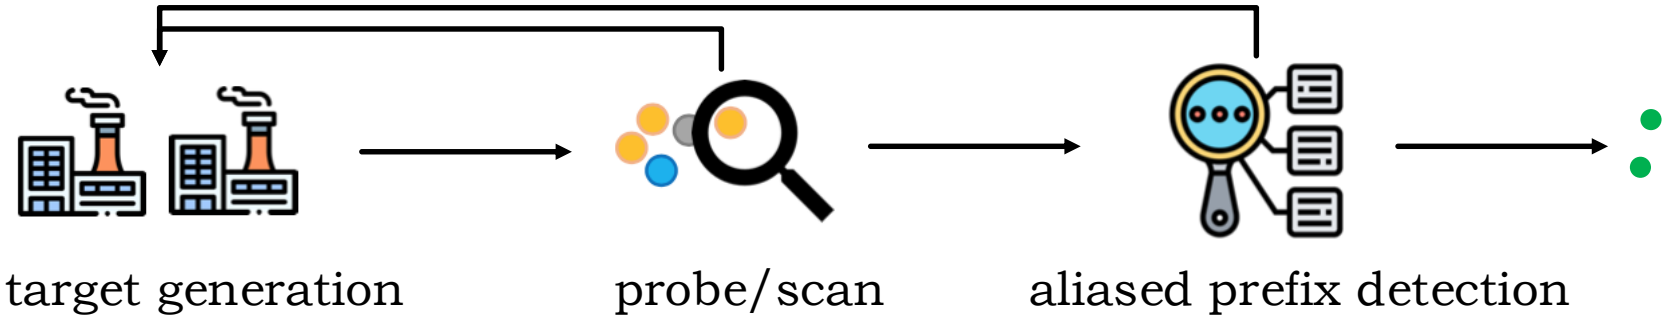

Aliased Prefixes in IPv6 Scanning

- Aliased Prefix is an entire prefix that is configured to map to a single network interface.
- Aliased prefix detection is Essential to prevent severely contaminated scan results and the massive waste of probing resources on these deceptive regions, ensuring more accurate and efficient IPv6 scanning.

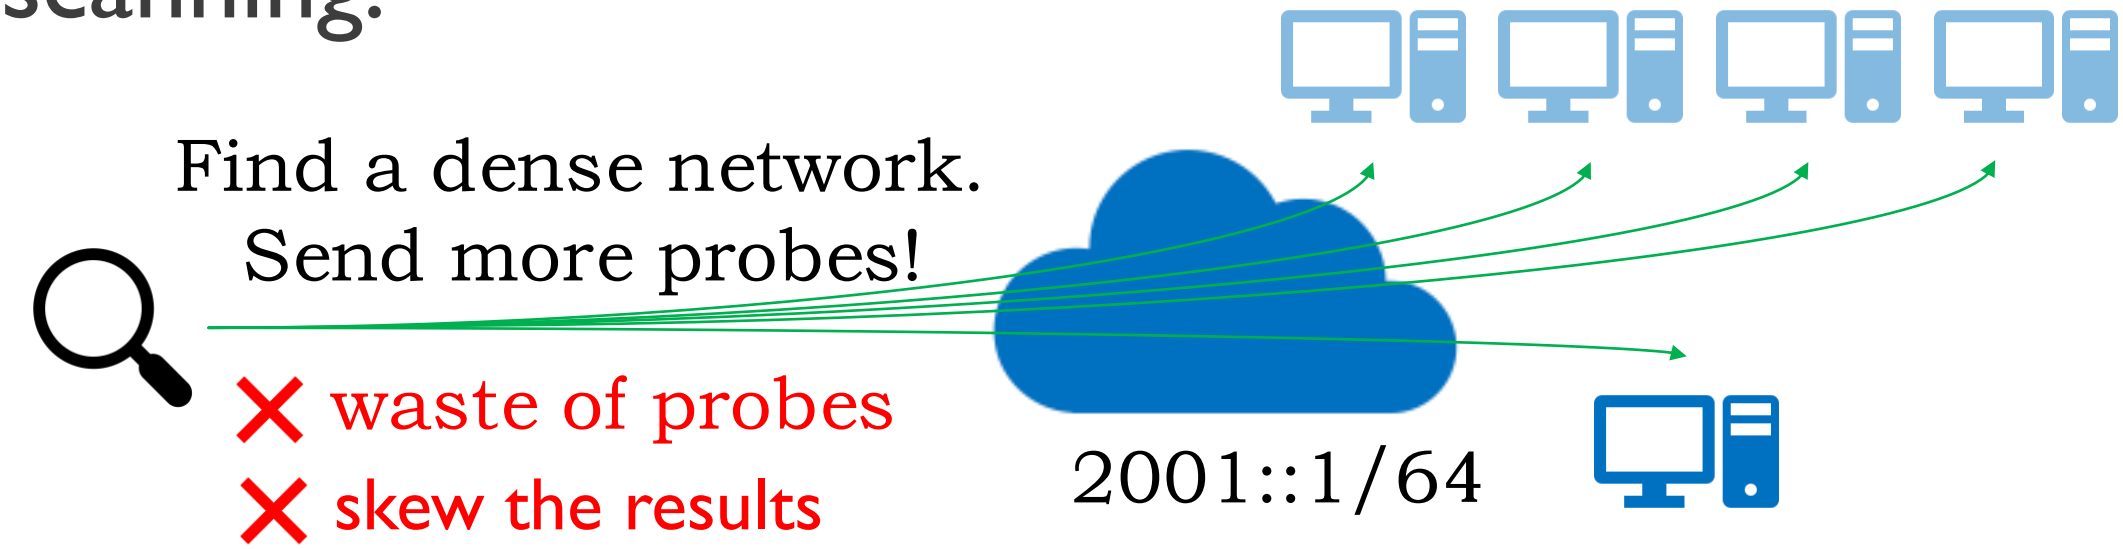

Advances in IPv6 Aliased Prefix Detection

Density-based methods

Core idea: probe a small number of random addresses within a given prefix. If the number of responses exceeds a certain threshold, the prefix is flagged as aliased.

- APD, 6Sense-APD
  - Fix length (/96) prefix
  - Probe 3 random addresses
- MAPD, HMap6-MAPD, Luori
  - Multi-Level length prefix
  - Probe 16 random addresses across 16 distinct sub-prefixes

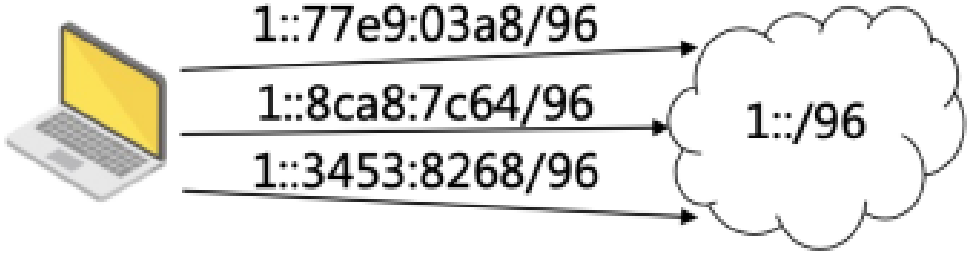


```
2001:0db8:0407:8000::/64
2001:0db8:0407:8000:0151:2900:77e9:03a8
2001:0db8:0407:8000:181c:4fcb:8ca8:7c64
2001:0db8:0407:8000:23d1:5e8e:3453:8268
                      :
2001:0db8:0407:8000:f693:2443:915e:1d2e
```

Fingerprint-based methods

- Core idea:
  - Used as a more precise follow up to validate the candidate aliased prefixes identified by density-based methods.
  - Instead of just counting probing responses, they analyze Fingerprints.
- Speedtrap, UAV6
- FAPD
  - Multi-Level Active prefix detection
  - Trick packet fragmentation
  - Check MTU caching behavior

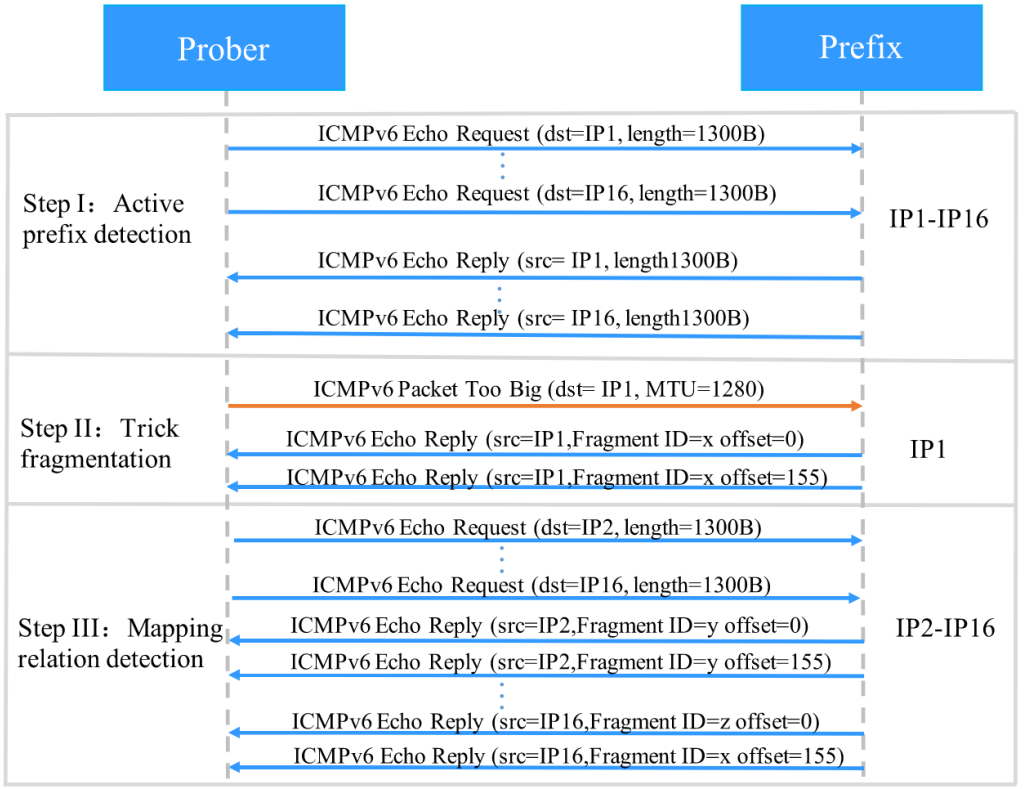

Summary of Existing Methods

- Flexible methods like MAPD or FAPD come with high overhead.
- Rely on active probing, accuracy is highly sensitive to packet loss.

Method		Prefix Length	Key Mechanism	Overhead
Density-based	APD 6Sense-APD	Fixed (/96)	Randomized testing within fixed range	3 Addresses
	6Tree-APD	Variable	Based on patterns discovered by 6Tree	160 Addresses
	MAPD HMap6-MAPD Luori	Multi-level	Recursive probing across sub-prefixes	16 Addresses
Fingerprint-based	Speedtrap UAV6 FBAR	N/A (Focus on Alias Addresses)	IP ID, PMTU, etc.	Pairwise comparison
	FAPD	Multi-level	MAPD+PMTU	MAPD cost + MTU induction

Limitations of Existing Methods

- Current methods are either limited in accuracy (like APD) or suffer from excessive overhead (like MAPD and FAPD).
- Our work aims to a solution that can achieve both high accuracy and low overhead.

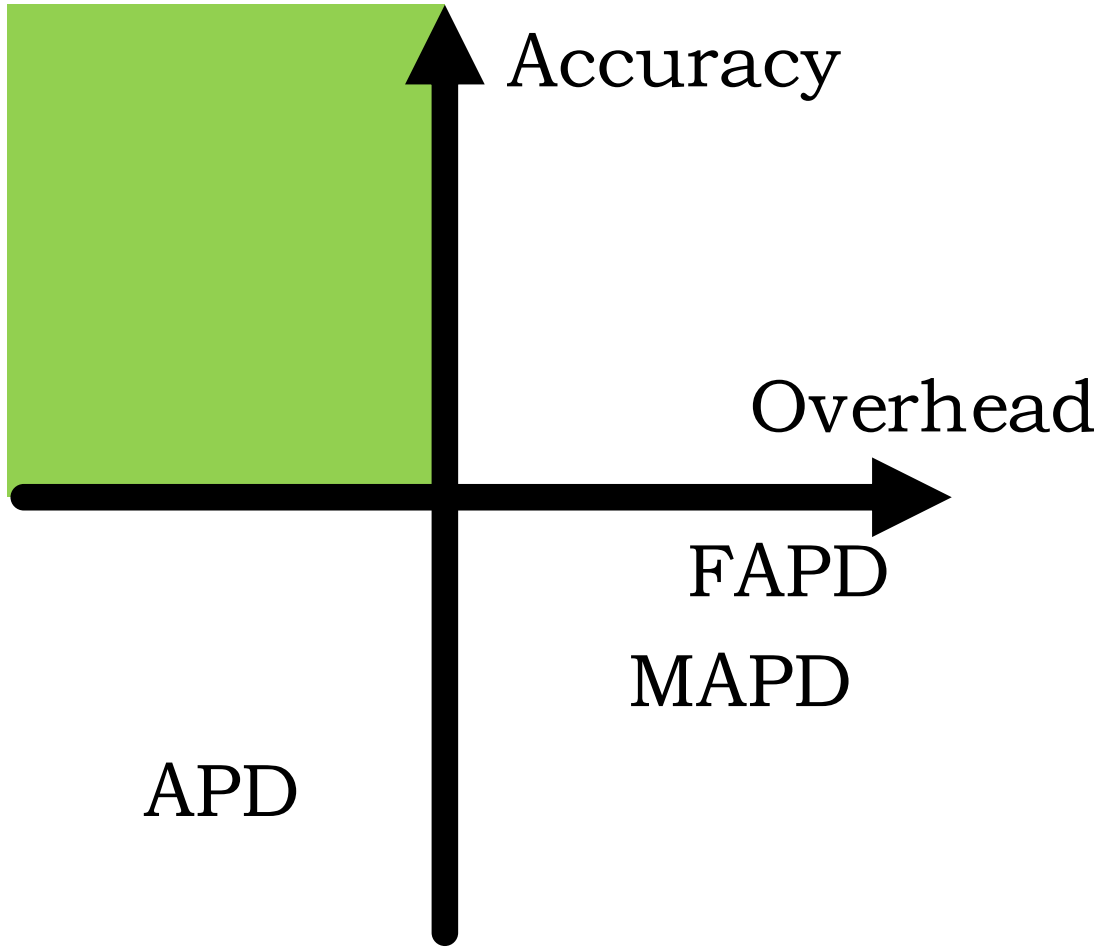

The PMAPD Approach

Overview

PMAPD: Passive-enhanced Multi-level Aliased Prefix Detection

- Integrate passive analysis with active probing for real-time detection within the ipv6 scanning loop.
- Each target first goes through a real-time filter. The results from the actual scan then feed back into our PMAPD module. This feedback is what triggers our "Passive and Active Components" to continuously refine and expand our knowledge of which prefixes are aliased.
- Flexibility: detecting prefix length across a wide and variable range, from BGP length up to /112 (in steps of 4 bits)

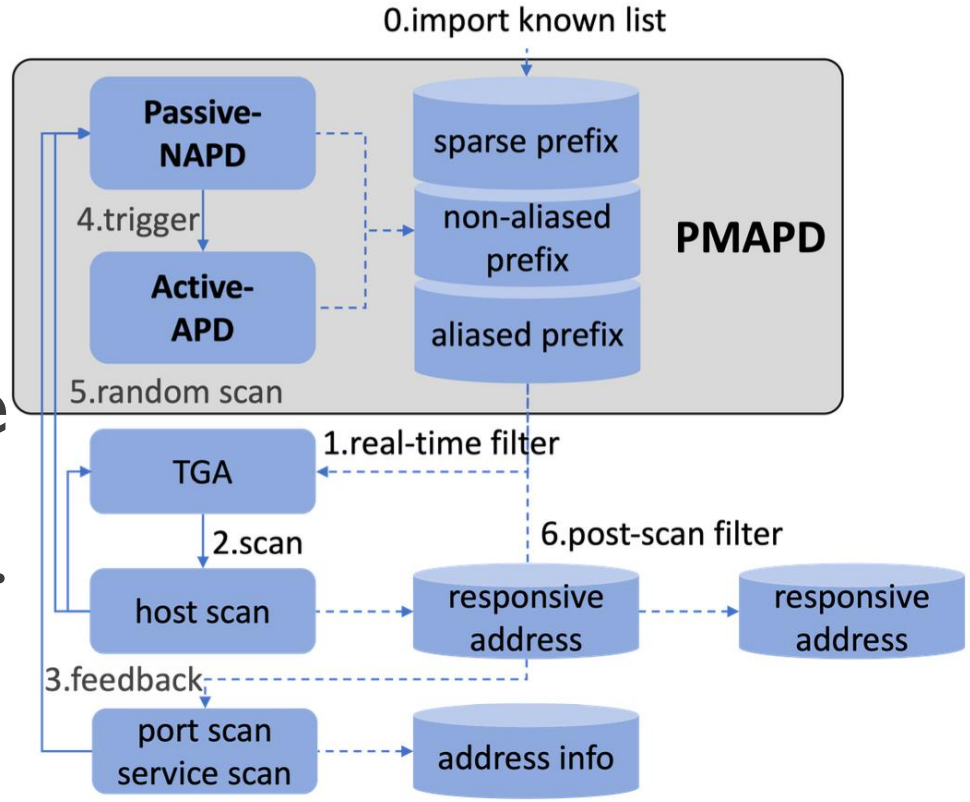

Passive-NAPD: Identifying Non-aliasing "For Free"

- Guiding Principle is simple
  - Addresses within an aliased prefix should exhibit consistent behavior.
- Passive Data, Multi-step filter
  - Service fingerprints
  - Host responsiveness
  - Port openness
- Method: Statistical tests
- Passive-first approach allows us to clear the vast majority of non-aliased prefixes without sending a single extra detection probe

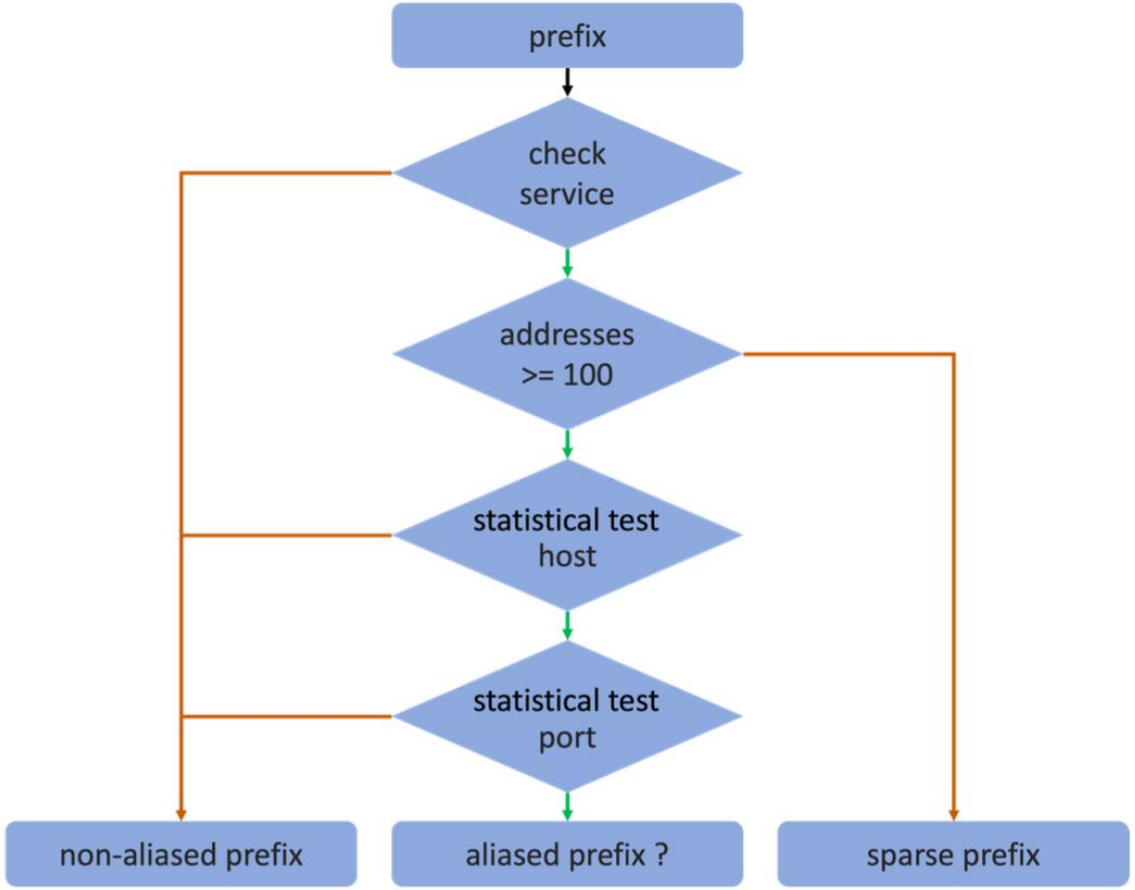

Non-Aliased Prefix Expansion

- When a prefix is identified as non-aliased, all its containing (i.e., shorter) parent prefixes are also marked as non-aliased.
  - Reduce detection overhead
  - Enhance robustness against interference from aliased sub prefixes

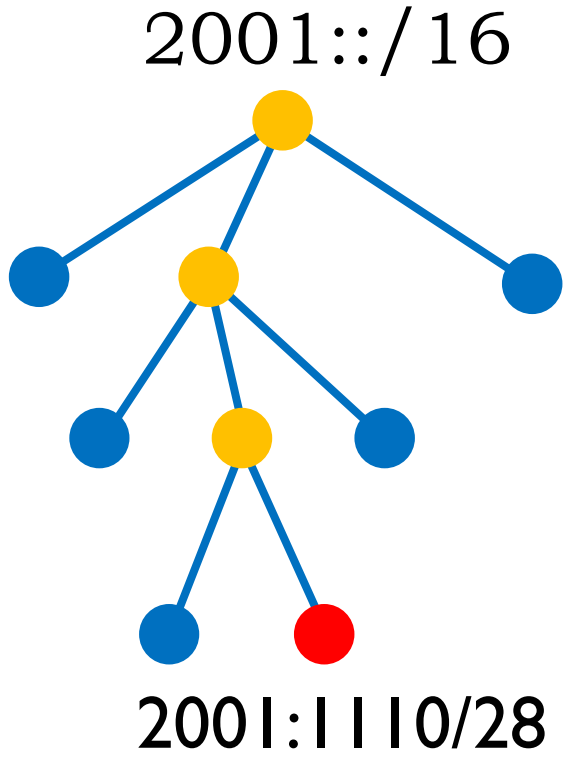

Active-APD: Optimized Probing for Suspicious Prefixes

- Sub prefixes selection
  - Instead of probing 3 completely random addresses, we select 3 distinct sub prefixes and probe one random address within each
  - More robust than APD and less overhead and sensitive to packet loss than MAPD
- Length-dependent threshold
  - Shorter Prefixes ≤ /96: require 2 responses out of 3 probes
  - Longer Prefixes> /96: require 3 responses out of 3 probes

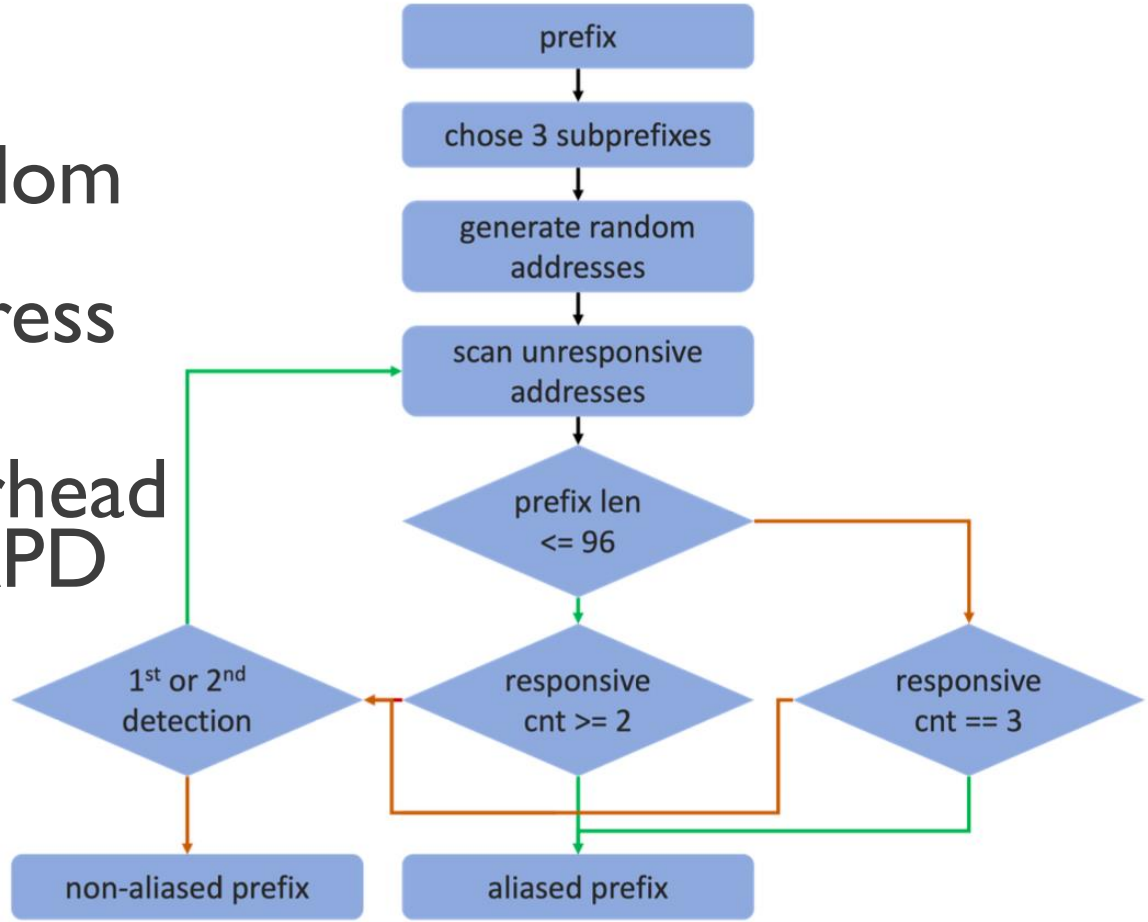

Theoretical Error Rate

- For a longer prefix like /112, a 2-out-of-3 rule would have a high false positive rate of 6.98% . That's why we use the stricter 3-out-of-3 rule, which drops the rate to 0.355%.
- For a shorter prefix like /96, the 2-out-of-3 rule has negligible false positive rate and false negative rate.
- Balancing accuracy and robustness

TABLE I: Theoretical false positive rate of Active-APD ($h = 10000$, $p = 0.05$)

Prefix Length	Threshold $r = 1$	$r = 2$	$r = 3$
/96	6.98×10^{-6}	1.63×10^{-11}	1.26×10^{-17}
/112	0.46	6.98×10^{-2}	3.55×10^{-3}

TABLE II: Theoretical false negative rate of Active-APD ($p = 0.05$)

Threshold $r = 1$	$r = 2$	$r = 3$
1.95×10^{-12}	4.69×10^{-8}	3.75×10^{-4}

Evaluation Setup

- TGA Selection
  - 6Tree, DET, 6Sense
- Comparison Methods: all running in a real-time detection mode
  - APD, 6Tree-APD, MAPD, HMap6-MAPD, and 6SENSE-APD
- Evaluation Metrics
  - De-aliased Hits, which measures how many real, unique hosts we discover. Higher is better.
  - Aliased Ratio in Responsive Addresses, which measures the pollution in our final results. Lower is better.
  - Detection Overhead, which measures the cost in probes. Lower is better.

Result 1: More De-aliased Hits than MAPD

- 6Tree
  - 4.64M vs 3.61M
  - Improved the hit count by 29% compared to MAPD
- DET
  - 16.85M vs15.02M
  - Improved the hit count by 12% compared to MAPD
- 6Sense
  - 3.53M vs 2.25M
  - Improved the hit count by 57% compared to MAPD
- Substantial increase in scanning efficiency

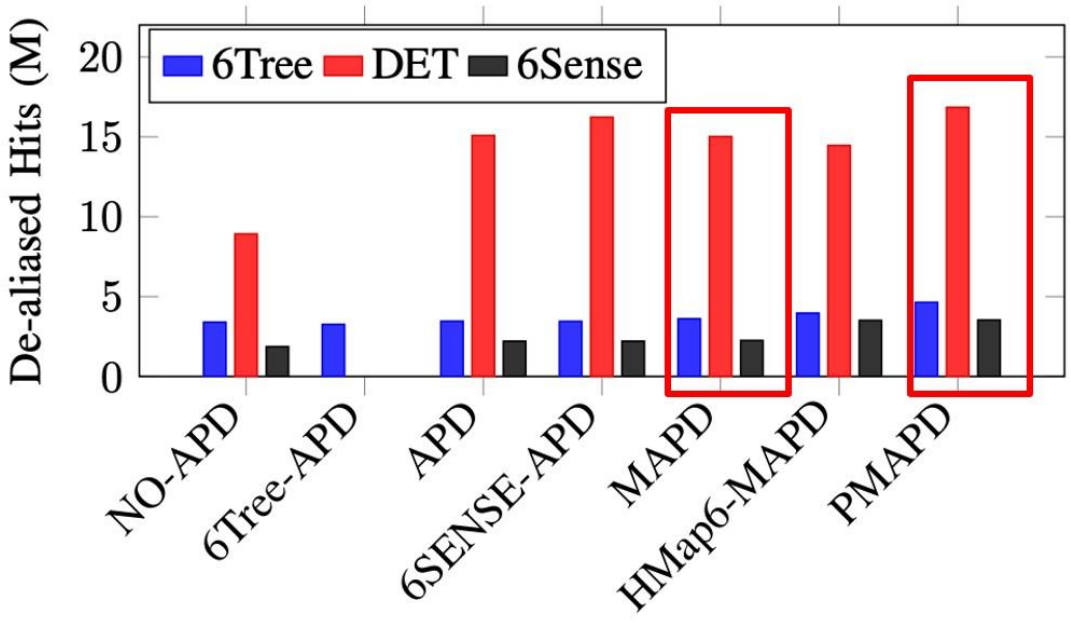

Result 2: Lower Aliased ratio than MAPD

- 6Tree
  - PMAPD Aliased ratio was only 16.84%, compared to 95.96% for MAPD
- DET
  - PMAPD Aliased ratio was 45.46 %, compared to 50.24 % for MAPD
- 6Sense
  - PMAPD Aliased ratio was only 4.3%, compared to 76% for MAPD
- A massive improvement in the quality and reliability of the scan data.

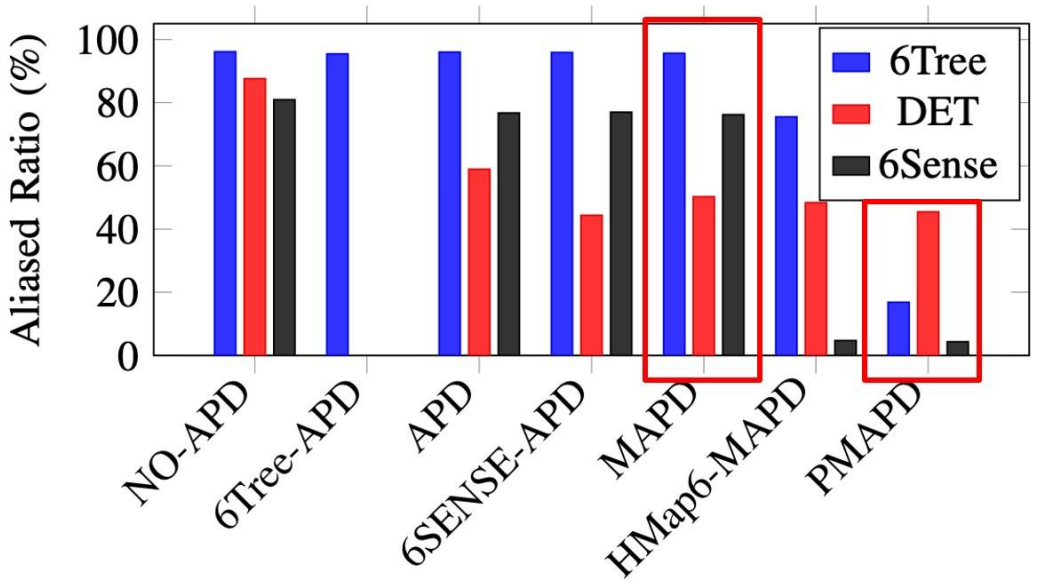

Result 3: Less Detection Overhead than MAPD

- PMAPD achieves these superior results with far less overhead because of highly effective passive-first strategy.
- 6Tree
  - PMAPD overhead was 1.58 million probes, while MAPD required 9.33 million.
- DET
  - PMAPD overhead was 0.99 million probes, while MAPD required 16.74 million.
- 6Sense
  - PMAPD overhead was 0.15 million probes, while MAPD required 4.56 million.

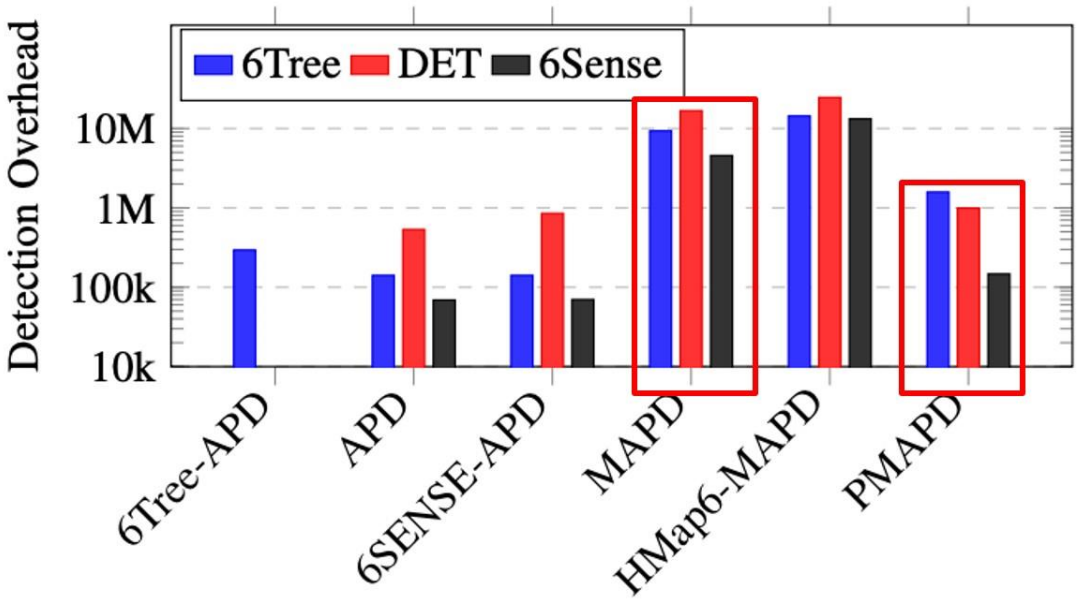

How Does PMAPD Achieve Such High Efficiency?

- Passive Component with free host Responsiveness data
  - 98.3% of all non-aliased prefixes are identified by Passive-NAPD
  - 80.7% of all non-aliased prefixes are identified with host Responsiveness

TABLE IV: Performance of PMAPD in identifying non-aliased prefixes (DET)

Method Component	Non-Aliased	%	Errors[a]	%
Passive-NAPD	798,878	98.3	230	0.03
Service Information[b]	138,286	17.0	/	/
Port Status	4,823	0.6	25	0.52
Responsiveness	655,769	80.7	205	0.03
Active-APD	14,084	1.7	13	0.09
Total PMAPD	**812,962**	**100.0**	**243**	**0.03**

Revealing the Landscape of Aliased Prefixes

- Only 1.60% of the ASes contains aliased prefixes
- Length distribution of aliased prefixes
  - Over 67% of the aliased prefixes we found are /64 or shorter
  - /96 prefix accounts for less than 0.7%
- Fixed-length detection methods are inadequate and highlights the absolute necessity for a flexible, wide-range detection approach like PMAPD.

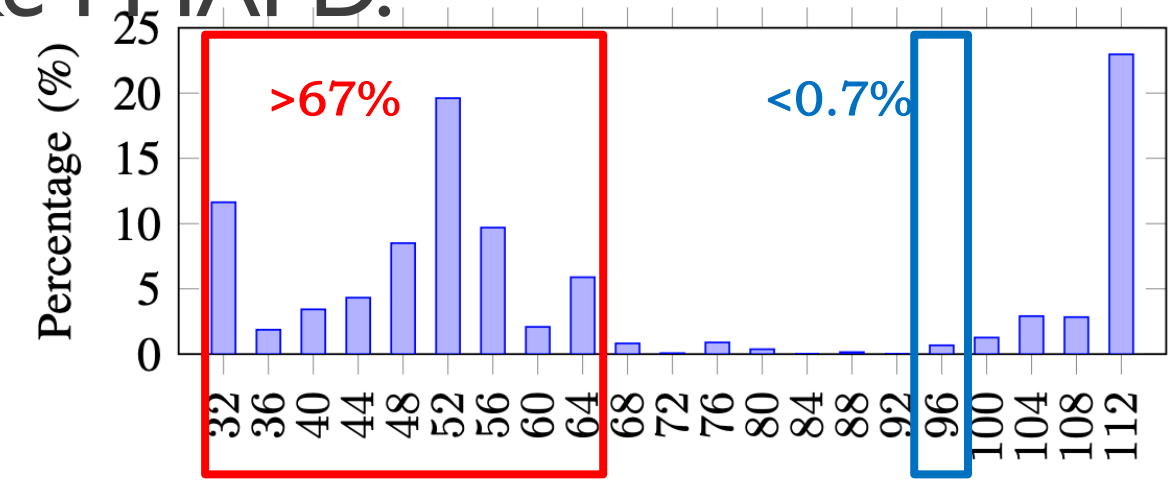

Fig. 9: Distribution of aliased prefix lengths (bits)

Discussion & Future Work

Discussion & Future Work

- Discussion
  - We acknowledge the challenge of establishing a perfect ground truth and the assumptions in our statistical model.
- Future work
  - Enhancing Passive Analysis with Richer Fingerprints
  - Developing "Aliasing-Aware" TGAs
- The related research has been published at the IEEE International Conference on Network Protocols (ICNP) 2025.

Thank you for listening !

Q & A

rengang@cernet.edu.cn

Supported by Open Project of National Engineering Laboratory for Technology of Internet Domain Name (NO.KF202502)