NAVAL
POSTGRADUATE
SCHOOL

Discovering the IPv6 Network Periphery

Erik Rye (CMAND)
Robert Beverly (NPS)

Passive and Active Measurement Conference
March 30, 2020

CMAND

Background

- IPv6:
  - Large address space + sparsity
  - Ephemeral and dynamic addressing
  - No need for address translation
- Implication:
  - IPv6 is deployed differently than IPv4!

IPv6 "Periphery"

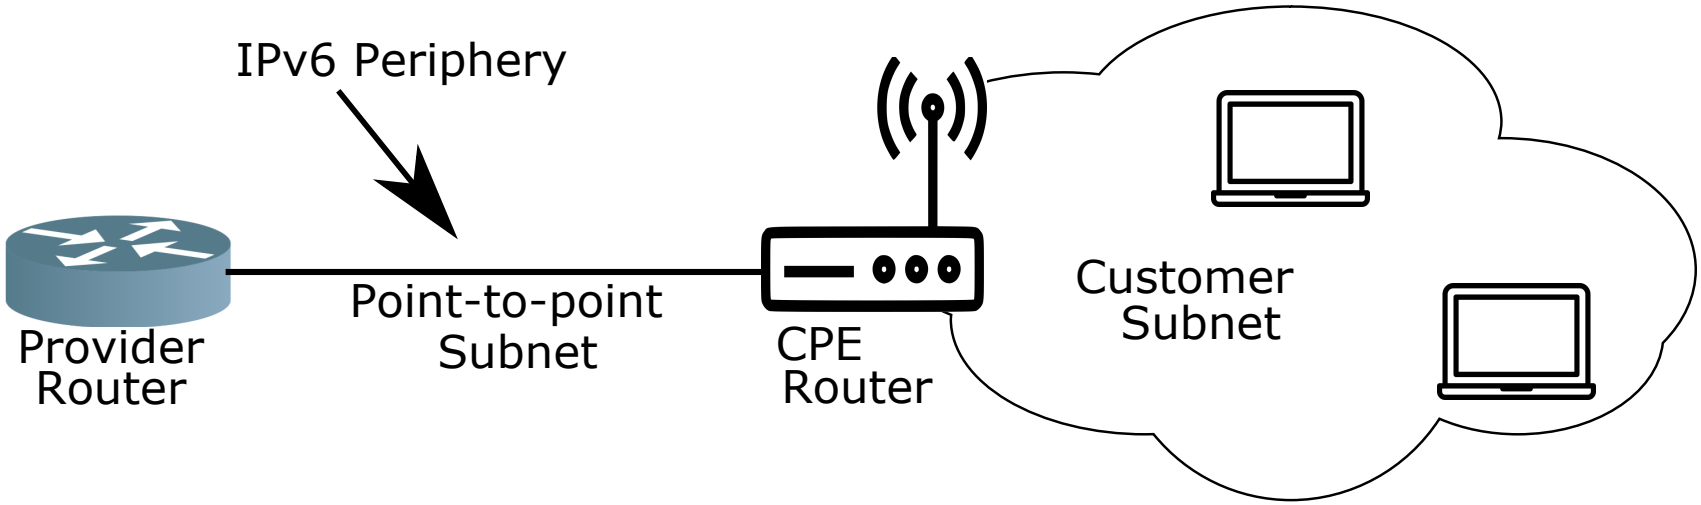

- Device at customer premises (CPE) is a routed hop!
- Subnet allocated to link between provider's router and CPE
- *Different* subnet allocated to customer, on other side of CPE

IPv4 Periphery Discovery

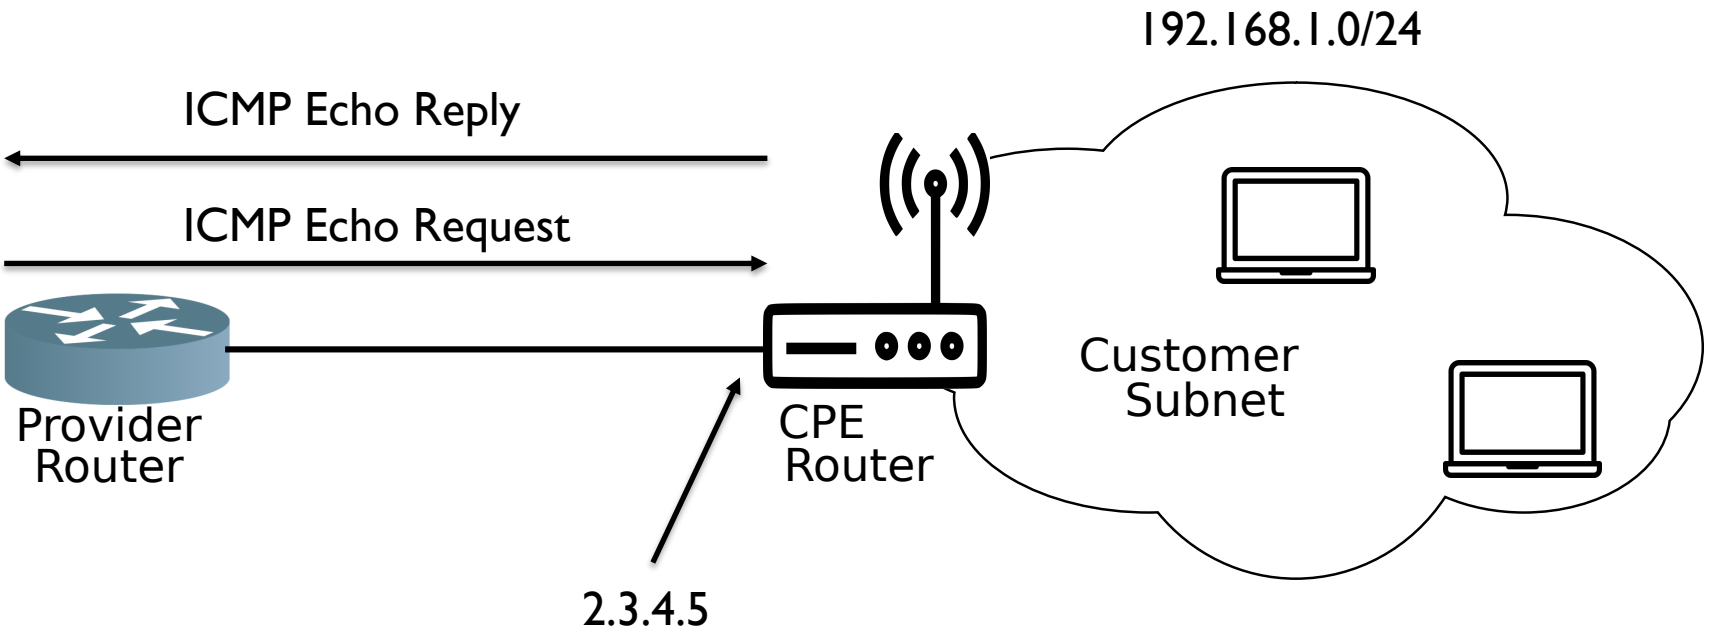

IPv4 address space can be exhaustively probed, so CPE do (or don't) respond to echo requests like every other public IPv4 host. Customer RFC1918 subnet isn't reachable

IPv6 Periphery Discovery

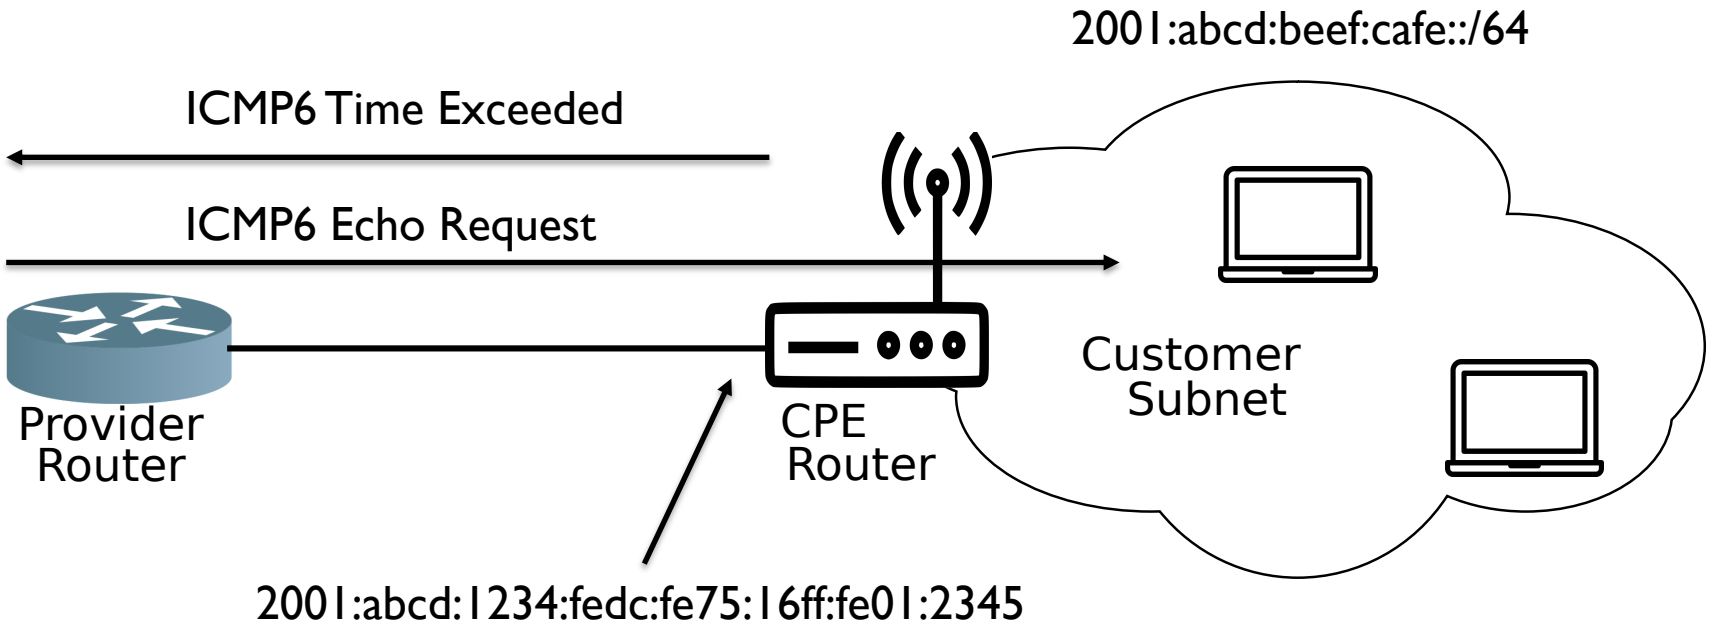

CPE device is a routed hop to on the path to the customer subnet. Traceroute echo request unlikely to hit a customer device – but doesn't need to in order to discover periphery.

The Reality of IPv6 Traceroutes

- Many mapping systems trace to a random address within advertised BGP prefixes:
  - Unlikely to reach a prefix allocated to a customer's CPE or her network
  - Even less likely to reach a responsive host
- Results are therefore ambiguous

The Reality of IPv6 Traceroutes…

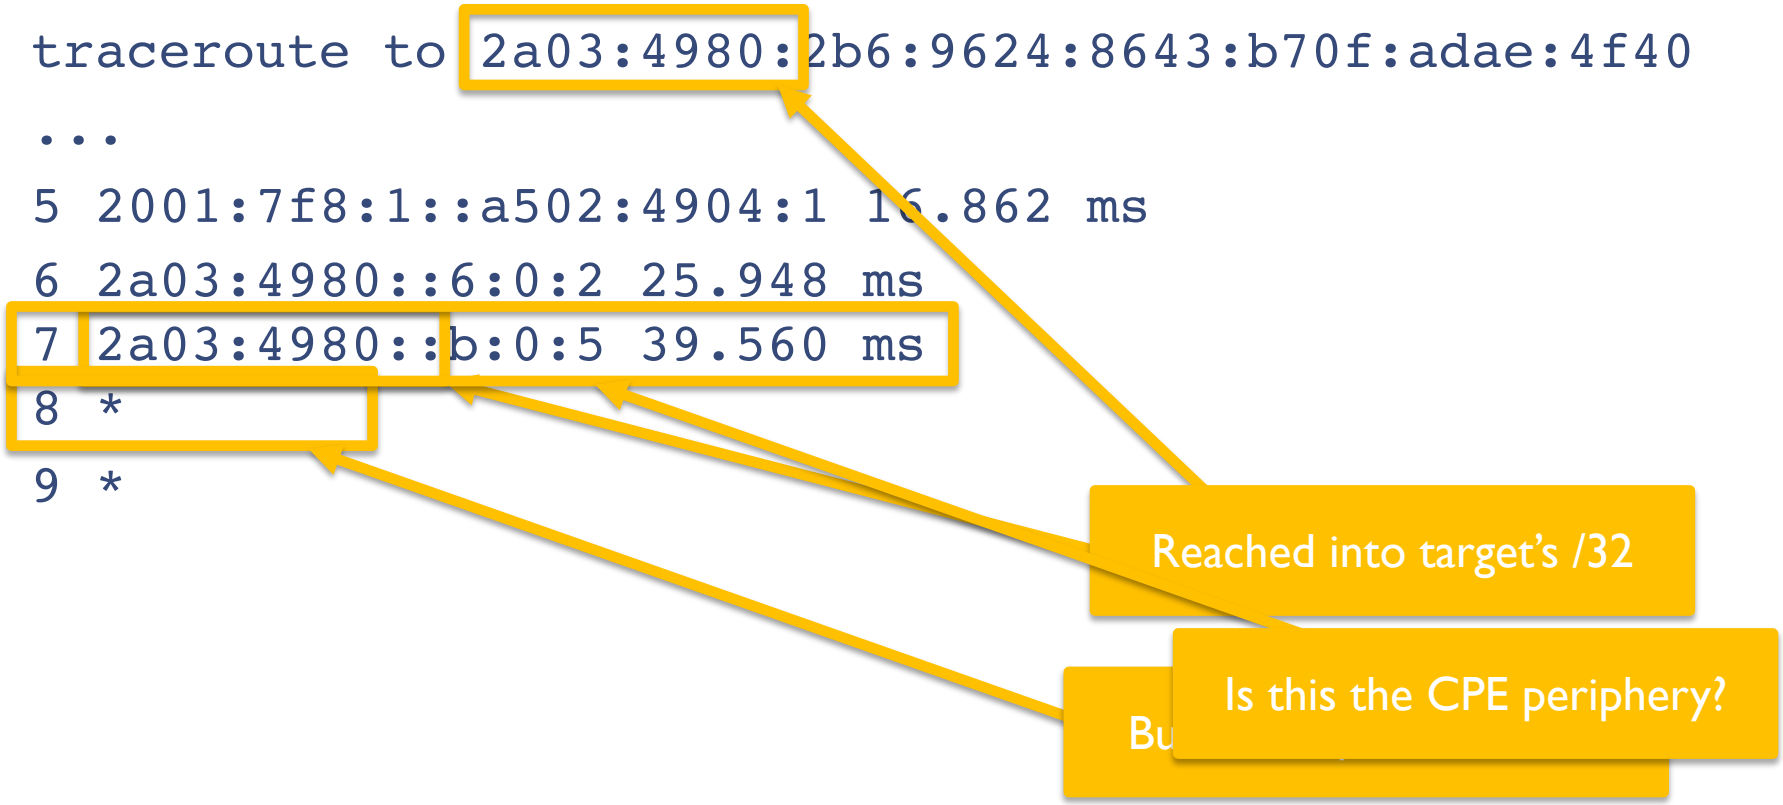

Contributions

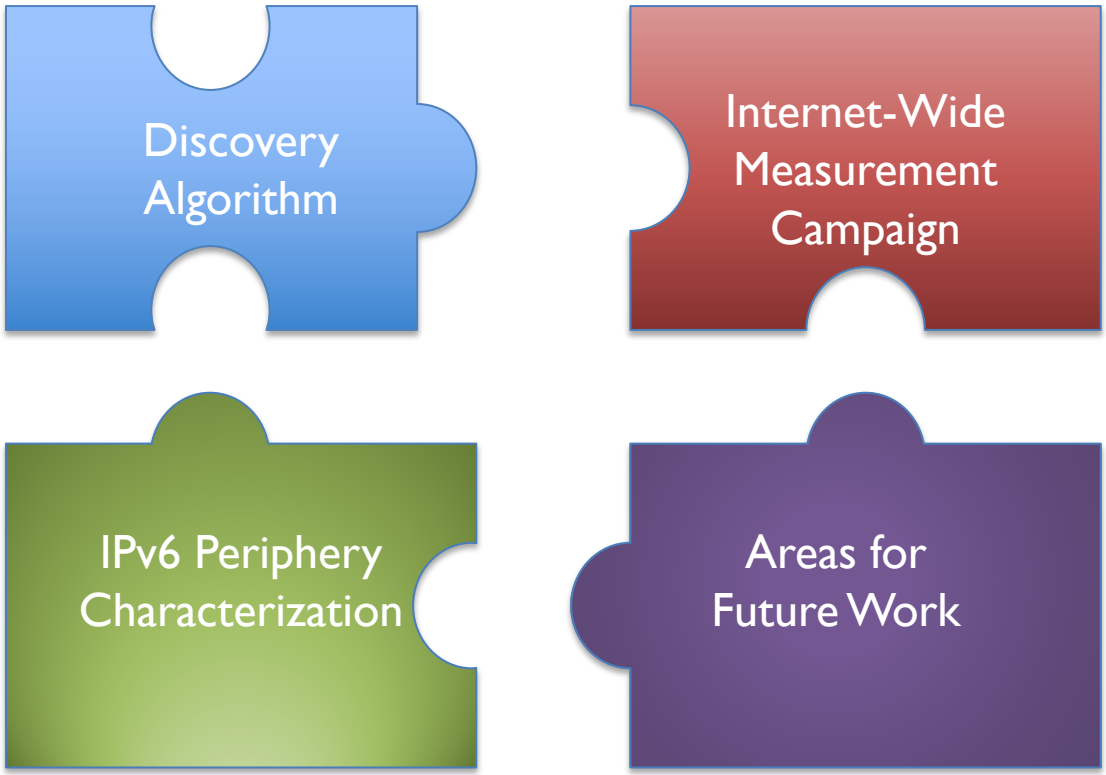

Contributions

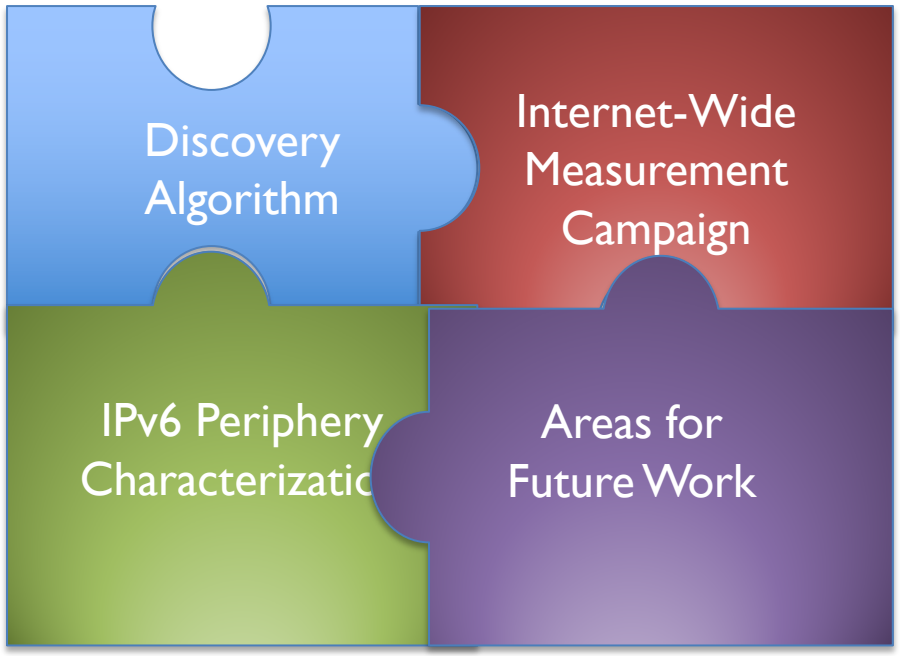

- More complete IPv6 topologies for:
  - Tracking adoption
  - Census
  - Reliability
  - Outages
  - Security

Discovery Algorithm: Edgy

- Two phases:
  - Initialization: find “interesting” /48s
  - Discovery: iteratively decompose /48 to find periphery

Edgy: Initialization

- Examine previous traceroute campaigns:
  - **BGP-Informed seed**
    - CAIDA trace to every routed /48, Aug 2018
  - **Hitlist-Informed seed**
    - Traces to targets in IPv6 hitlist
    - “IP of the Beholder”, IMC 2018
- Find “interesting” target /48 prefixes:
  - Last hops unique to one /48

Initialization: Scenario 1

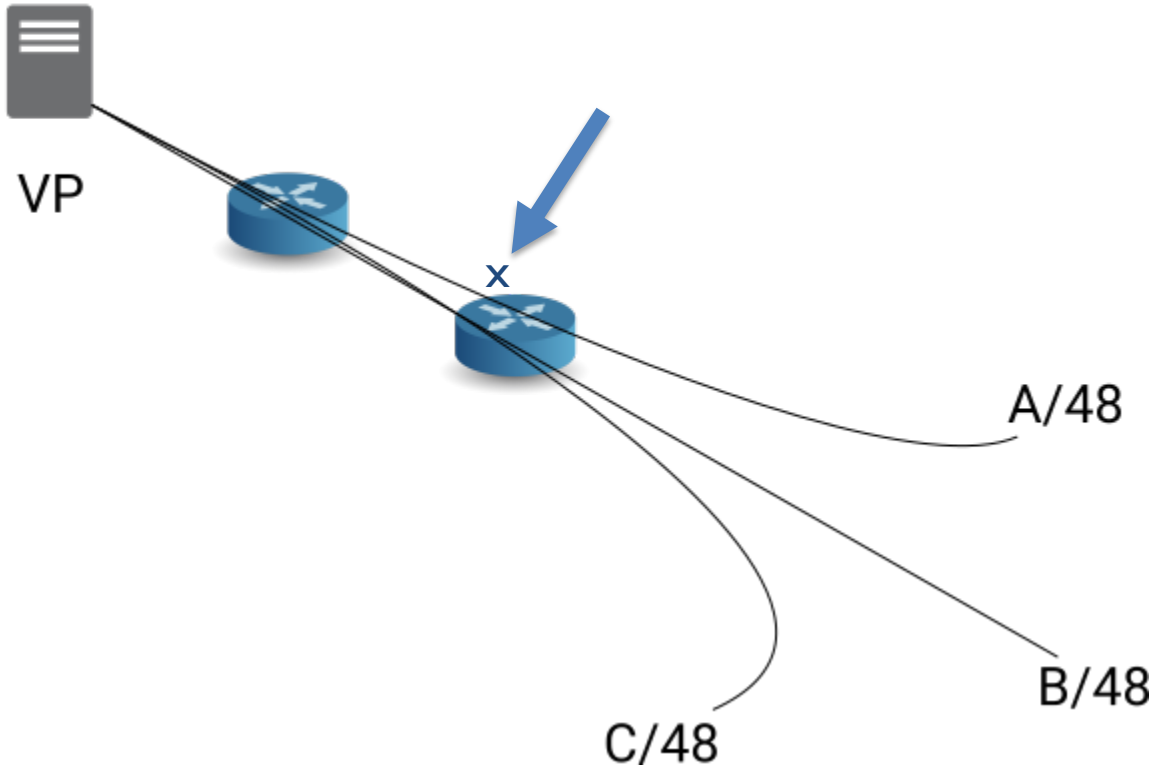

Traces to three different /48s share last hop address

- ICMP6 filtering
- A, B, and C in same subnet

Initialization: Scenario 2

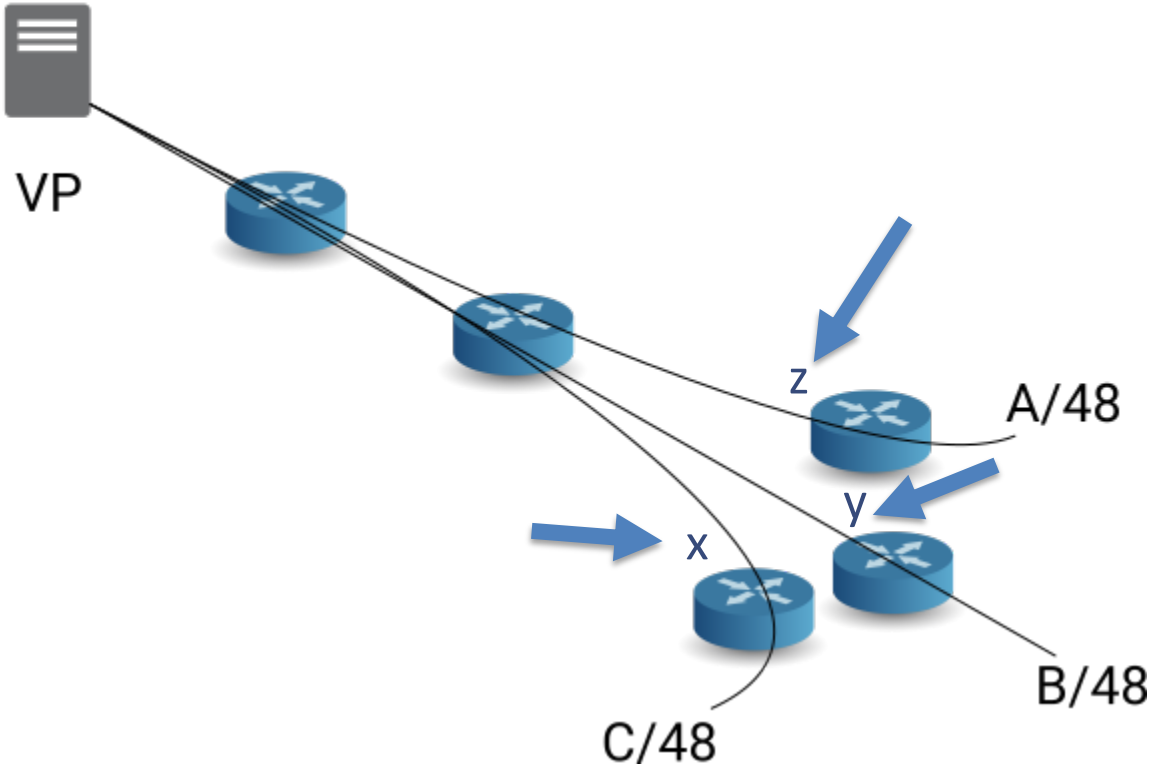

Last hops x, y, z appear only in traces to single /48s. These /48s become target prefixes for discovery probing

Edgy: Discovery Phase

- Probe target /48 prefixes in rounds with yarrp
  - Each with different probe granularity
    - All /56, /60, /62, and /64 subnets of target /48
- Continuation threshold
  - Number of new addresses > n
- Intuition – Probe prefixes that produce new periphery addresses at progressively finer granularities

Discovery Algorithm: Edgy

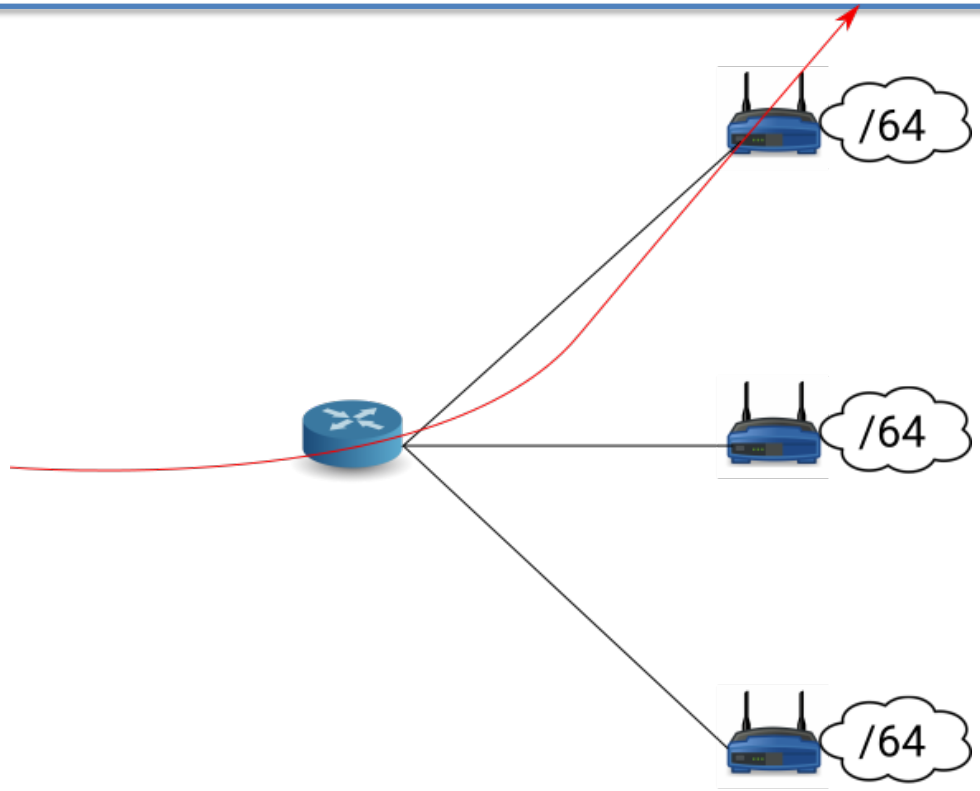

Coarse-grained discovery finds some periphery topology, but misses significant portions if small subnets allocated

Discovery Algorithm: Edgy

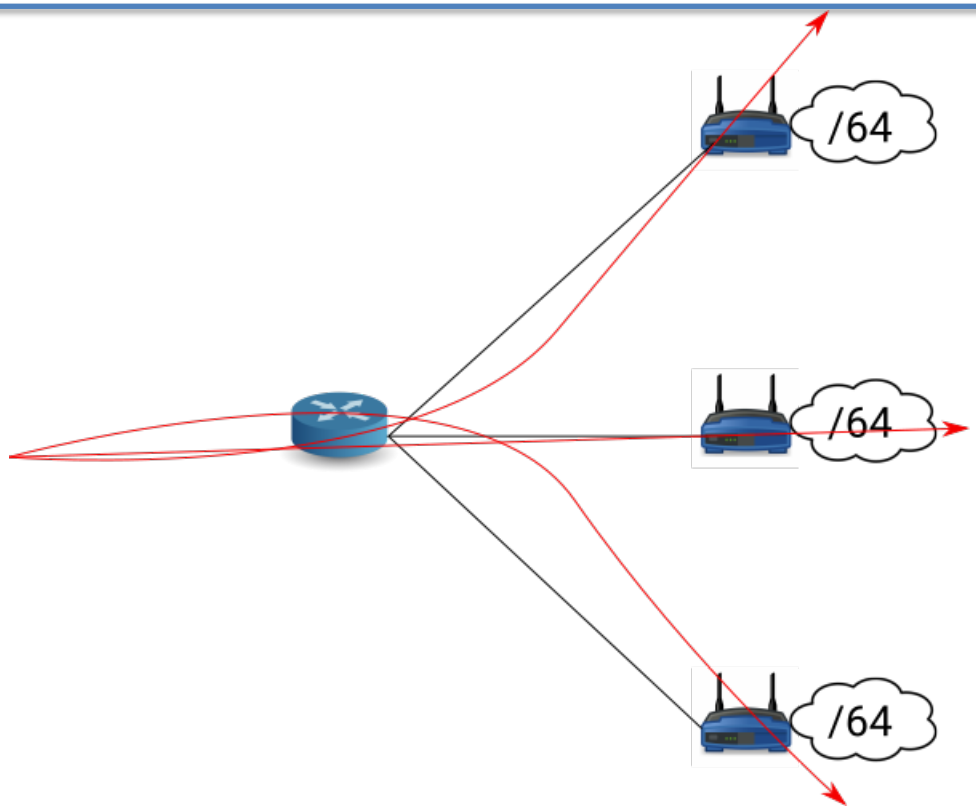

/48 prefixes that pass discovery thresholds are reprobed
at progressively finer granularities, uncovering more periphery structure

Measurement Campaign

- Sept – Oct 2019
- Probed 130k (BGP-Informed) and 111k (Hitlist-Informed) /48 prefixes
  - Single VP in Lausanne, Switzerland
- Followed ethical probing best practices
  - Received no opt-out requests
- Discover ~64M unique router interface addresses
- Nearly entirely disjoint from input seed
- Results from two different seeds largely disjoint
  - Edgy discovers new topology
  - Different seeds discover different new topology

Periphery Characterization

IPv6 Periphery Characterization

	BGP-Informed				Hitlist-Informed			
Round	Prefixes Probed	Unique Last Hops	Unique Last Hop /48s	Cum. Unique Last Hops	Prefixes Probed	Unique Last Hops	Unique Last Hop /48s	Cum. Unique Last Hops
1 (/56)	130,447	4,619,692	33,831	4,619,692	111,670	9,217,137	89,268	9,217,137
2 (/60)	34,520	12,228,916	26,082	13,410,601	67,107	11,021,329	74,302	11,365,910
3 (/62)	12,014	14,770,061	11,675	24,832,391	4,462	5,428,992	19,942	15,569,221
4 (/64)	2,641	15,326,298	7,833	37,169,357	1,531	15,340,591	32,718	29,248,703

Begin with ~same # prefixes

~25% vs ~60% pass threshold

Round / probe granularity

Only 1.5M in set intersection

Last hop addressing characteristics differ

EUI-64

IPv6 Periphery Characterization

- EUI-64 addresses are *still* pervasive
  - RFC4941 Privacy Extensions for SLAAC published in 2007
- 30M EUI-64 addresses seen (~50% total discovered)
- 16M unique MAC addresses (prefix cycling in select providers)

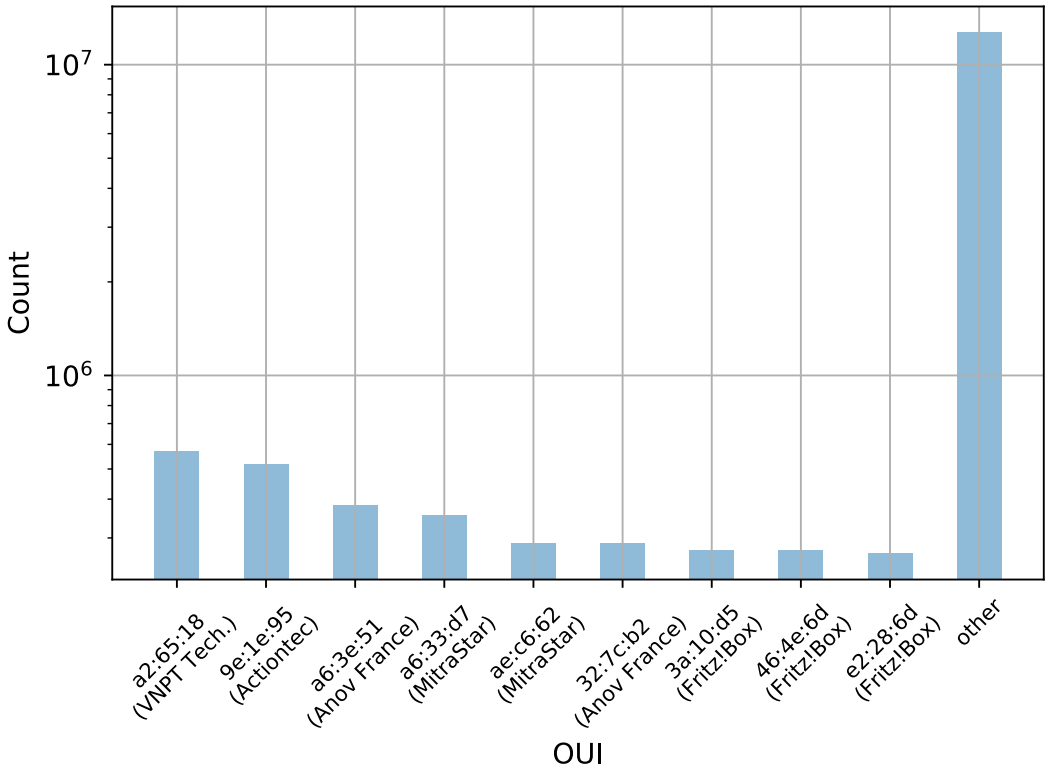

IID Entropy

IPv6 Periphery Characterization

- **Edgy-Discovered Addresses**
  - **Higher entropy IIDs**
    - (BGP-, Hitlist-Informed seed plot lines)
  - EUI-64 SLAAC, SLAAC w/P.E.
  - Suggests periphery (eg CPE, unmanaged devices)
  - Ex: 429b:cdff:fe1e:c5e0, 8871:14ad:4cf4:50a2
- **Previous Studies**
  - **Lower entropy IIDs**
  - Often manually assigned
  - Easy to recall
  - Suggests managed devices (eg provider infrastructure, servers)
  - Ex: ::1, ::beef

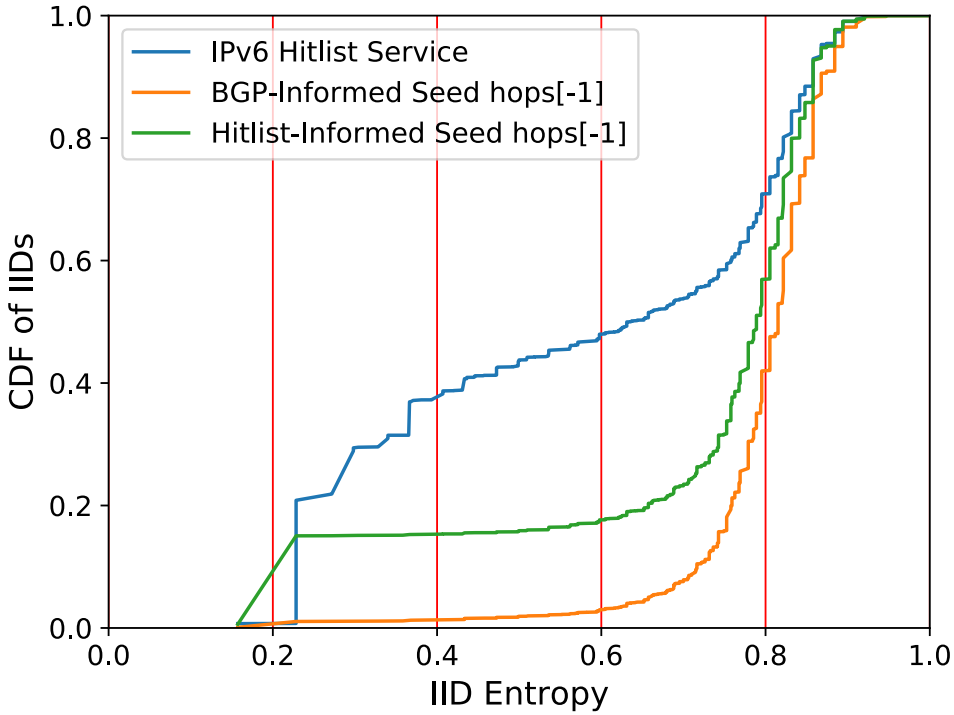

Edgy/Ark Comparison

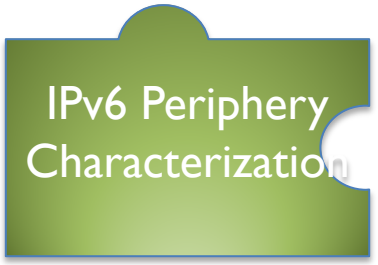

- Edgy traces reach destination more often, and farther into destination prefix
  - But, 2 orders of magnitude more probes, so not directly comparable
- Day of CAIDA Ark IPv6 traces vs edgy results
  - **40% Ark traces vs 87% edgy reach target AS**
- Median common bitmask length between target and last hop address:
  - **Ark – /13**
  - **Edgy - /32**

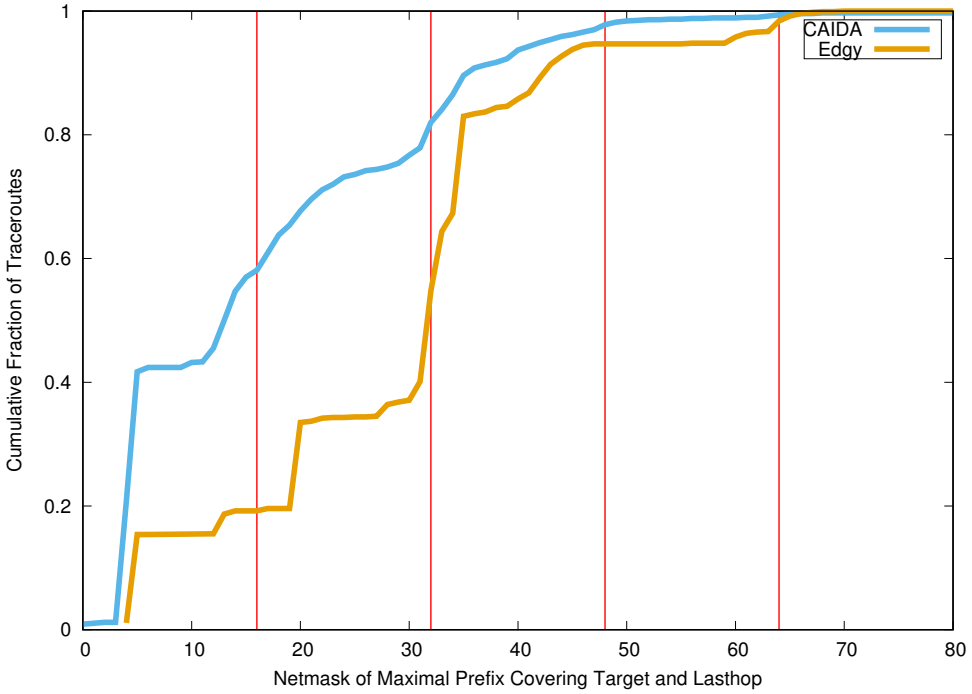

Pathologies: Prefix Cycling

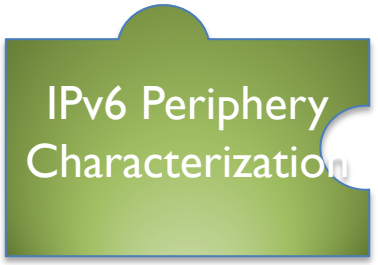

- Observe high frequency prefix cycling in some providers
  - 1und1.net (Versatel), Vietnam Posts and Telecommunications Group (VNPT)
  - ~24 hour lifetime before new prefix issued
  - Track EUI-64 addresses across prefix rotations

A week in the life of a MAC address

IPv6 Periphery Characterization

Multiple addresses seen (in single day)

Varies each day

/32

Lower 3 bytes anonymized

1 Feb 2020	2001:16b8	0100:10b3:	3a10:d5ff:feaa:bbcc
1 Feb 2020	2001:16b8	0101:c256:	3a10:d5ff:feaa:bbcc
2 Feb 2020	2001:16b8	0101:c256:	3a10:d5ff:feaa:bbcc
2 Feb 2020	2001:16b8	0103:74fe:	3a10:d5ff:feaa:bbcc
3 Feb 2020	2001:16b8	0101:1f20:	3a10:d5ff:feaa:bbcc
4 Feb 2020	2001:16b8	0102:d3c4:	3a10:d5ff:feaa:bbcc
5 Feb 2020	2001:16b8	0102:d3c4:	3a10:d5ff:feaa:bbcc
5 Feb 2020	2001:16b8	0100:98a5:	3a10:d5ff:feaa:bbcc
6 Feb 2020	2001:16b8	0100:98a5:	3a10:d5ff:feaa:bbcc
6 Feb 2020	2001:16b8	0102:5360:	3a10:d5ff:feaa:bbcc
7 Feb 2020	2001:16b8	0100:0cac:	3a10:d5ff:feaa:bbcc

Address carries over between days

All within same /46

Pathologies: MAC reuse

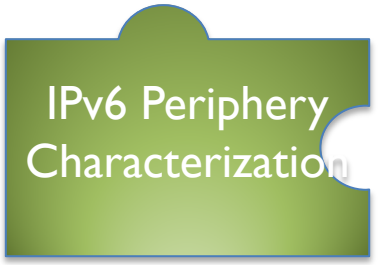

- Of 16M unique MAC addresses in EUI-64 IPv6 addresses,
  - 12.5M only observed once
  - 2.8M observed less than 10 times
    - Likely prefix rotation during study
  - 66 seen more than 1000 times
- 58:02:03:04:05:06 > **750,000 times!**
  - Observed in the LTE WAN interface IPv6 address on Huawei hotspots
  - Maybe others?
- f0:7d:68:15:a2:a2 > **186,000 times!**
  - D-Link address, but unclear what
  - Maybe another default LTE interface address?

Provider Allocation Policies

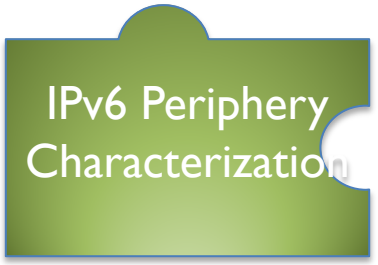

- Edgy sends probes into customer subnets
- Based on last hop responsive addresses, can:
  - Infer how providers allocate subnets to customers
    - Size, eg /48, /52, smaller
    - Uniform vs non-uniform allocations

- Use edgy results to visualize three distinct deployments
  - Uniform /56s
  - "Binary Tree" allocation
  - Uniform /64s

Uniform /56 Allocation

IPv6 Periphery Characterization

- Send probe to random IID in each /64 of a /48
- Plot target /48
  - y-axis: 7th byte of IPv6 address
  - x-axis: 8th byte of IPv6 address
  - Each color represents different responsive address

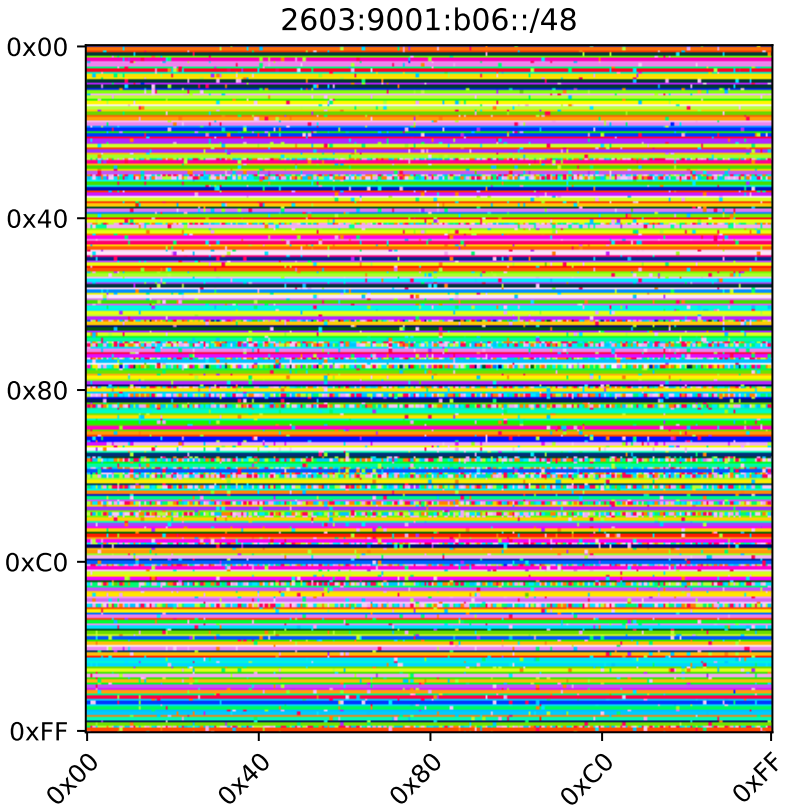

Charter Communications /48 divided evenly into 256 /56s

Binary Tree Allocation

IPv6 Periphery Characterization

- Send probe to random IID in each /64 of a /48
- Plot target /48
  - y-axis: 7th byte of IPv6 address
  - x-axis: 8th byte of IPv6 address
  - Each color represents different responsive address

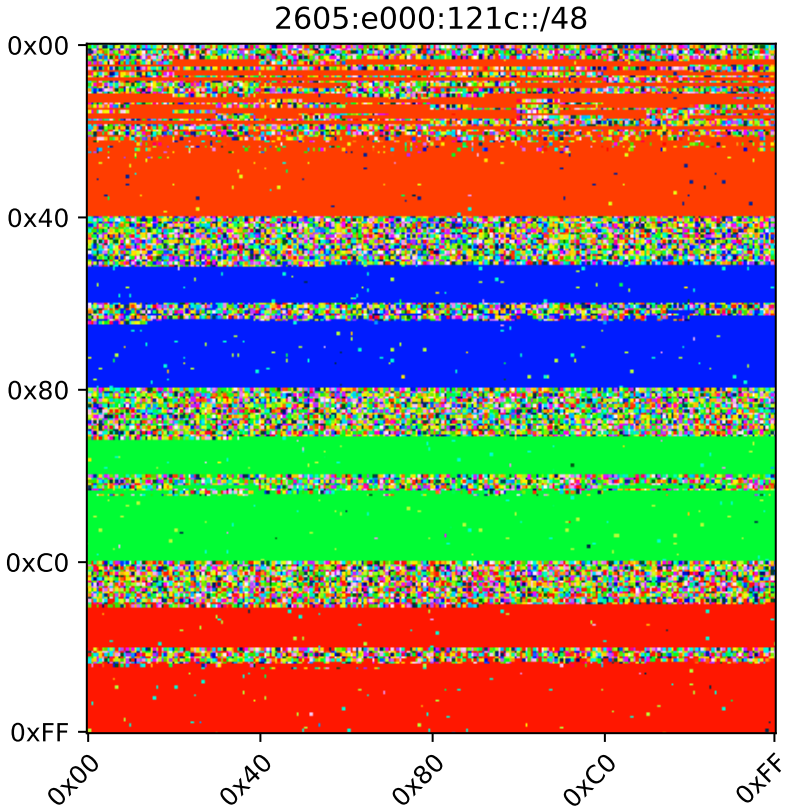

Time Warner /48 split into 4 /52s, which are then split into /64s for customers. Banding pattern suggests a binary tree approach. Significant portions of each /52 remain unallocated.

Uniform /64 Allocation

IPv6 Periphery Characterization

- Send probe to random IID in each /64 of a /48
- Plot target /48
  - y-axis: 7th byte of IPv6 address
  - x-axis: 8th byte of IPv6 address
  - Each color represents different responsive address

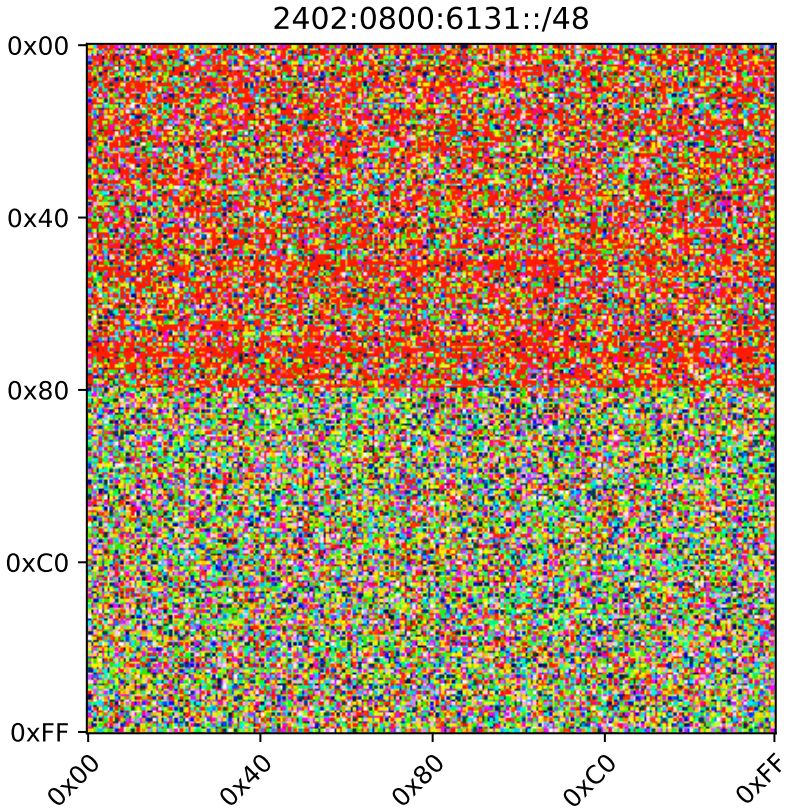

Viettel Group (VN) /48 split into two /49s, which are then split into /64s for customers. Majority of the /48 is subnetted into /64s.

Future Work

Areas for Future Work

- Longitudinal study of prefix cycling
  - Can we predict/quantify:
    - Exactly when prefixes change?
    - The next prefix for an IID?
    - How addresses move in relation to one another?
- Couple edgy discovery with other measurements
  - ICMPv6 Echoes, banner grabs

Conclusions

- Introduce edgy, a technique to discover IPv6 periphery
  - Probe prefixes at increasingly finer granularities while address discovery meets threshold
  - More of the IPv6 periphery is discoverable than previously mapped
  - Step toward more complete IPv6 topology mapping
- Deeper insights into the IPv6 periphery:
  - Prefix cycling, EUI-64s, MAC reuse
  - Per-provider allocations and deployment

Thanks!

Backup

Address ASN Distribution

IPv6 Periphery Characterization

- 5,109 unique ASNs
- Well-known providers contribute significant #s of addresses to total
  - 1und1.net (8881)
  - Deutsche Telekom (3220)
  - VNPT (45899)
  - Sky (5607)
  - Cox (22773)
- Provider prefix churn dynamics inflate totals of some ISPs
  - In particular, 8881 and 45899

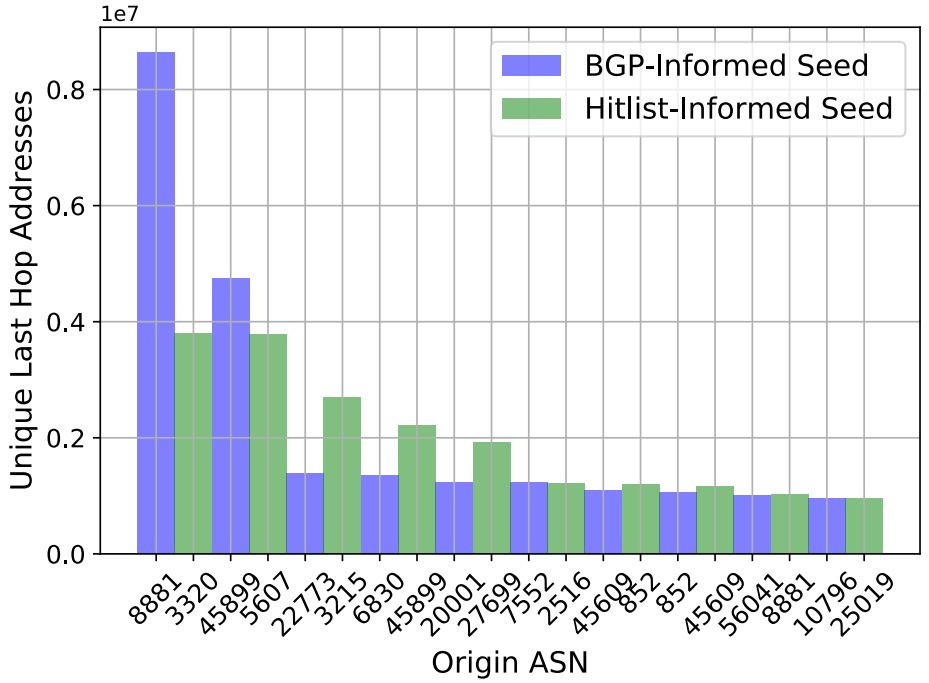

Address Country Distribution

IPv6 Periphery Characterization

- 153 countries represented
- Distribution of countries uneven between seed data sources
  - US second in BGP-Informed, but 14th in Hitlist-Informed
- Again, prefix cycling over-represents some countries
  - BGP-informed DE and VN, especially