Leveraging Internet Background Radiation for Opportunistic Network Analysis

Karyn Benson,
Alberto Dainotti, kc claffy, Alex C. Snoeren, Michael Kallitsis

Quiz: How do you measure _____

Quiz: How do you measure _____

- The number of open resolvers on the Internet?

- The uptime of many machines?

- How often routes to a destination change?

Quiz: How do you measure _____

- The number of open resolvers on the Internet?
  - With specially crafted DNS packets [Schomp *et al.* IMC ’13]
- The uptime of many machines?
  - By connecting to port 80 every hour [Huang *et al*. IMC ’08]
- How often routes to a destination change?
  - Through analysis of publicly available BGP messages [Rexford *et al.* IMW ’02]

Quiz: How do you measure _____

- The number of open resolvers on the Internet?
    - With Internet Background Radiation
- The uptime of many machines?
    - With Internet Background Radiation
- How often routes to a destination change?
    - With Internet Background Radiation

Quiz: How do you measure _____

- The number of open resolvers on the Internet?
    - With Internet Background Radiation
- The uptime of many machines?
    - With Internet Background Radiation
- How often routes to a destination change?
    - With Internet Background Radiation

Goal: understand when IBR is an appropriate datasource for Internet-wide measurement

Internet Background Radiation (IBR)

How is IBR applicable to Internet-wide measurement?

Relevant properties and limitations of IBR

What can we expect in terms of coverage, traffic components, and repeated contact?

Case studies

When is using IBR beneficial?

Internet Background Radiation (IBR)

How is IBR applicable to Internet-wide measurement?

Relevant properties and limitations of IBR

What can we expect in terms of coverage, traffic components, and repeated contact?

Case studies

When is using IBR beneficial?

Internet Background Radiation (IBR)

- Many hosts send unsolicited traffic, called Internet Background Radiation
- Also known as network telescopes, darknets are a collection of unused IP addresses that capture unsolicited traffic

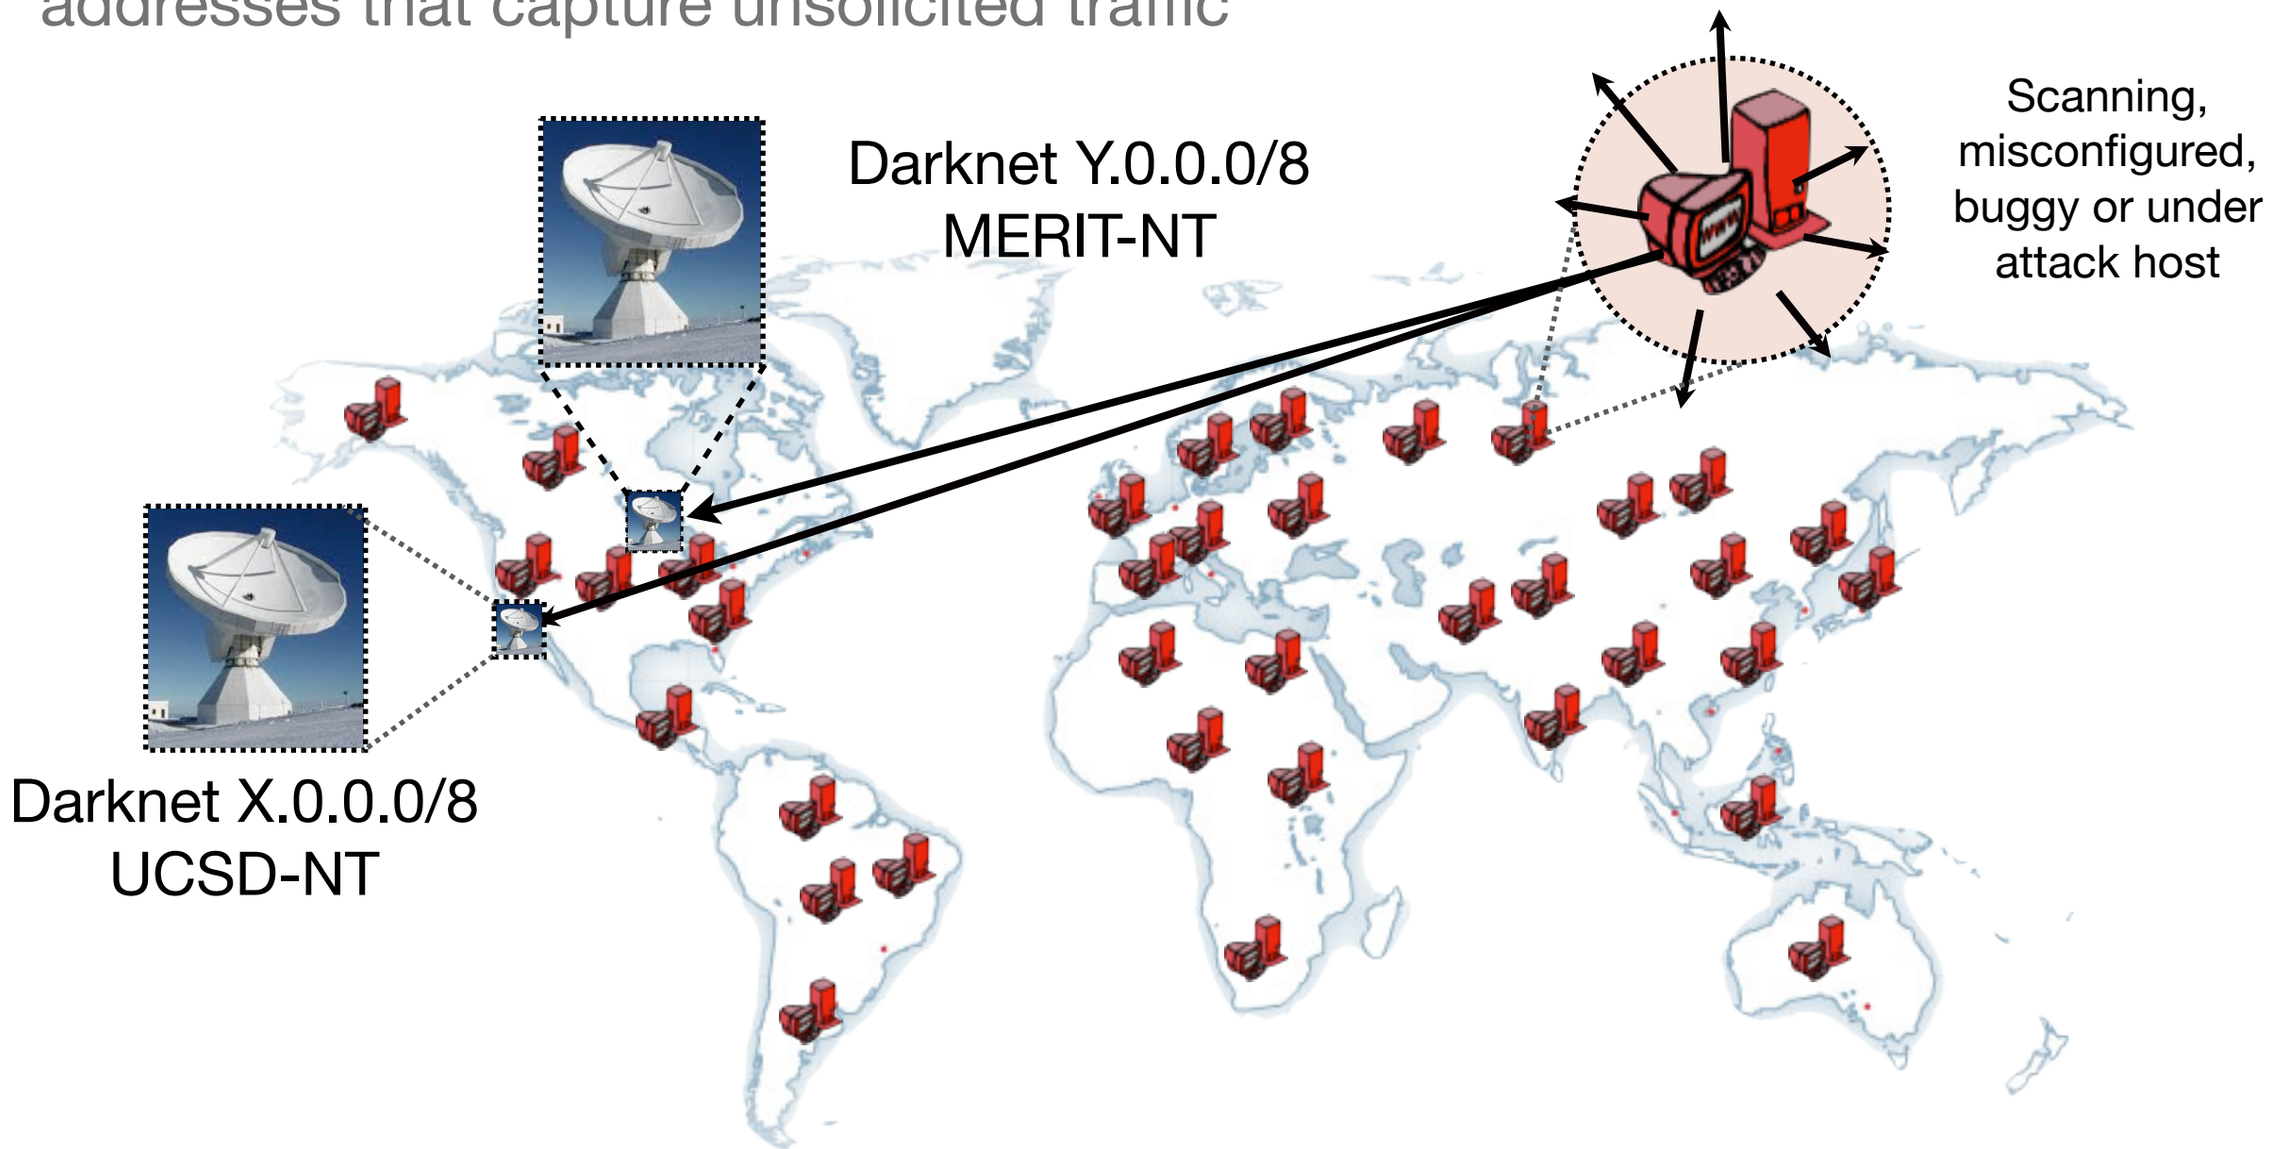

Pros
of using IBR for measurement data

PROS

- **P**ersistent
- **P**ervasive:
  - From sources Internet-wide
  - From many different types of sources
- Few **p**rivacy concerns

Pros and Cons of using IBR for measurement data

PROS

- **P**ersistent
- **P**ervasive:
  - From sources Internet-wide
  - From many different types of sources
- Few **p**rivacy concerns

CONS

- No **c**ontrol over who sends and when
  - Mix of traffic **c**hanges frequently
- Unidirectional data source
- **C**leaning: Need to remove spoofed traffic

IBR Datasets

IBR Datasets

- Pcap data from 34 days (around July)
  - **UCSD-12**
  - **UCSD-13**
  - **MERIT-13**

IBR Datasets

- Pcap data from 34 days (around July)
  - **UCSD**-12
  - **UCSD**-13
  - **MERIT**-13

Where Collected

IBR Datasets

- Pcap data from 34 days (around July)
  - **UCSD-12**
  - **UCSD-13**
  - **MERIT-13**

Where Collected

Year Collected

IBR Datasets

- Pcap data from 34 days (around July)
  - **UCSD-12**
  - **UCSD-13**
  - **MERIT-13**

Where Collected

Year Collected

- 7+ years of packet-header information

IBR Datasets

- Pcap data from 34 days (around July)
  - **UCSD-12**
  - **UCSD-13**
  - **MERIT-13**

Where Collected

Year Collected

- 7+ years of packet-header information
- Use method in [CCR ’14] to remove spoofed traffic

We can study background radiation *and* the networks generating IBR

Malicious Activity:

- Worm propagation [Moore *et al.* ’02, Moore *et al.* ’03, Kumar *et al.* ’05, Bailey *et al.* ’05, Aben ’08]
- DDoS targets [Moore *et al.* ’06]
- Scanning techniques [Dainotti *et al.* ’12, Durumeric *et al*. ’14]

We can study background radiation *and* the networks generating IBR

Malicious Activity:

- Worm propagation [Moore *et al.* ’02, Moore *et al.* ’03, Kumar *et al.* ’05, Bailey *et al.* ’05, Aben ’08]
- DDoS targets [Moore *et al.* ’06]
- Scanning techniques [Dainotti *et al.* ’12, Durumeric *et al*. ’14]

What can we learn about the hosts/ networks sending unsolicited traffic?

We can study background radiation *and* the networks generating IBR

Malicious Activity:

- Worm propagation [Moore *et al.* ’02, Moore *et al.* ’03, Kumar *et al.* ’05, Bailey *et al.* ’05, Aben ’08]
- DDoS targets [Moore *et al.* ’06]
- Scanning techniques [Dainotti *et al.* ’12, Durumeric *et al*. ’14]

Existing Examples:

- Uptime of Witty-infected machines [Kumar *et al.*’05]
- Country-wide outages [Dainotti *et al.*’11]
- Packet loss during BGP-leaks [Benson *et al.*’13]
- Filtering policy in 2011 [Sargent *et al.* ’15]

What can we learn about the hosts/ networks sending unsolicited traffic?

We can study background radiation *and* the networks generating IBR

Malicious Activity:

- Worm propagation [Moore *et al.* ’02, Moore *et al.* ’03, Kumar *et al.* ’05, Bailey *et al.* ’05, Aben ’08]
- DDoS targets [Moore *et al.* ’06]
- Scanning techniques [Dainotti *et al.* ’12, Durumeric *et al*. ’14]

Existing Examples:

- Uptime of Witty-infected machines [Kumar *et al.*’05]
- Country-wide outages [Dainotti *et al.*’11]
- Packet loss during BGP-leaks [Benson *et al.*’13]
- Filtering policy in 2011 [Sargent *et al.* ’15]

We can study background radiation *and* the networks generating IBR

Malicious Activity:

- Worm propagation [Moore *et al.* ’02, Moore *et al.* ’03, Kumar *et al.* ’05, Bailey *et al.* ’05, Aben ’08]
- DDoS targets [Moore *et al.* ’06]
- Scanning techniques [Dainotti *et al.* ’12, Durumeric *et al*. ’14]

Existing Examples:

- Uptime of Witty-infected machines [Kumar *et al.*’05]
- Country-wide outages [Dainotti *et al.*’11]
- Packet loss during BGP-leaks [Benson *et al.*’13]
- Filtering policy in 2011 [Sargent *et al.* ’15]

Is the analysis repeatable for other networks? other types of traffic? other time periods?

We can study background radiation *and* the networks generating IBR

Malicious Activity:

- Worm propagation [Moore *et al.* ’02, Moore *et al.* ’03, Kumar *et al.* ’05, Bailey *et al.* ’05, Aben ’08]
- DDoS targets [Moore *et al.* ’06]
- Scanning techniques [Dainotti *et al.* ’12, Durumeric *et al*. ’14]

Existing Examples:

- Uptime of Witty-infected machines [Kumar *et al.*’05]
- Country-wide outages [Dainotti *et al.*’11]
- Packet loss during BGP-leaks [Benson *et al.*’13]
- Filtering policy in 2011 [Sargent *et al.* ’15]

This Paper:

- What properties of IBR make it amenable to Internet-wide measurement?
- When should we use IBR to learn about the Internet?

Internet Background Radiation (IBR)

How is IBR applicable to Internet-wide measurement?

Relevant properties and limitations of IBR

What can we expect in terms of coverage, traffic components, and repeated contact?

Case studies

When is using IBR beneficial?

IBR-based inferences

IBR-based inferences

	IP
One observation	Ascertaining IPv4 Space Utilization

IBR-based inferences

	IP	TCP/UDP	Application
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers

IBR-based inferences

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage		

IBR-based inferences

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages			

IBR-based inferences

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages	Recognizing Packet loss	Number of Disk (Witty)	

IBR-based inferences

For a given analysis technique:
success, or coverage, depends on the traffic type and number of required observations.

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages	Recognizing Packet loss	Number of Disk (Witty)	

IBR-based inferences

For a given analysis technique:
success, or coverage, depends on the traffic type and number of required observations.

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages	Recognizing Packet loss	Number of Disk (Witty)	

IBR-based inferences

For a given analysis technique:
success, or **coverage**, depends on the traffic type and number of required observations.

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages	Recognizing Packet loss	Number of Disk (Witty)	

IBR-based inferences

For a given analysis technique:
success, or **coverage**, depends on the **traffic type** and number of required observations.

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages	Recognizing Packet loss	Number of Disk (Witty)	

IBR-based inferences

For a given analysis technique:
success, or **coverage**, depends on the **traffic type** and **number of** required **observations**.

	IP	TCP/UDP	Application	
One observation	Ascertaining IPv4 Space Utilization	Discovering Services	Locating Open Resolvers	Inferring Filtering Policy
Two observations	Identifying Path Changes	Determining Host Uptime	Evaluating Security Improvements	
Many observations	Deducing Packet Sending Rate	Detecting NAT Usage	Assessing BitTorrent Client Popularity	
Predictable observations	Detecting Outages	Recognizing Packet loss	Number of Disk (Witty)	

IBR-based inferences

coverage

traffic type

number of observations

Dimensions relevant to Internet-wide network analysis

coverage

traffic type

number of observations

Desire: to observe traffic from many diverse sources

coverage

traffic type

number of observations

Desire: to observe traffic from many diverse sources

coverage

traffic type

number of observations

	Percent BGP Announced	Total UCSD-13
IP addresses	5%	133M
/24 blocks	30%	3.15M
Prefixes	45%	205k
ASes	54%	24.2k
Countries	99%	233

- We observe traffic from most countries, large ASes

Desire: to have a rich signal composed of many different types of traffic

coverage

traffic type

number of observations

Desire: to have a rich signal composed of many different types of traffic

coverage

traffic type

number of observations

Misconfigurations & Bugs	UCSD-13 /24s
P2P (BitTorrent, eDonkey, QQlive)	2.5M
Qihoo 360 Safe	1.4M
Encapsulated IPv6 (6in4, Teredo)	0.7M
Gaming (Xbox, Steam)	0.5M
Botnet C&C (Sality, ZeroAccess)	0.2M
Scanning	
Conficker	0.6M
Bro Scanner	0.2M
Backscatter	
Backscatter	0.4M
Unclassified	
Encrypted	1.5M
Other	1.9M
TOTAL	3.15M

- Many applications with significant number of sources
- Usefulness of traffic depends on the type of network analysis (e.g., packets vs sources)

Desire: to have a rich signal composed of many different types of traffic

coverage

traffic type

number of observations

Misconfigurations & Bugs	UCSD-13 /24s
P2P (BitTorrent, eDonkey, QQlive)	2.5M
Qihoo 360 Safe	1.4M
Encapsulated IPv6 (6in4, Teredo)	0.7M
Gaming (Xbox, Steam)	0.5M
Botnet C&C (Sality, ZeroAccess)	0.2M
Scanning	
Conficker	0.6M
Bro Scanner	0.2M
Backscatter	
Backscatter	0.4M
Unclassified	
Encrypted	1.5M
Other	1.9M
TOTAL	3.15M

- Many applications with significant number of sources
- Usefulness of traffic depends on the type of network analysis (e.g., packets vs sources)

Desire: to have a rich signal composed of many different types of traffic

coverage

traffic type

number of observations

Most /24 blocks

Misconfigurations & Bugs	UCSD-13 /24s
P2P (BitTorrent, eDonkey, QQlive)	2.5M
Qihoo 360 Safe	1.4M
Encapsulated IPv6 (6in4, Teredo)	0.7M
Gaming (Xbox, Steam)	0.5M
Botnet C&C (Sality, ZeroAccess)	0.2M
Scanning	
Conficker	0.6M
Bro Scanner	0.2M
Backscatter	
Backscatter	0.4M
Unclassified	
Encrypted	1.5M
Other	1.9M
TOTAL	3.15M

- Many applications with significant number of sources
- Usefulness of traffic depends on the type of network analysis (e.g., packets vs sources)

Desire: to have a rich signal composed of many different types of traffic

coverage

traffic type

number of observations

Misconfigurations & Bugs	UCSD-13 /24s
P2P (BitTorrent, eDonkey, QQlive)	2.5M
Qihoo 360 Safe	1.4M
Encapsulated IPv6 (6in4, Teredo)	0.7M
Gaming (Xbox, Steam)	0.5M
Botnet C&C (Sality, ZeroAccess)	0.2M
Scanning	
Conficker	0.6M
Bro Scanner	0.2M
Backscatter	
Backscatter	0.4M
Unclassified	
Encrypted	1.5M
Other	1.9M
TOTAL	3.15M

- Many applications with significant number of sources
- Usefulness of traffic depends on the type of network analysis (e.g., packets vs sources)

Desire: to have a rich signal composed of many different types of traffic

coverage

traffic type

number of observations

83% of packets

Misconfigurations & Bugs	UCSD-13 /24s
P2P (BitTorrent, eDonkey, QQlive)	2.5M
Qihoo 360 Safe	1.4M
Encapsulated IPv6 (6in4, Teredo)	0.7M
Gaming (Xbox, Steam)	0.5M
Botnet C&C (Sality, ZeroAccess)	0.2M
Scanning	
Conficker	0.6M
Bro Scanner	0.2M
Backscatter	
Backscatter	0.4M
Unclassified	
Encrypted	1.5M
Other	1.9M
TOTAL	3.15M

- Many applications with significant number of sources
- Usefulness of traffic depends on the type of network analysis (e.g., packets vs sources)

Desire: to be able to repeatedly infer properties of remote networks

coverage

traffic type

number of observations

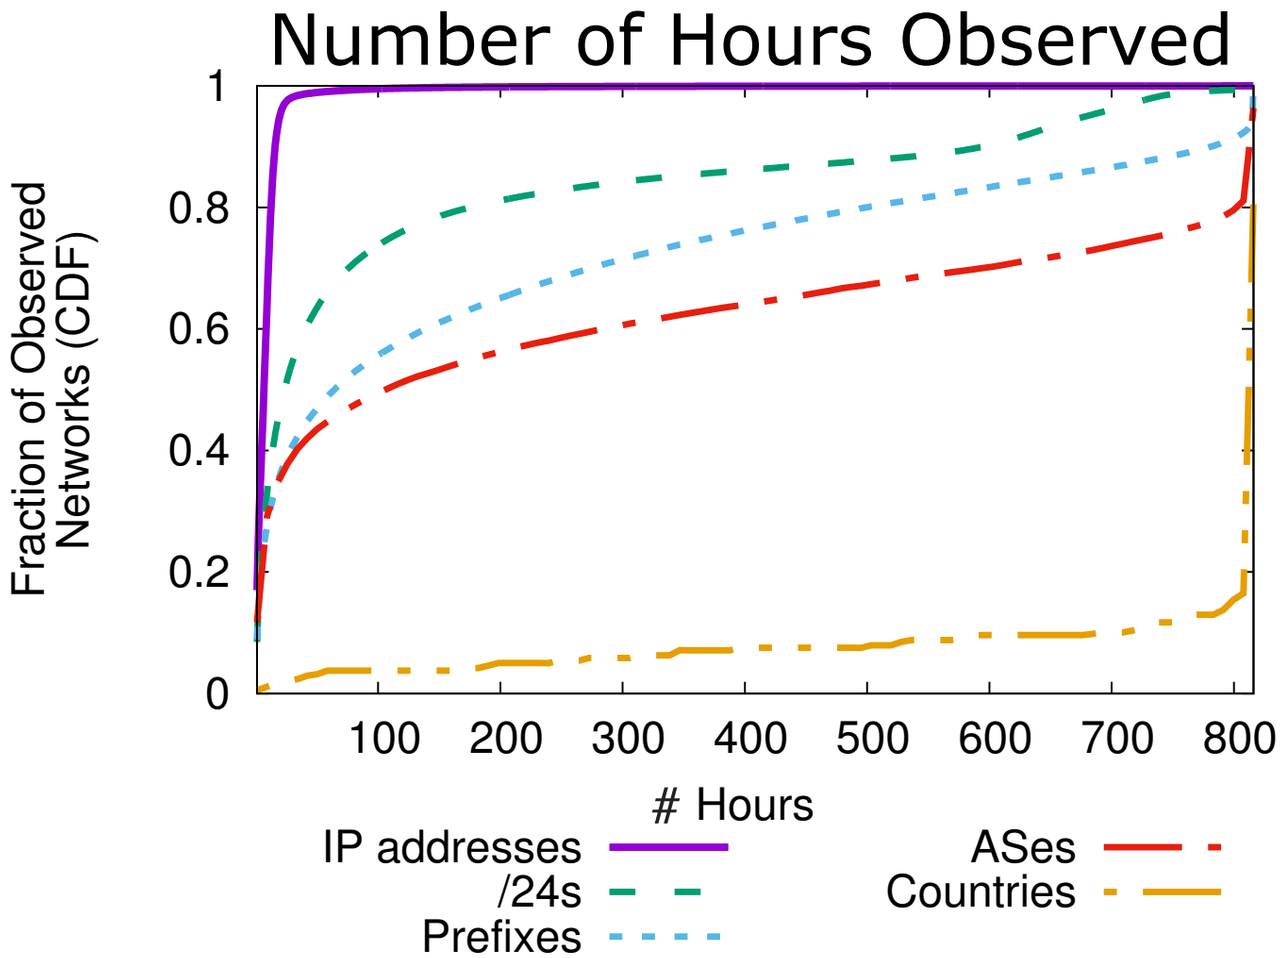

- Analyses requiring observations every hour: only possible for countries and some ASes

Desire: to be able to repeatedly infer properties of remote networks

coverage

traffic type

number of observations

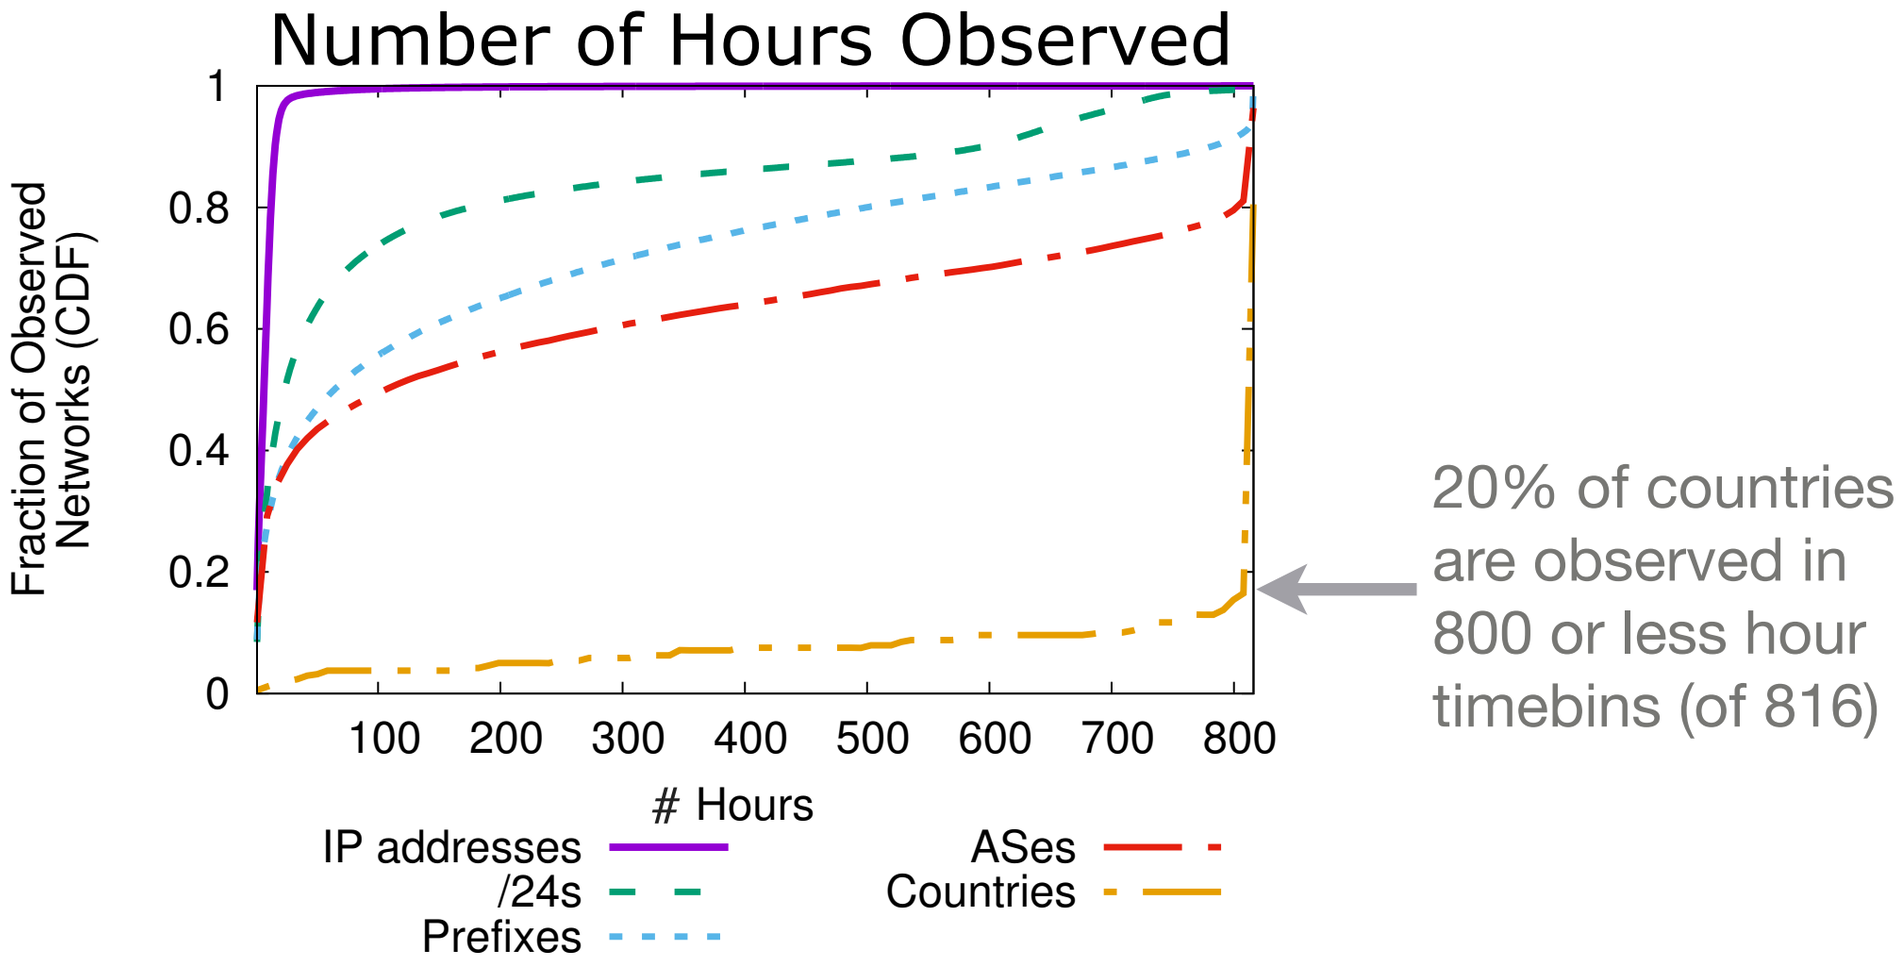

- Analyses requiring observations every hour: only possible for countries and some ASes

Desire: to be able to repeatedly infer properties of remote networks

coverage

traffic type

number of observations

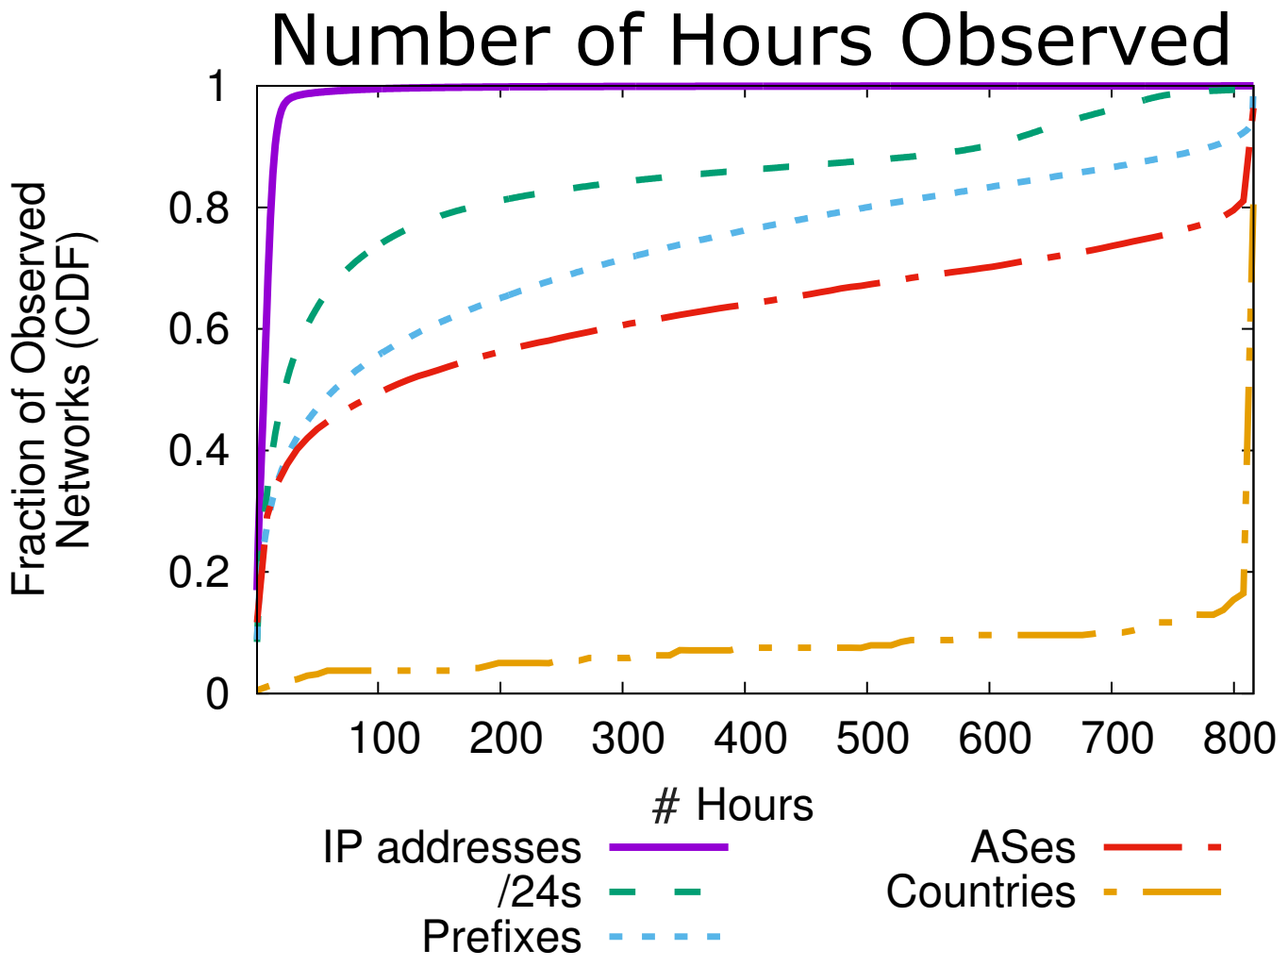

- Analyses requiring observations every hour: only possible for countries and some ASes

Desire: to be able to repeatedly infer properties of remote networks

coverage

traffic type

number of observations

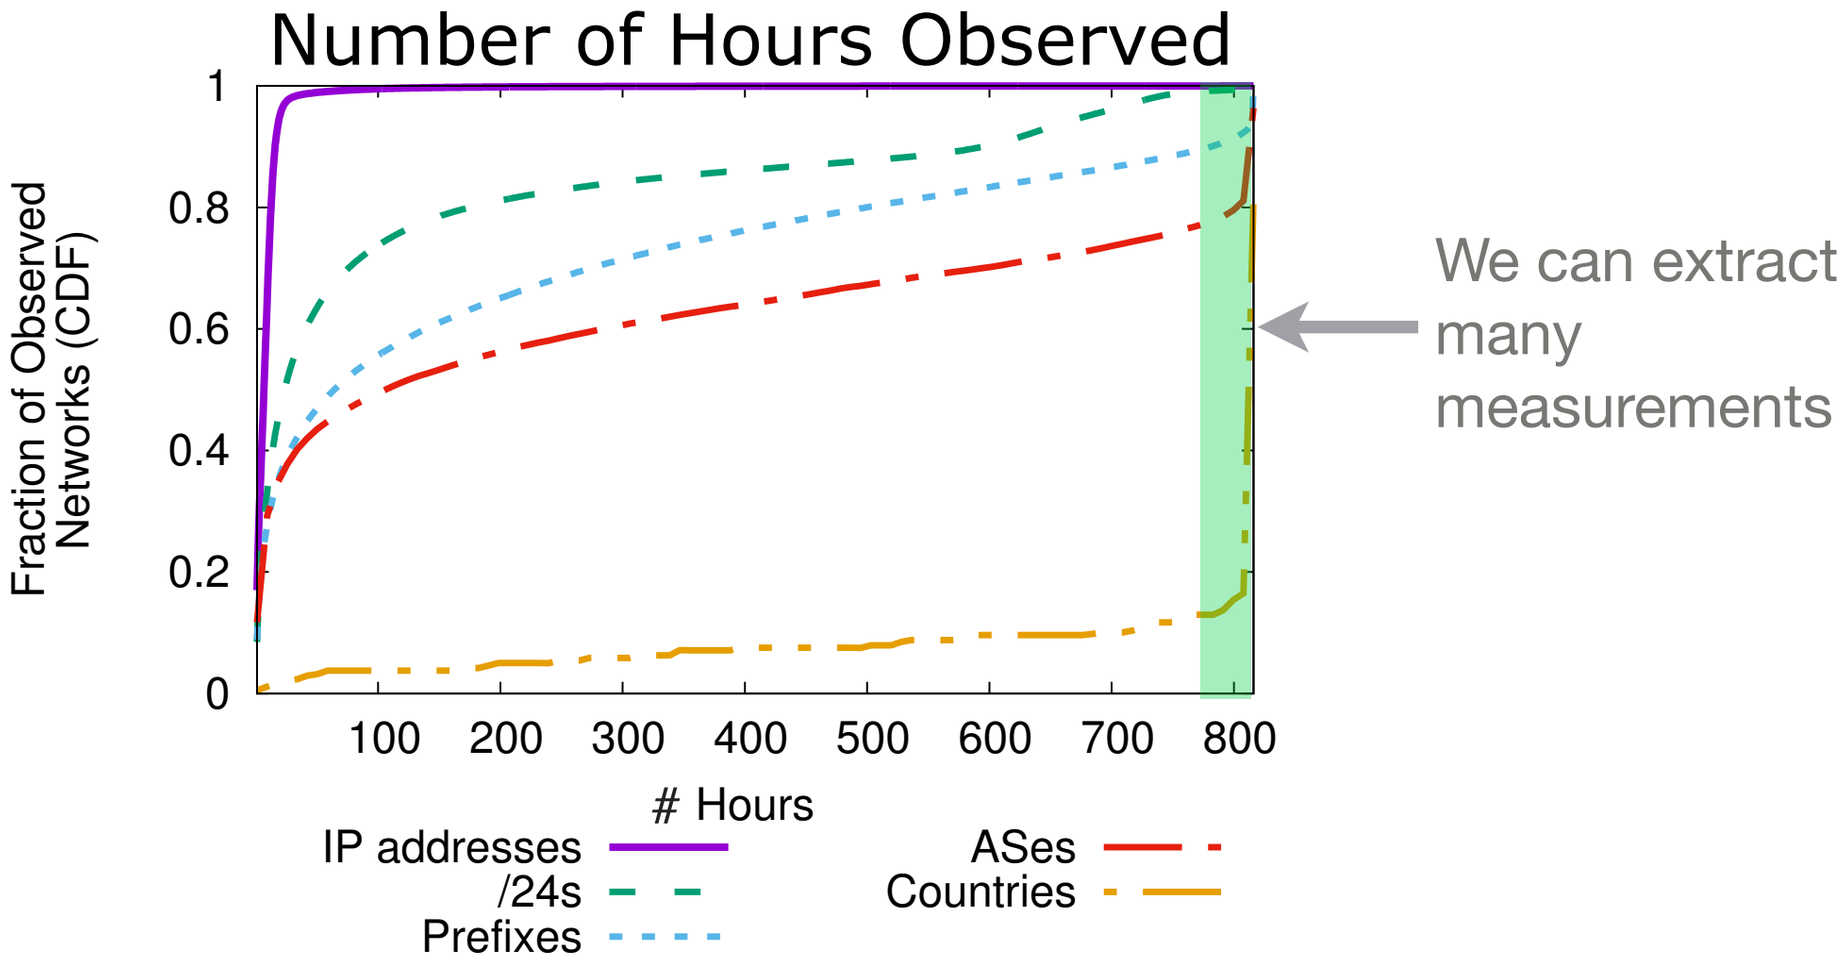

- Analyses requiring observations every hour: only possible for countries and some ASes

Desire: to be able to repeatedly infer properties of remote networks

coverage

traffic type

number of observations

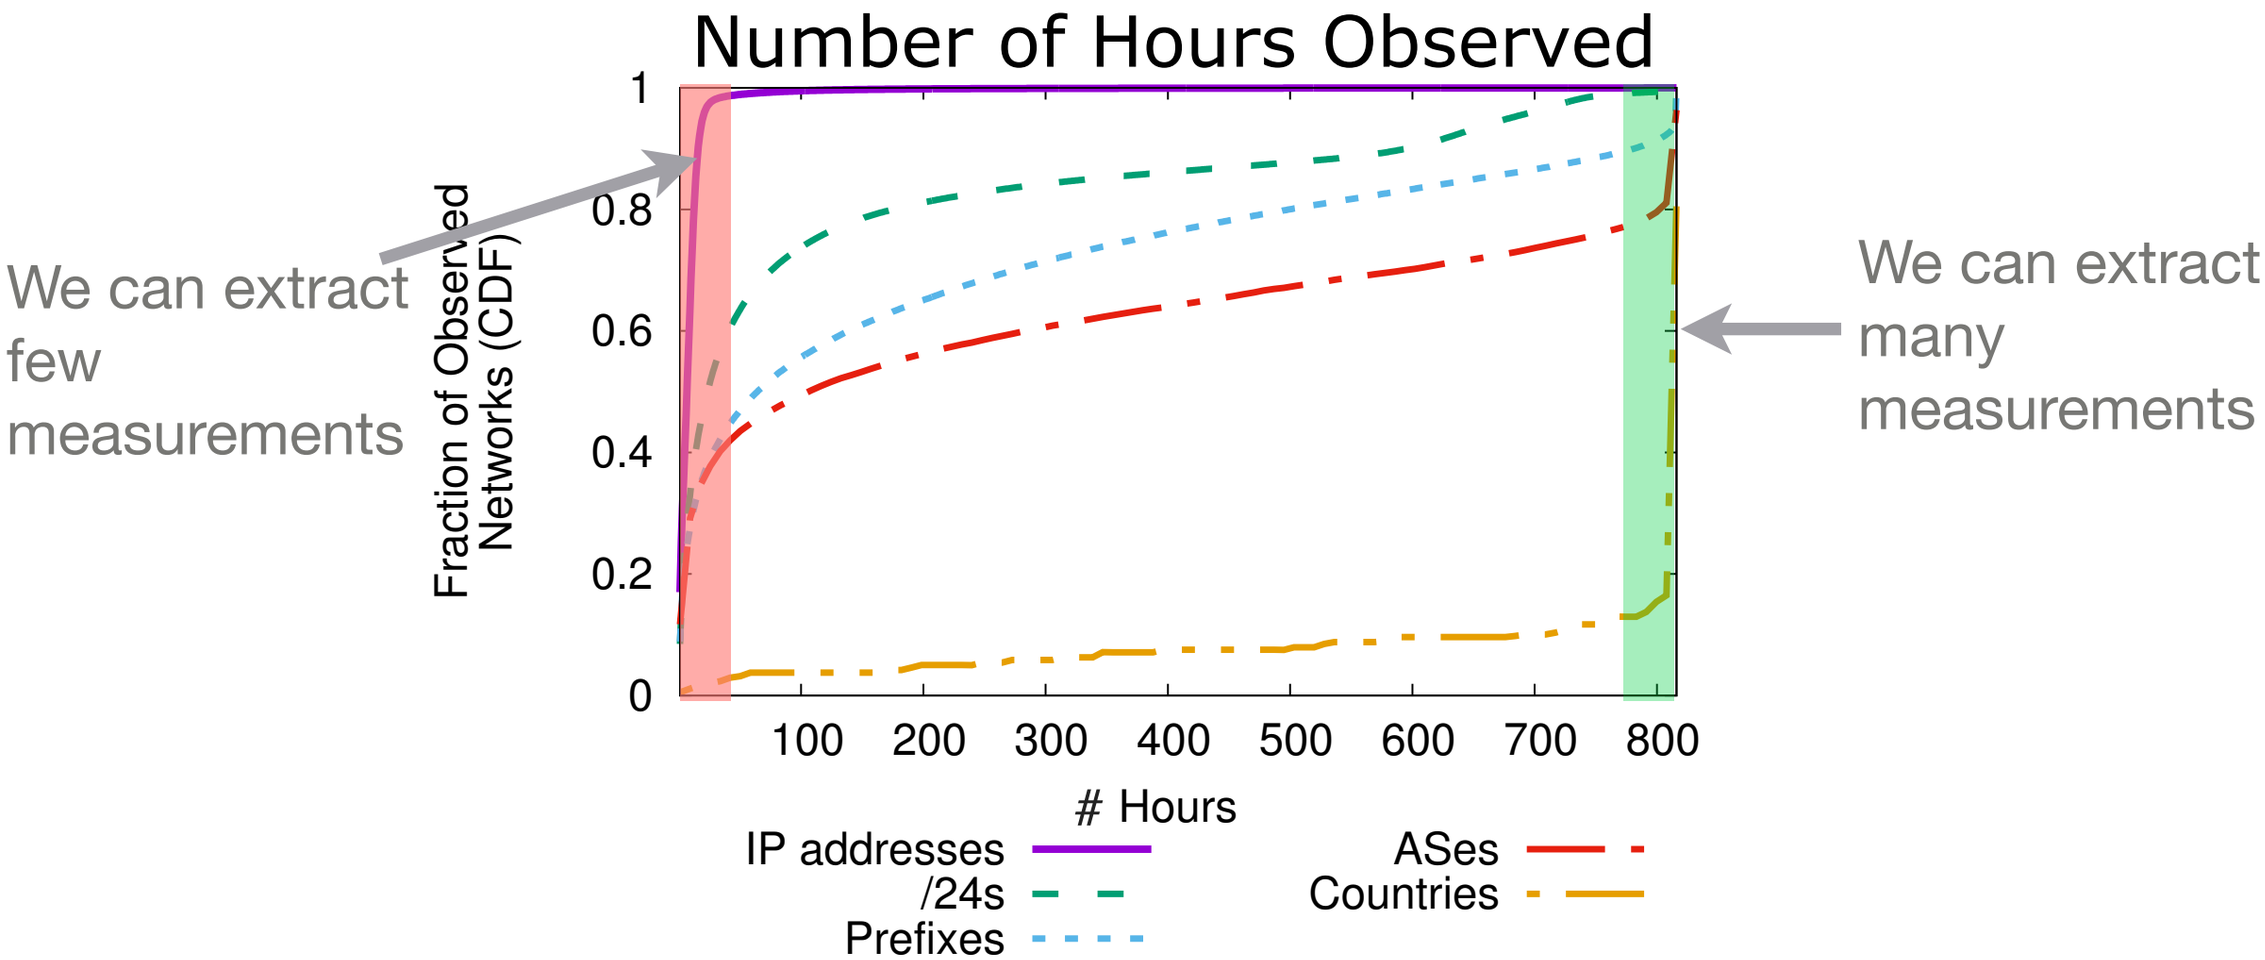

- Analyses requiring observations every hour: only possible for countries and some ASes

Dimensions relevant to Internet-wide network analysis

coverage

traffic type

number of observations

Wait. You used really large darknets!

Wait. You used really large darknets!

- We analyze how darknet size influences number of observed sources
  - Consider contiguous subnets of UCSD-NT as “mini-darknet”

Wait. You used really large darknets!

- We analyze how darknet size influences number of observed sources
  - Consider contiguous subnets of UCSD-NT as “mini-darknet”

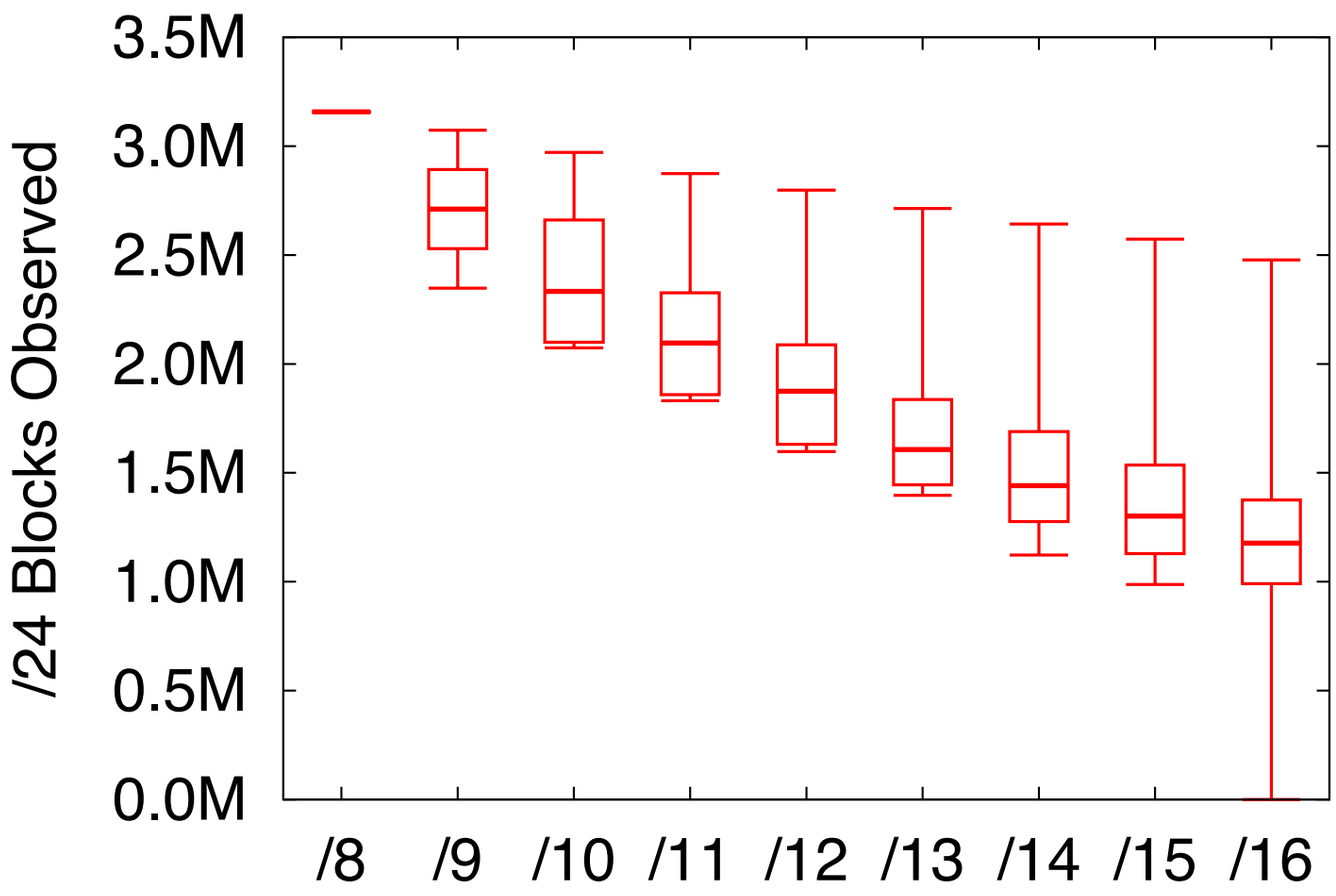

Internet Background Radiation (IBR)

How is IBR applicable to Internet-wide measurement?

Relevant properties and limitations of IBR

What can we expect in terms of coverage, traffic components, and repeated contact?

Case studies

When is using IBR beneficial?

Results of three case studies

	Coverage	Traffic Type	Number of Observations
Locating Open Resolvers	~1.5M IPs	Application: DNS	One
Determining Host Uptime	~200k IPs	Transport: TCP w/ timestamps	Two
Identifying Path Changes	Always analyzable: ~1.5k ASes	IP: TTL field	Continual: >= 1 IP in each set of consecutive time bins

Determining the existence of a resource requires almost no effort with IBR

	Coverage	Traffic Type	Number of Observations
Locating Open Resolvers	~1.5M IPs	Application: DNS	One
Determining Host Uptime	~200k IPs	Transport: TCP w/ timestamps	Two
Identifying Path Changes	Always analyzable: ~1.5k ASes	IP: TTL field	Continual: >= 1 IP in each set of consecutive time bins

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

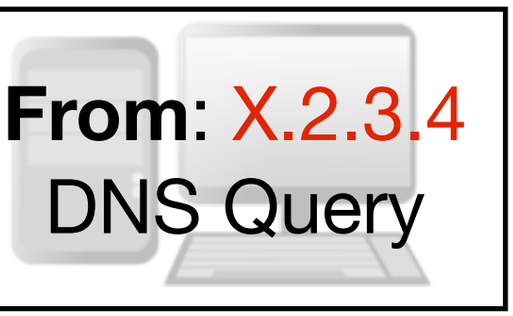

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

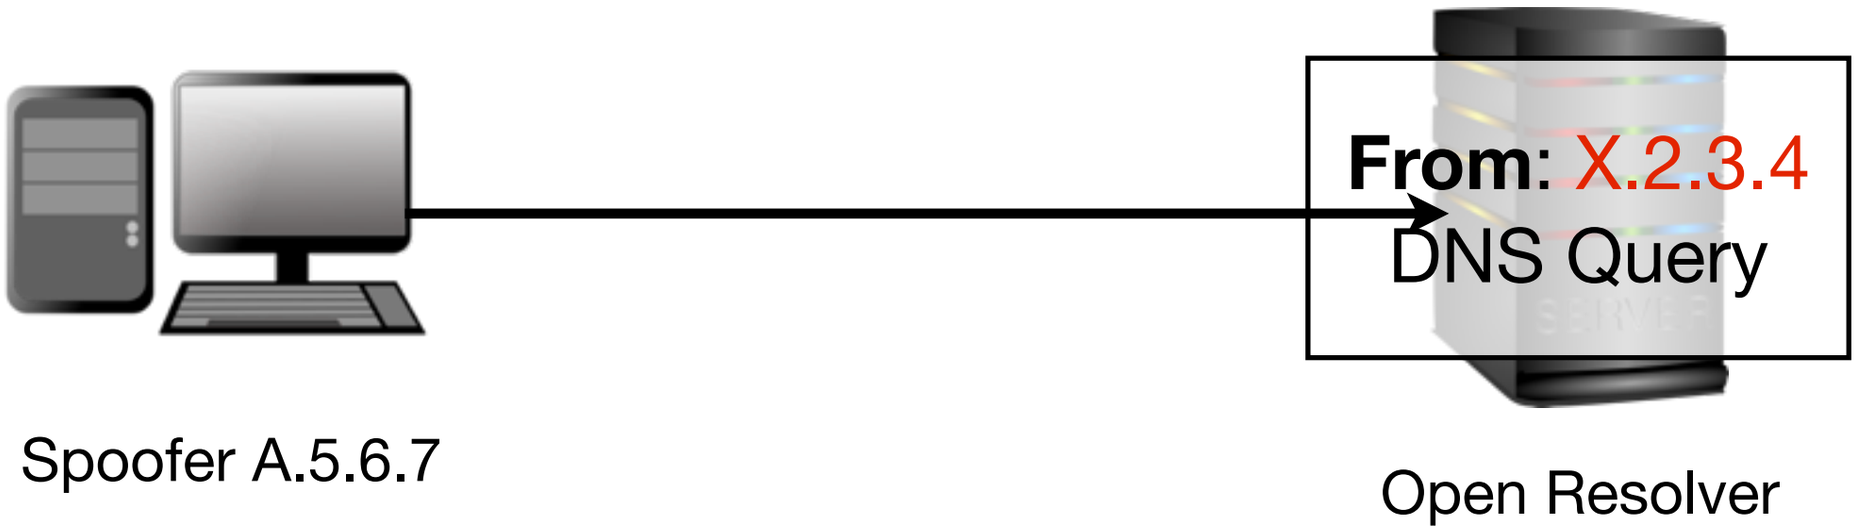

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

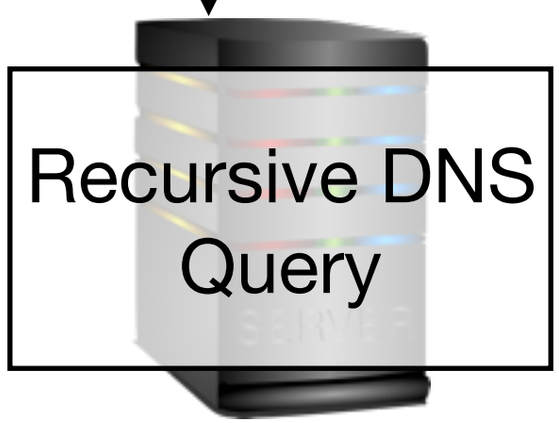

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

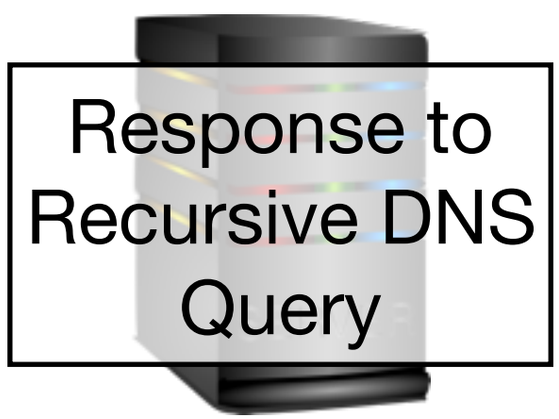

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

Darknet X.0.0.0/8

Authoritative NS

Why we receive traffic from open resolvers: spoofed queries

Spoofer A.5.6.7

Open Resolver

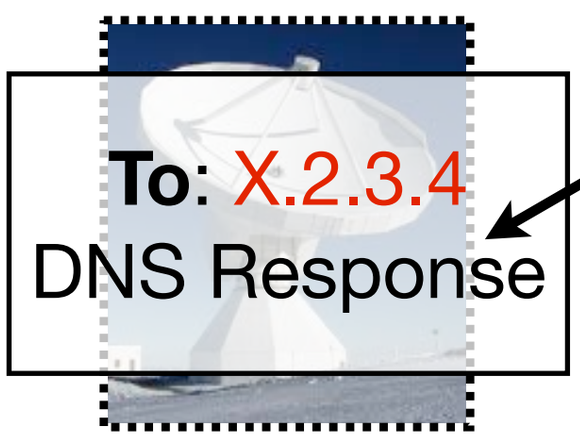

Darknet X.0.0.0/8

Authoritative NS

We see more open resolvers as a result of a change in traffic composition

We see more open resolvers as a result of a change in traffic composition

	IPs
IBR UCSD-13	3.4k

We see more open resolvers as a result of a change in traffic composition

	IPs
IBR UCSD-13	3.4k
IBR ~Feb. 2014	1.56M

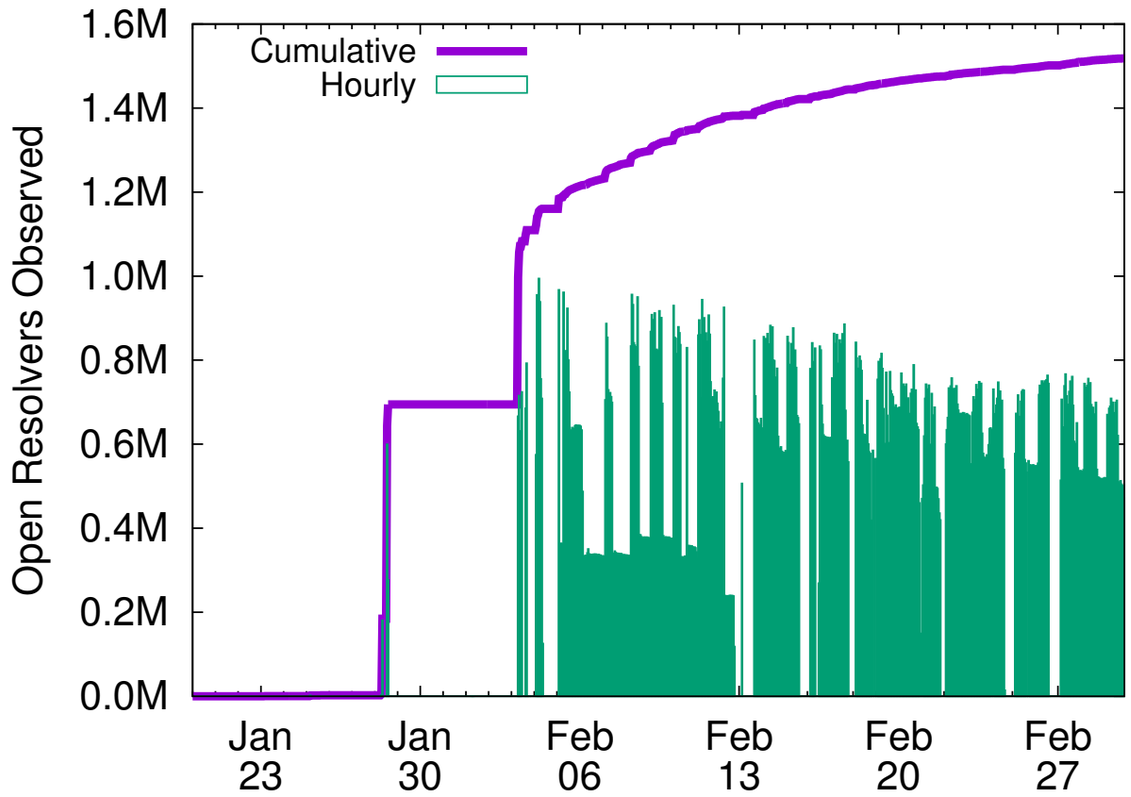

Very few open resolvers before Jan 29, 2014

We see more open resolvers as a result of a change in traffic composition

	IPs
IBR UCSD-13	3.4k
IBR ~Feb. 2014	1.56M

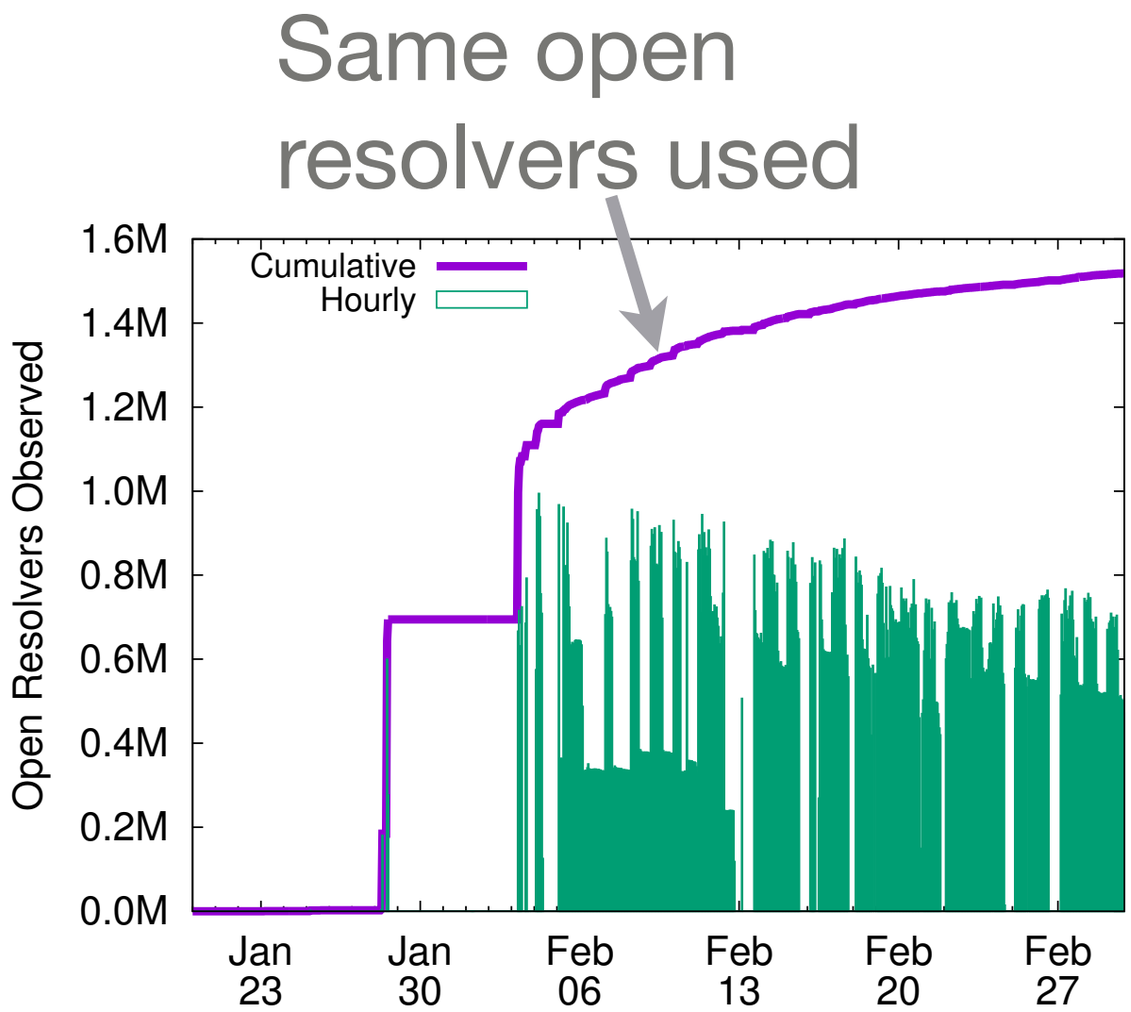

But the number of open resolvers we see is much less than active probing

	IPs
IBR ~July 2013	3.4k
IBR ~Feb. 2014	1.56M
Open Resolver Project ~Feb. 2014	37.6M

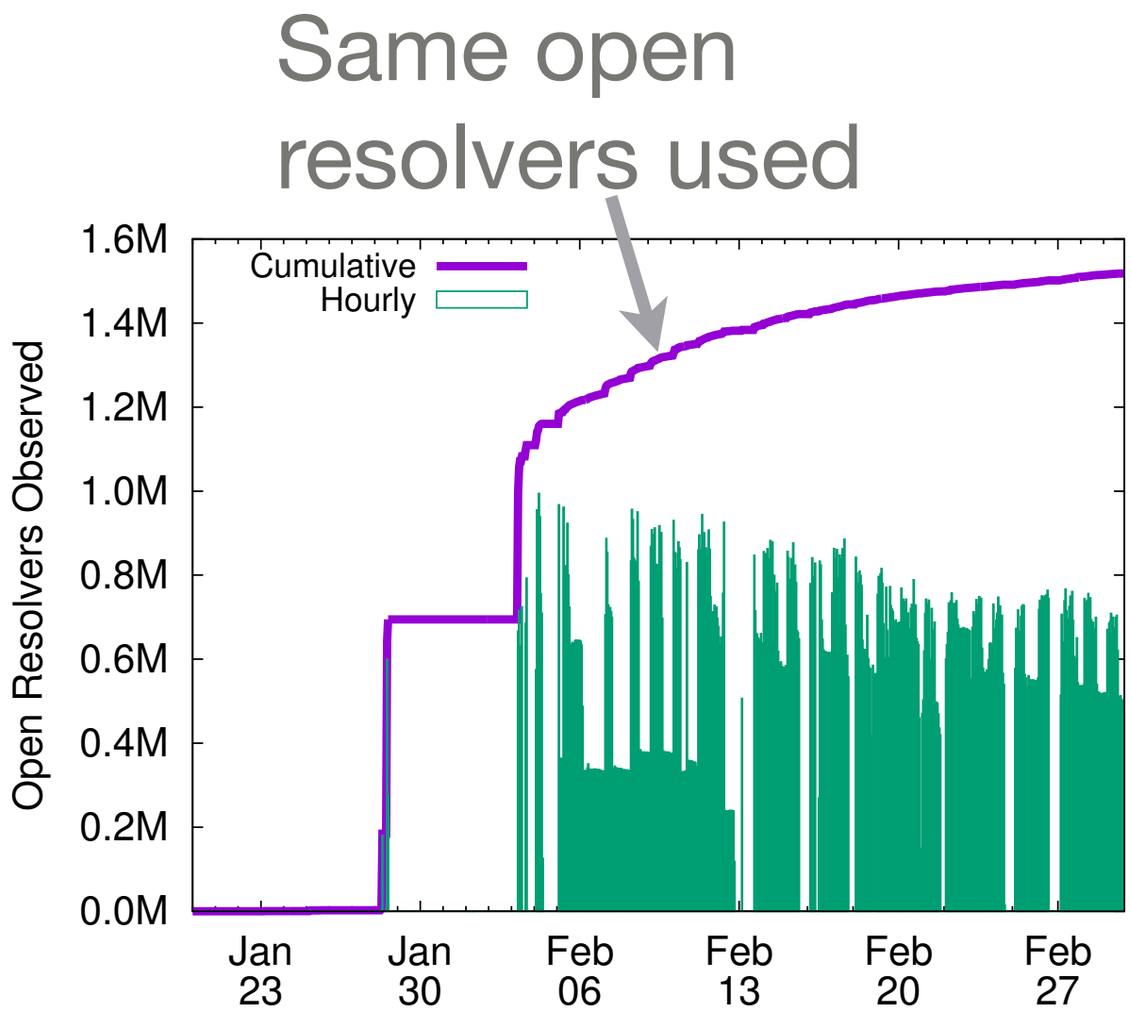

The open resolvers we observe are used in DoS attacks

	IPs	OPCODE: OK	OPCODE: SERVFAIL	OPCODE: NAMEFAIL
IBR ~July 2013	3.4k	3.0k	148	200
IBR ~Feb. 2014	1.56M	1.44M	1.45M	1.35M
Open Resolver Project ~Feb. 2014	37.6M	32.6M	0.92M	0.15M

The open resolvers we observe are used in DoS attacks

	IPs	OPCODE: OK	OPCODE: SERVFAIL	OPCODE: NAMEFAIL
IBR ~July 2013	3.4k	3.0k	148	200
IBR ~Feb. 2014	1.56M	1.44M	1.45M	1.35M
Open Resolver Project ~Feb. 2014	37.6M	32.6M	0.92M	0.15M

Low number of errors

The open resolvers we observe are used in DoS attacks

	IPs	OPCODE: OK	OPCODE: SERVFAIL	OPCODE: NAMEFAIL
IBR ~July 2013	3.4k	3.0k	148	200
IBR ~Feb. 2014	1.56M	1.44M	1.45M	1.35M
Open Resolver Project ~Feb. 2014	37.6M	32.6M	0.92M	0.15M

[VanNice Nanog '15]

Explanation: Authoritative NS is under attack

Low number of errors

Value of IBR

- Without sending any special probes we gain insight into **new phenomena.**

- IBR can provide **additional context**. E.g., starting point of hosts to investigate.

Determining host uptime requires a specific type of traffic

	Coverage	Traffic Type	Number of Observations
Locating Open Resolvers	~1.5M IPs	Application: DNS	One
Determining Host Uptime	~200k IPs	Transport: TCP w/ timestamps	Two
Identifying Path Changes	Always analyzable: ~1.5k ASes	IP: TTL field	Continual: >= 1 IP in each set of consecutive time bins

There are multiple ways to determine host uptime

- [Kumar *et al.* IMC ‘05] used IBR to infer uptime of machines infected with the Witty Worm

	IPs
Witty-inferable March 2004	800

There are many more hosts sending packets with TCP timestamps than Witty payload

- [Kumar *et al.* IMC '05] used IBR to infer uptime of machines infected with the Witty Worm

- p0f, Nmap use TCP timestamps to infer uptime (2 packets required)

- Many IBR packets have TCP timestamps

	IPs
Witty-inferable March 2004	800

There are many more hosts sending packets with TCP timestamps than Witty payload

- [Kumar *et al.* IMC '05] used IBR to infer uptime of machines infected with the Witty Worm
- p0f, Nmap use TCP timestamps to infer uptime (2 packets required)
- Many IBR packets have TCP timestamps

	IPs
Witty-inferable March 2004	800
Sent TCP UCSD-13	16M
Sent TCP timestamps UCSD-13	1.7M

Not all TCP timestamps are usable for uptime inferences

- [Kumar *et al.* IMC ‘2005] used IBR to infer uptime of machines infected with the Witty Worm
- p0f, Nmap use TCP timestamps to infer uptime (2 packets required)
- Many IBR packets have TCP timestamps
- We find that the TCP timestamp method is inaccurate for several operating systems

	IPs
Witty-inferable March 2004	800
Sent TCP UCSD-13	16M
Sent TCP timestamps UCSD-13	1.7M
Sent usable TCP timestamps UCSD-13	208k

Many types of traffic send packets with TCP timestamps

- [Kumar *et al.* IMC ‘2005] used IBR to infer uptime of machines infected with the Witty Worm
- p0f, Nmap use TCP timestamps to infer uptime (2 packets required)
- Many IBR packets have TCP timestamps
- We find that the TCP timestamp method is inaccurate for several operating systems
- With TCP timesamps we can infer uptime consistently over time

	IPs
Witty-inferable March 2004	800
Sent TCP UCSD-13	16M
Sent TCP timestamps UCSD-13	1.7M
Sent usable TCP timestamps UCSD-13	208k
Sent usable TCP timestamps UCSD-12	291k

Value of IBR

- Ability to make inferences for **hard to measure hosts** (E.g., behind NAT).

- Easy to get a **large sample** size.

Continually identifying path changes requires multiple packets and repeated contact

	Coverage	Traffic Type	Number of Observations
Locating Open Resolvers	~1.5M IPs	Application: DNS	One
Determining Host Uptime	~200k IPs	Transport: TCP w/ timestamps	Two
Identifying Path Changes	Continual: ~1.5k ASes	IP: TTL field	Continual: >= 1 IP in each set of consecutive time bins

Using multiple packets to infer path changes

- TTL field reflects number of hops from remote host to destination
  - A change in TTL implies that the path has changed
    (we miss path changes when new and old route are the same number of hops)

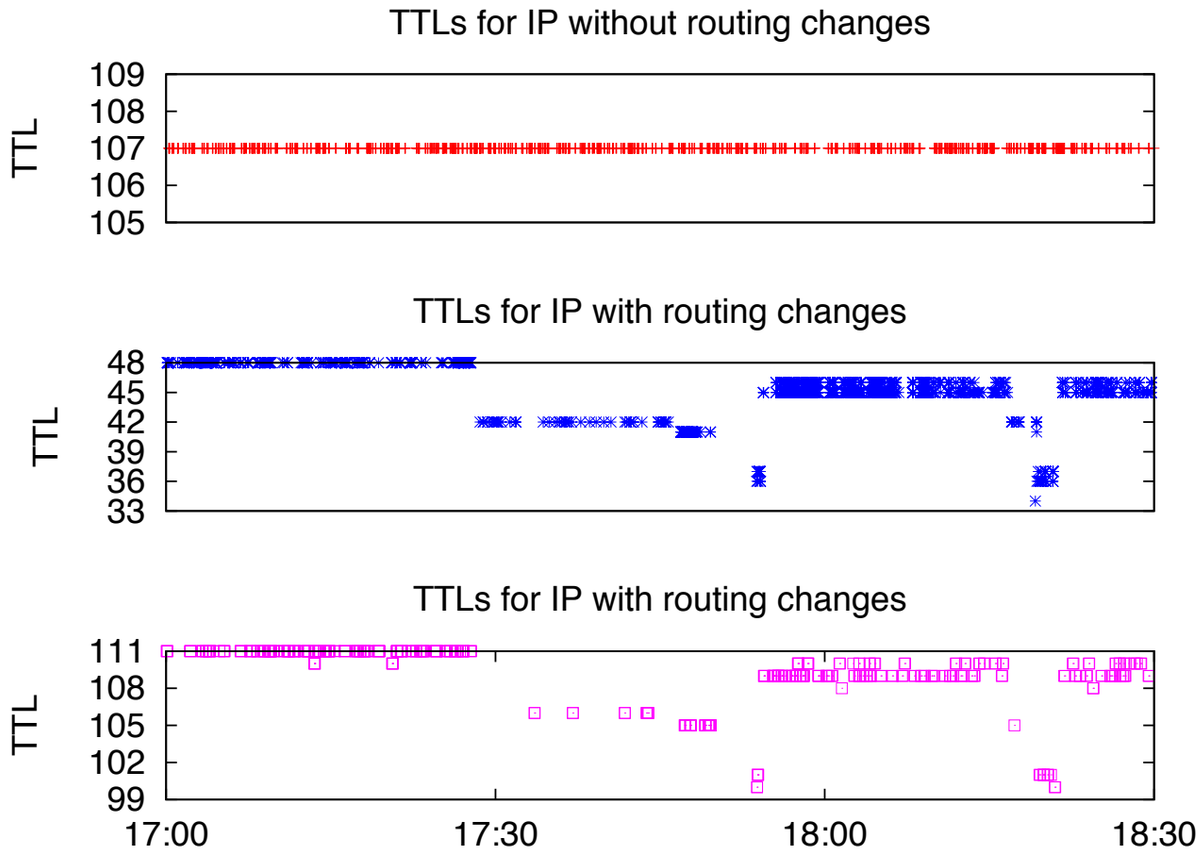

At aggregate levels, we can constantly monitor many path changes

- Technique:
  - Divide dataset into 5 minute bins
  - Compare TTL values of each IP address appearing in consecutive time bins

At aggregate levels, we can constantly monitor many path changes

- Technique:
  - Divide dataset into 5 minute bins
  - Compare TTL values of each IP address appearing in consecutive time bins

Granularity	Always-Analyzable
IP addresses	2.8k
/24 blocks	2.6k
BGP announced prefixes	3.6k
ASes	1.7k
Countries	155

At aggregate levels, we can constantly monitor many path changes

- Technique:
  - Divide dataset into 5 minute bins
  - Compare TTL values of each IP address appearing in consecutive time bins

Granularity	Always-Analyzable
IP addresses	2.8k
/24 blocks	2.6k
BGP announced prefixes	3.6k
ASes	1.7k
Countries	155

0.02% of IPs sending IBR

At aggregate levels, we can constantly monitor many path changes

- Technique:
  - Divide dataset into 5 minute bins
  - Compare TTL values of each IP address appearing in consecutive time bins

Granularity	Always-Analyzable
IP addresses	2.8k
/24 blocks	2.6k
BGP announced prefixes	3.6k
ASes	1.7k
Countries	155

0.02% of IPs sending IBR

7% of ASes sending IBR

Our method provides similar insight as traceroute based methods

- Heuristic for inferring when a path change occurs yields similar results to a traceroute based method (Ark)

Our method provides similar insight as traceroute based methods

- Heuristic for inferring when a path change occurs yields similar results to a traceroute based method (Ark)

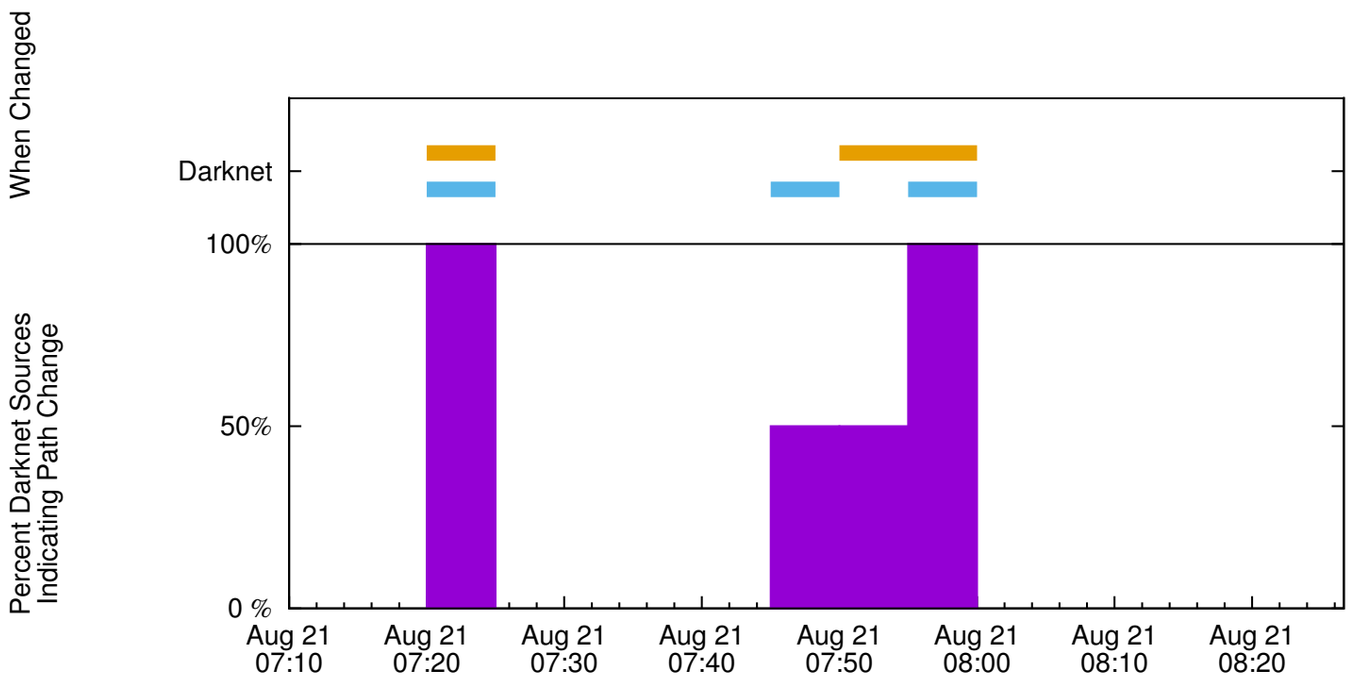

Our method provides similar insight as traceroute based methods

- Heuristic for inferring when a path change occurs yields similar results to a traceroute based method (Ark)

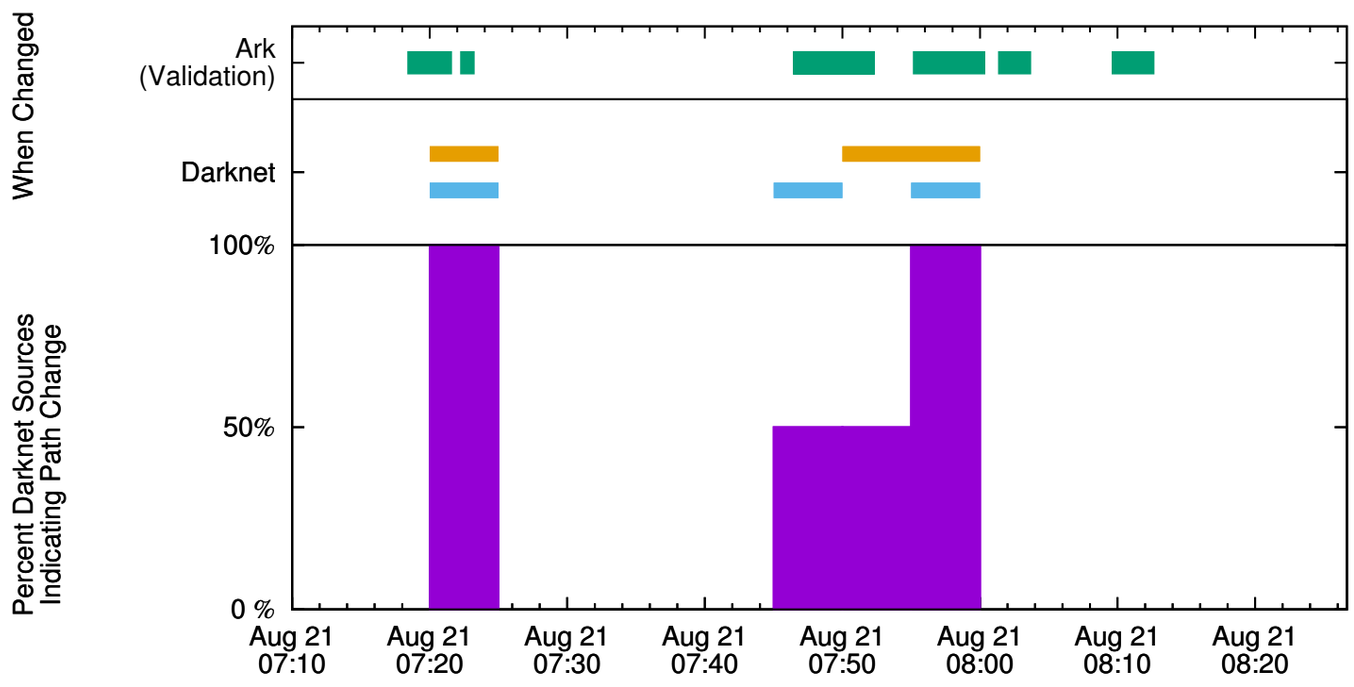

Value of IBR

- For many ASes: **aggregating** the signal from many hosts provides diverse and continuous visibility.

- IBR can **reduce the overhead** of active methods by providing guidance on when and where to probe.

Summary

- IBR is useful for a variety of network analysis tasks

- A technique's coverage is a function of:
  - the traffic components (network layer, fields) used
  - the frequency of observations
  - the infrastructure to capture IBR

- Guidelines for IBR is useful: to provide additional context, for hard-to-measure hosts, for large samples, and to reduce measurement overhead

- Interested in using IBR for your measurement study?
  Contact us: karyn@caida.org