

Understanding the Efficacy of Deployed Internet Source Address Validation Filtering



Robert Beverly, Arthur Berger (MIT),
Young Hyun, k claffy (UCSD/CAIDA)

ACM Internet Measurement Conference 2009

Spoofers Project

- Background
- Recent Relevance
- Project Methodology
- Results
- Parting Thoughts

Spoofed-Source IP Packets

- Circumvent host network stack to forge or “spoof” *source address* of an IP packet
- Lack of source address accountability a basic Internet weakness:
 - Anonymity, indirection [VP01], amplification
- Security issue for more than *two-decades* [RTM85, SB89]
- Still an attack vector?


0	4	8	16	19	31
Version	HLen	Tos	Length		
Ident		Flags	Offset		
TTL	Protocol	Checksum			
Source Address					
Destination Address					
Options (Variable)				Padding (Variable)	
Data					

Circa 2004...

IP source spoofing
doesn't matter!

- a) All providers **filter**
- b) All modern attacks use **botnets**
- c) Compromised hosts are behind **NATs**

- Strong opinions from many:
 - Academic
 - Operational
 - Regulatory
- ...but only anecdotal data



spoofers.csail.mit.edu

- Internet-wide active measurement effort:
 - Quantify the extent and nature of Internet source address filtering
 - Understand real-world efficacy of available best-practice defenses
 - Validate common assumption of edge filtering
- Began Feb. 2005
 - Understand how filtering has evolved
 - Basis for driving design of more secure architectures



Spoofers Project

- Background
- **Recent Relevance**
- Project Methodology
- Results
- Parting Thoughts

Prediction: spoofing increasingly a problem in the future

- **Spoofed traffic complicates a defenders job**
- Tracking spoofs is operationally difficult:
 - [Greene, Morrow, Gemberling NANOG 23]
 - Hash-based IP traceback [Snoeren01]
 - ICMP traceback [Snoeren01]
- Consider **Slide from SRUTI 2005**
 - Today (worst case scenario), non-spoofing zombies are widely distributed, a network operator must defend against attack packets from 5% of routeable netblocks.
 - Future: if 25% of zombies capable of spoofing significant volume of the traffic could appear to come any part of the IPv4 address space
- Adaptive programs that make use of all local host capabilities to amplify their attacks

The Spoofing Problem (2009)

- DNS Amplifier Attacks
- DNS Cache Poisoning
- In-Window TCP Reset Attacks
- Bots that probe for ability to spoof
- Spam Filter Circumvention
- UW reverse traceroute
- etc, etc...

Can't anticipate next attack employing IP spoofing

The Operational Side


- Arbor 2008 Infrastructure Survey:
 - *“Reflective amplification attacks responsible for the largest attacks (40Gbps) exploit IP spoofing”*
 - *“No bots were used in this attack. The attacker had a small number of compromised Linux boxes from which he’d launch the spoofed source DNS query”*
- What’s an operator to do?

Operational View

- IETF BCP38 best filtering practice
- But, not all sources created equal:

<u>Example Source IP</u>	<u>Description</u>	<u>Possible Defense</u>
192.168.1.1	RFC1918 private	Static ACL
1.2.3.4	Unallocated	Bogon Filters
6.1.2.3	Valid (In BGP table)	uRPF (loose/strict)
Client IP \oplus (2^N)	Neighbor Spoof	Switch, DOCSIS

harder



Operational View

- We have defenses, what's the problem?
- BCP38 suffers from:
 - Lack of hardware support (see NANOG)
 - Global participation requirement
 - Management nightmare (edge filters)
 - Multi-homing, asymmetry, etc implies loose uRPF, implies little protection
- This work: understand the real-world efficacy of these best practices

Spoofers Project

- Background
- Recent Relevance
- **Project Methodology**
- Results
- Parting Thoughts

Spoofers Test



- Willing participants run “spoofers” client to test policy, perform inference, etc.
 - Binaries, source publicly available
 - Useful diagnostic tool for many
 - Runs once, not an agent
- Clients self-selecting
 - Understand population and potential bias

Spoofing Test


- Testing vulnerability of Internet to source spoofing, not prevalence of source spoofing (e.g. backscatter analysis)
- Uses CAIDA's Ark infrastructure to test many paths
- Aggregate results, tomography, etc to form global picture of best-practices (BCP38) efficacy

Archipelagio

- Tied into CAIDA's distributed measurement infrastructure (Ark)
- ~40 nodes, globally distributed
- Ark nodes act as IPv4/v6 spoof probe receivers

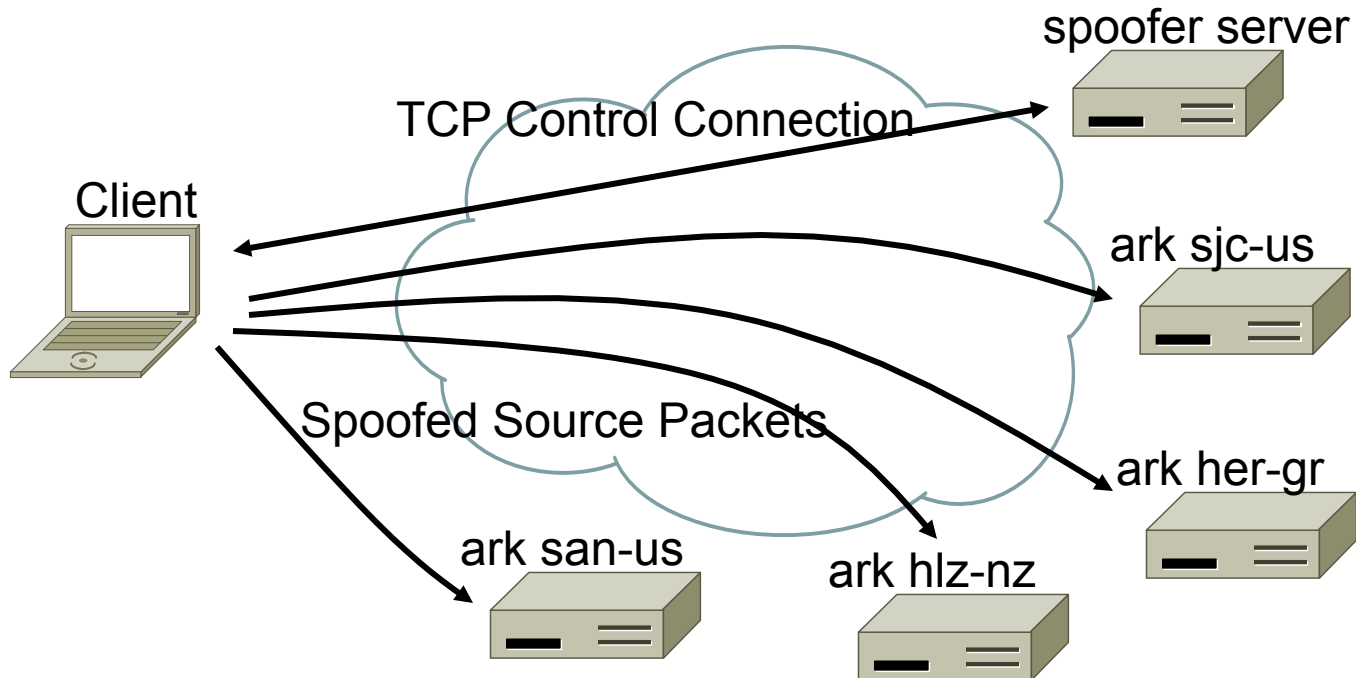


Spoofing Operation



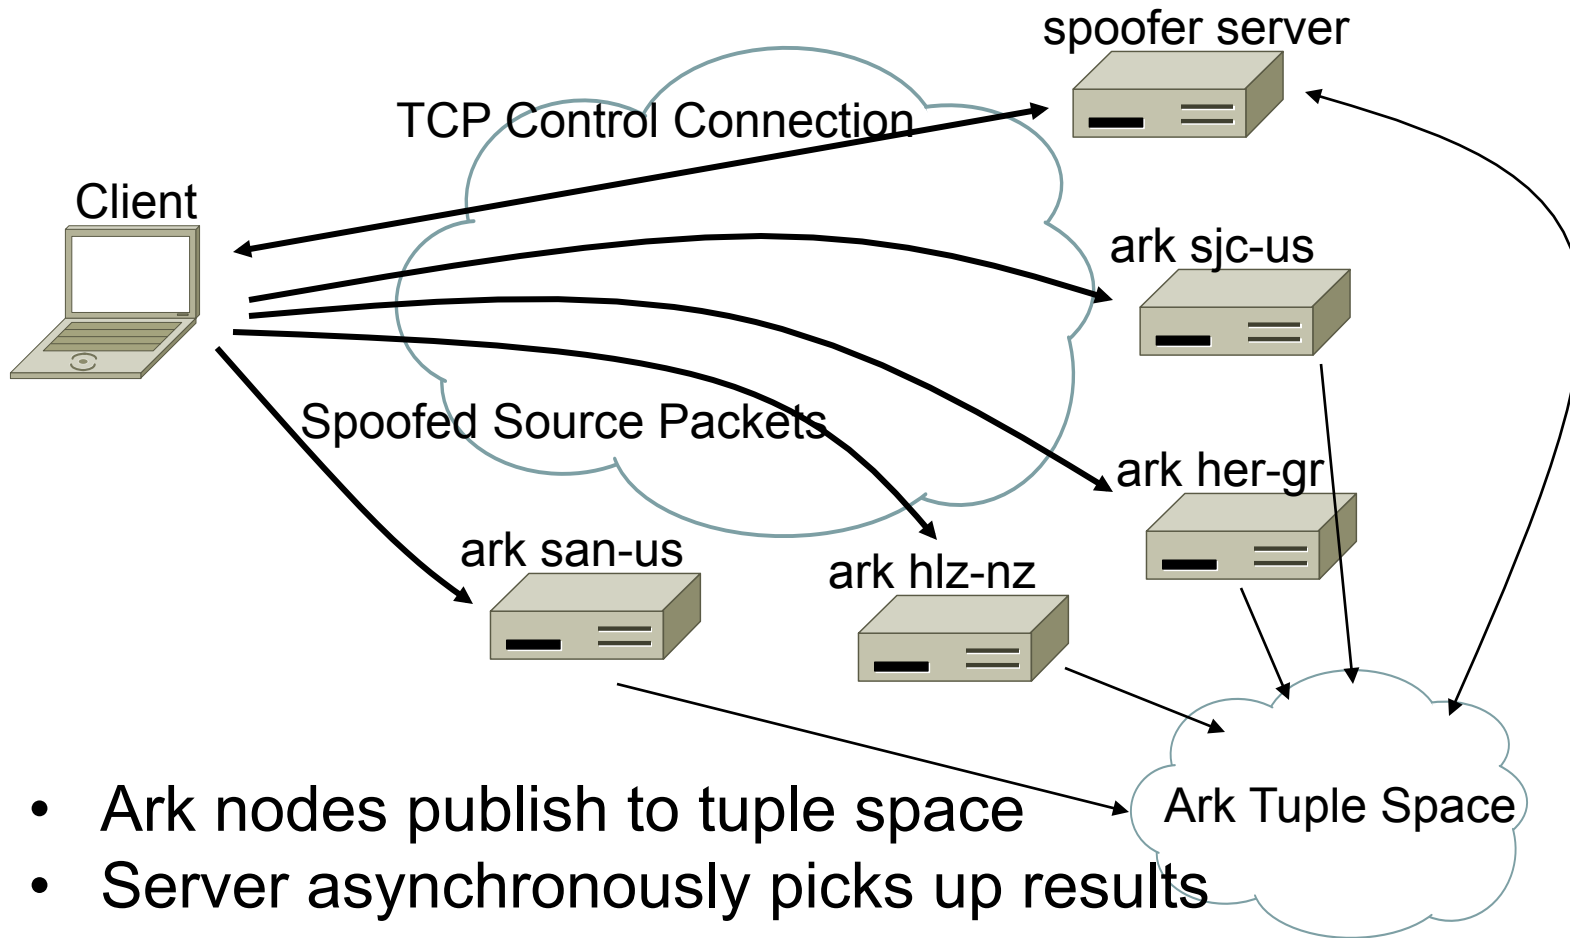
- Client confers with control server, receives test
- (SRC, DST, HMAC, SEQ) probe tuples
- Use TCP destination port 80 to avoid secondary filtering

Distributed Probing



- Client sends HMAC keyed spoof probes to ark nodes
- Includes ground-truth validation (non-spoofed) probes
- UDP port 53 + random delay to avoid secondary filtering
- Client runs traceroute to each ark node in parallel

Distributed Probing

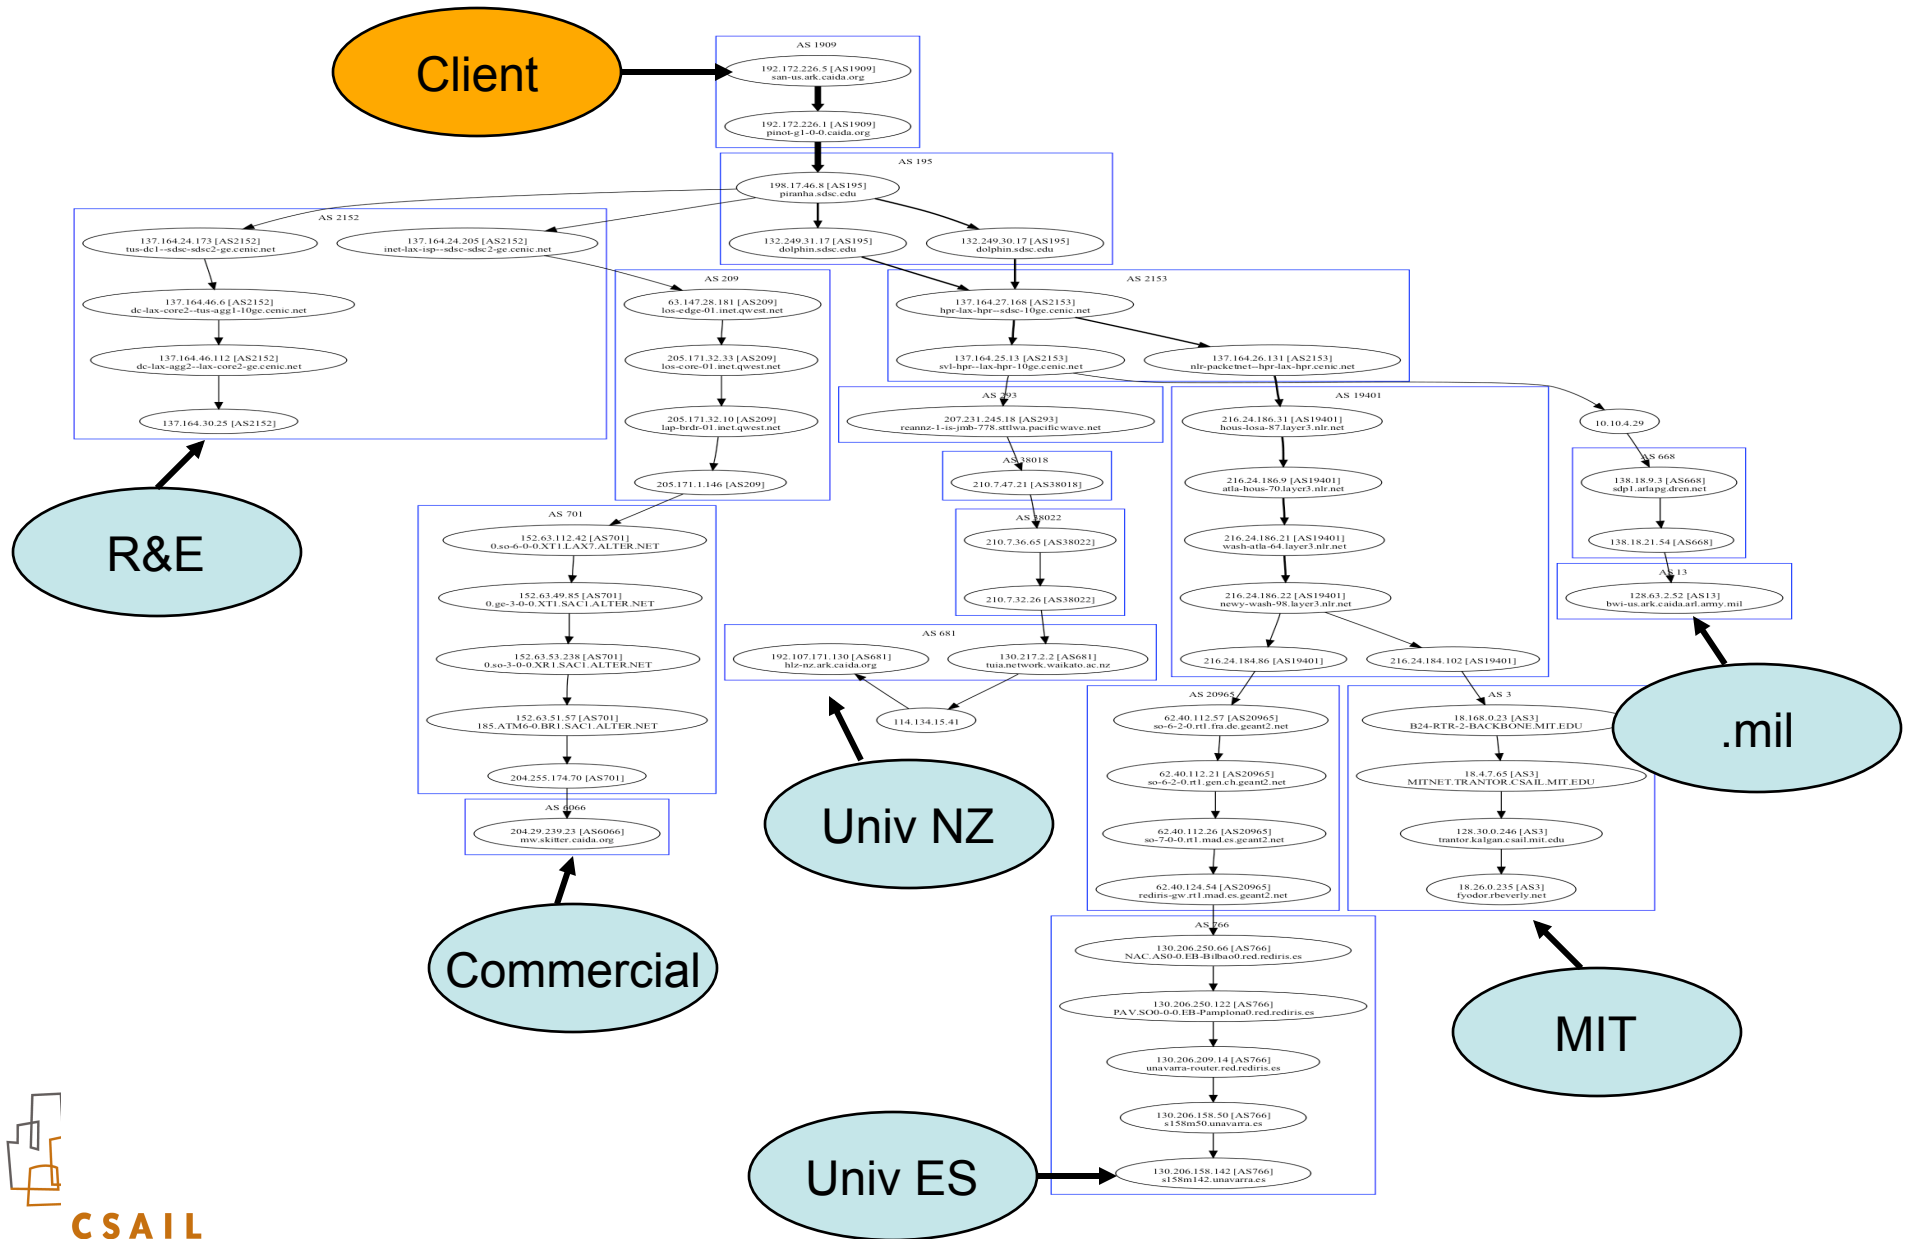


- Ark nodes publish to tuple space
- Server asynchronously picks up results
- Run tracefilter (described next)
- Return results to user via web report

Outcome of a Probe

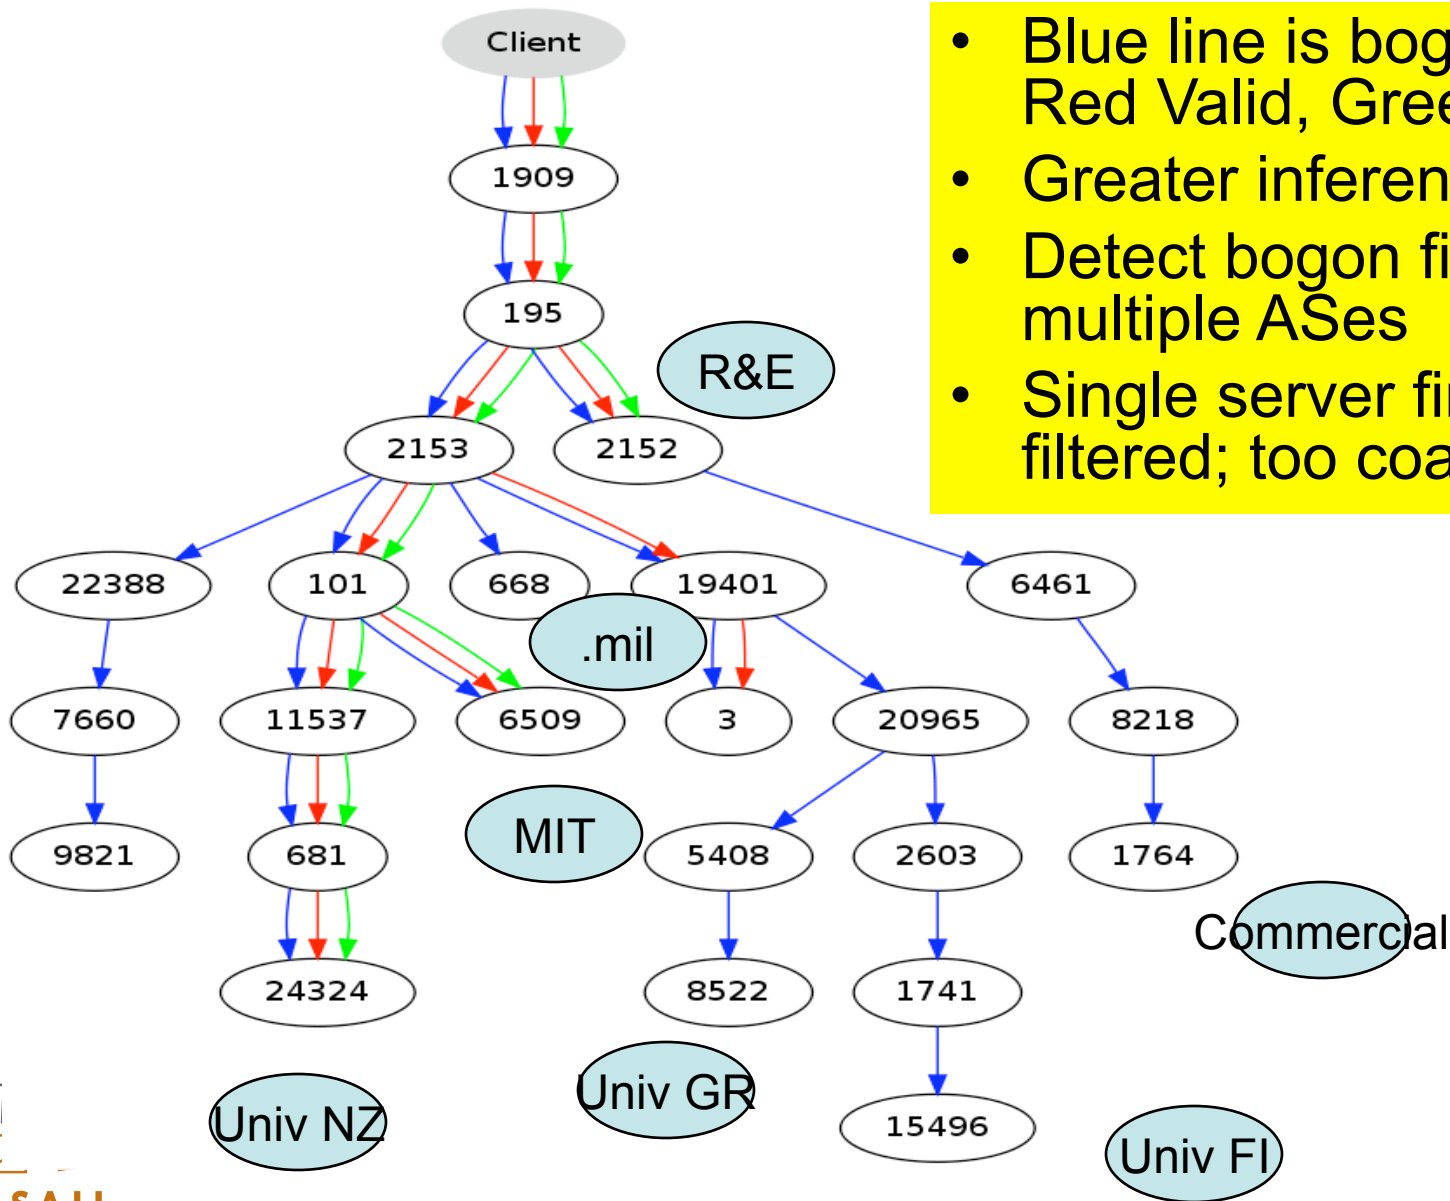
- Blocked by OS:
 - Detect, revert to raw Ethernet
- Hits a NAT along path:
 - Detect, exclude from results
- Other blocking (proxy, congestion):
 - Detect, exclude from results
- Blocked by source validation filter
- Successfully received at Ark node

Ark Enables Better Inferences



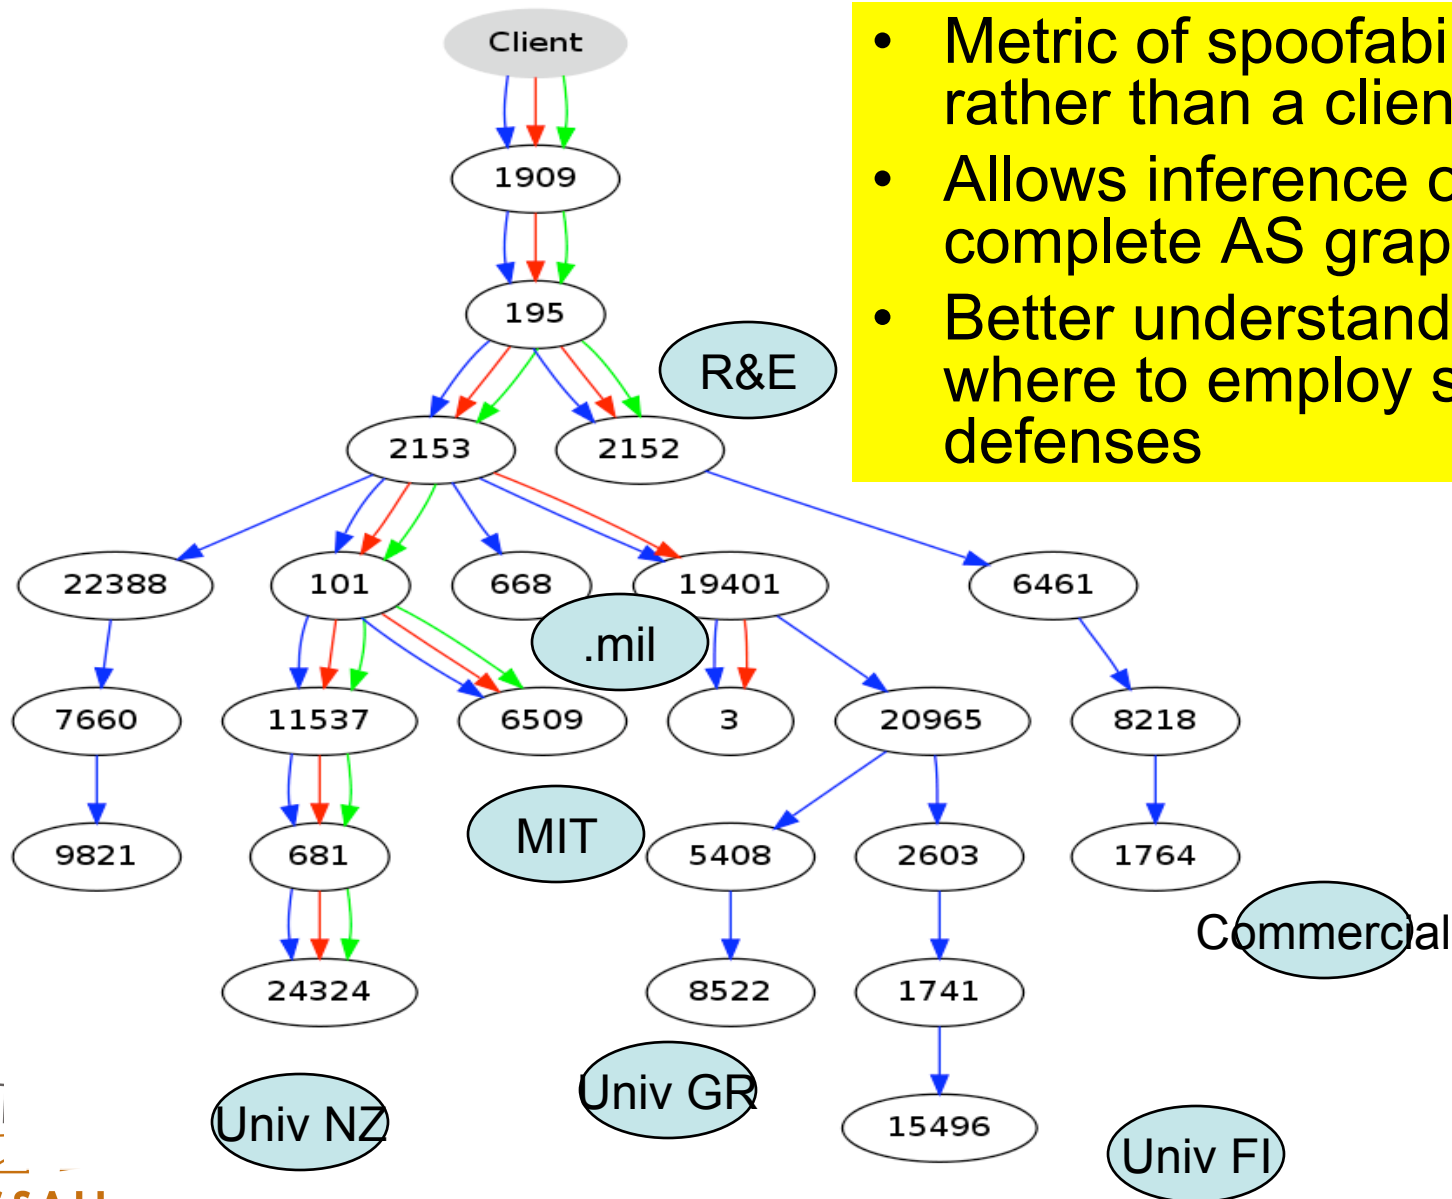
Multiple Destinations

- Blue line is bogon traffic, Red Valid, Green private
- Greater inference power
- Detect bogon filtering at multiple ASes
- Single server finds valid filtered; too coarse!



Multiple Destinations


- Metric of spoofability a path rather than a client
- Allows inference on the complete AS graph
- Better understanding of where to employ spoofing defenses



tracefilter


- A tool for *locating* source address validation (anti-spoofing) filters along path
- “traceroute for BCP38”
- Better understand at finer granularity (router) who is/is not filtering

tracefilter



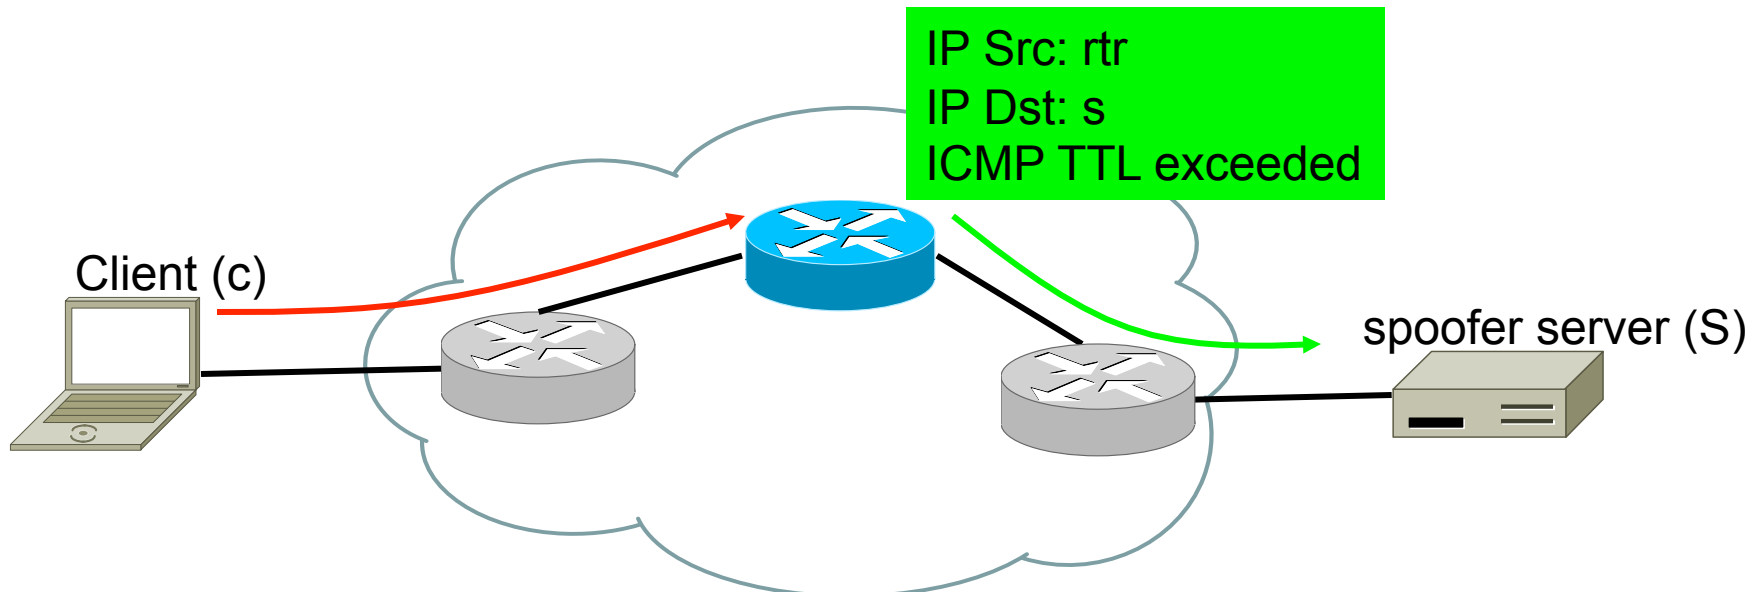
- Client c works in conjunction with our server S

tracefilter




- c sends spoofed packet with:
- $t\text{tl}=x$, $\text{src}=\text{S}$, $\text{dst}=\text{S}+1$ for $0 < x < \text{pathlen}$

tracefilter



- S receives ICMP expiration messages from routers along path
- For each decoded TTL, S records which spoofed packets are received

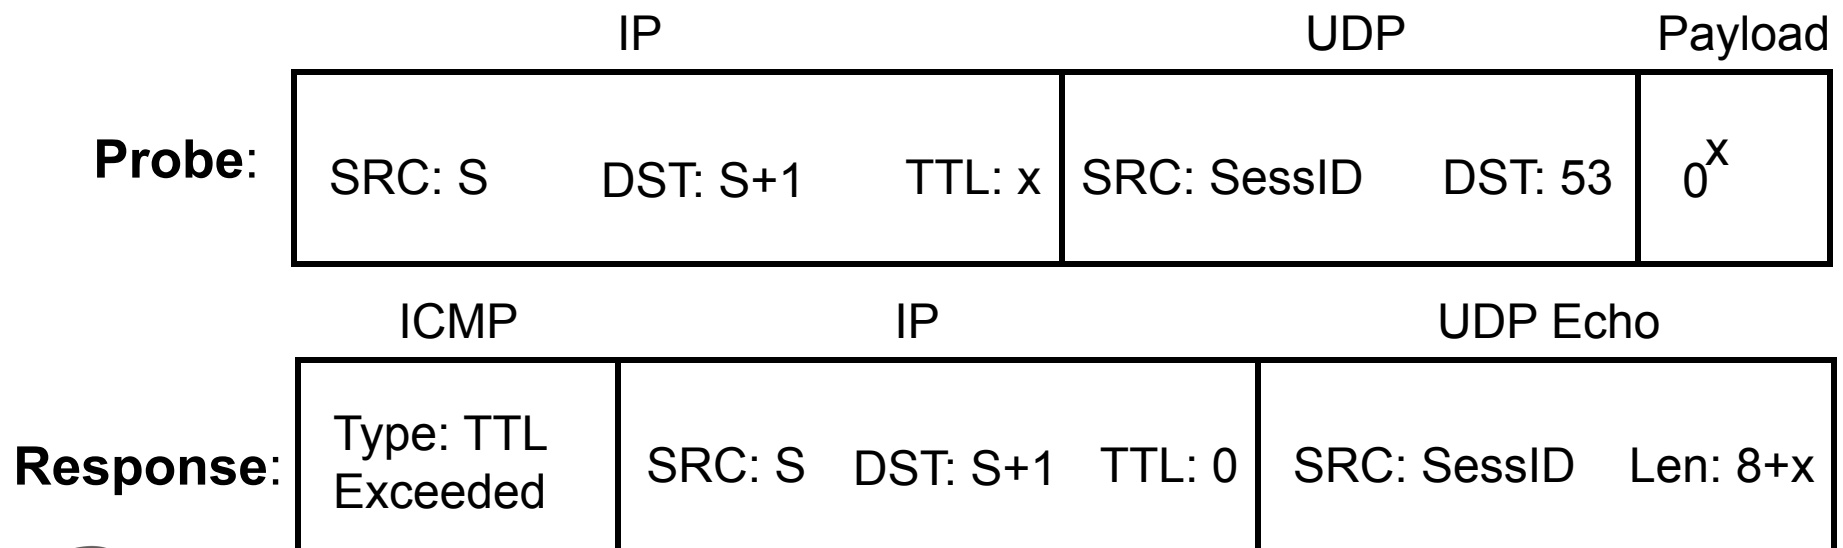
tracefilter



- Increase TTL, repeat
- Largest TTL indicates filtering point

tracefilter

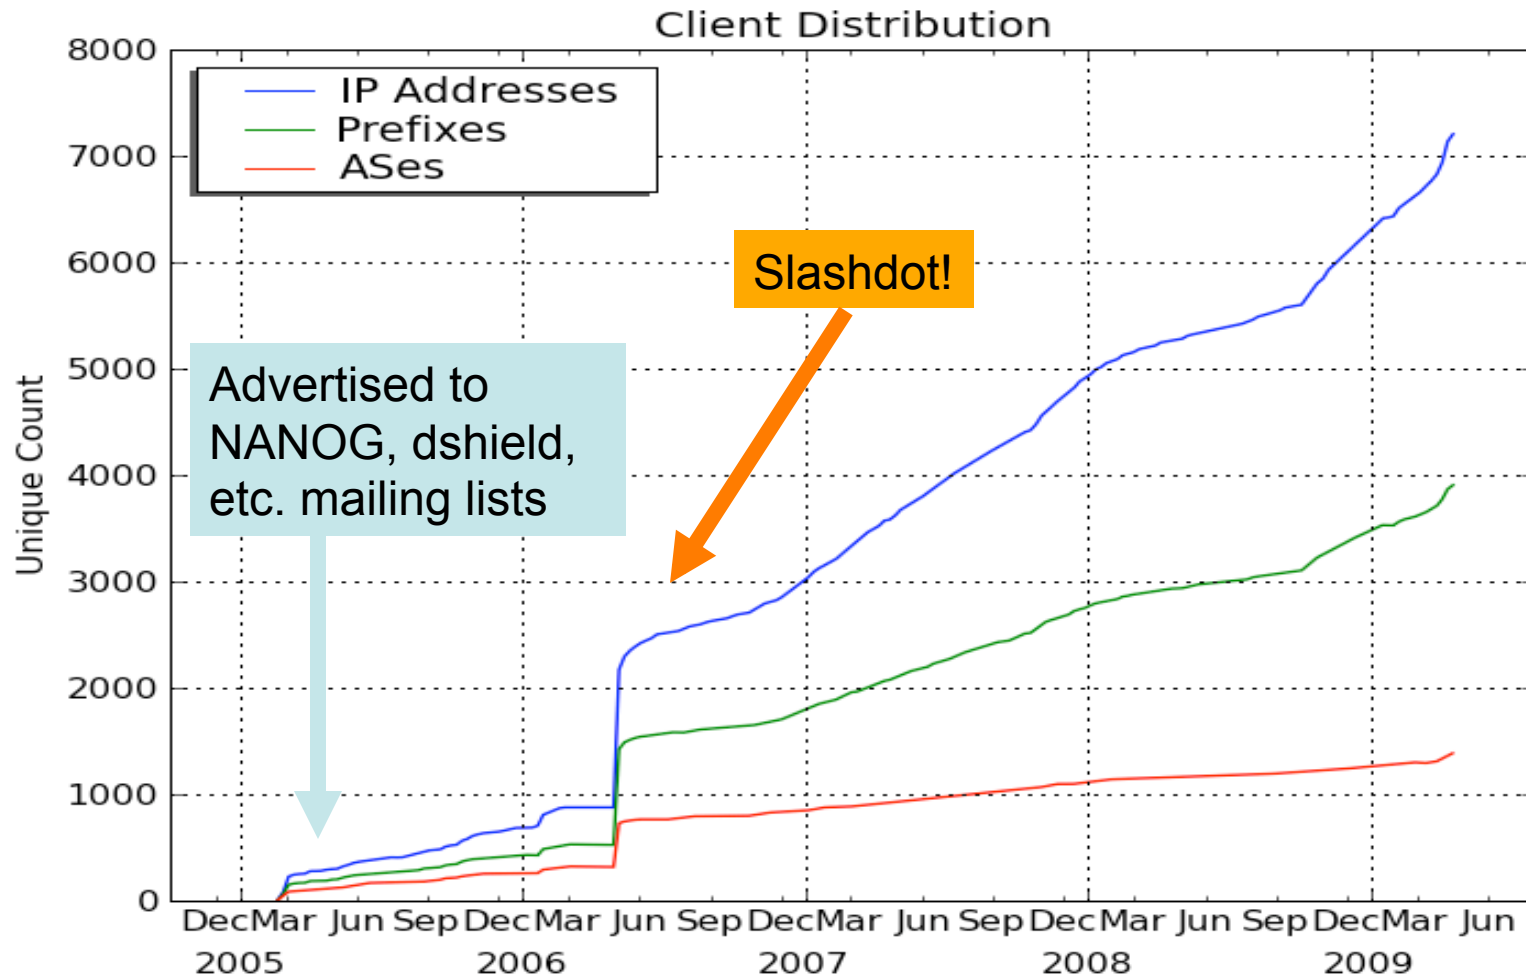
- How can *S* determine *originating* TTL of *c*'s packets?
- ICMP echo includes only 28 bytes of expired packet
- *c* encodes TTL by padding payload with zeros



Spoofers Project

- Background
- Recent Relevance
- Project Methodology
- **Results**
- Parting Thoughts

Client Population



Sample Bias

- Obtain general population using 20.8M unique IPs from random topology traces
- Use NetAcuity for geolocation, descriptive statistics
- Aggregate general population into /24s to eliminate non-homogenous properties

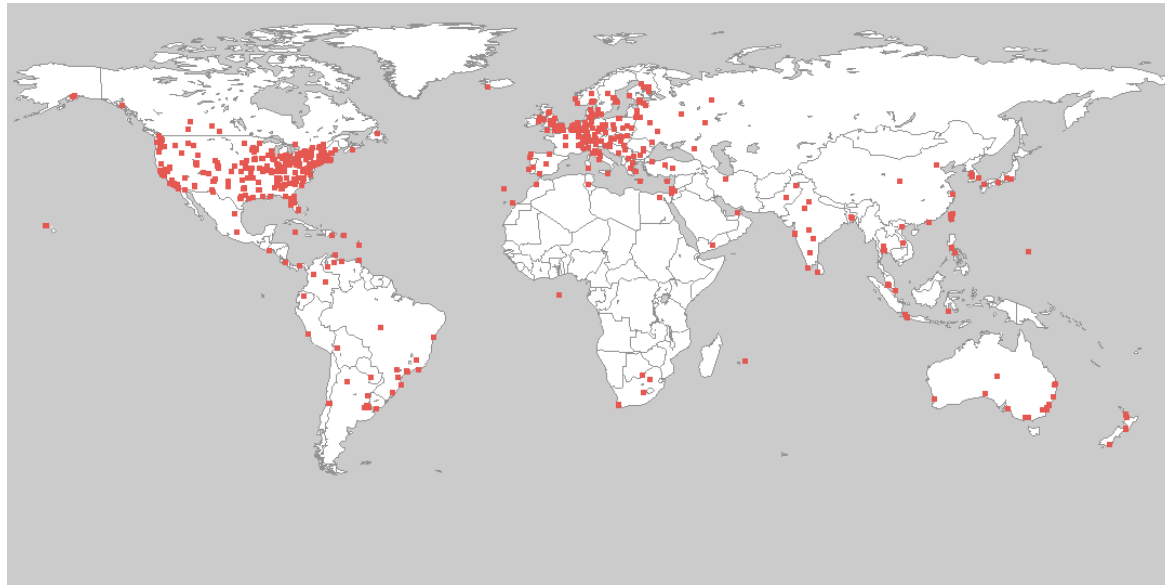
Comparing Populations

- Evaluate Bias:
 - Country, speed, organization type, etc.
- Continent Analysis

Continent	Population	Measurement Set
N. America	37%	36%
Europe	29%	33%
Asia	28%	17%
S. America	4%	4%
Oceania	1%	2%
Africa	0.5%	6%

Client Population Distribution

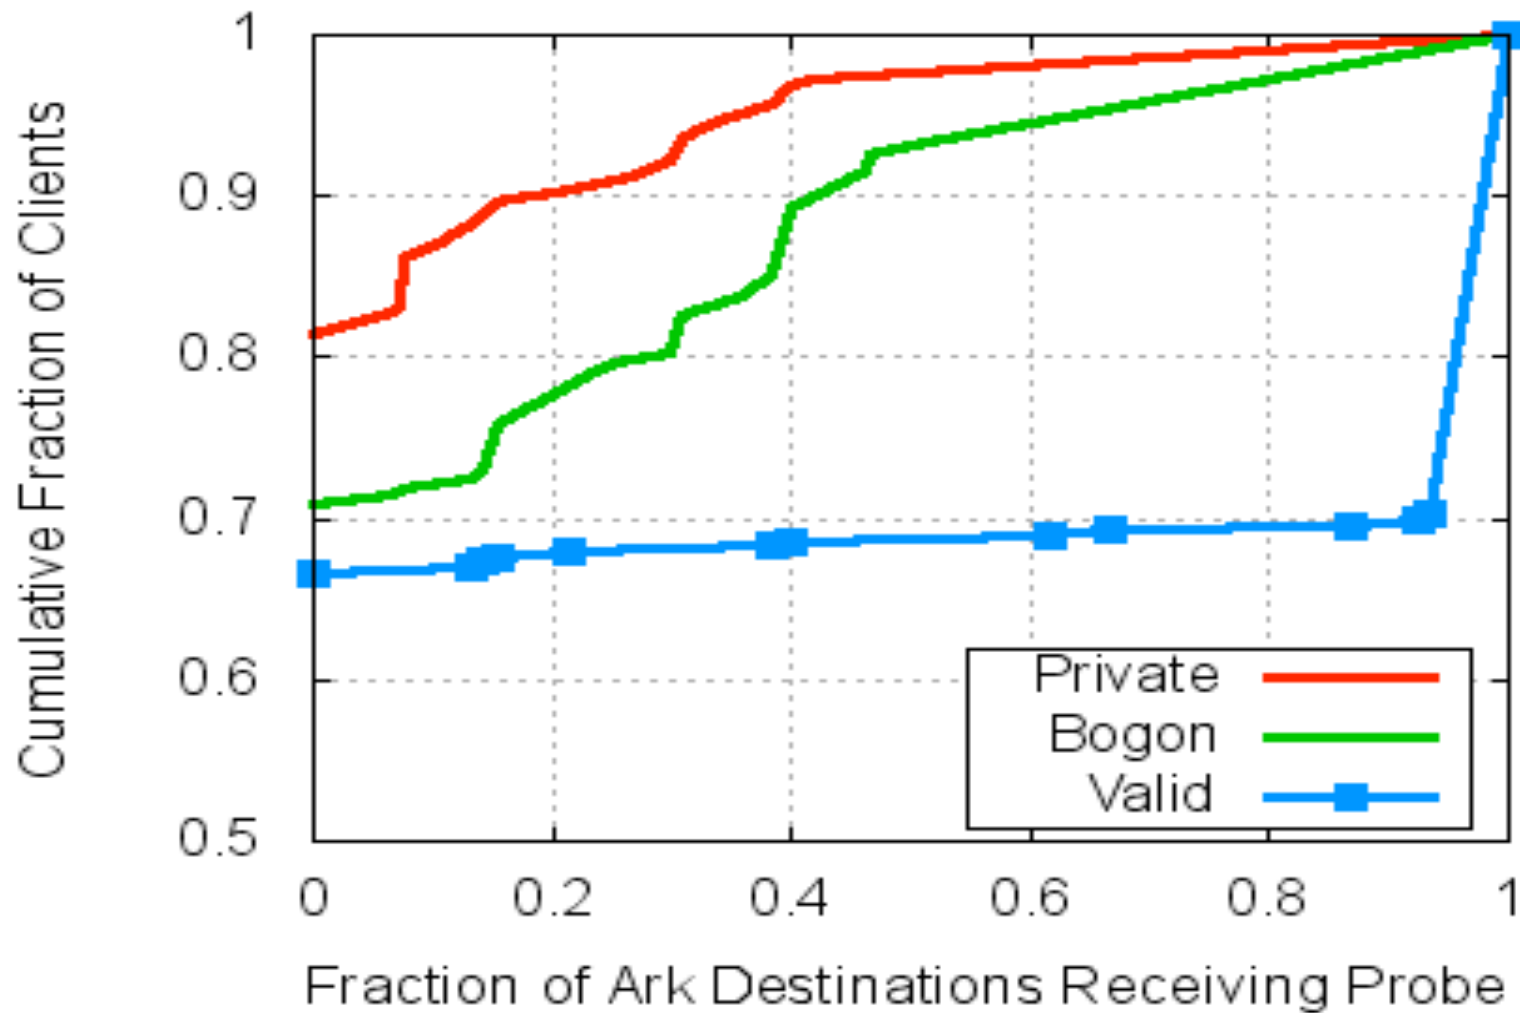
- ~12,000 unique tests
- 132 countries present in data set
- Don't claim zero bias
- Do claim diverse and representative



Questions

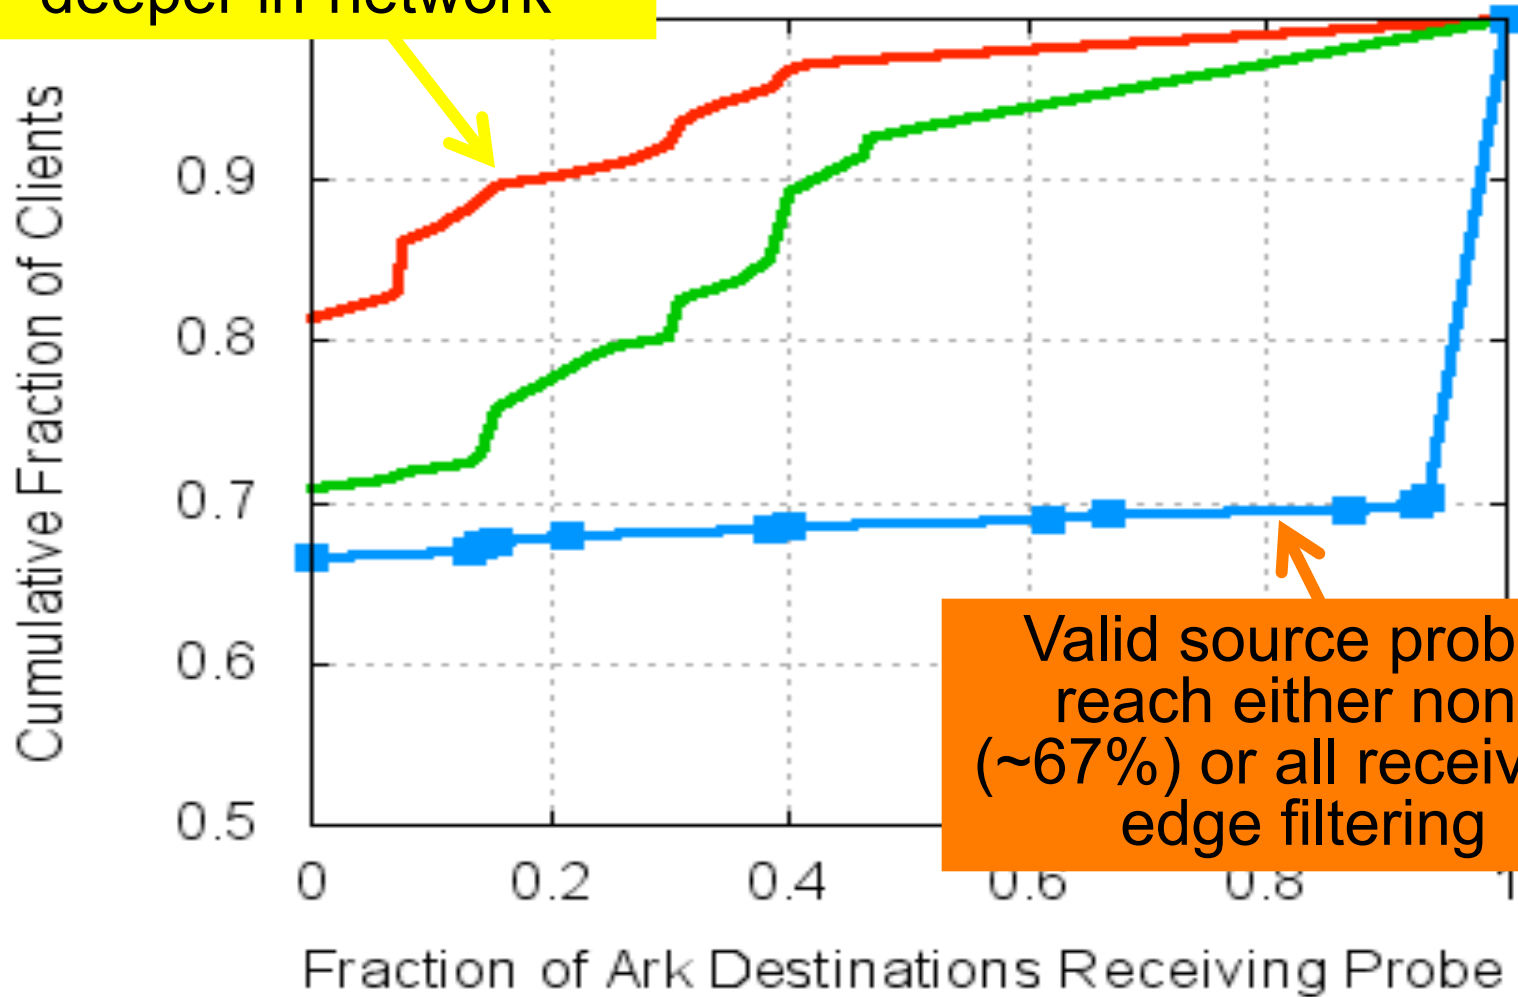
- Are there filtering variations among paths?
- What filtering methods are used?
- Where in network is source validation?
- Granularity of filtering?
- How vulnerable is the Internet?
- How has filtering evolved over >4 years?

Path-based Filtering Variation?



Surprising variation among bogon and private sources: filtering deeper in-network

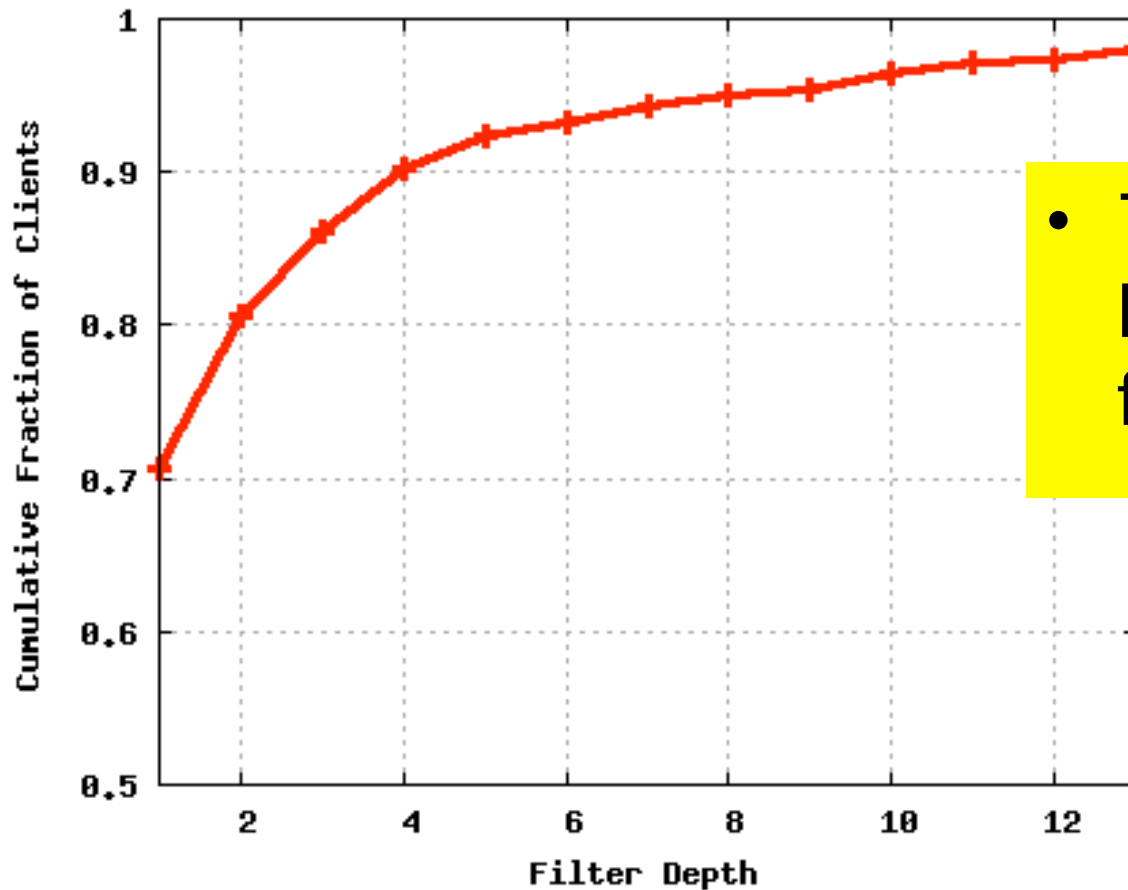
Variation



Valid source probes reach either none (~67%) or all receivers: edge filtering

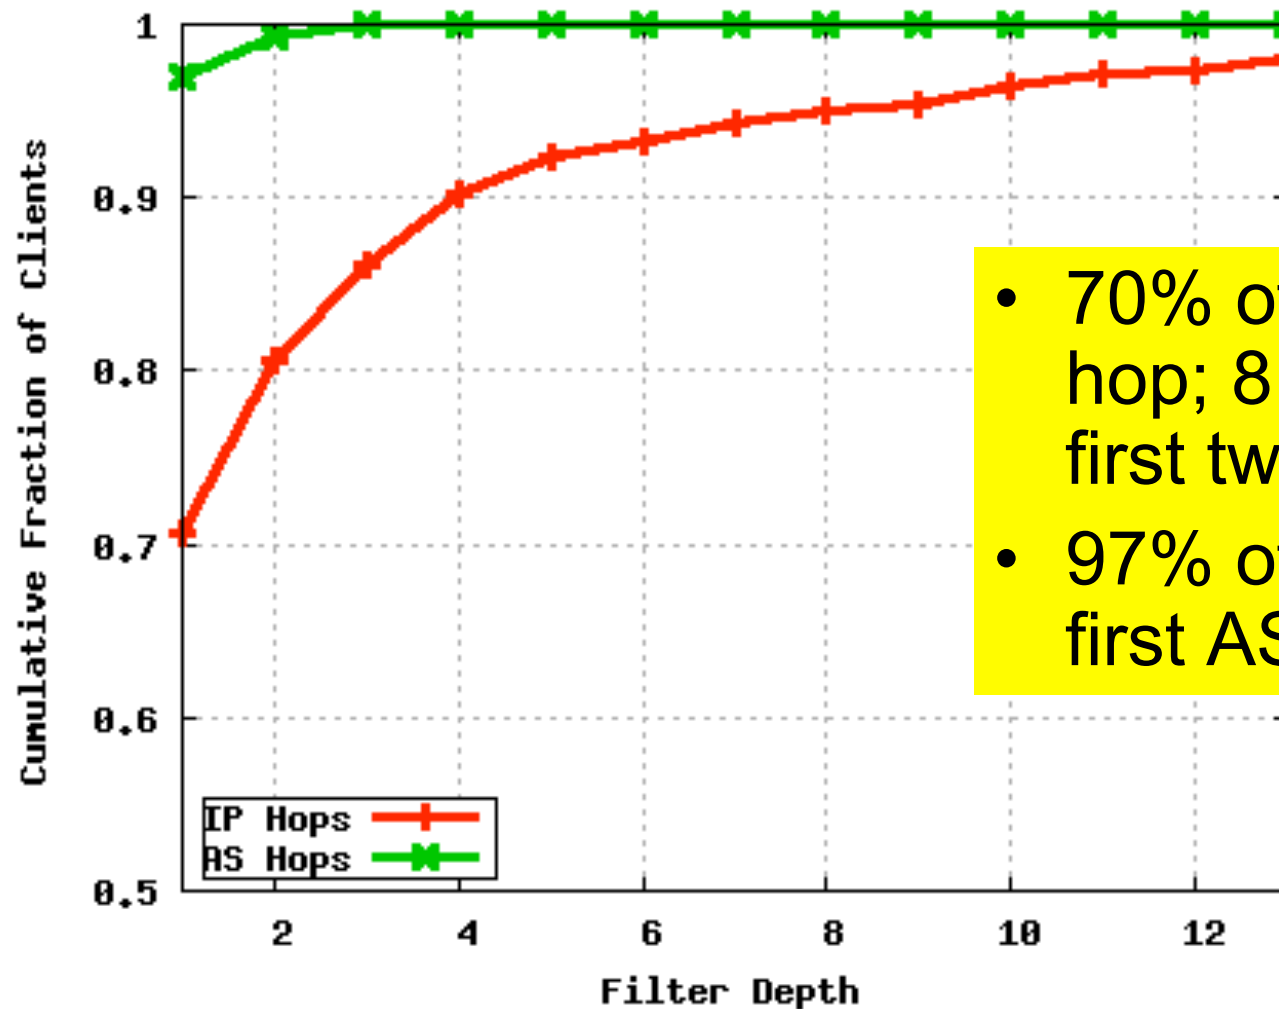
Where is source validation?

- tracefilter results:




- 70% of filters at 1st hop; 81% within first two hops

tracefilter Results



- 70% of filters at 1st hop; 81% within first two hops
- 97% of filters within first AS

tracefilter Results




- 70% of filters at 1st hop; 81% within first two hops
- 97% of filters within first AS

If a spoofed packet passes through first two hops, likely to travel unimpeded to destination

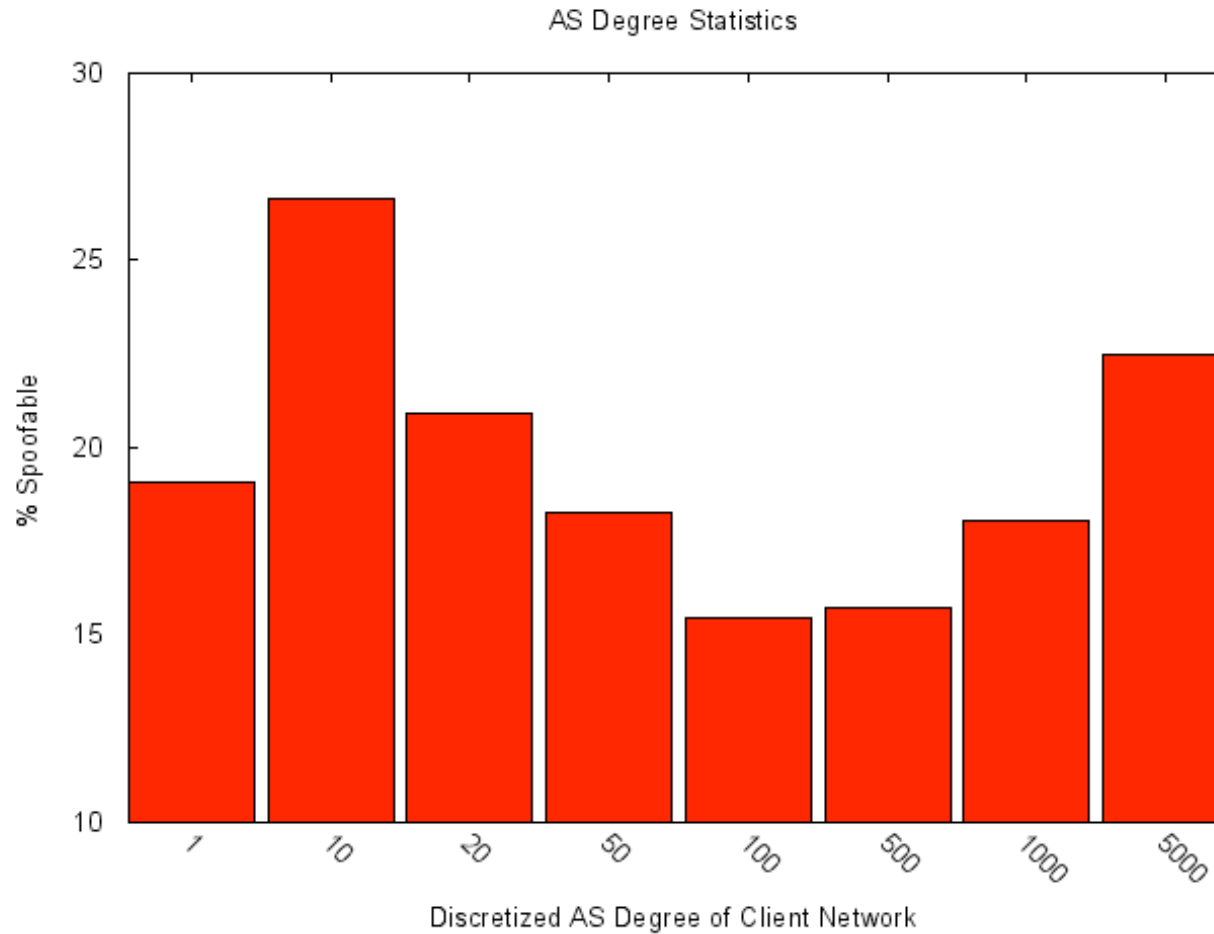
Filtering Granularity

- Clients test own IP $\oplus (2^n)$ for $0 < n < 24$
- Filtering on a /8 boundary enables a client within that network to spoof ~16M addresses



~70% of clients unable to spoof test sources can spoof neighbors

AS Degree



- Small or large providers filtering?
- Surprisingly, no clear trend
- Work required across the board (or a new solution)

Evolution of Spoofability

- Find two three-month periods with large and comparable sample sizes

	Proportion Spoofable		
Metric	2005 (single dest)	2009 (single dest)	2009 (all dests)
Sessions	18.8%	29.9%	31.2%
Netblocks	20.0%	30.2%	31.7%
Addresses	5.0%	11.0%	11.1%
ASes	23.4%	31.8%	34.1%

Evolution of Spoofability

- Find two three-month periods with large and comparable sample sizes

	Proportion Spoofable		
Metric	2005 (single dest)	2009 (single dest)	2009 (all dests)
Sessions	18.8%	29.9%	31.2%
Netblocks	20.0%	30.2%	31.7%
Addresses	5.0%	11.0%	11.1%
ASes	23.4%	31.8%	34.1%

Less filtering
four years later

Change not attributable
to increasing number
of destinations

Spoofers Project

- Background
- Recent Relevance
- Project Methodology
- Results
- Parting Thoughts

Parting Thoughts

- Even after all these years, source spoofing problem not solved. It's the incentives:
 - Provider can follow BCP38 and still receive anonymous, spoofed traffic
 - Others can spoof a provider's address space
 - Disincentive in form of accidental blocking
- Single unfiltered ingress can compromise entire Internet system
 - Can we plug every hole?
 - Regulatory Response? ... but multinational?
 - Spoofer page for public provider flogging?

Parting Thoughts

- Tracefilter exposes operational tension between filtering incentives and managing edge filters
- If a spoofed packet isn't filtered at edge, will travel unimpeded to destination
- Needed?
 - Filtering in the core
 - Clean slate design
- Think (seriously) about alternate techniques?
 - StackPI [Yaar, Perrig, Song 2006]
 - Passport [Liu, Li, Yang, Wetherall 2008]
 - Others?

Parting Thoughts

- Tracefilter exposes operational tension between filtering incentives and managing edge filters
- If a spoofed packet isn't filtered at edge, will travel unimpeded to destination
- Ne
 - F
 - C
- Think (seriously) about alternate techniques?
 - StackPI [Yaar, Perrig, Song 2006]
 - Passport [Liu, Li, Yang, Wetherall 2008]
 - Others?

Thanks!