IP Spoofer Project

Observations on four-years of data

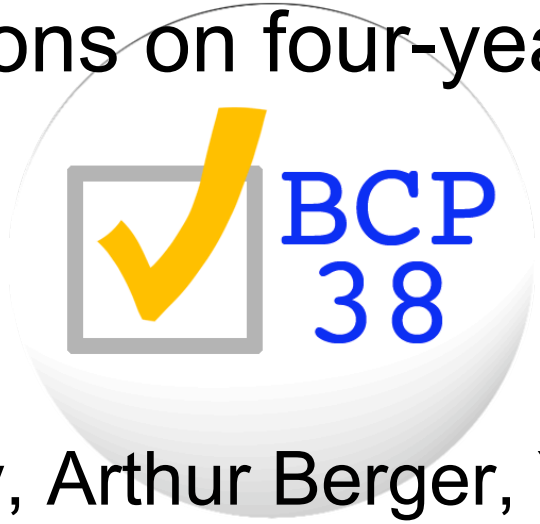

Rob Beverly, Arthur Berger, Young Hyun

{rbeverly,awberger}@csail.mit, youngh@caida

ISMA AIMS 2009

February 12, 2009

CSAIL

MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Project

- Background
- Recent Relevance
- Project Description
- What's New: Methodology
- What's New: Data
- Parting Thoughts

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

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, Bellovin89]
- Still an attack vector?

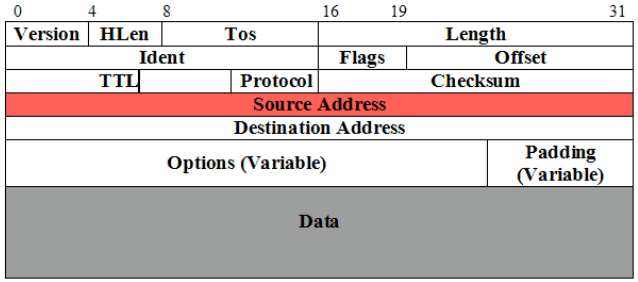

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Circa 2004…

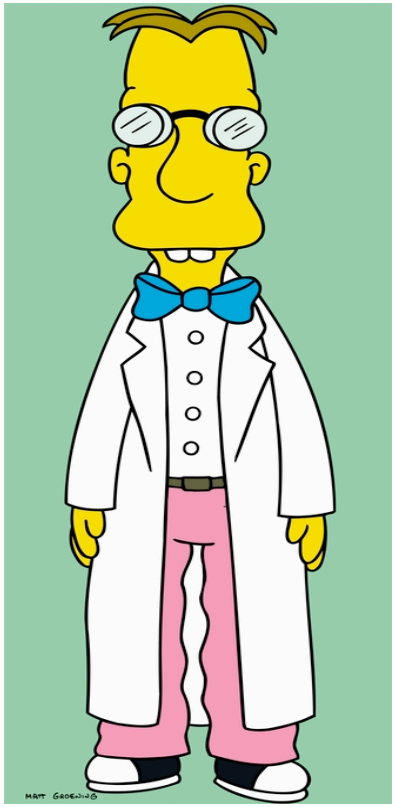

IP Source Spoofing doesn't matter!

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

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Circa 2004…

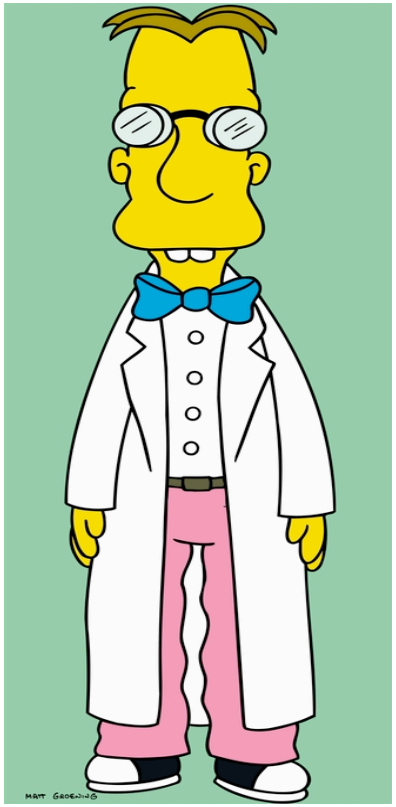

IP Source Spoofing doesn't matter!

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

!?!?!

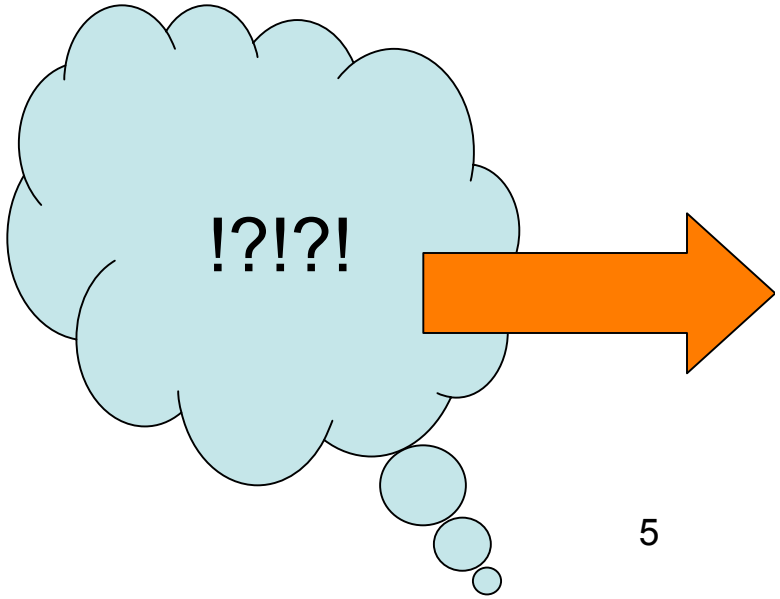

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

The Spoofer Project

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

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

spoofer.csail.mit.edu

- Internet-wide active measurement effort:
  - Quantify the extent and nature of Internet source address filtering
- We learn and form inferences over:
  - Filtering policies/currently employed defenses
  - Filtering specificity, locations, providers, etc.
  - Distribution of filtering
- Began Feb. 2005

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Project

- Background
- Recent Relevance
- Project Description
- What's New: Methodology
- What's New: Data
- Parting Thoughts

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

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 tr[illegible]
- Consider[illegible] Slide from SRUTI 2005
  - Today ([illegible]) [illegible] 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

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Prominent 2008 Example: DNS Amplifier Attack

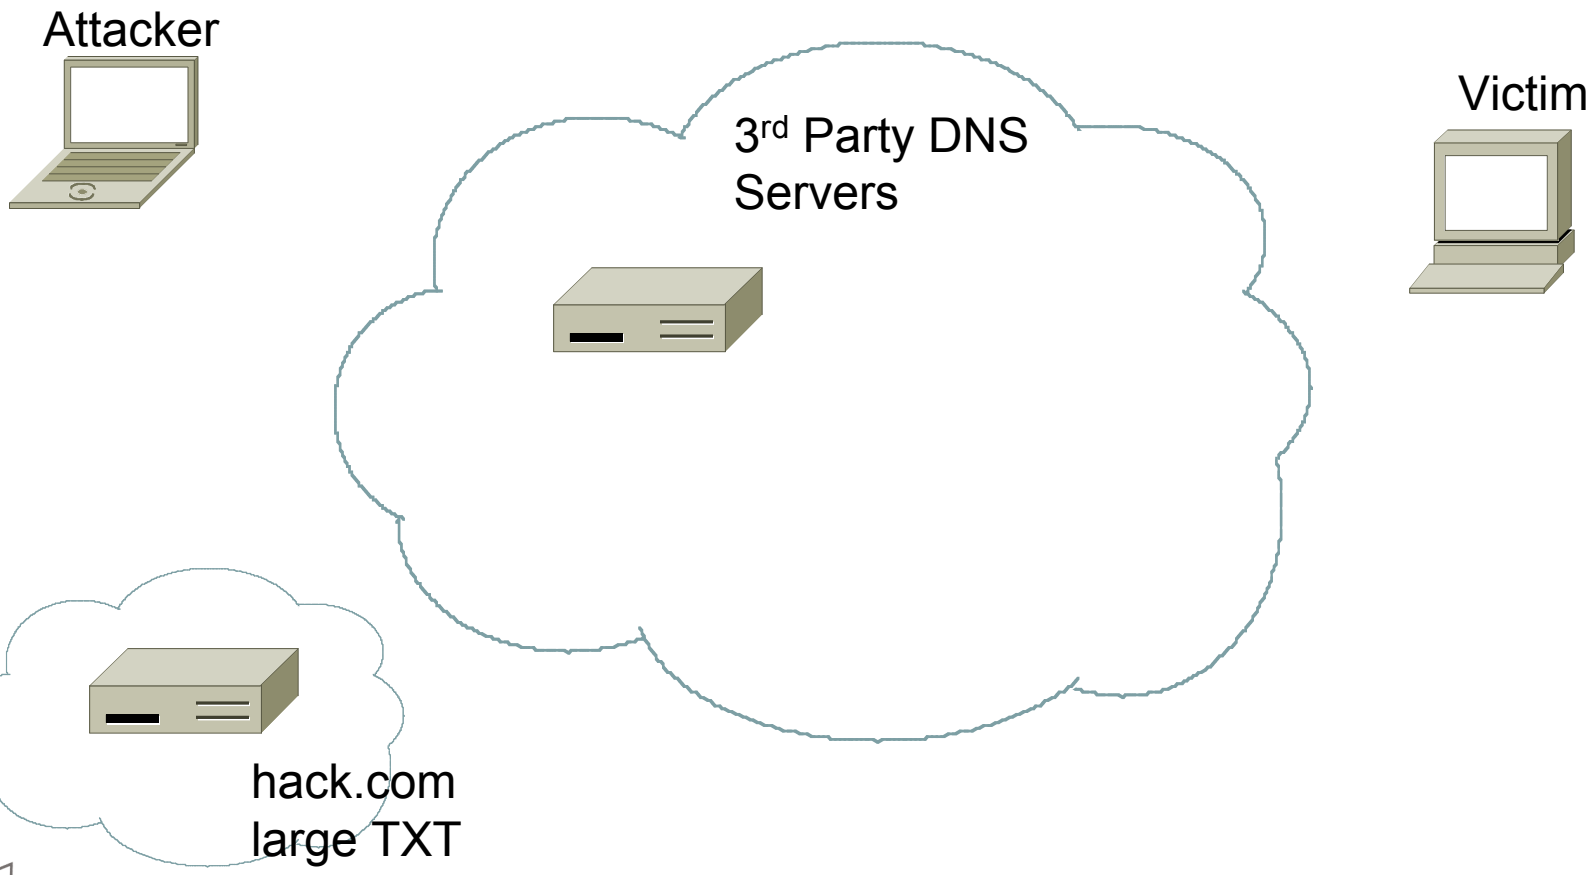

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Prominent 2008 Example: DNS Amplifier Attack

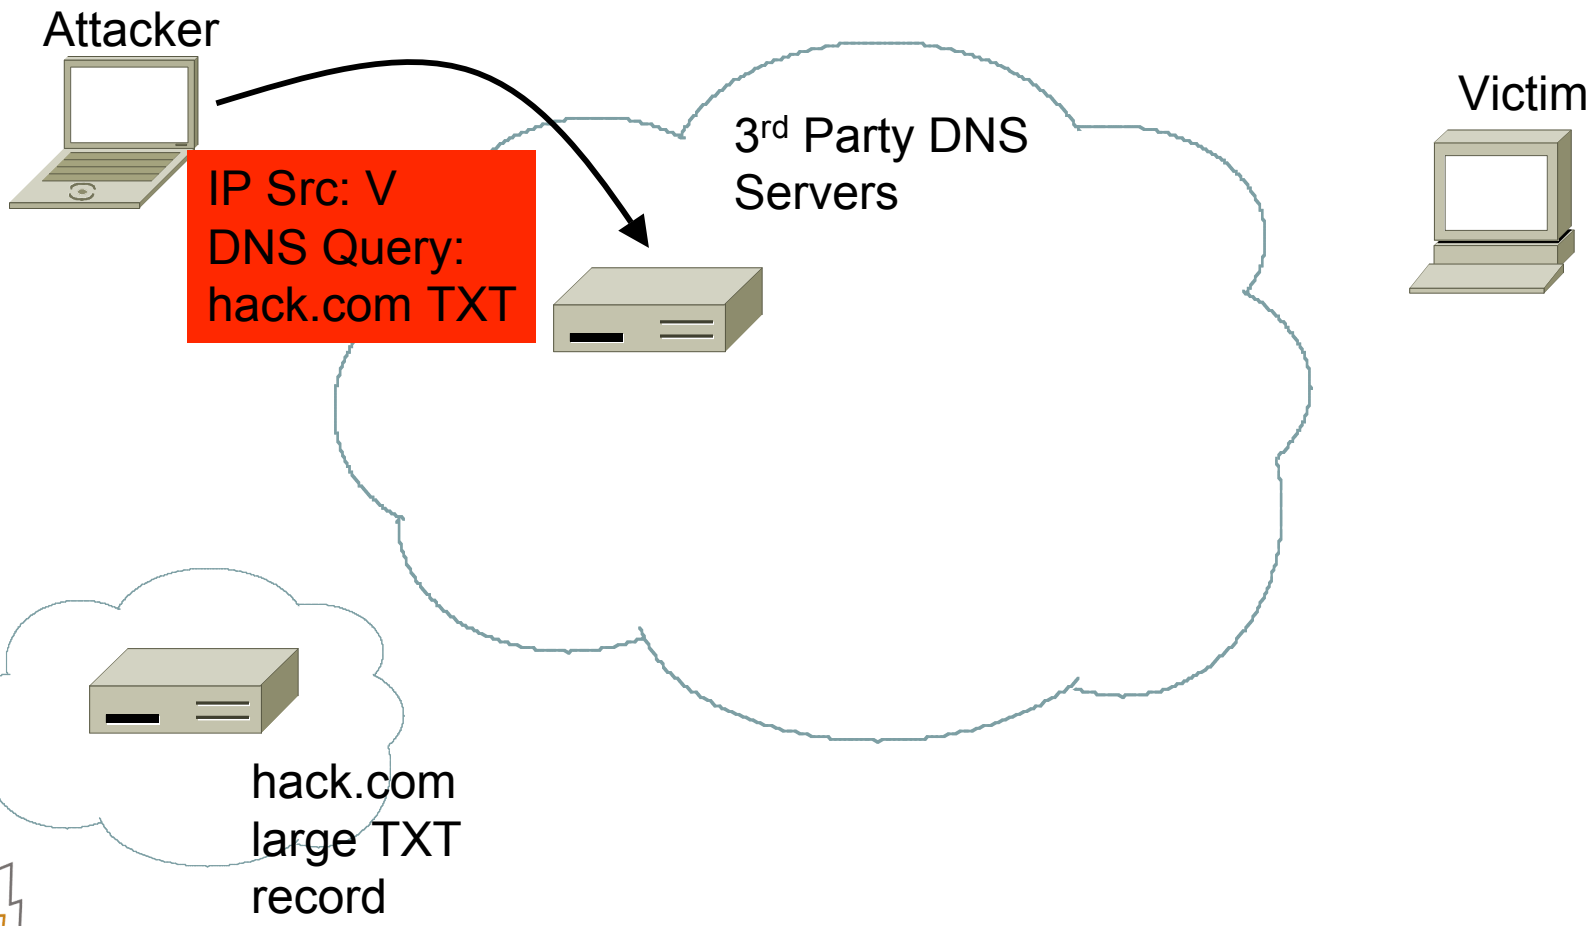

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Prominent 2008 Example: DNS Amplifier Attack

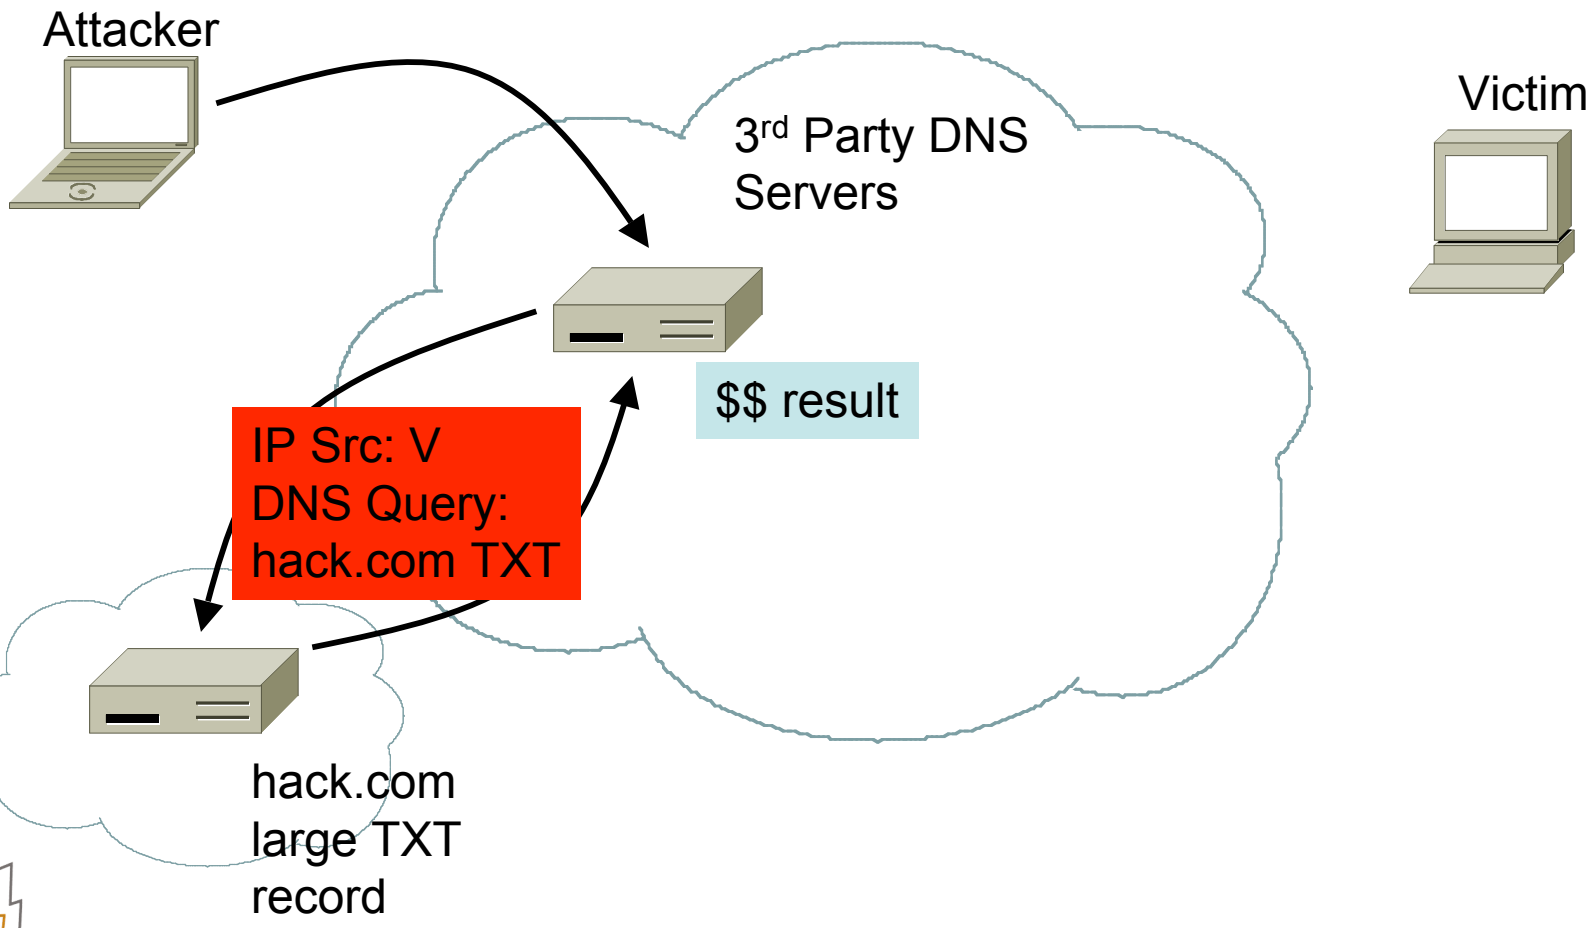

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Prominent 2008 Example: DNS Amplifier Attack

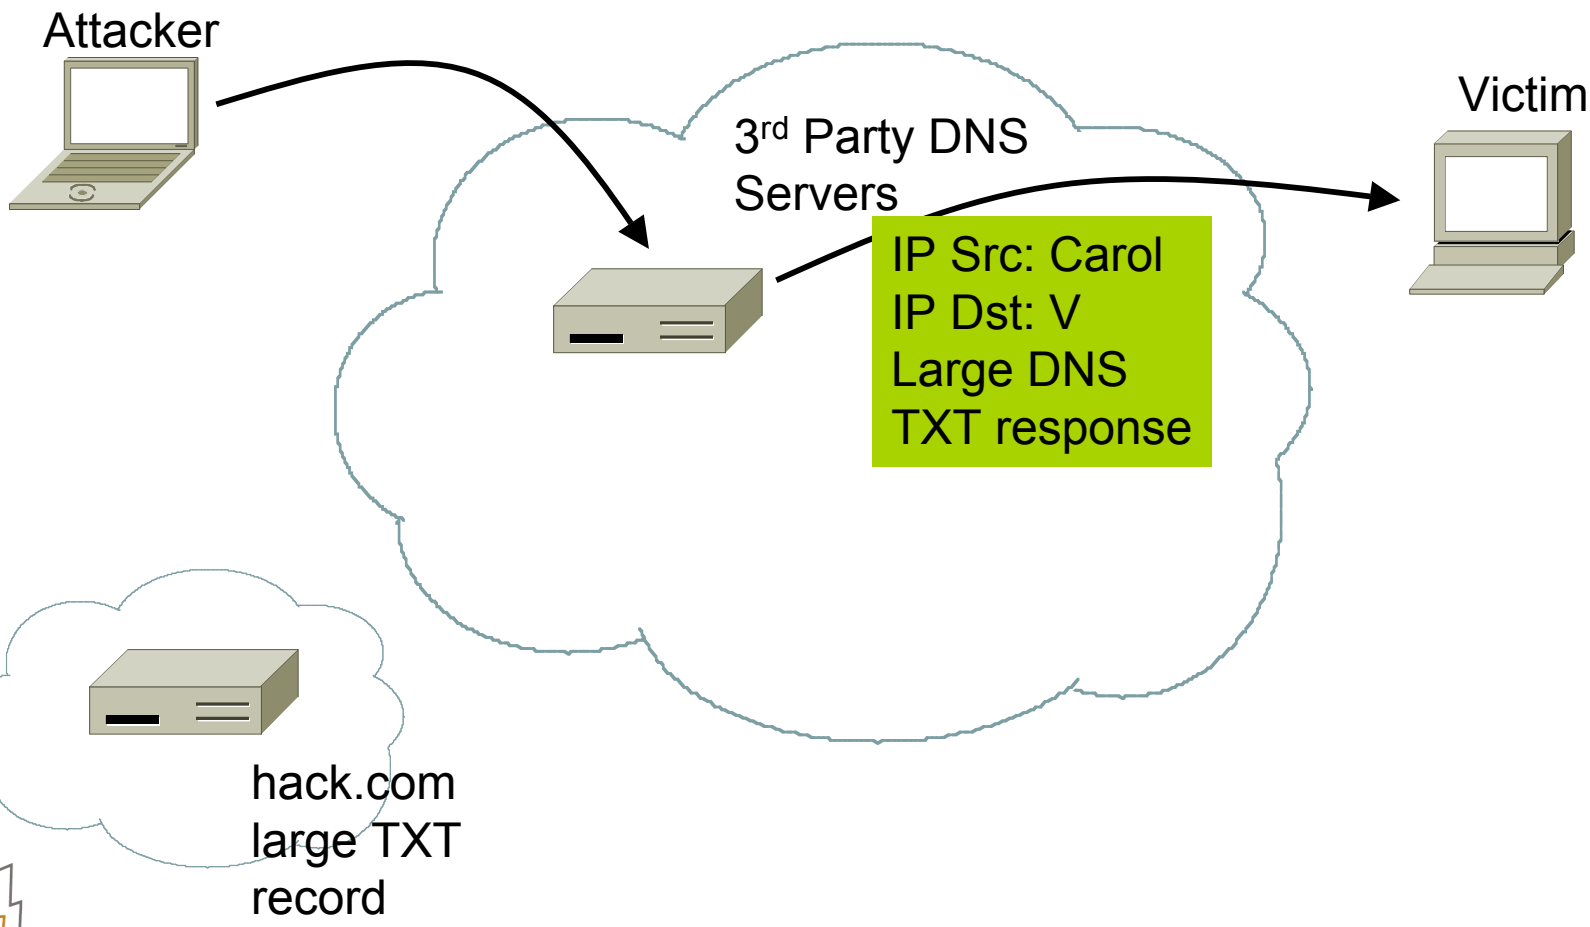

Prominent 2008 Example: DNS Amplifier Attack

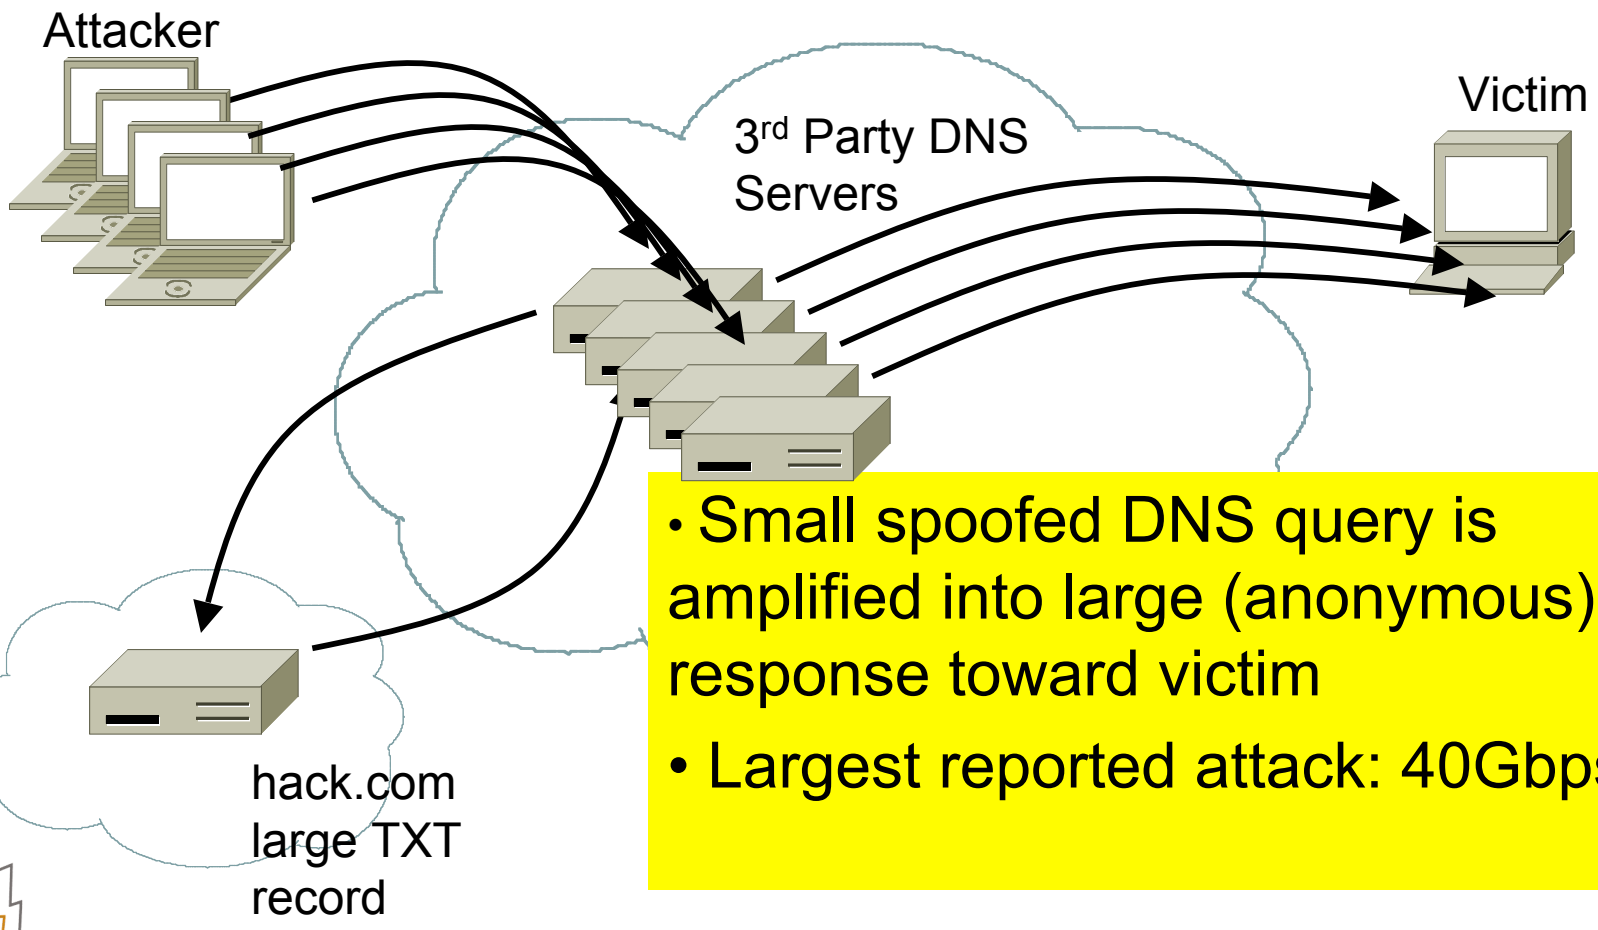

- Small spoofed DNS query is amplified into large (anonymous) response toward victim
- Largest reported attack: 40Gbps*

*Arbor networks 2008 infrastructure security survey

Reasons to Believe Spoofing Matters (2009)

- DNS Amplifier Attacks
- In-Window TCP Reset Attacks
- Spam Filter Circumvention
- DNS Cache Poisoning
- UW reverse traceroute
- Spoofer web site statistics

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

The Operational Side

- Arbor:
  - "Reflective amplification attacks responsible for the largest attacks 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?

*Arbor networks 2008 infrastructure security survey

Operational View

- Not all sources are created equal
- IETF BCP38 best filtering practice

Example Source IP	**Description**	**Possible Defense**
1.2.3.4	Unallocated	Bogon Filters
6.1.2.3	Valid (In BGP table)	uRPF
192.168.1.1	RFC1918 private	Static ACL
Client IP ⊕ (2^N)	Neighbor Spoof	Switch, DOCSIS

IPv4 Address Space

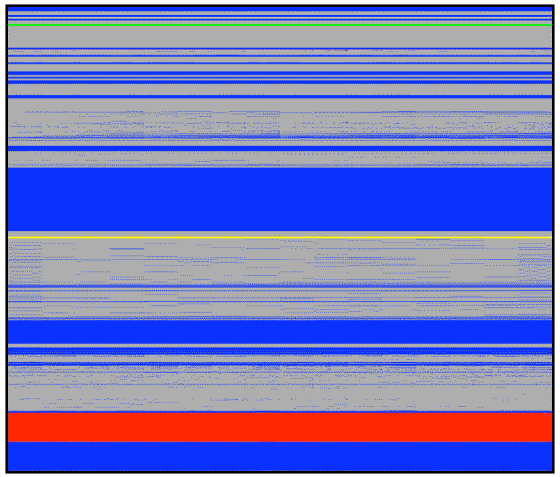

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Operational View

- We have defenses, what's the problem?
- BCP38 suffers from:
  - Incentive problem
  - Lack of hardware support (see NANOG)
  - Management nightmare (edge filters)

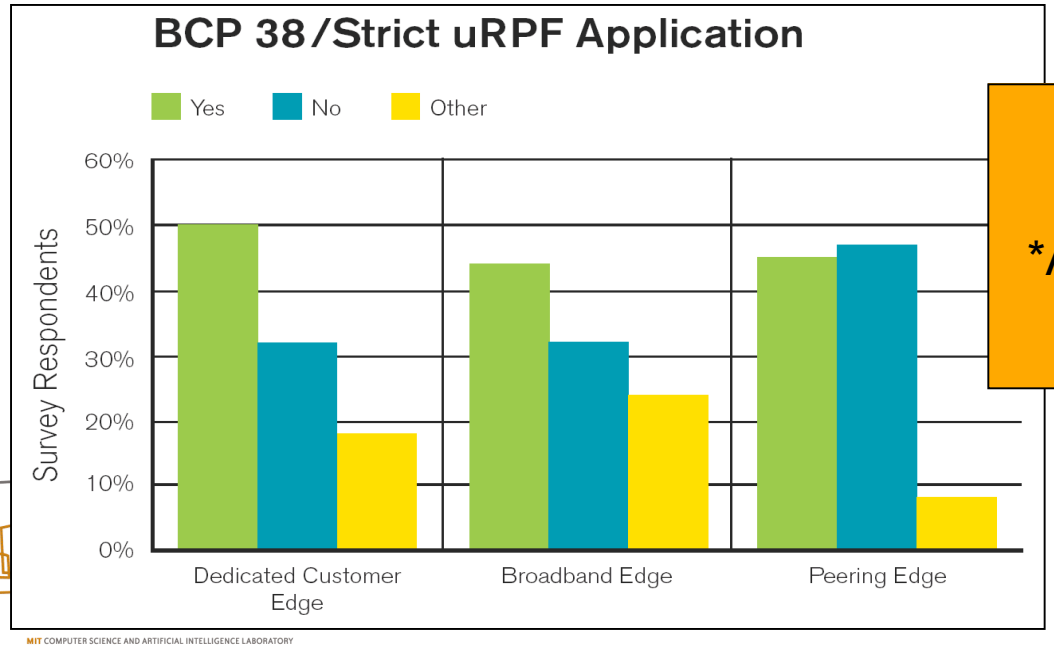

> 30% don't filter!

*Arbor networks 2008 infrastructure security survey

Spoofer Project

- Background
- Recent Relevance
- Project Description
- What's New: Methodology
- What's New: Data
- Parting Thoughts

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Test Client

- Willing participants run “spoofer” client to test policy, perform inference, etc.
  - Binaries, source publicly available

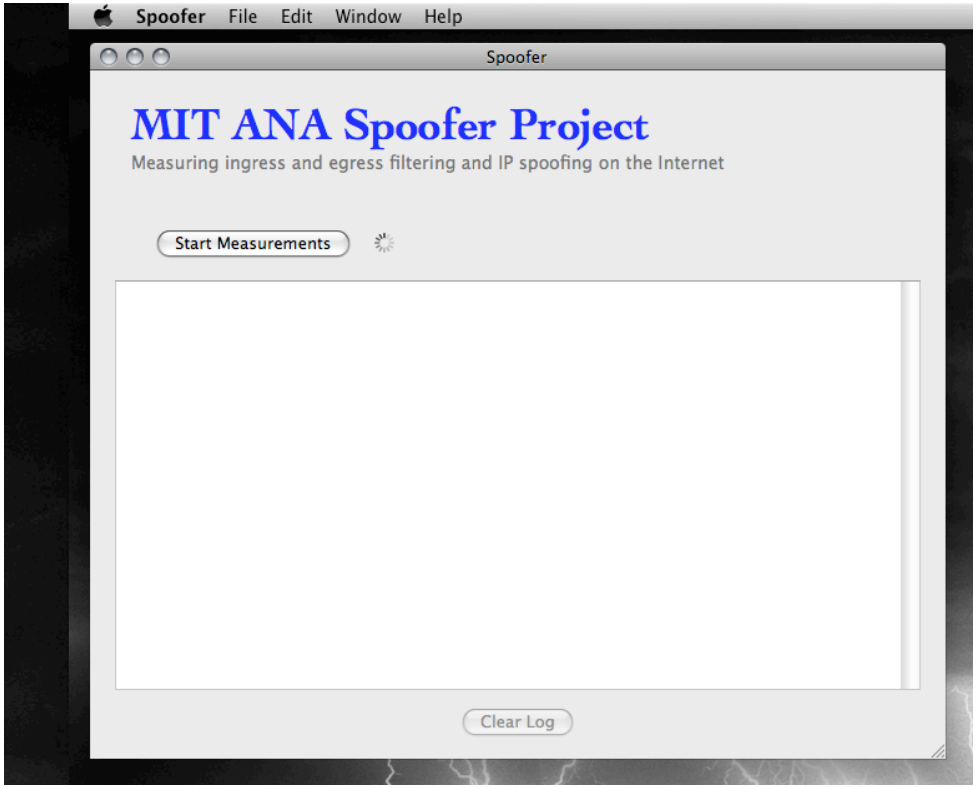

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Operation

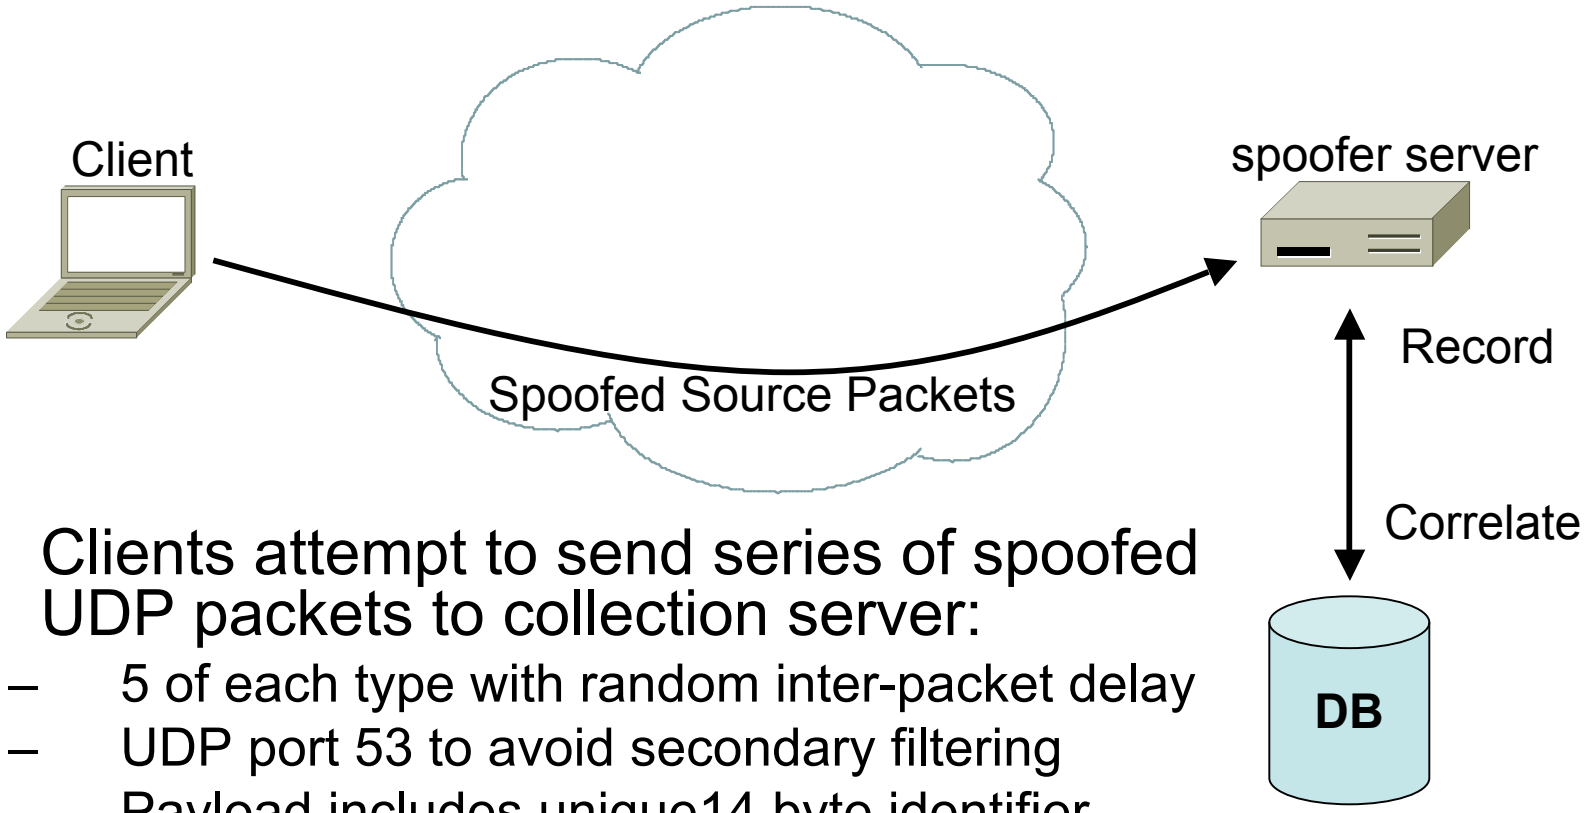

- Clients attempt to send series of spoofed UDP packets to collection server:
  - 5 of each type with random inter-packet delay
  - UDP port 53 to avoid secondary filtering
  - Payload includes unique14 byte identifier
- Server stores received packets in DB

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Operation

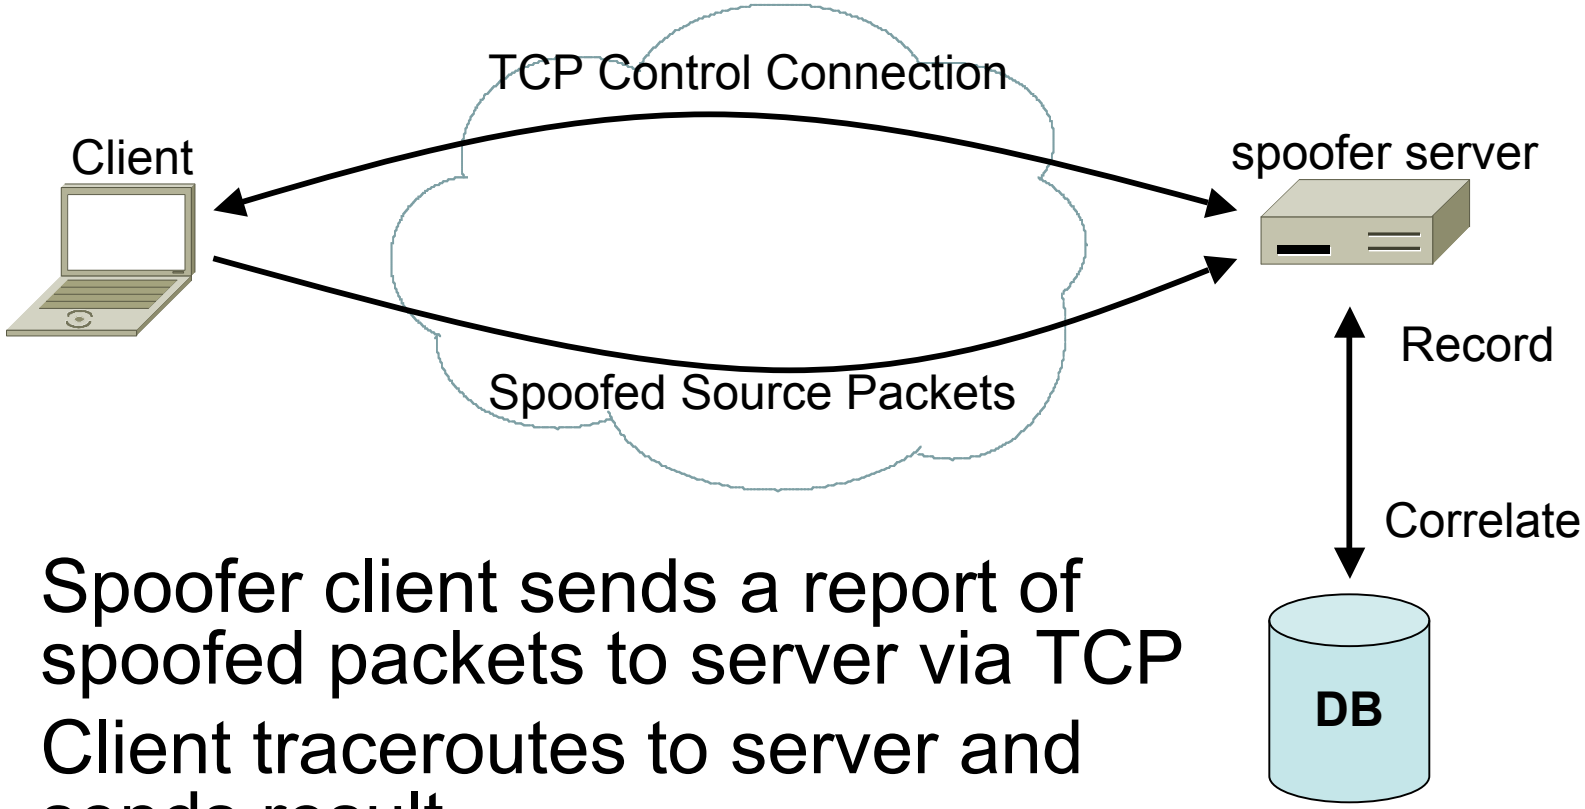

- Spoofer client sends a report of spoofed packets to server via TCP
- Client traceroutes to server and sends result
- TCP destination port 80 used to avoid secondary filtering effects

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Client Population

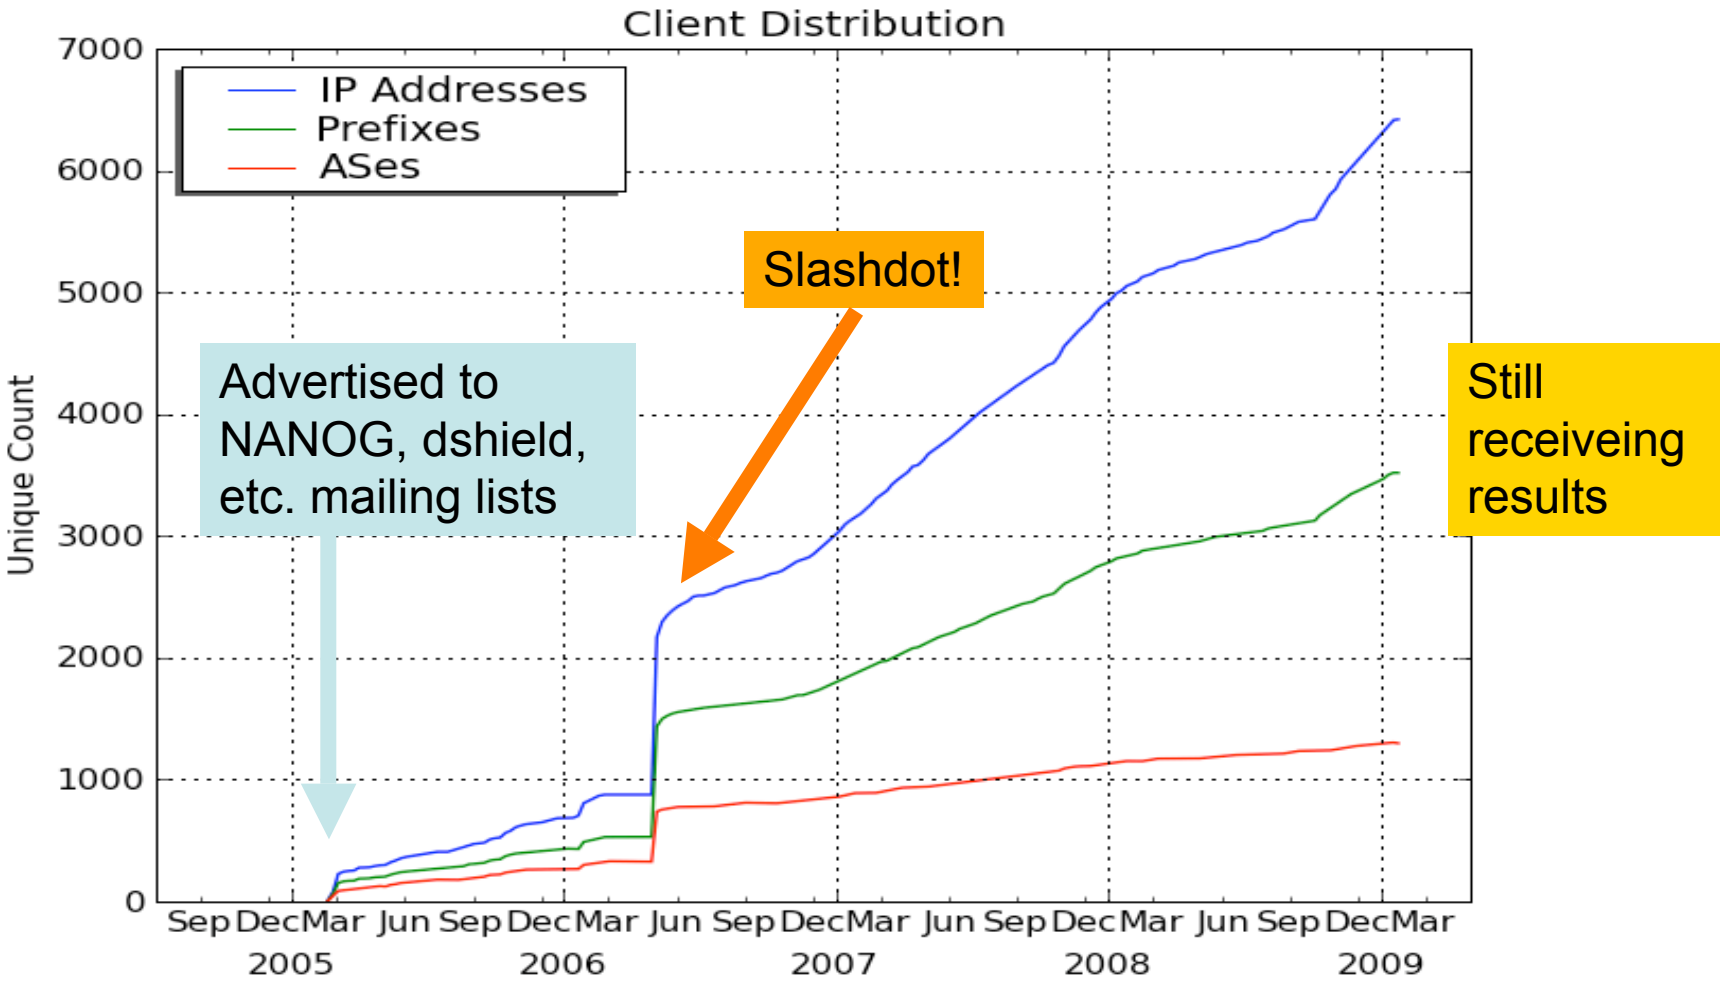

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Client Population Distribution

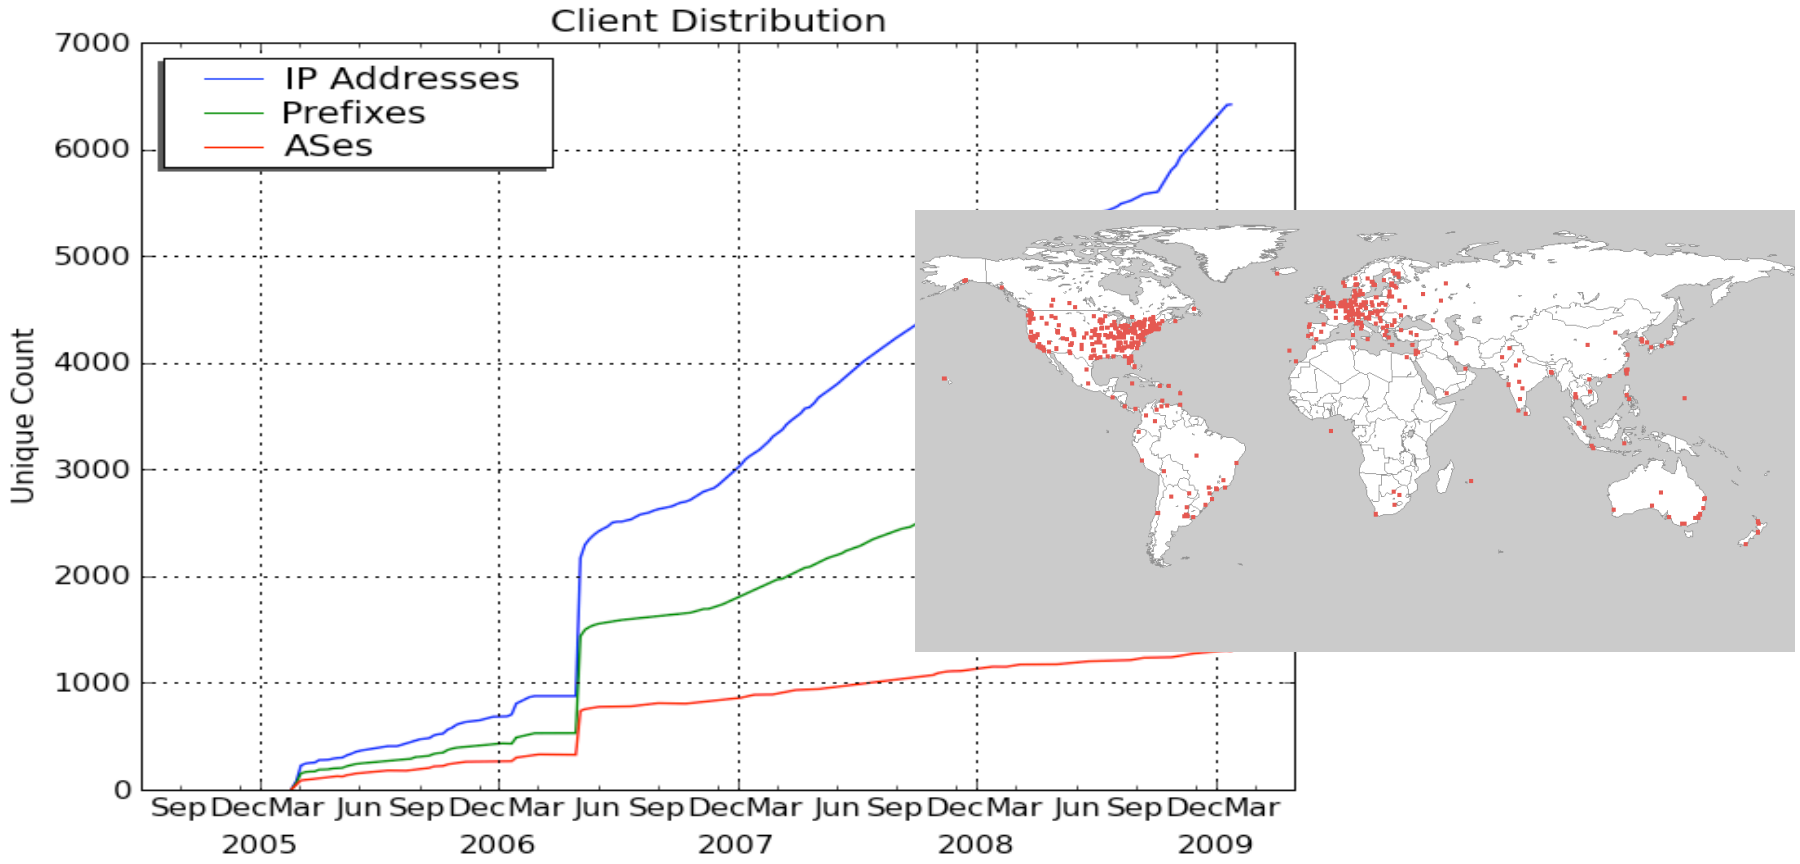

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Filtering Specificity

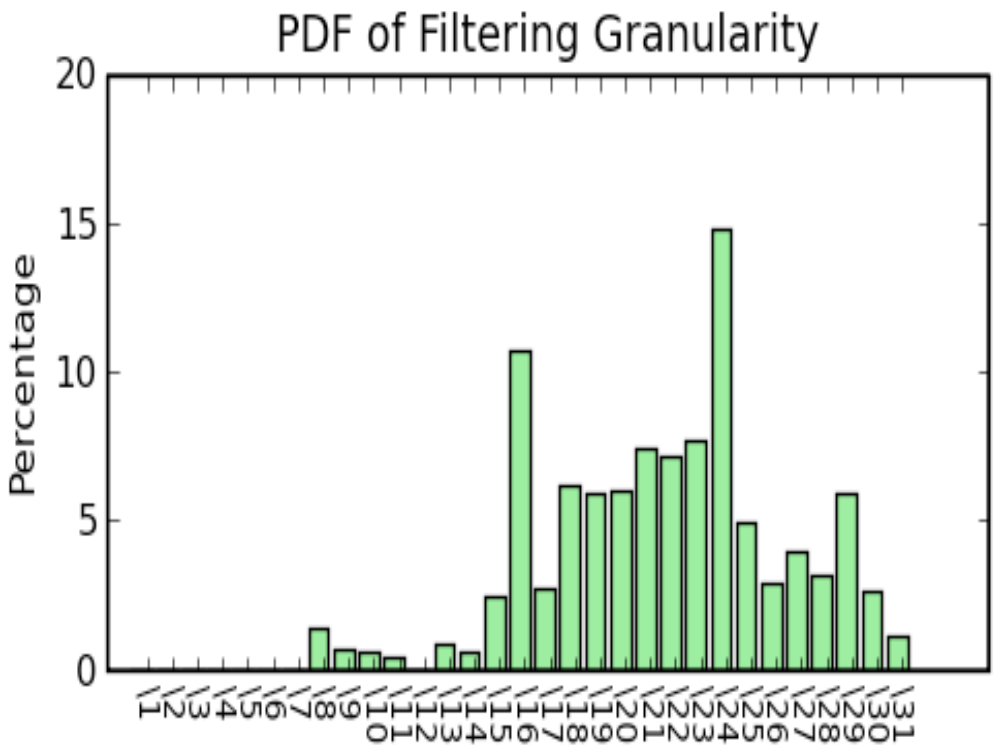

- Clients test own IP ⊕ (2^n) for 0<n<24
- Filtering on a /8 boundary enables a client within that network to spoof ~16M addresses
- >30% of clients "unable" to spoof can spoof neighbors
- Exclude "neighbor spoof" from macro results

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

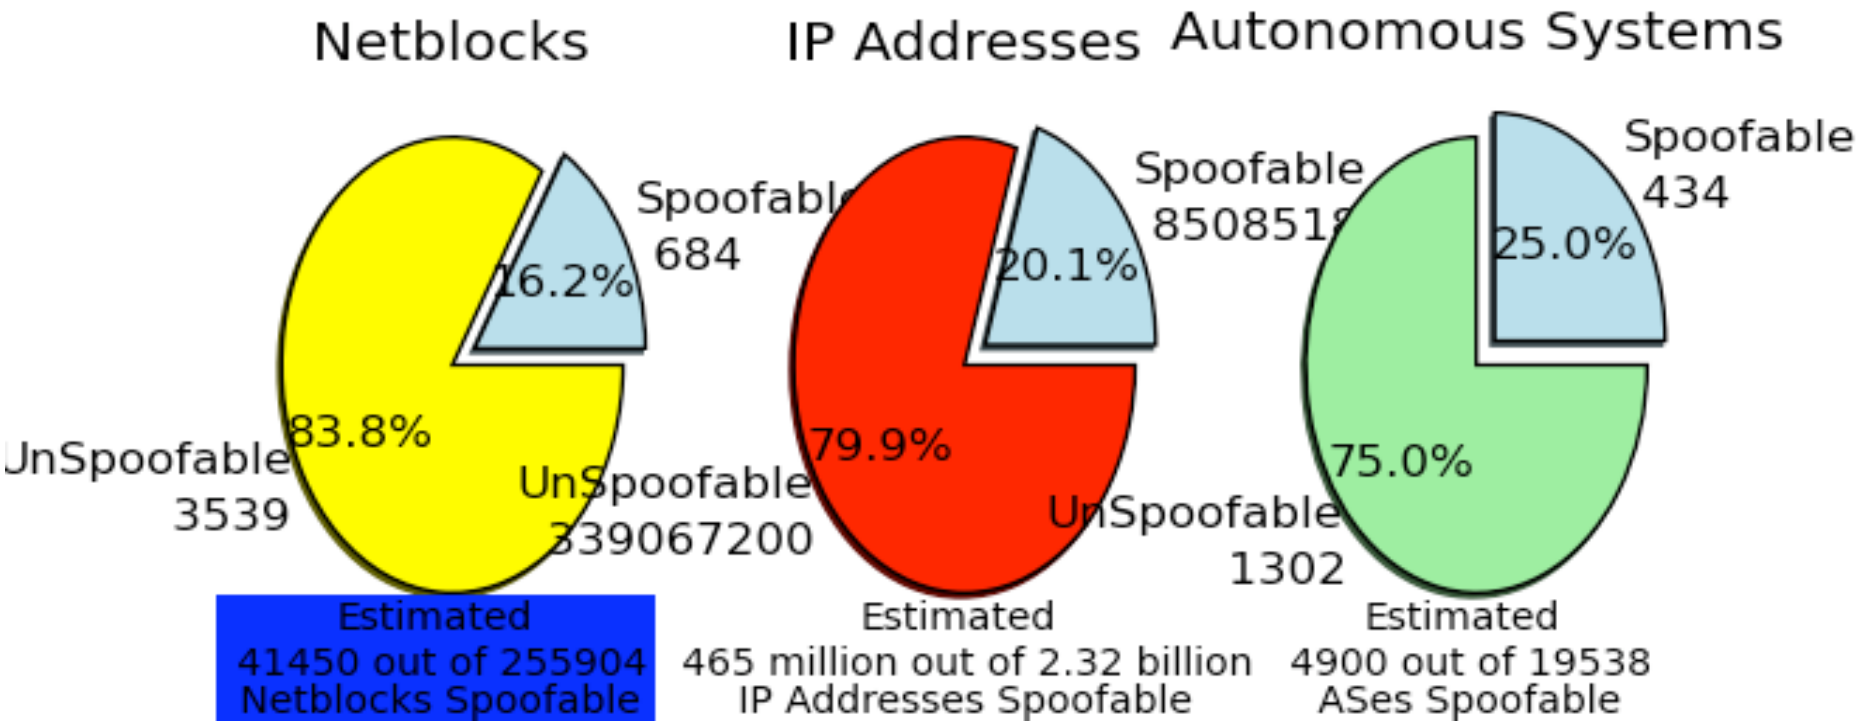

- **Spoofable:** spoofing of private, unallocated, or valid IP packets possible from these locations
- Agrees to a first-order with Arbor survey
- But… these numbers cause even more disagreement!

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Project

- Background
- Recent Relevance
- Project Description
- What's New: Methodology
- What's New: Data
- Parting Thoughts

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

What's New: Methodology

- Goal:
  - Resolve ambiguity
  - Increase confidence
- New:
  - tracefilter
  - Tied into CAIDA's ark distributed measurement infrastructure
  - More detailed analysis
  - Longitudinal analysis over four-years of data

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter

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

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter

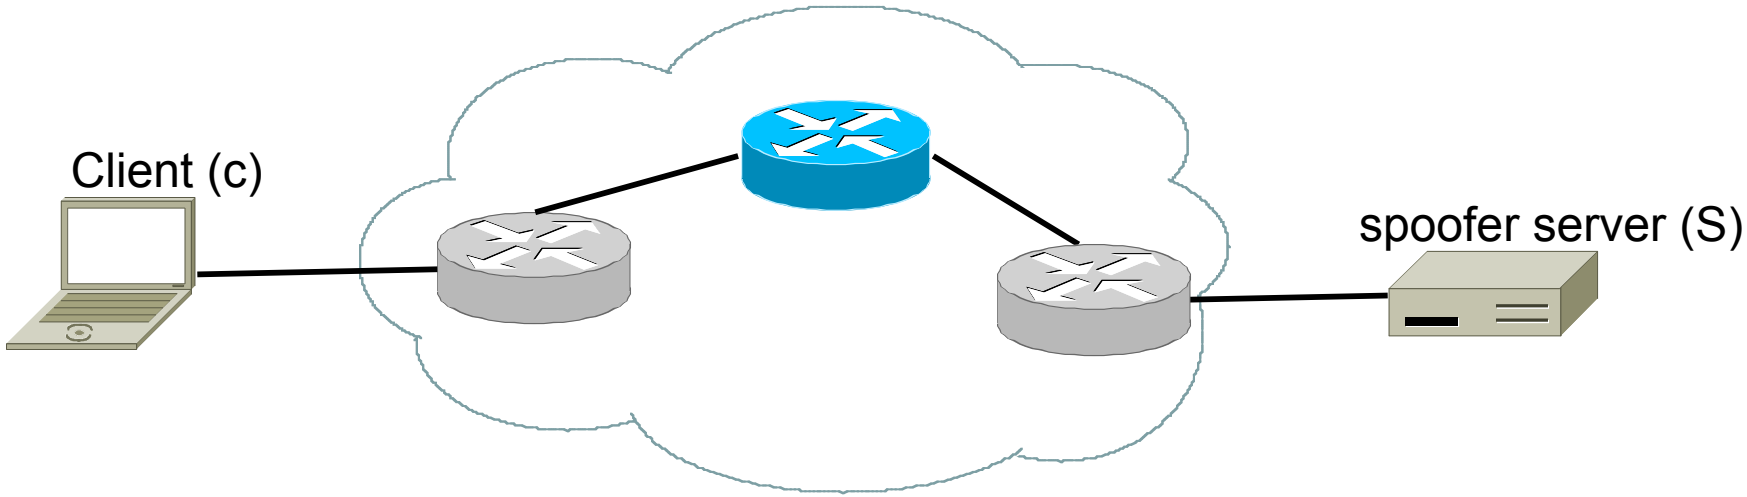

- Client c works in conjunction with our server S

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter

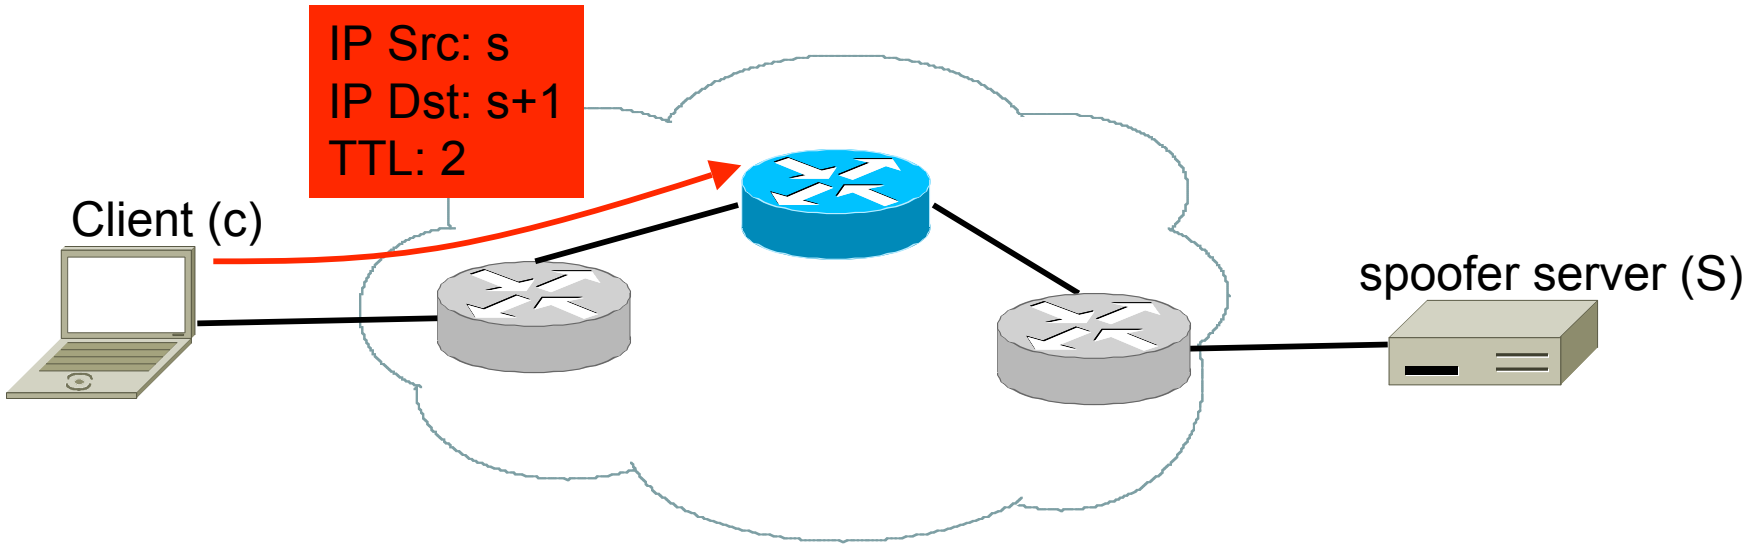

- c sends spoofed packet with:
- ttl=x, src=S, dst=S+1 for 0<x<pathlen

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter

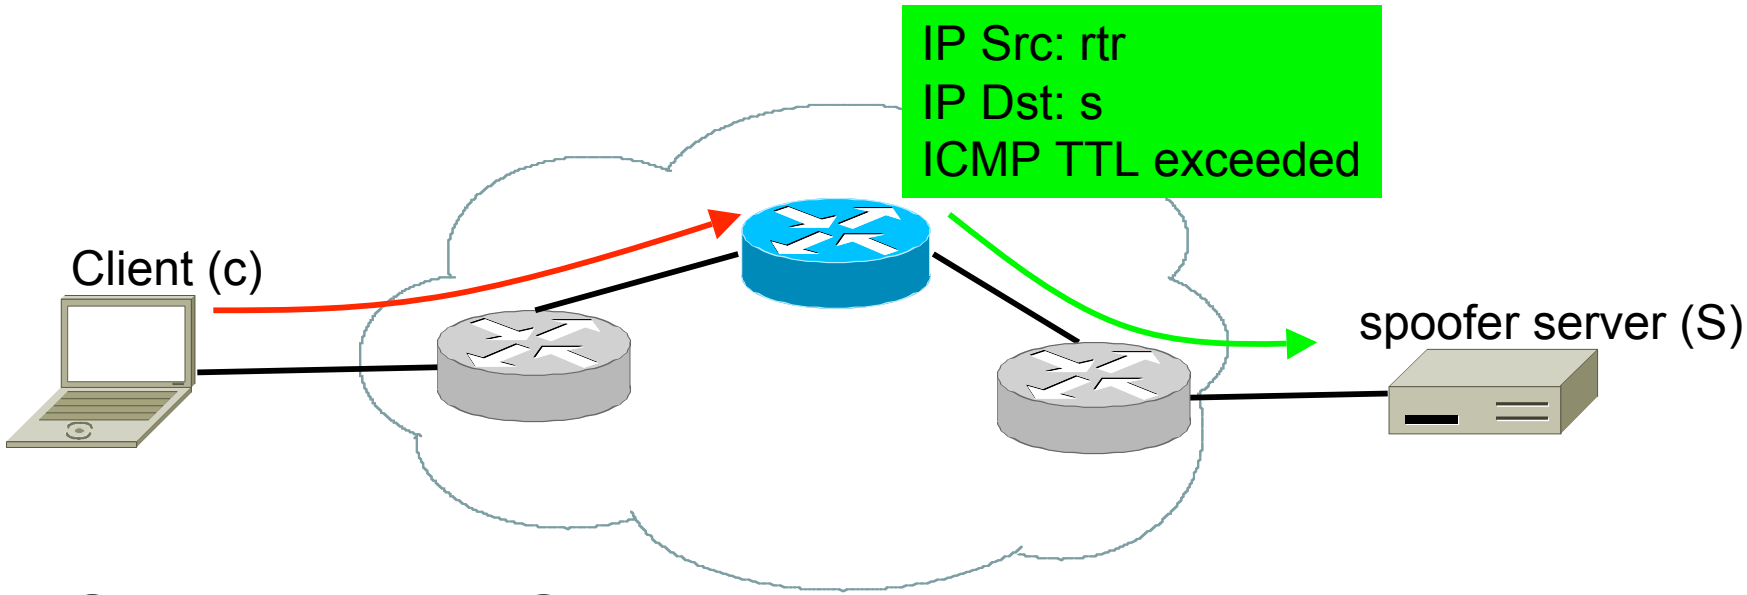

- S receives ICMP expiration messages from routers along path
- For each decoded TTL, S records which spoofed packets are received

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter

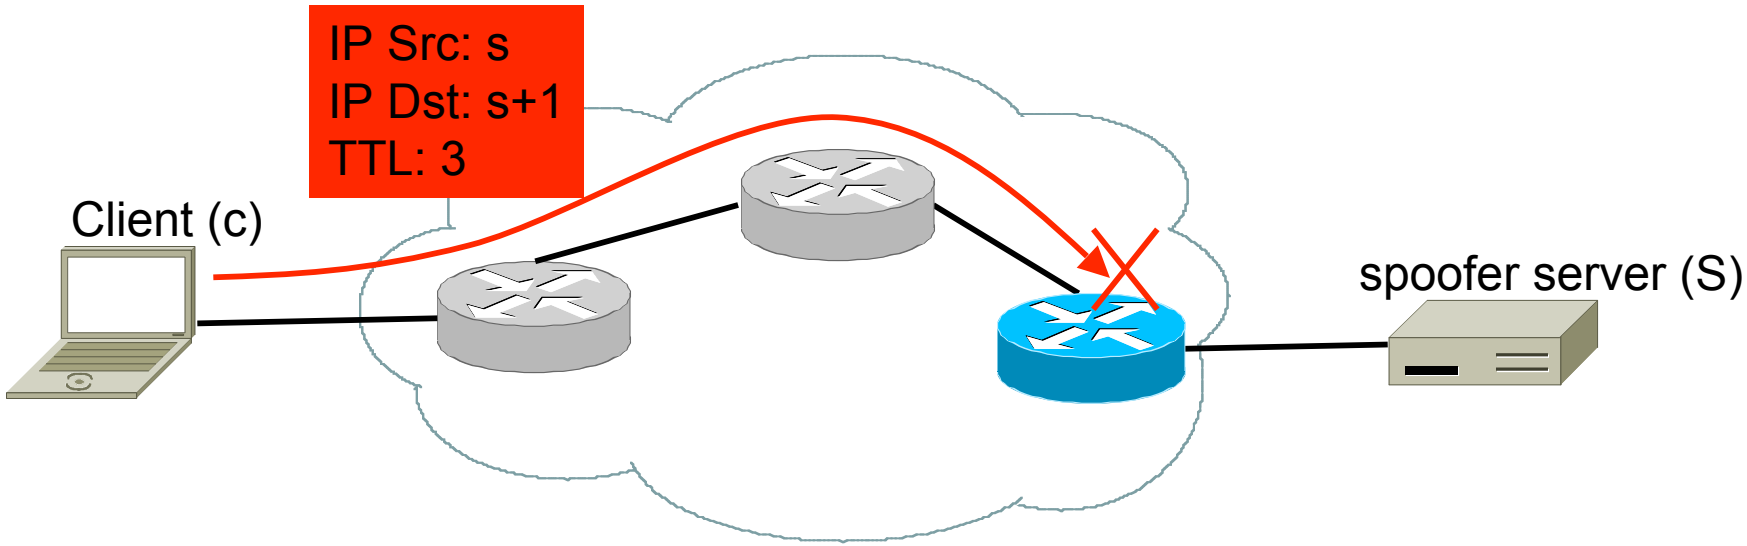

- Increase TTL, repeat
- Largest TTL indicates filtering point

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

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

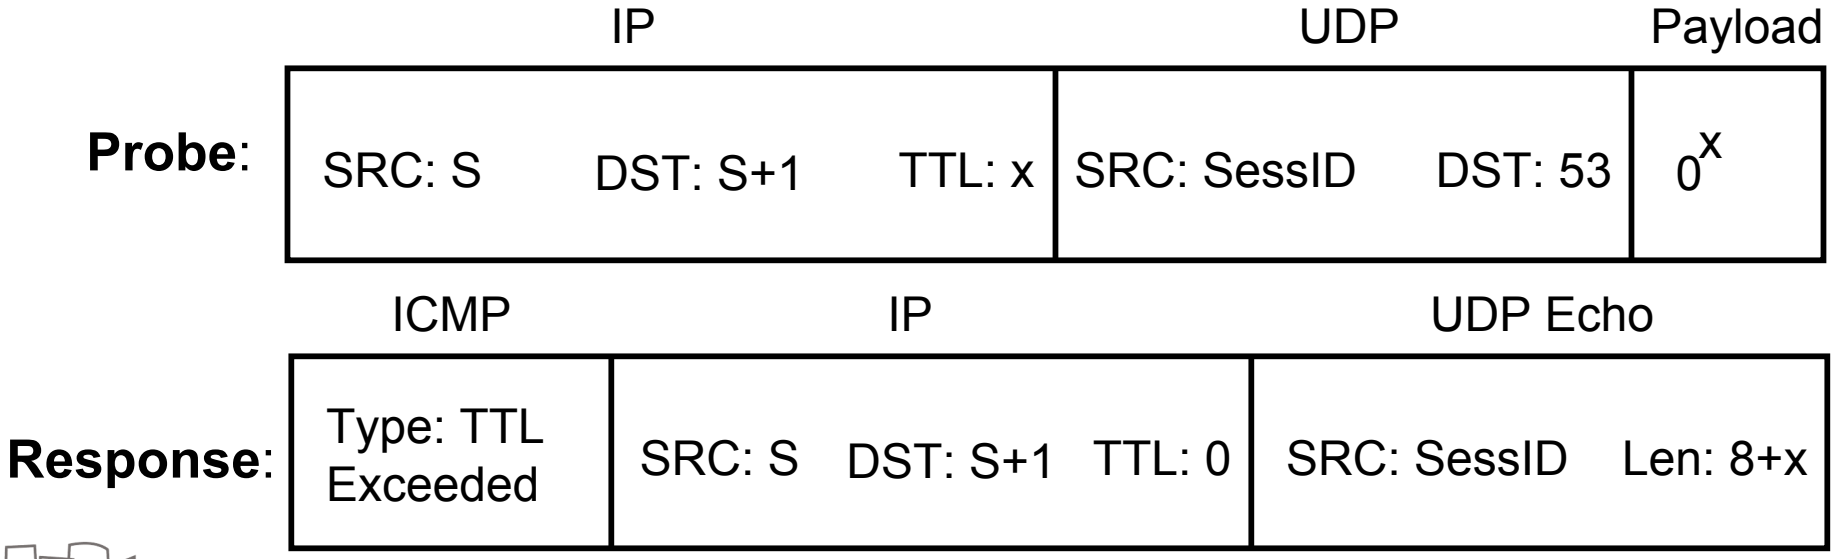

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter Results

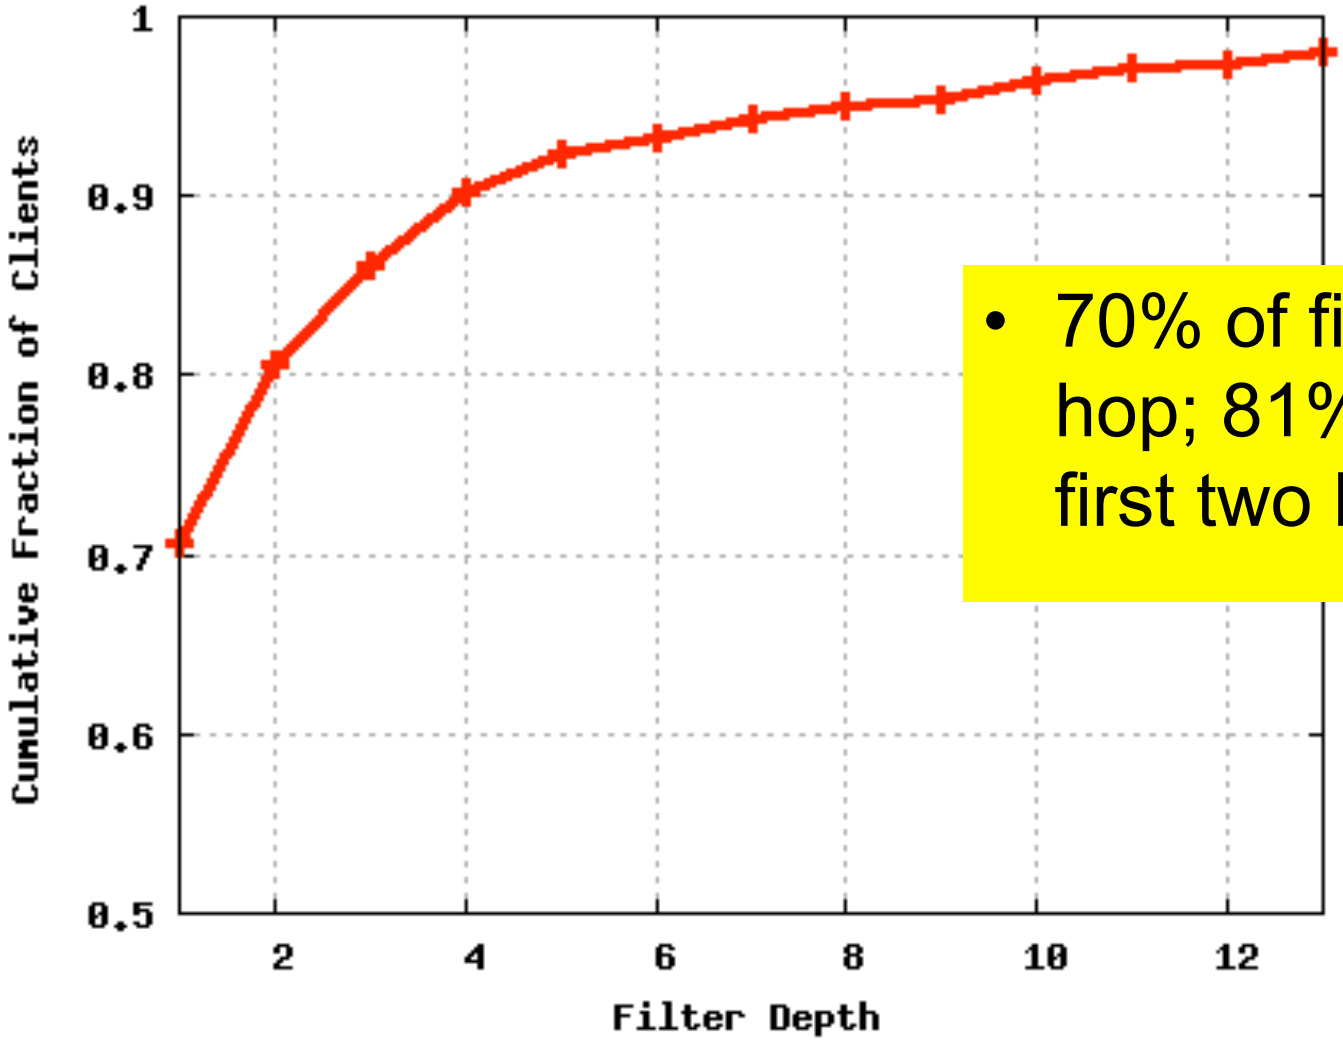

- 70% of filters at 1st hop; 81% within first two hops

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter Results

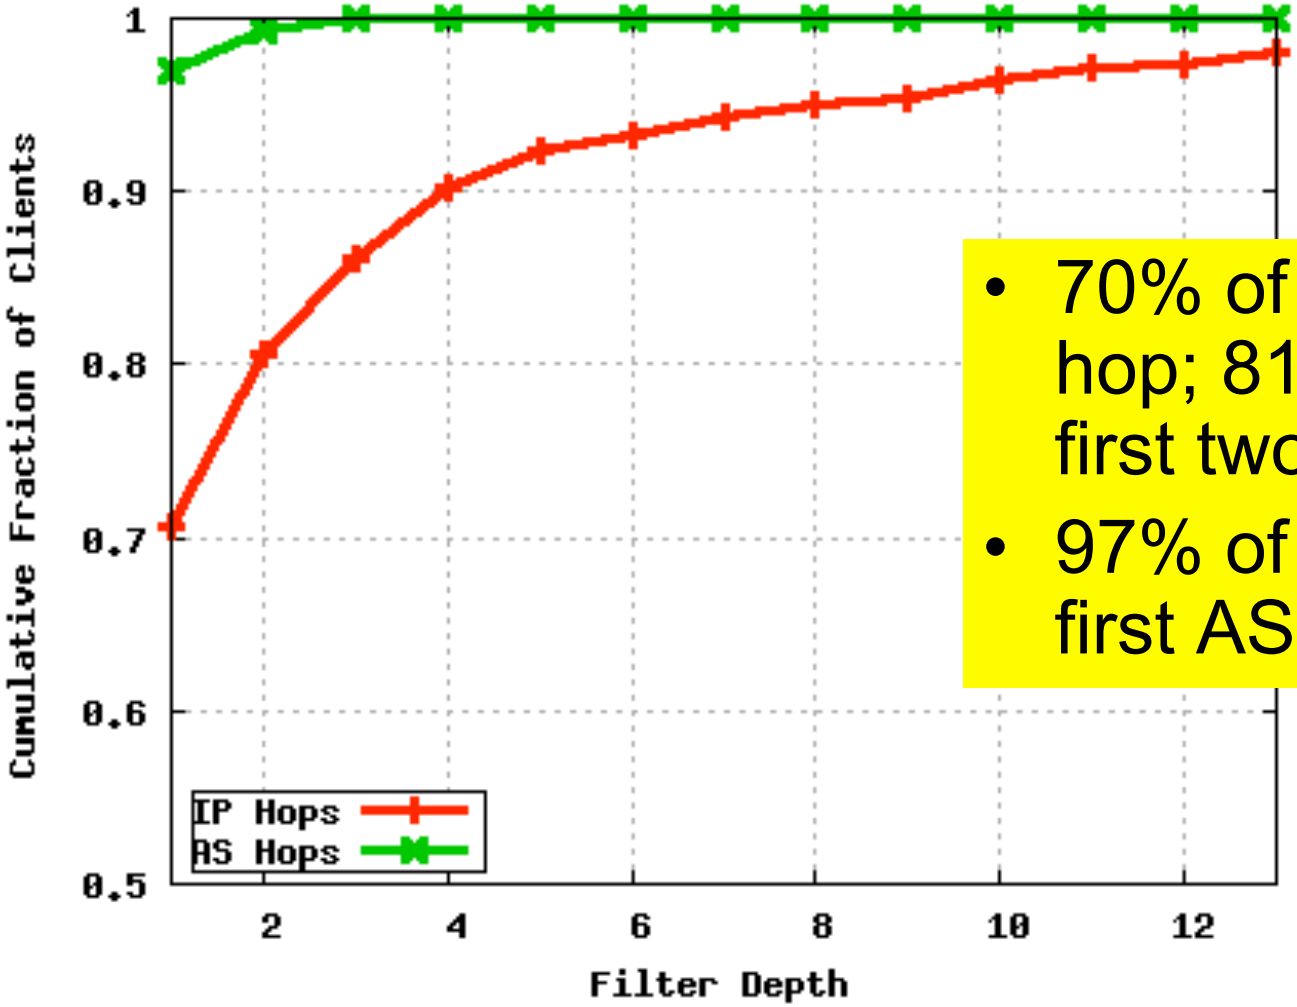

- 70% of filters at 1st hop; 81% within first two hops
- 97% of filters within first AS

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

tracefilter Results

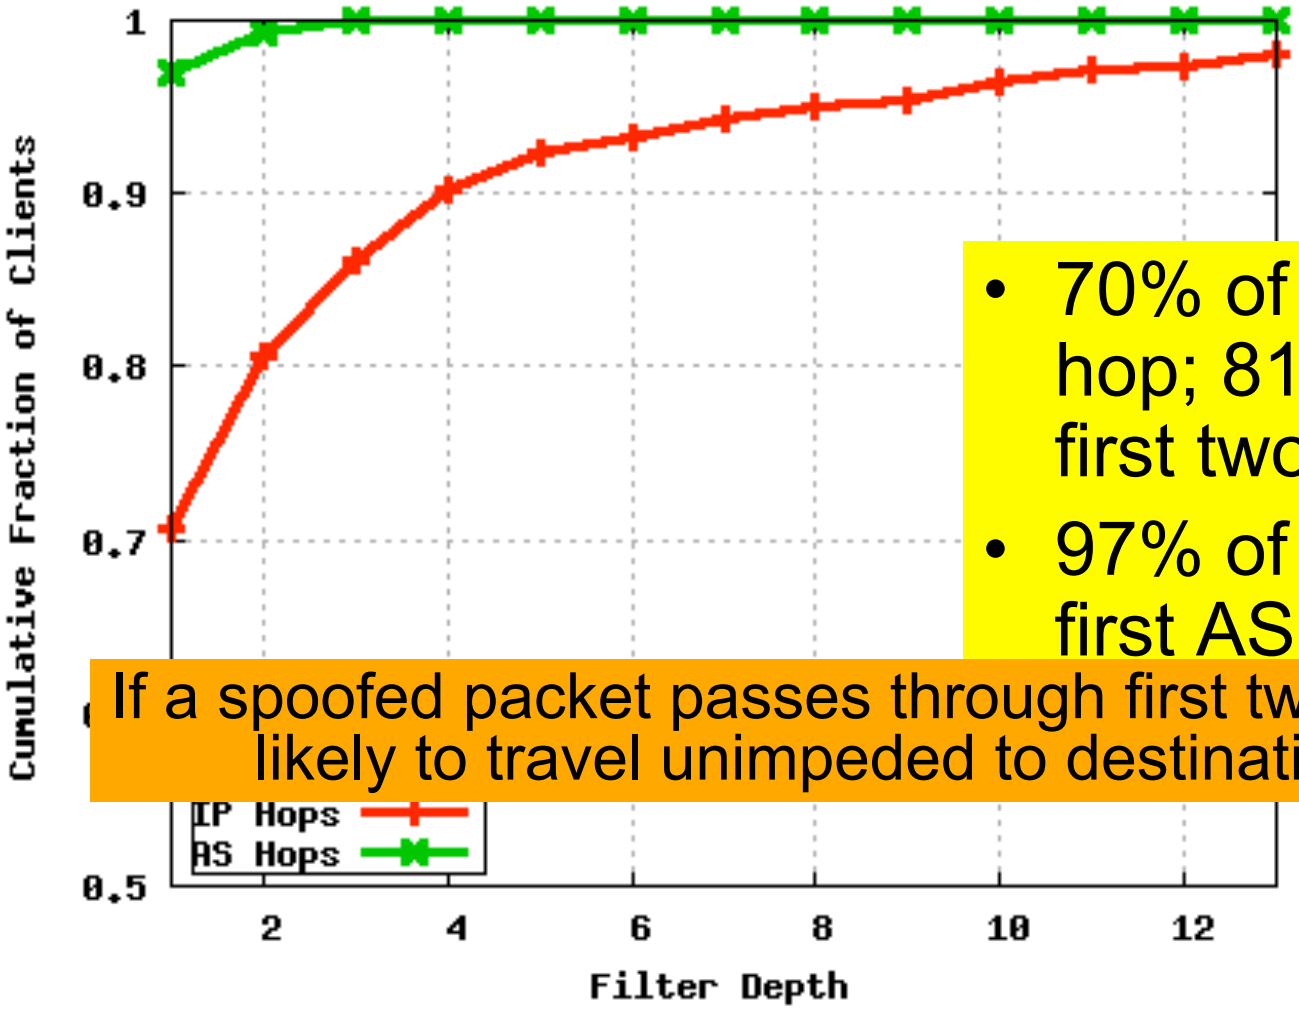

- 70% of filters at 1st hop; 81% within first two hops
- 97% of filters within first AS

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

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Ark Support

- Spoofer tester now tied into CAIDA's archipelago distributed measurement infrastructure (Ark)
- Provides invaluable additional inference capability
- Allows us to resolve aforementioned ambiguity

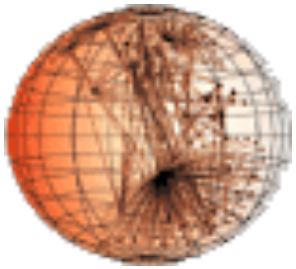

caida

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Utilizing Ark Infrastructure

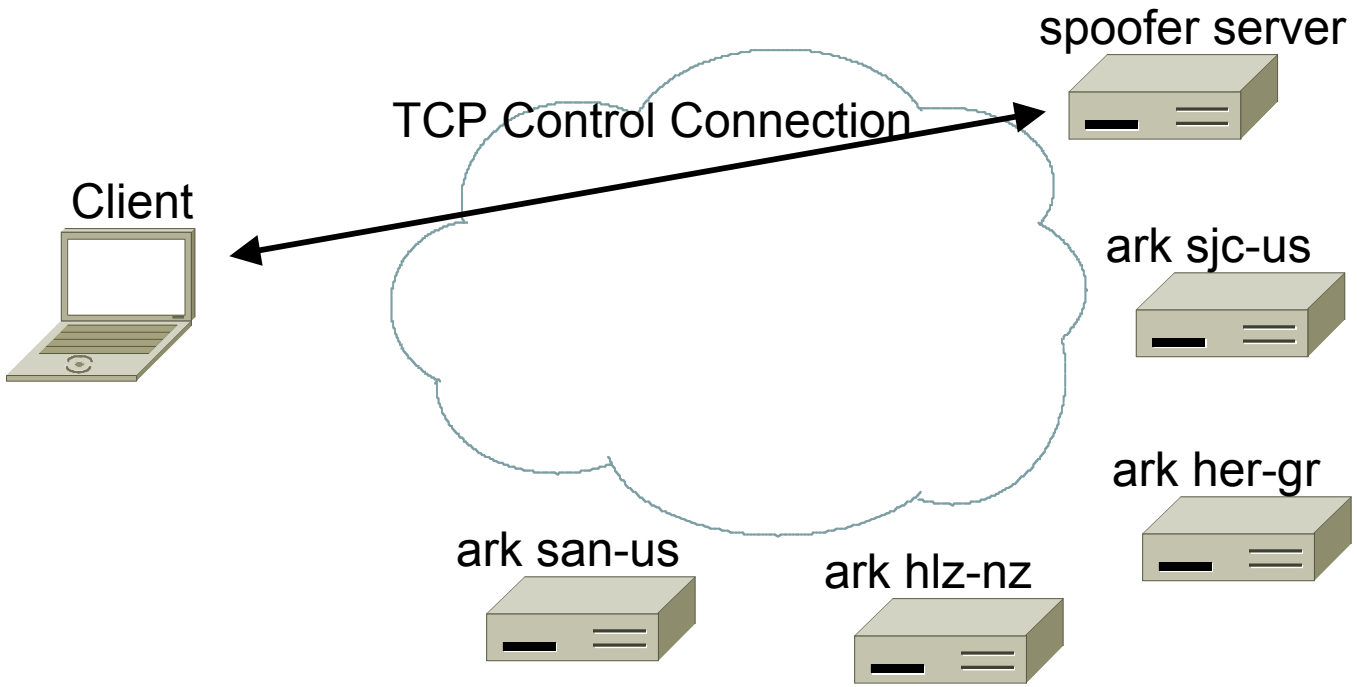

- Server and Ark nodes agree on common HMAC key
- Provide client with (SRC, DST, KEY, SEQ) tuples

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Utilizing Ark Infrastructure

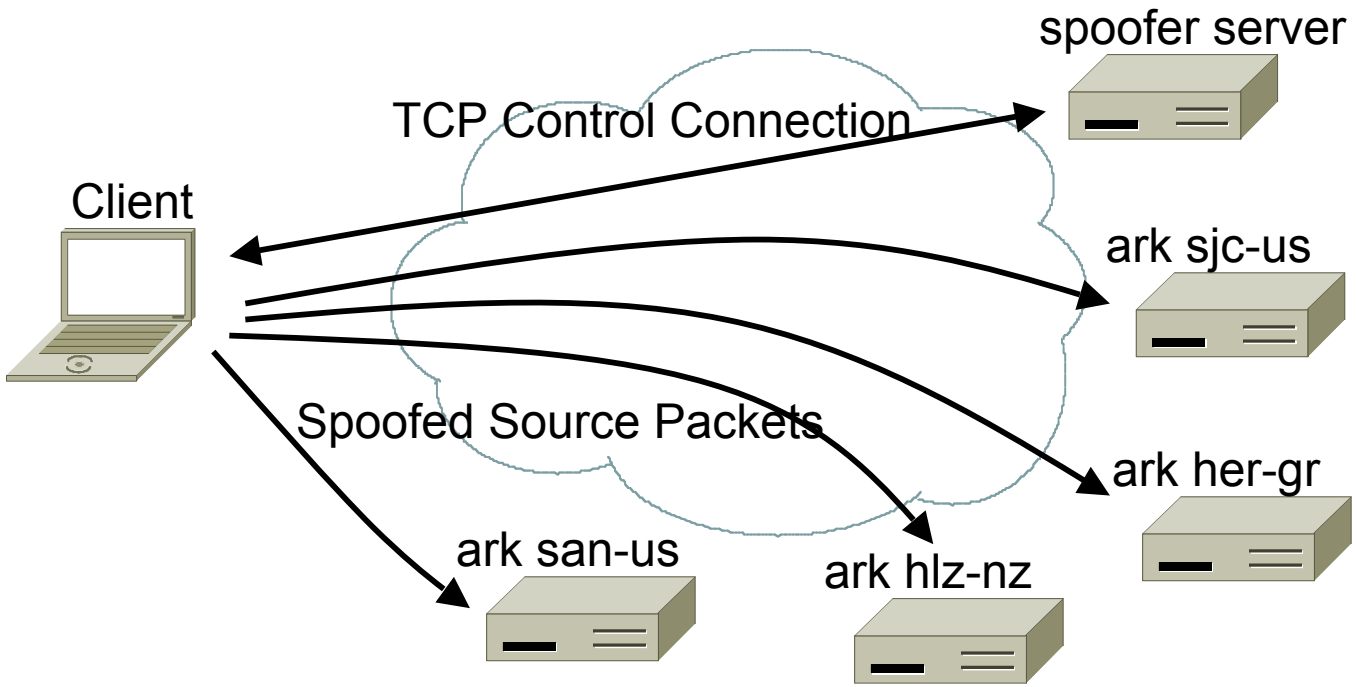

- Client sends HMAC keyed spoof probes to ark nodes
- Client runs traceroute to each ark node in parallel

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Utilizing Ark Infrastructure

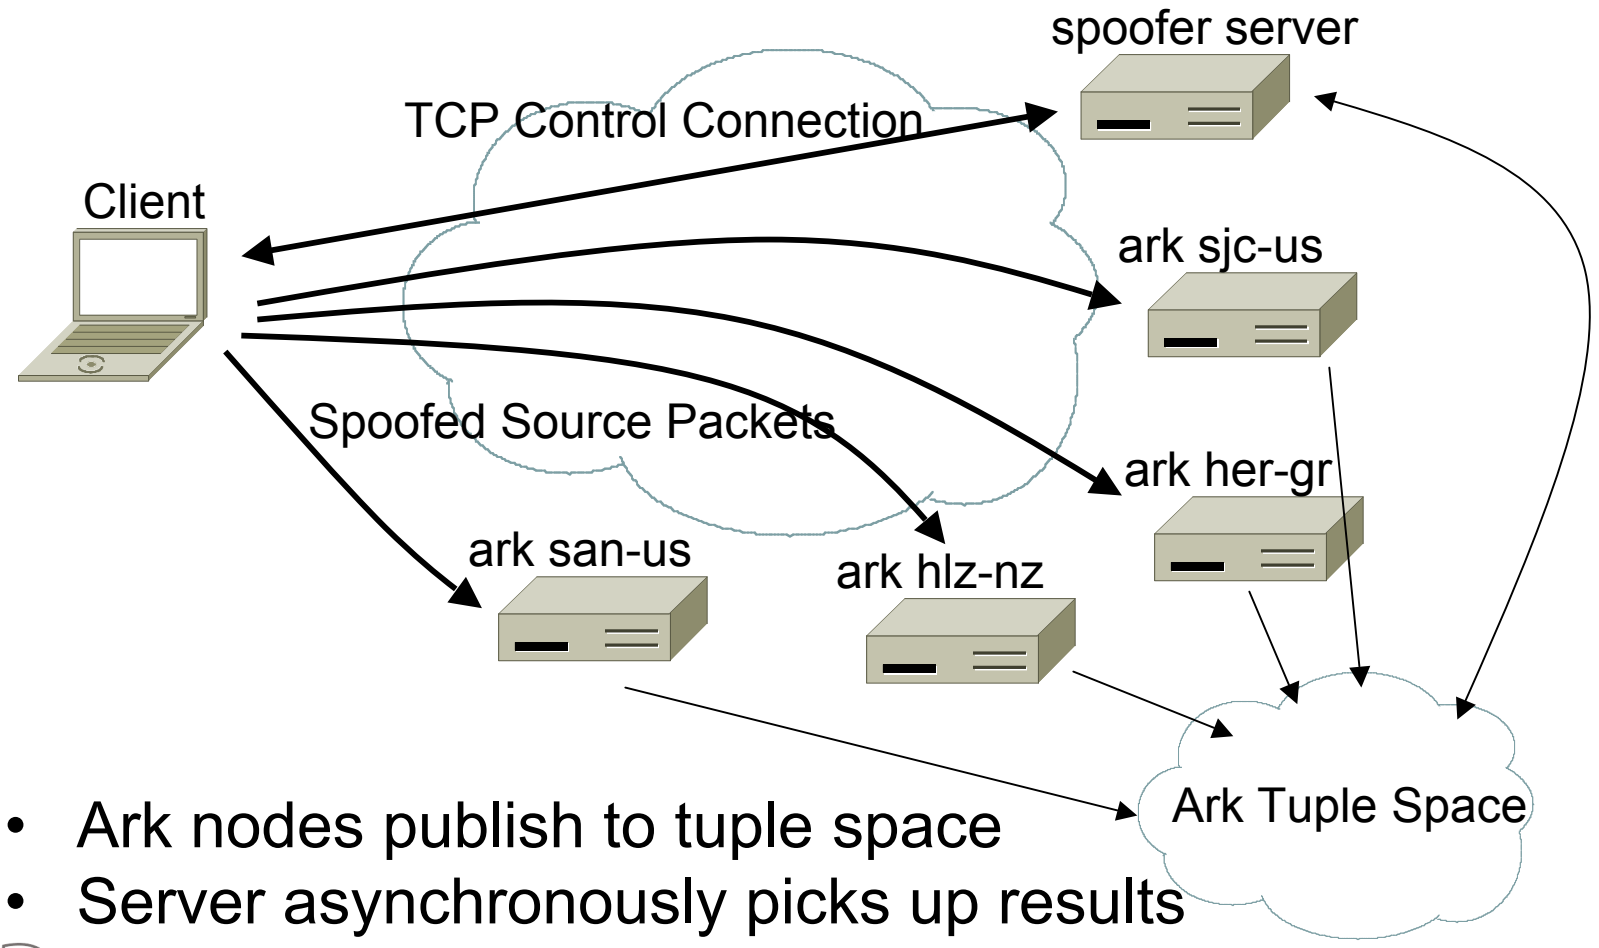

- Ark nodes publish to tuple space
- Server asynchronously picks up results

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Value of Ark

- How does Ark allow us better inference
- Example:

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Multiple Destinations

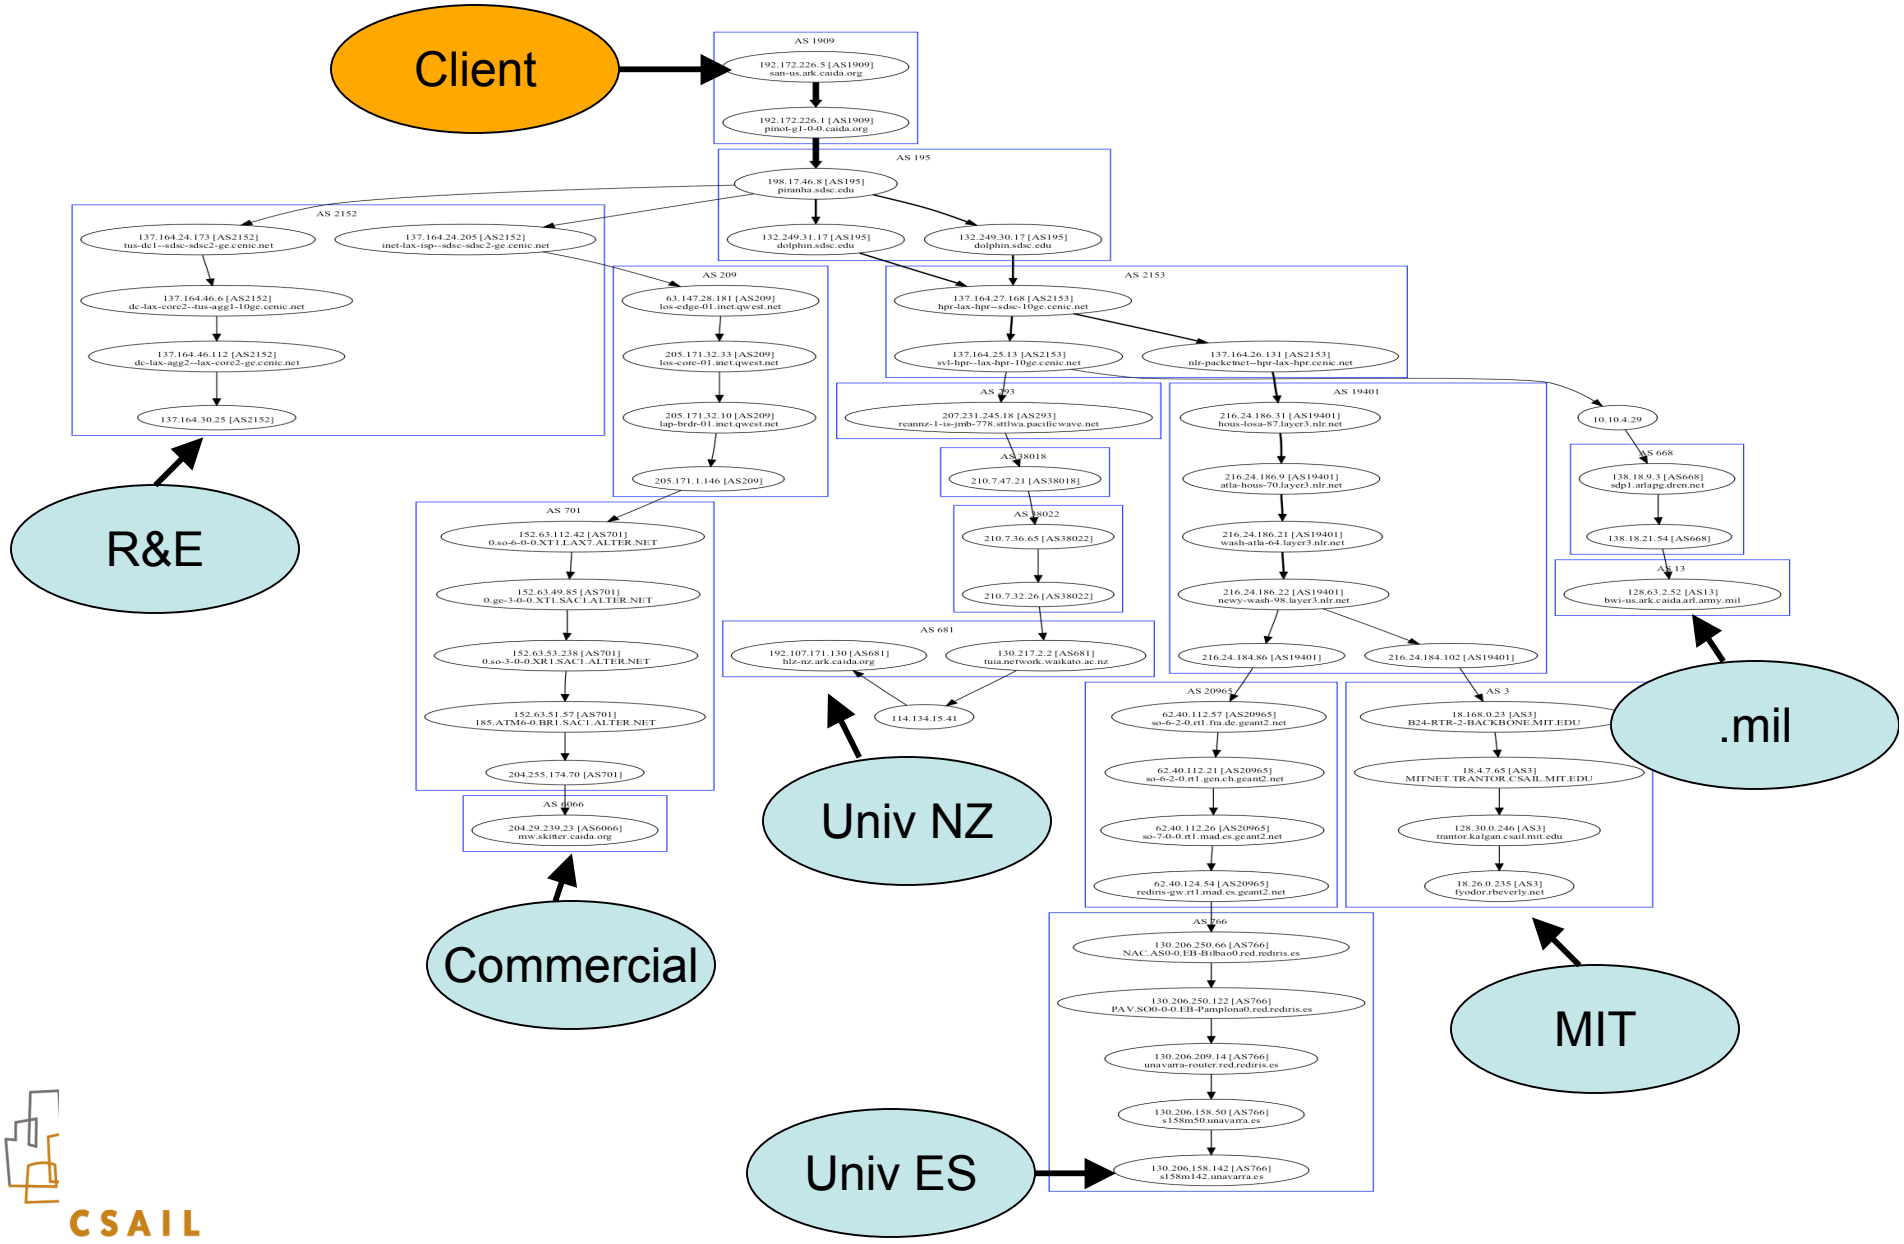

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Multiple Destinations

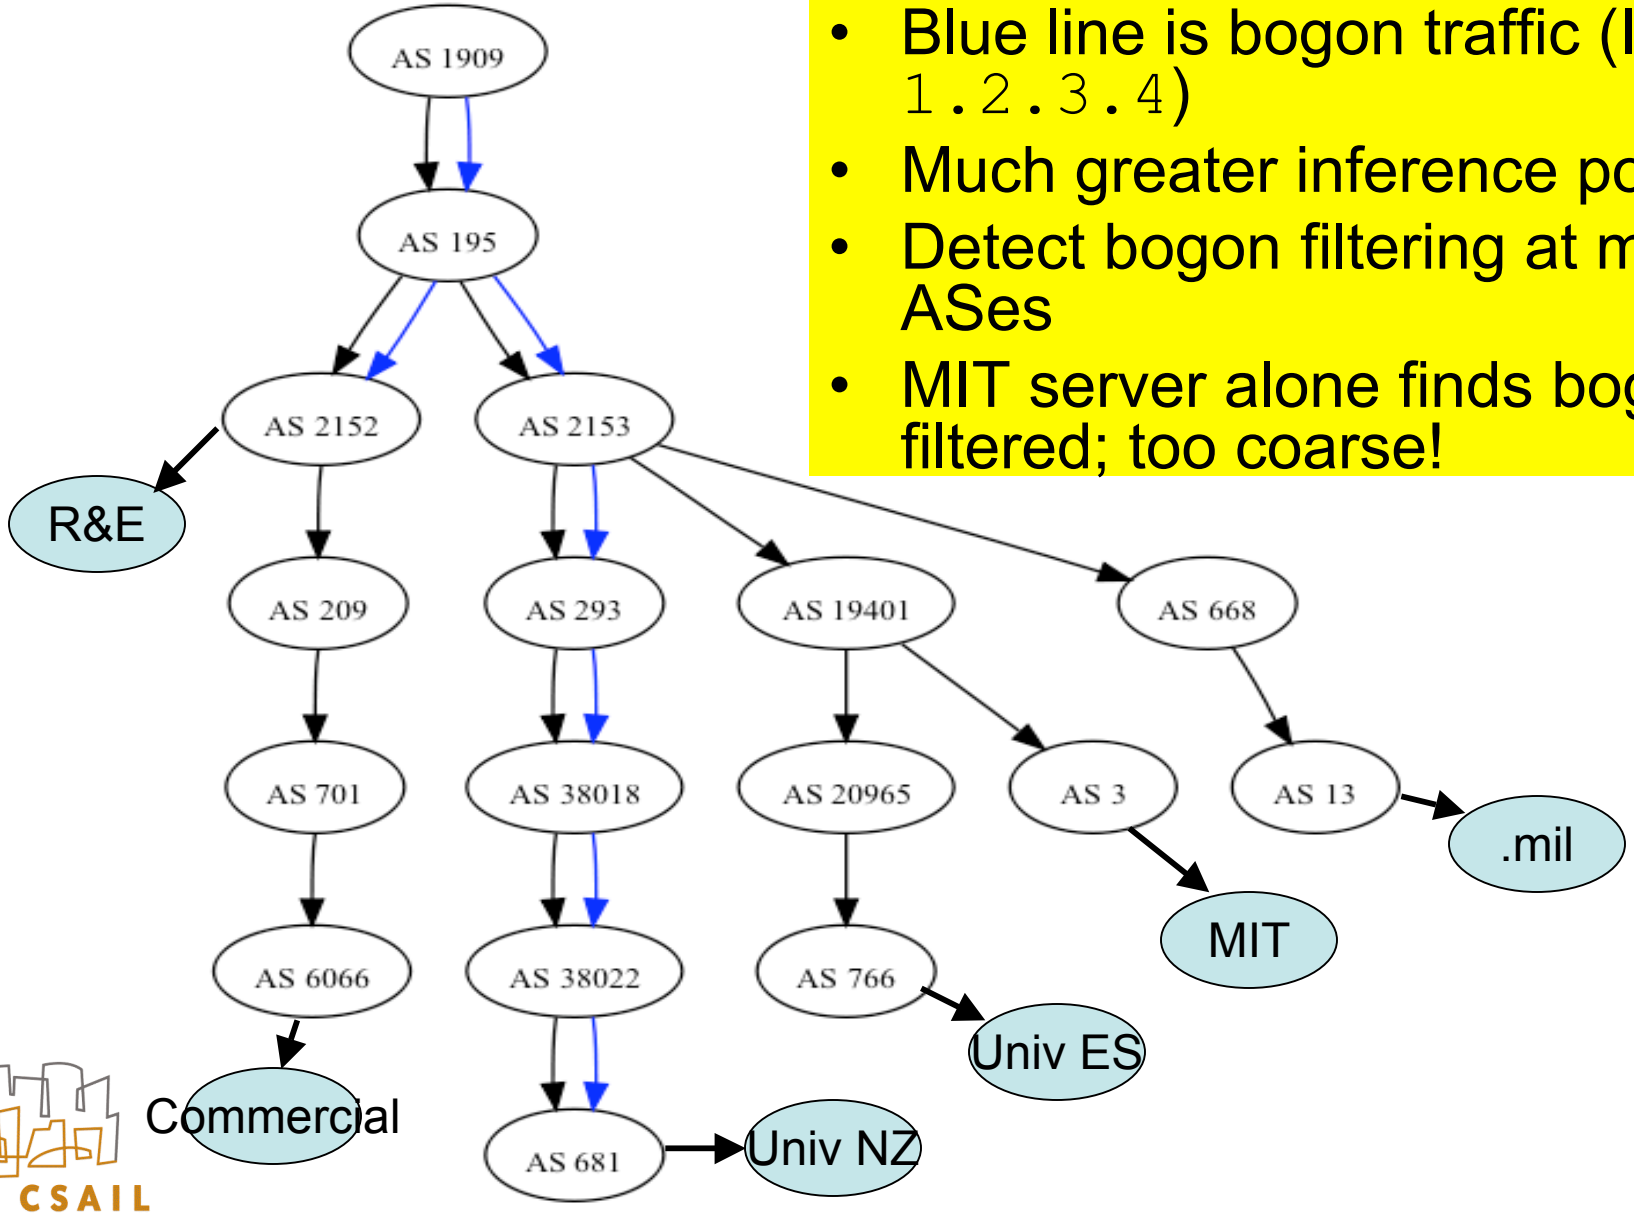

- Blue line is bogon traffic (IP: `1.2.3.4`)
- Much greater inference power
- Detect bogon filtering at multiple ASes
- MIT server alone finds bogons filtered; too coarse!

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Multiple Destinations

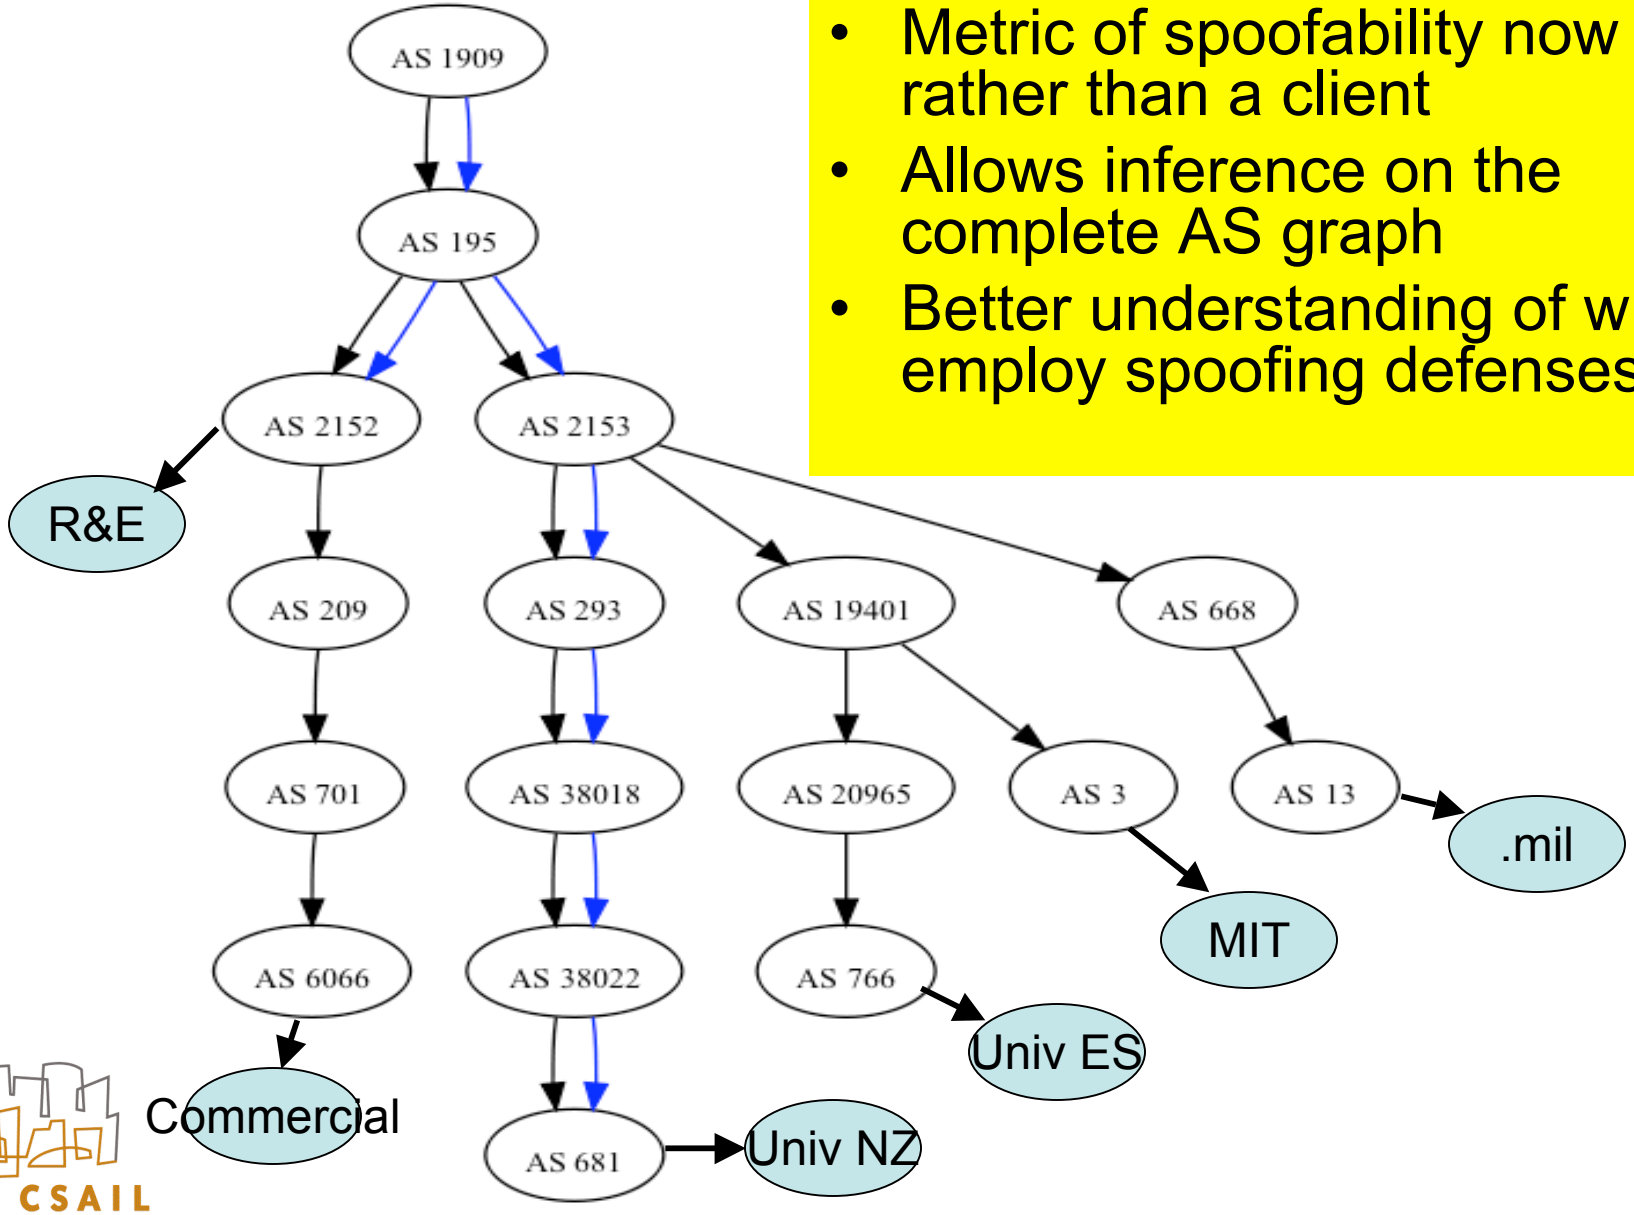

- Metric of spoofability now a path rather than a client
- Allows inference on the complete AS graph
- Better understanding of where to employ spoofing defenses

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Project

- Background
- Recent Relevance
- Project Description
- What's New: Methodology
- What's New: Data
- Parting Thoughts

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Deeper Analysis

- Question we want to answer:
  - Geographic analysis
  - Large or small providers filter?
  - What kinds of providers?

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Geographic (Tests)

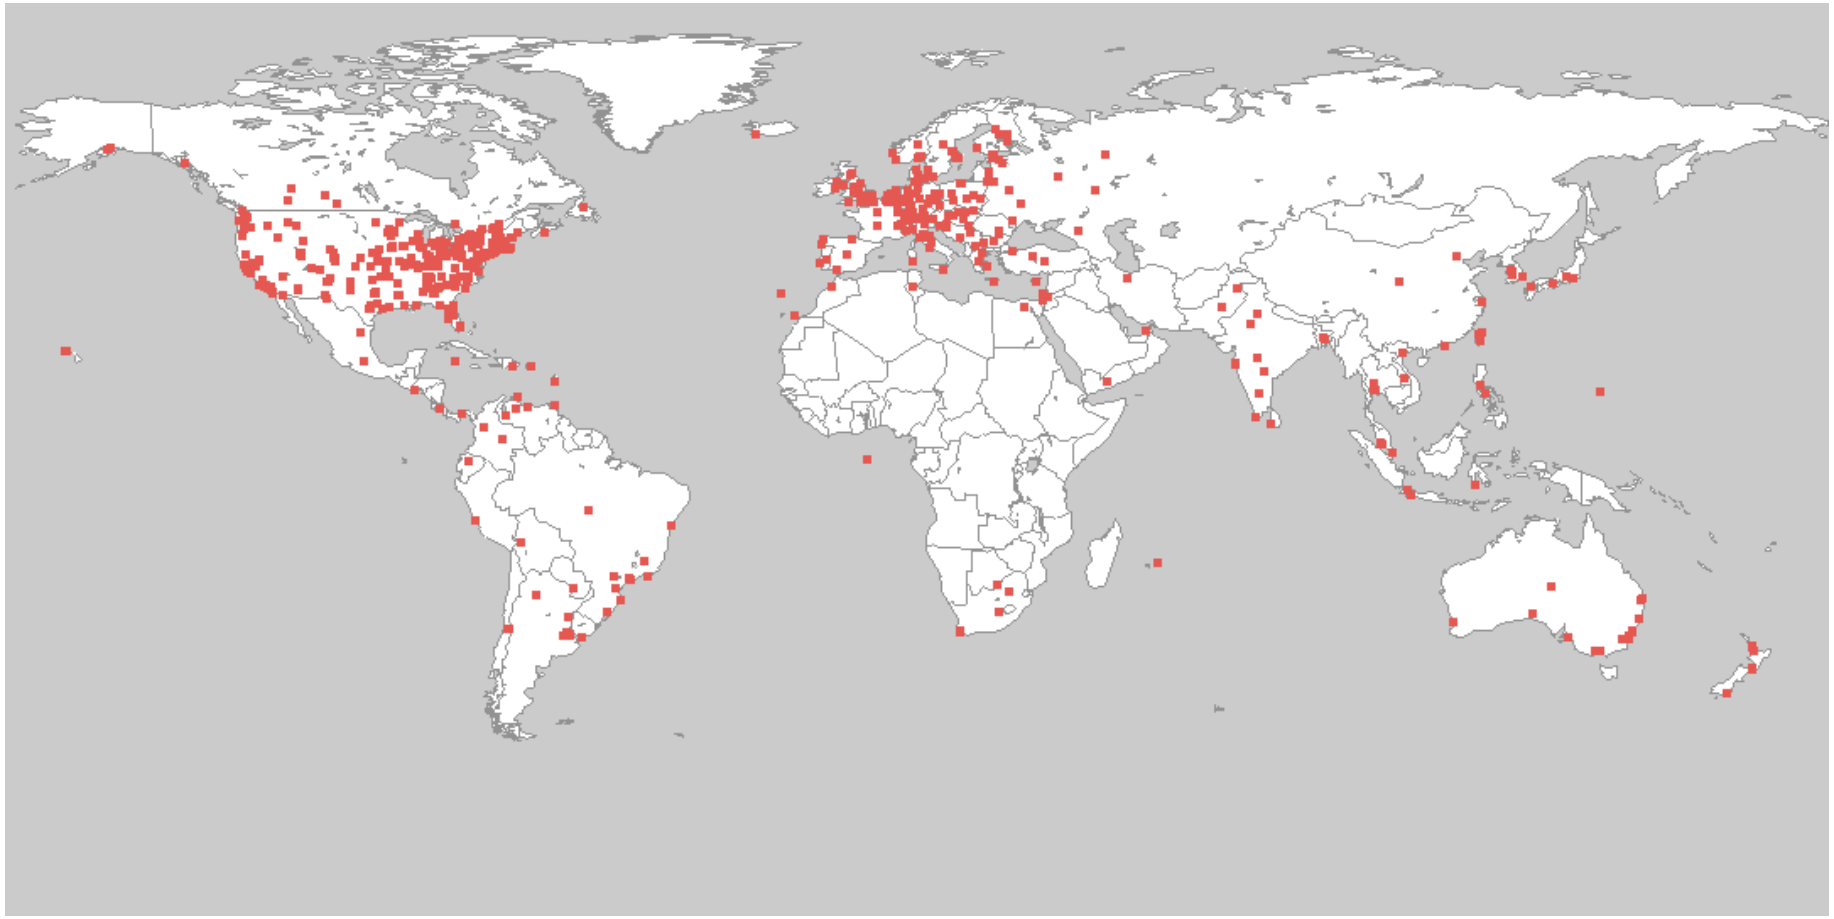

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Geographic (Spoofable)

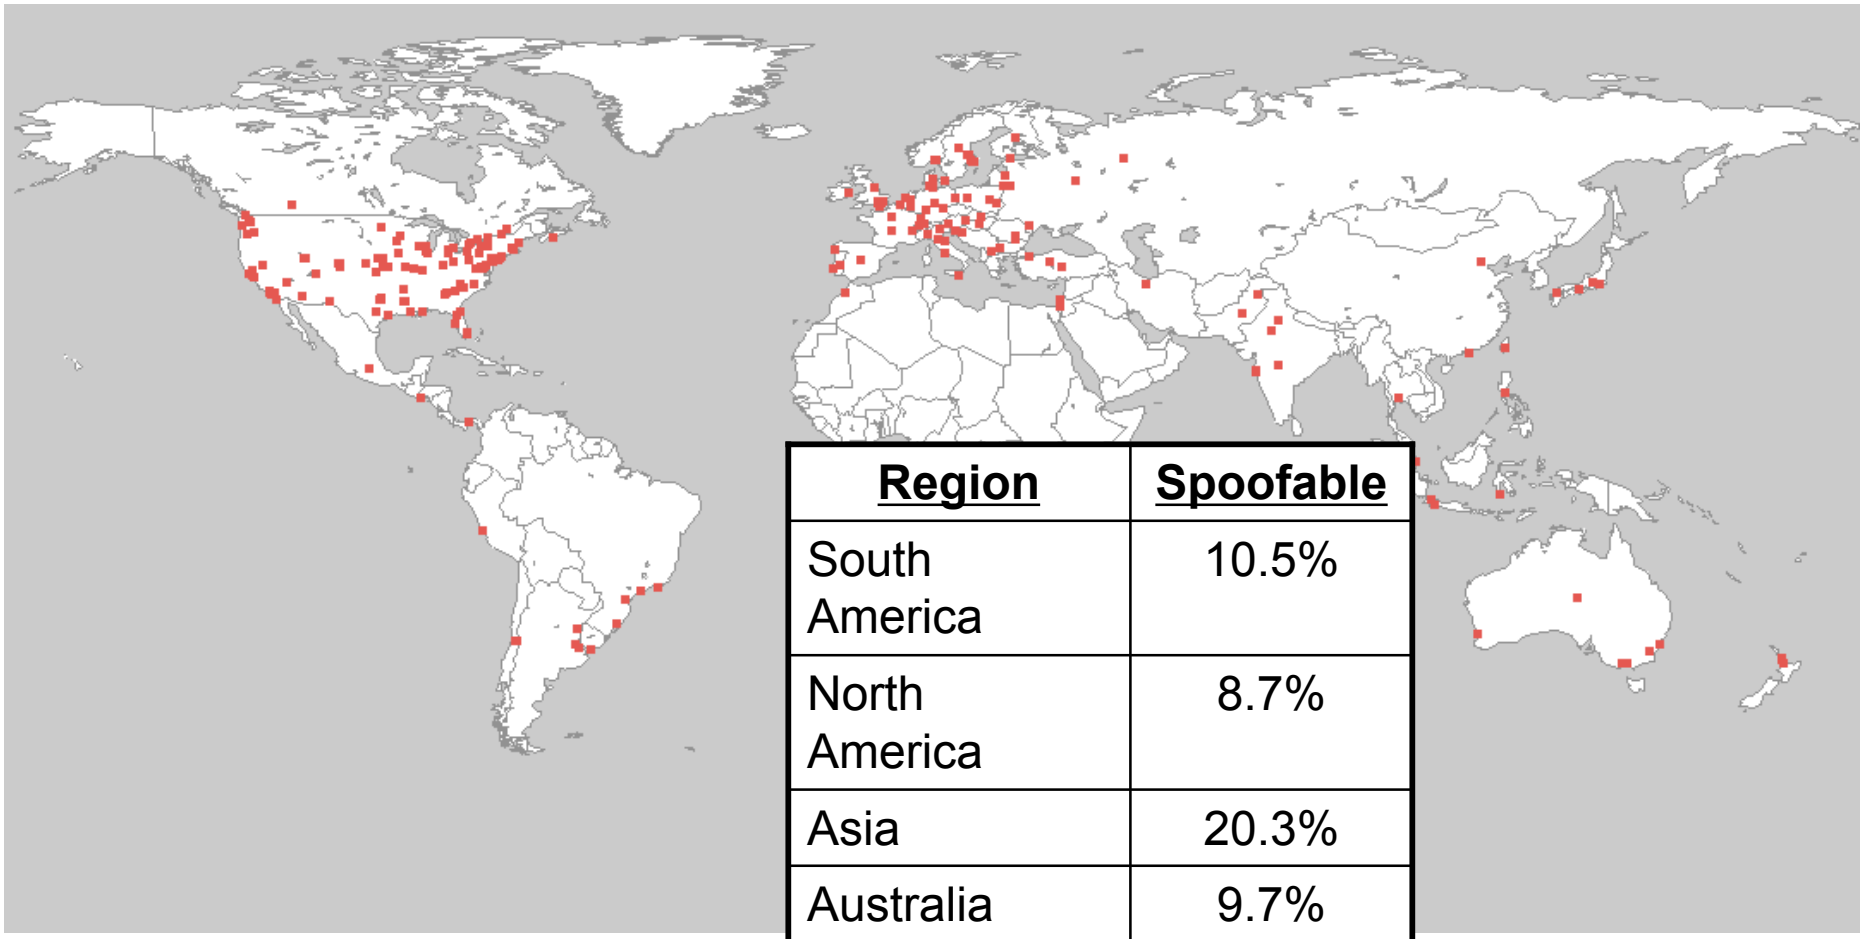

Region	Spoofable
South America	10.5%
North America	8.7%
Asia	20.3%
Australia	9.7%
Europe	9.2%
Africa	10.5%

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

DNS Stats

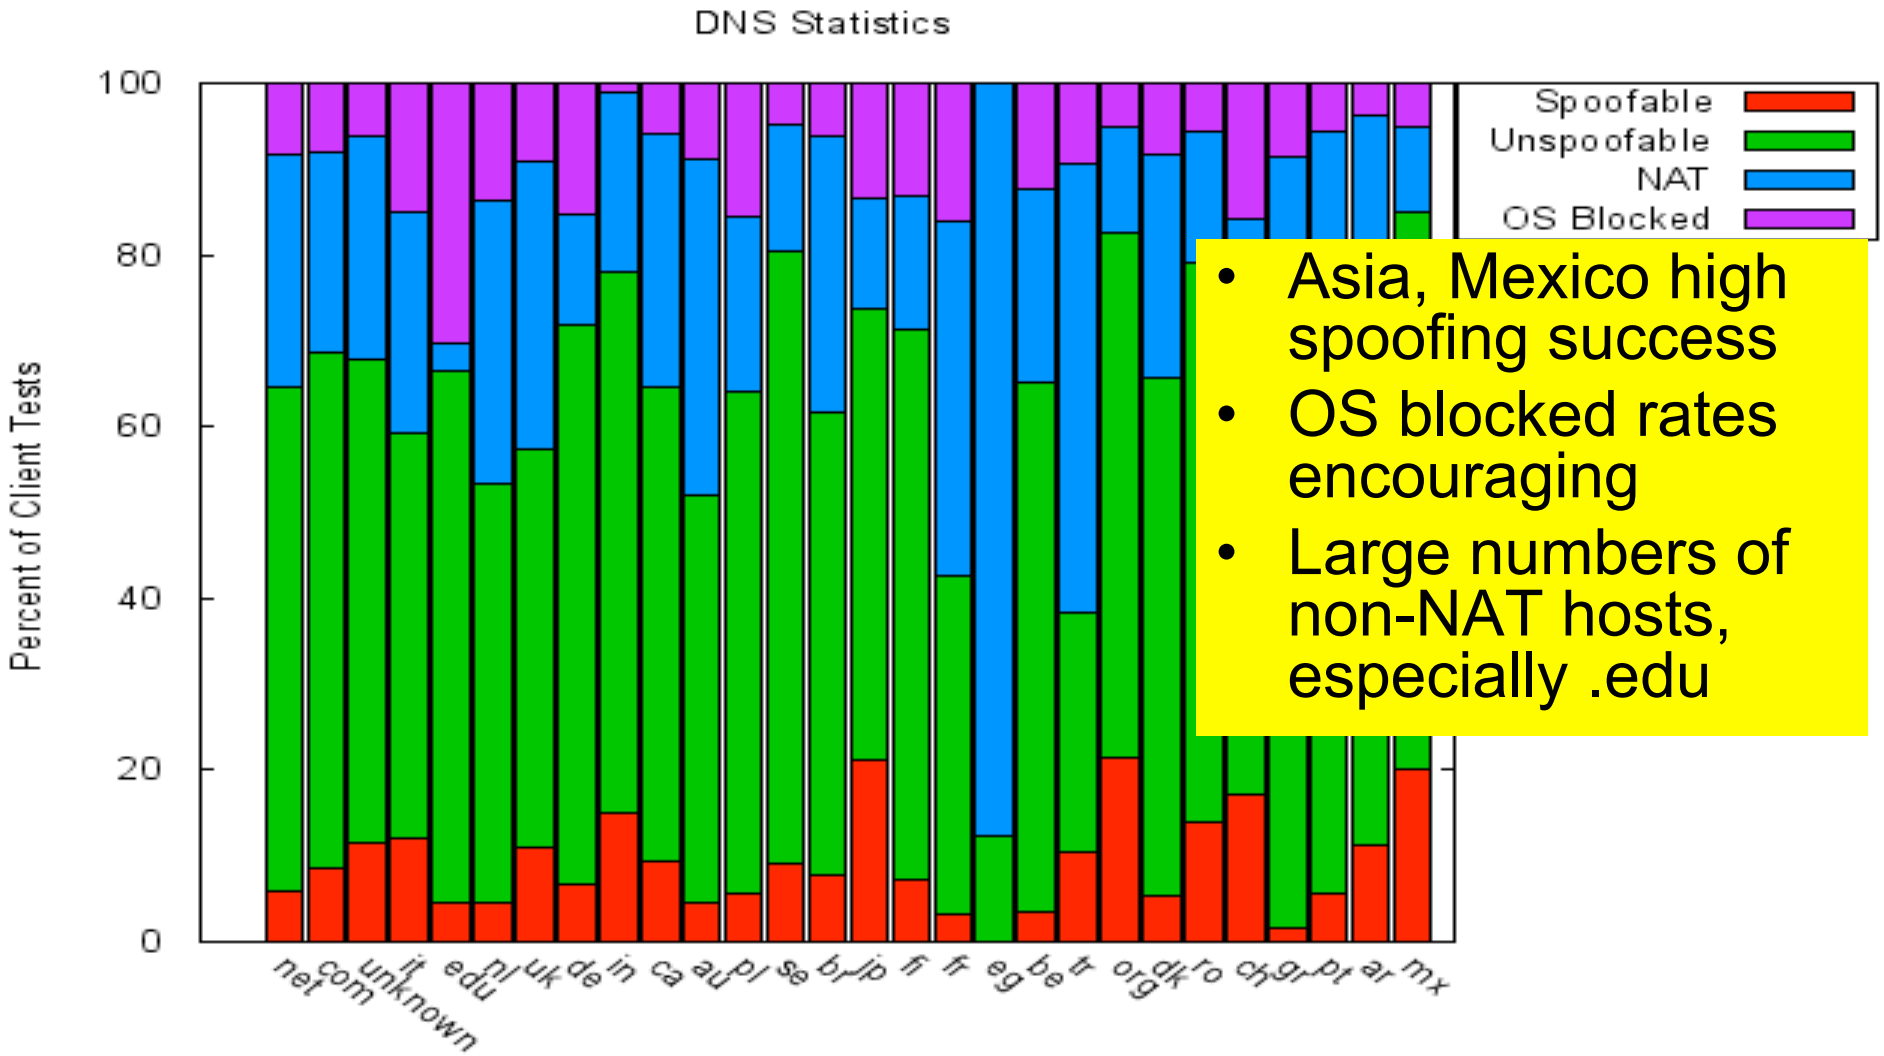

- Asia, Mexico high spoofing success
- OS blocked rates encouraging
- Large numbers of non-NAT hosts, especially .edu

Connection Classes

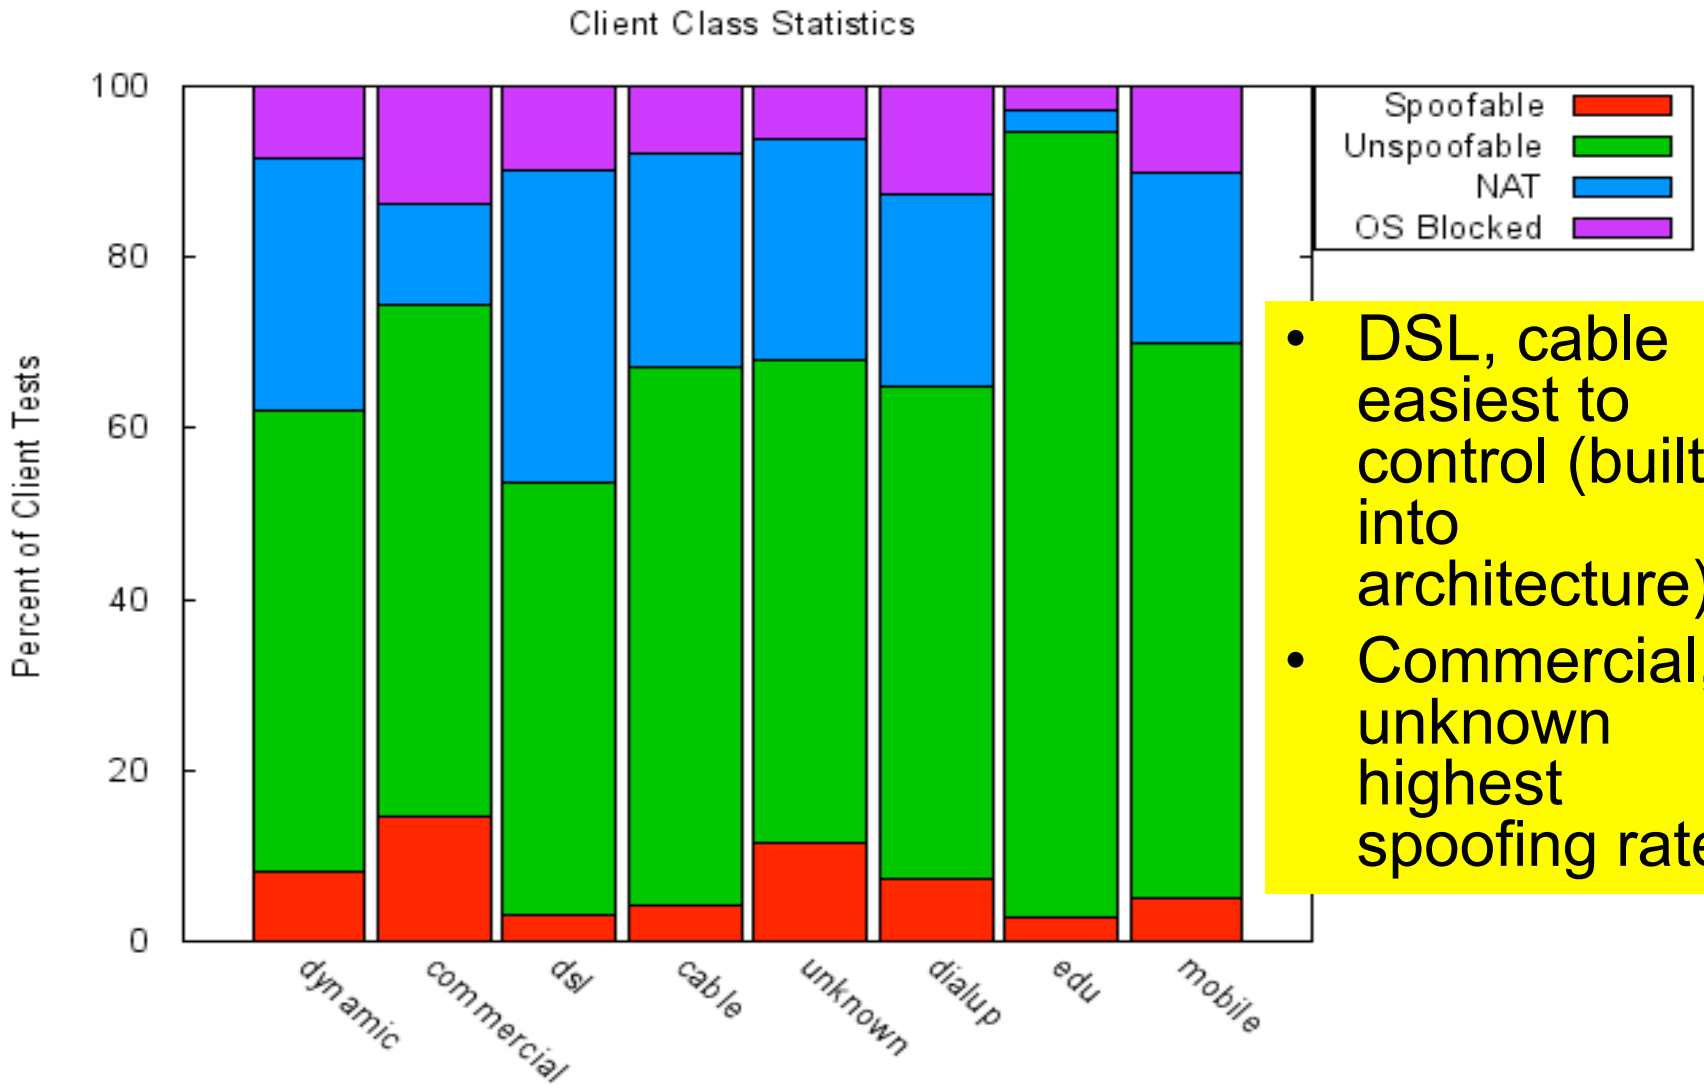

- DSL, cable easiest to control (built into architecture)
- Commercial, unknown highest spoofing rates

AS Degree

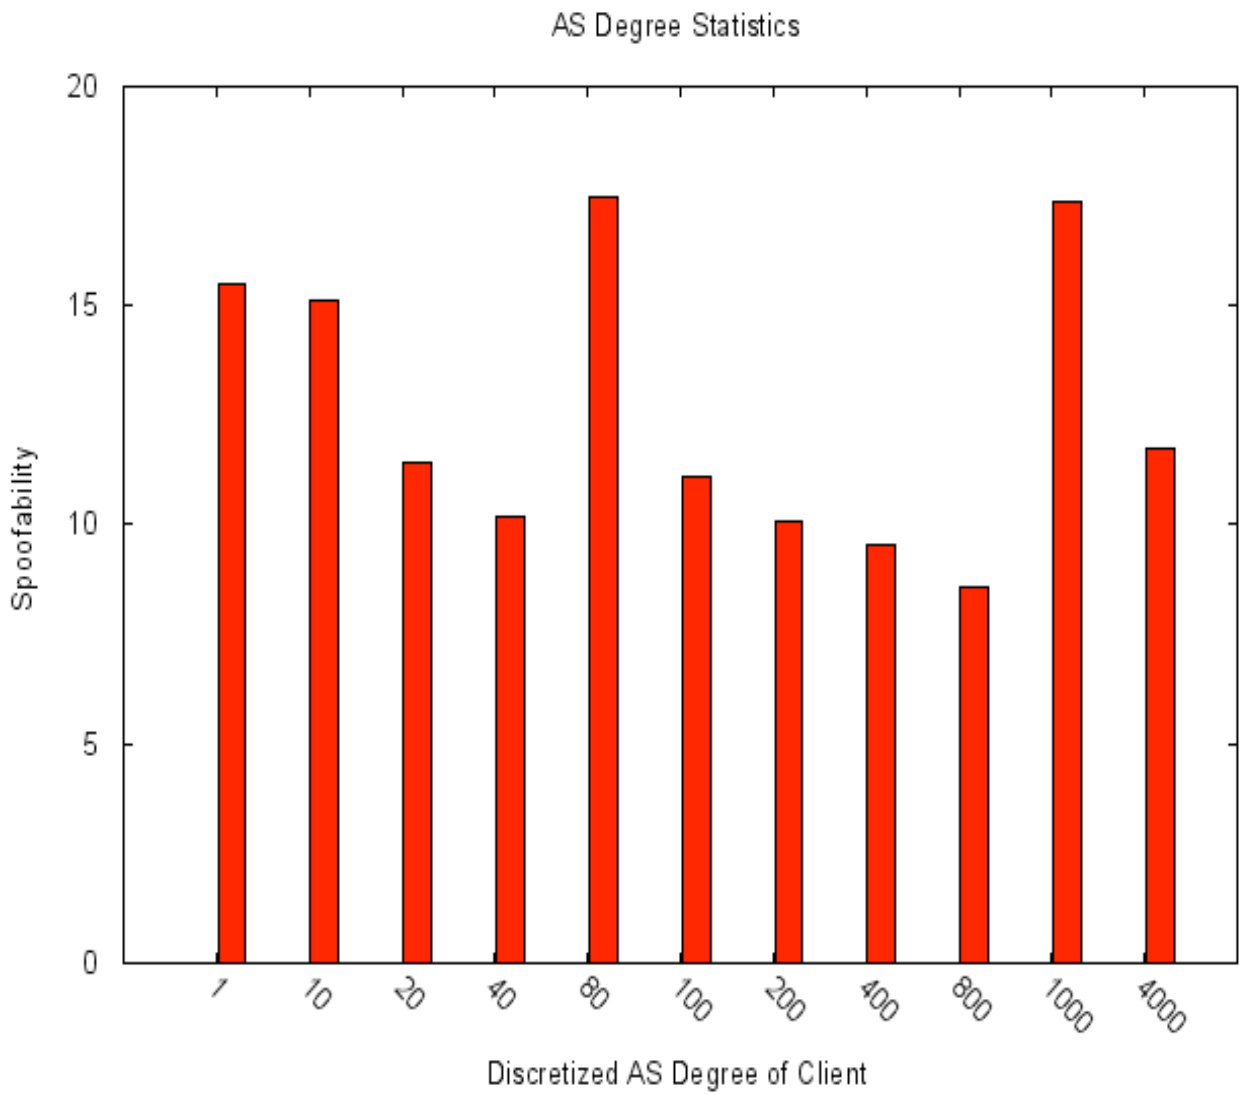

- Small or large providers filtering?
- Surprisingly, no clear trend
- Work required across the board (or a new solution)

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Spoofer Project

- Background
- Recent Relevance
- Project Description
- What's New: Methodology
- What's New: Data
- Longitudinal Analysis
- Parting Thoughts

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Parting Thoughts (1)

- Among clients able to spoof, what sources can they spoof?

Spoofed Source	Description	Defense	Percent
6.1.2.3	Valid (In BGP table)	uRPF	90%
1.2.3.4	Unallocated	Bogon Filters	58%
172.16.1.100	RFC1918 private	Static ACL	1%

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Parting Thoughts (1)

Spoofed Source	Description	Defense	Percent
6.1.2.3	Valid (In BGP table)	uRPF	90%
1.2.3.4	Unallocated	Bogon Filters	58%
172.16.1.100	RFC1918 private	Static ACL	1%

Low hanging fruit already employed, problem is harder!

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Parting Thoughts (2)

- Tracefilter exposes operational tension between current filtering incentives and difficulty managing edge filters
- If a spoofed packet isn't filtered at edge, will travel unimpeded to destination
- Should we think about core filtering techniques?
  - StackPI
  - ML approaches with soft response (rbeverly thesis work)
  - Others

Parting Thoughts

- Even after all these years, source spoofing problem not solved
  - BCP38 has been around for 9 years
  - BCP38 great, but incentives wrong
- Single unfiltered ingress can compromise entire Internet system
  - Can we plug every hole?
  - Regulatory Response? … but multinational?
  - Spoofer page for public provider flogging?
- What's needed (biased opinion!):
  - Clean slate design
  - Filtering in the core

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY

Parting Thoughts

- Even after all these years, source spoofing problem not solved
  - BCP38 has been around for x years
  - BCP38 great, but incentives wrong
- Single unfiltered ingress can compromise entire Inte
  - C
  - F
  - S
- What's needed (biased opinion!):
  - Clean slate design
  - Filtering in the core

Thanks!

`http://spoofer.csail.mit.edu`

CSAIL
MIT COMPUTER SCIENCE AND ARTIFICIAL INTELLIGENCE LABORATORY