The Internet Measurement Toolbox

Justine Sherry, University of Washington

@ University of Puget Sound | April 12, 2010

Research does not occur in an ivory tower of monastic isolation…

This work is the effort of a diverse group:

- Ethan Katz-Bassett
- Harsha Madhyastha
- Vijay Adhikiri
- Colin Scott
- Steve Rutherford
- Mary Pimenova
- Ashoat Tevosyan
- Peter van Wesep
- Arvind Krishnamurthy
- Tom Anderson

Research does not occur in an ivory tower of monastic isolation…

This work is the effort of a diverse group:

- Ethan Katz-Bassett
- Harsha Madhyastha
- Vijay Adhikiri
- Colin Scott
- Steve Rutherford
- Mary Pimenova
- Ashoat Tevosyan
- Peter van Wesep
- Arvind Krishnamurthy
- Tom Anderson

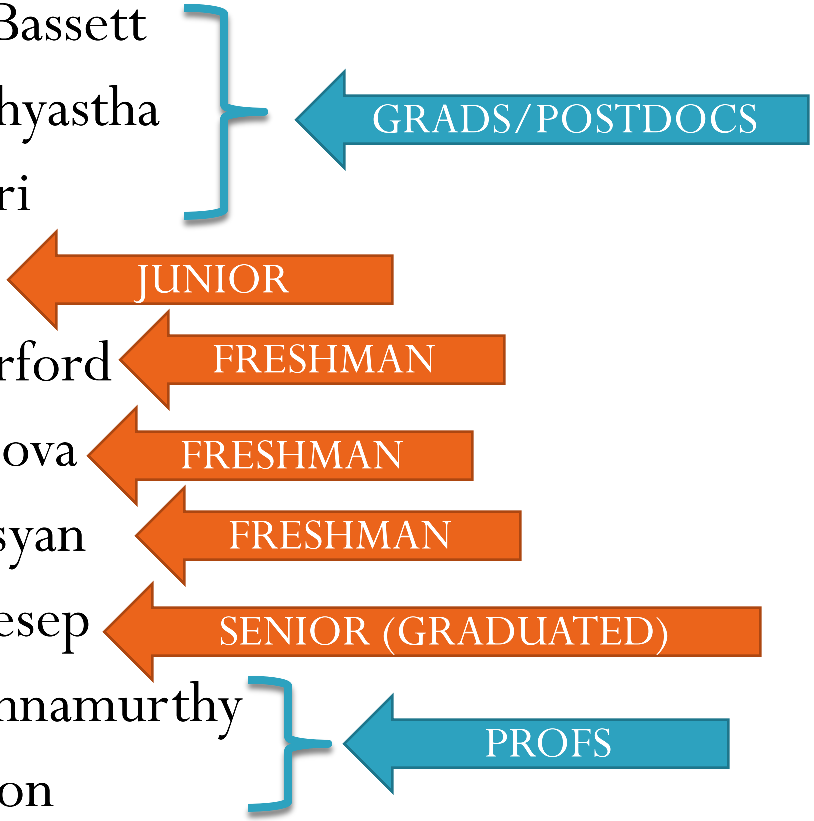

GRADS/POSTDOCS

JUNIOR

FRESHMAN

FRESHMAN

FRESHMAN

SENIOR (GRADUATED)

PROFS

A Really Fast Intro to the Internet

Words you should know:

Packet

IP Address

Packet Header

Router

A Really Fast Intro to the Internet

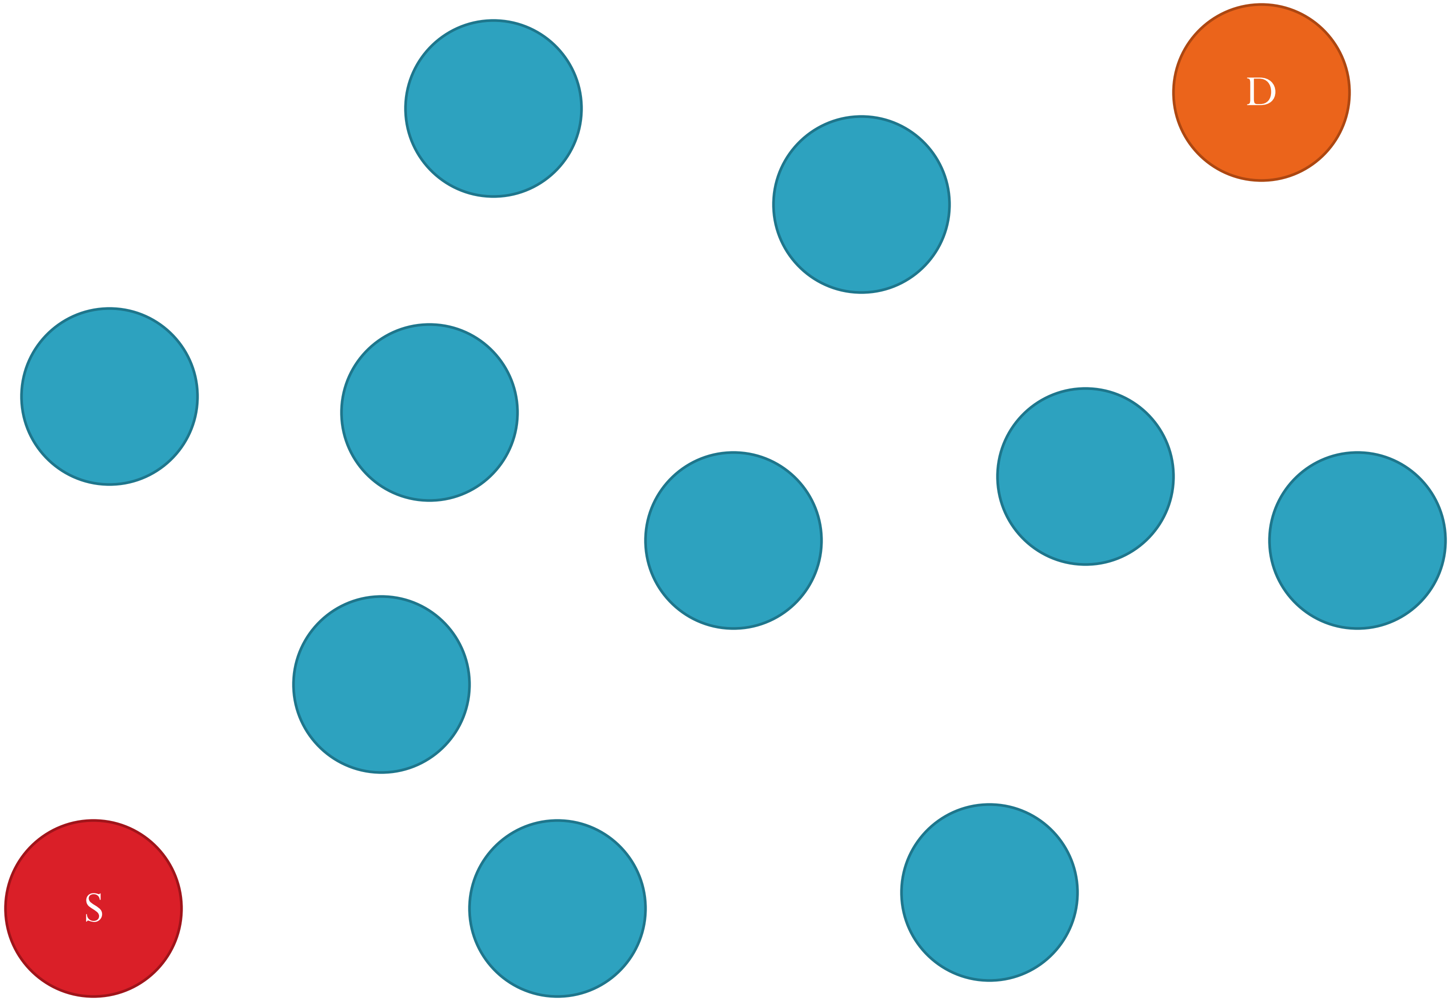

A Really Fast Intro to the Internet

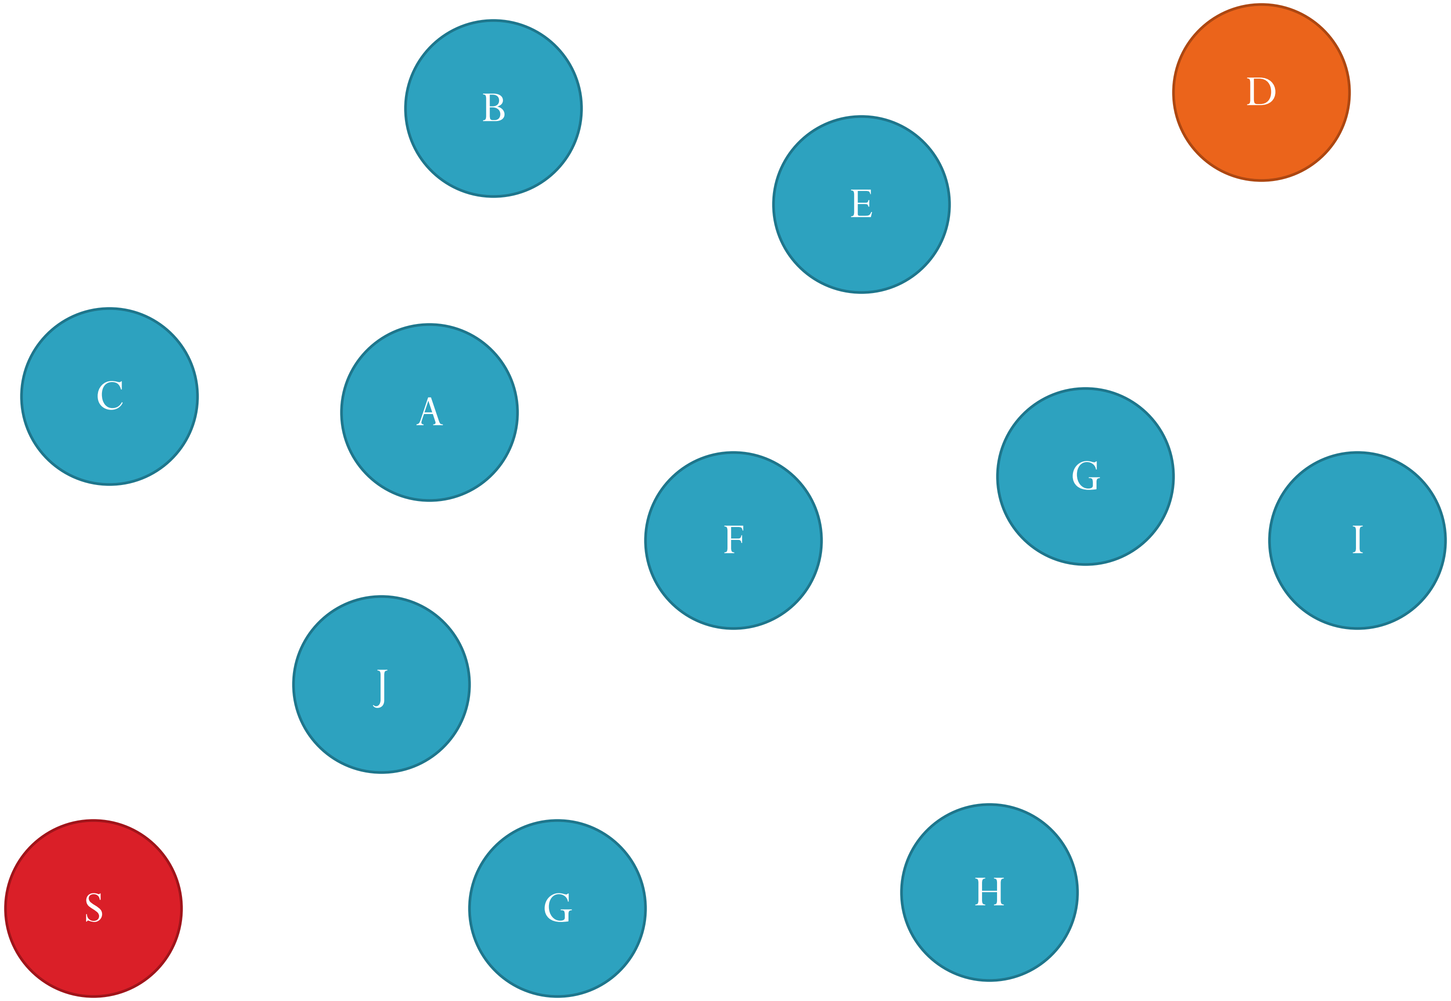

A Really Fast Intro to the Internet

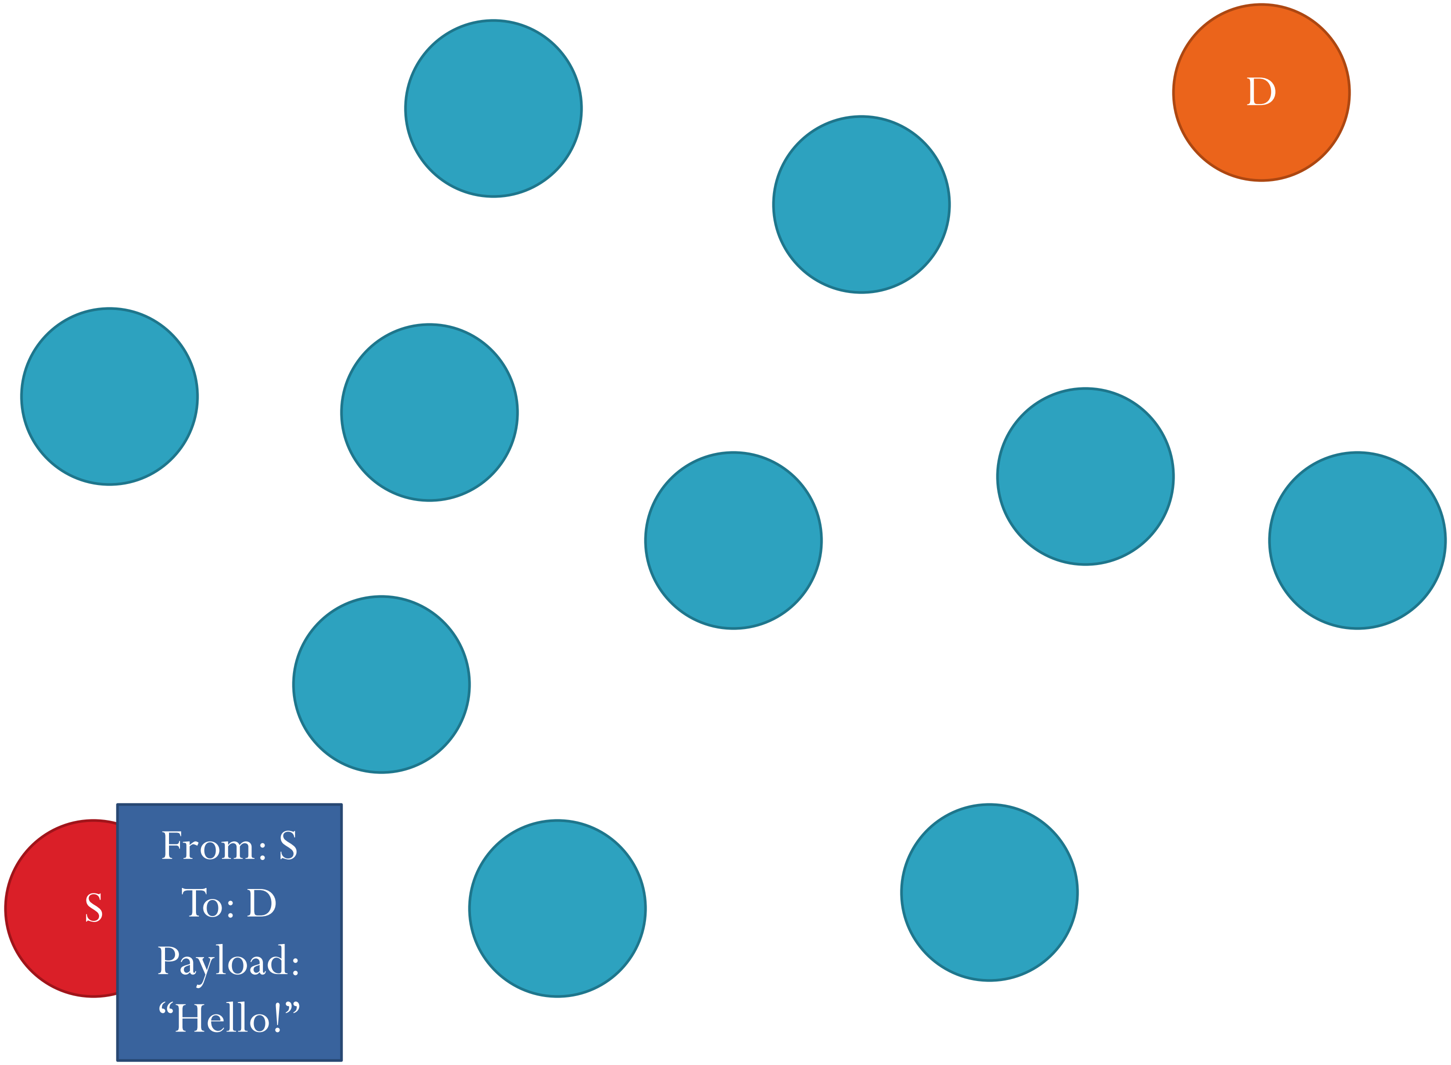

From: S
To: D
Payload:
"Hello!"

A Really Fast Intro to the Internet

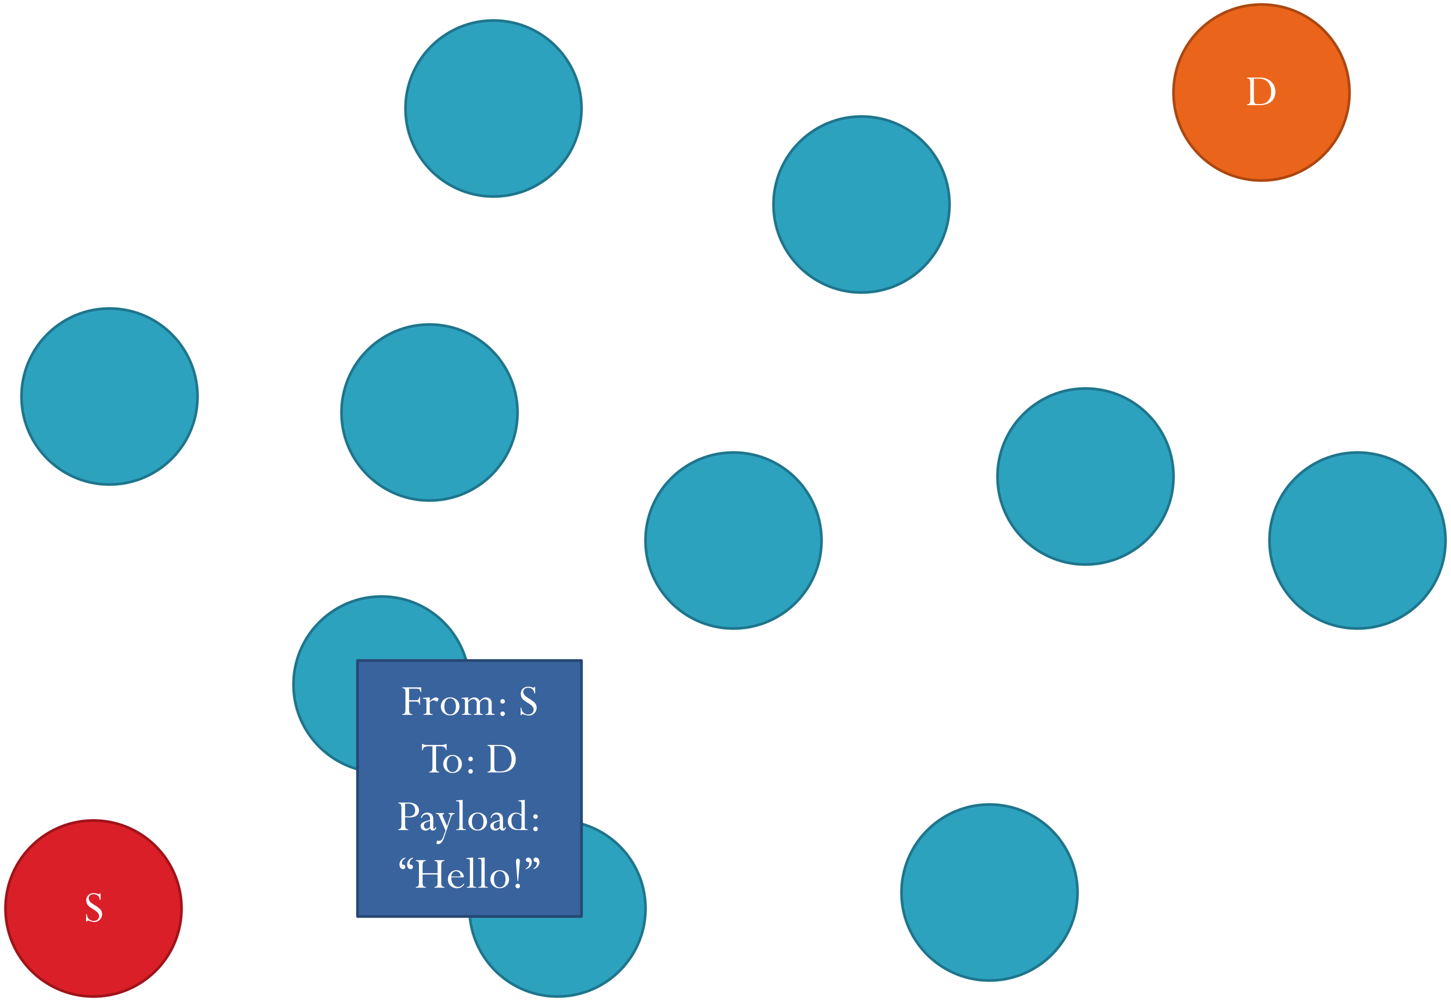

A Really Fast Intro to the Internet

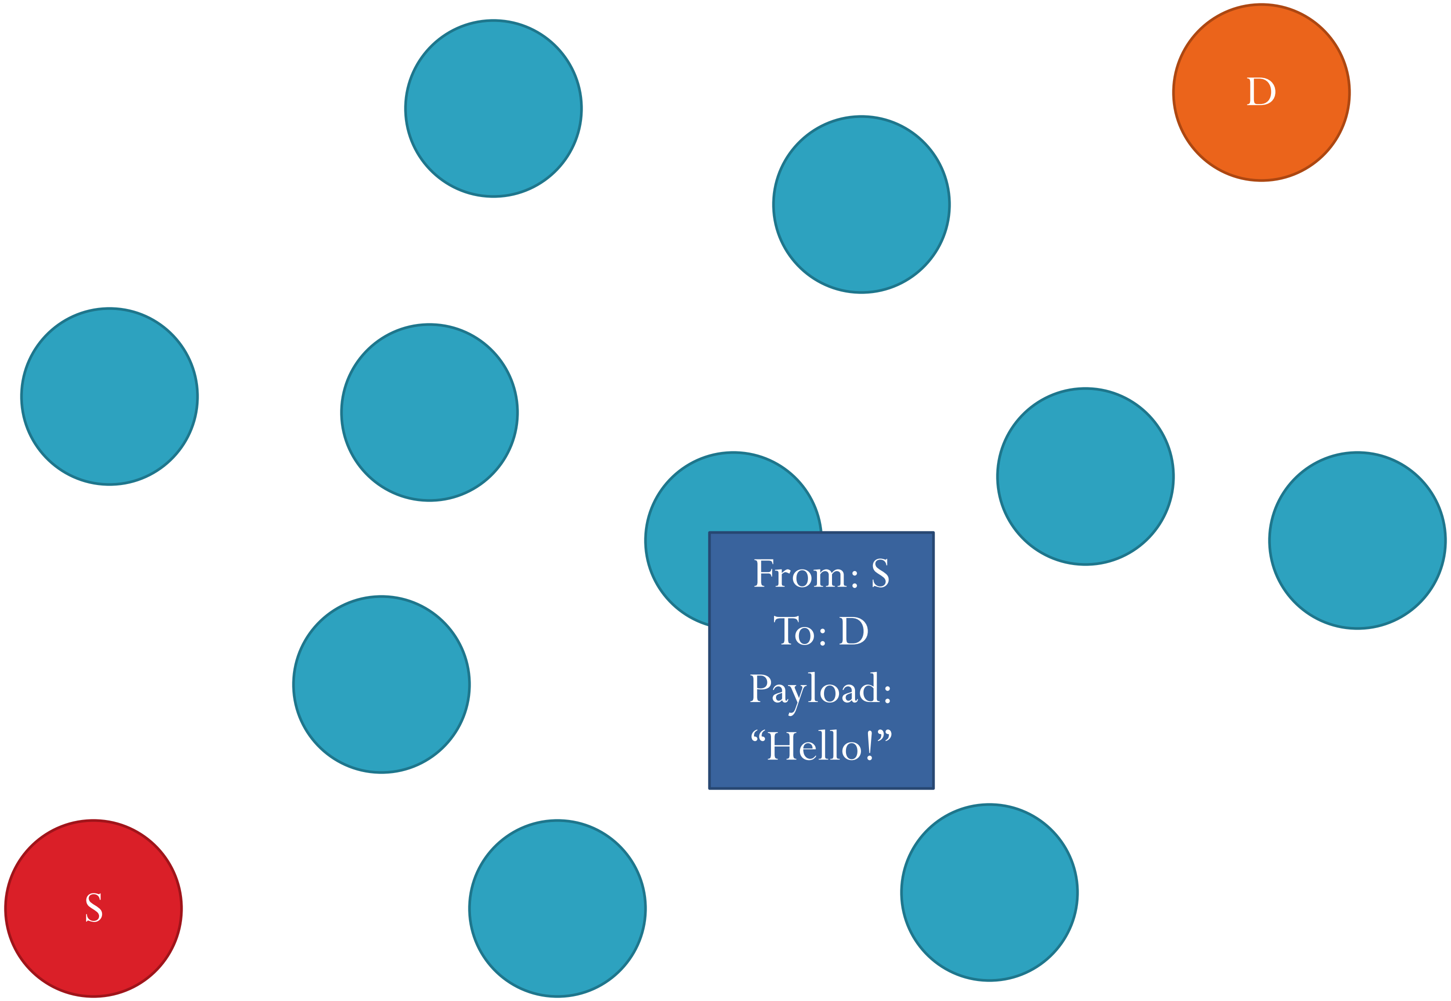

D

From: S
To: D
Payload:
“Hello!”

S

A Really Fast Intro to the Internet

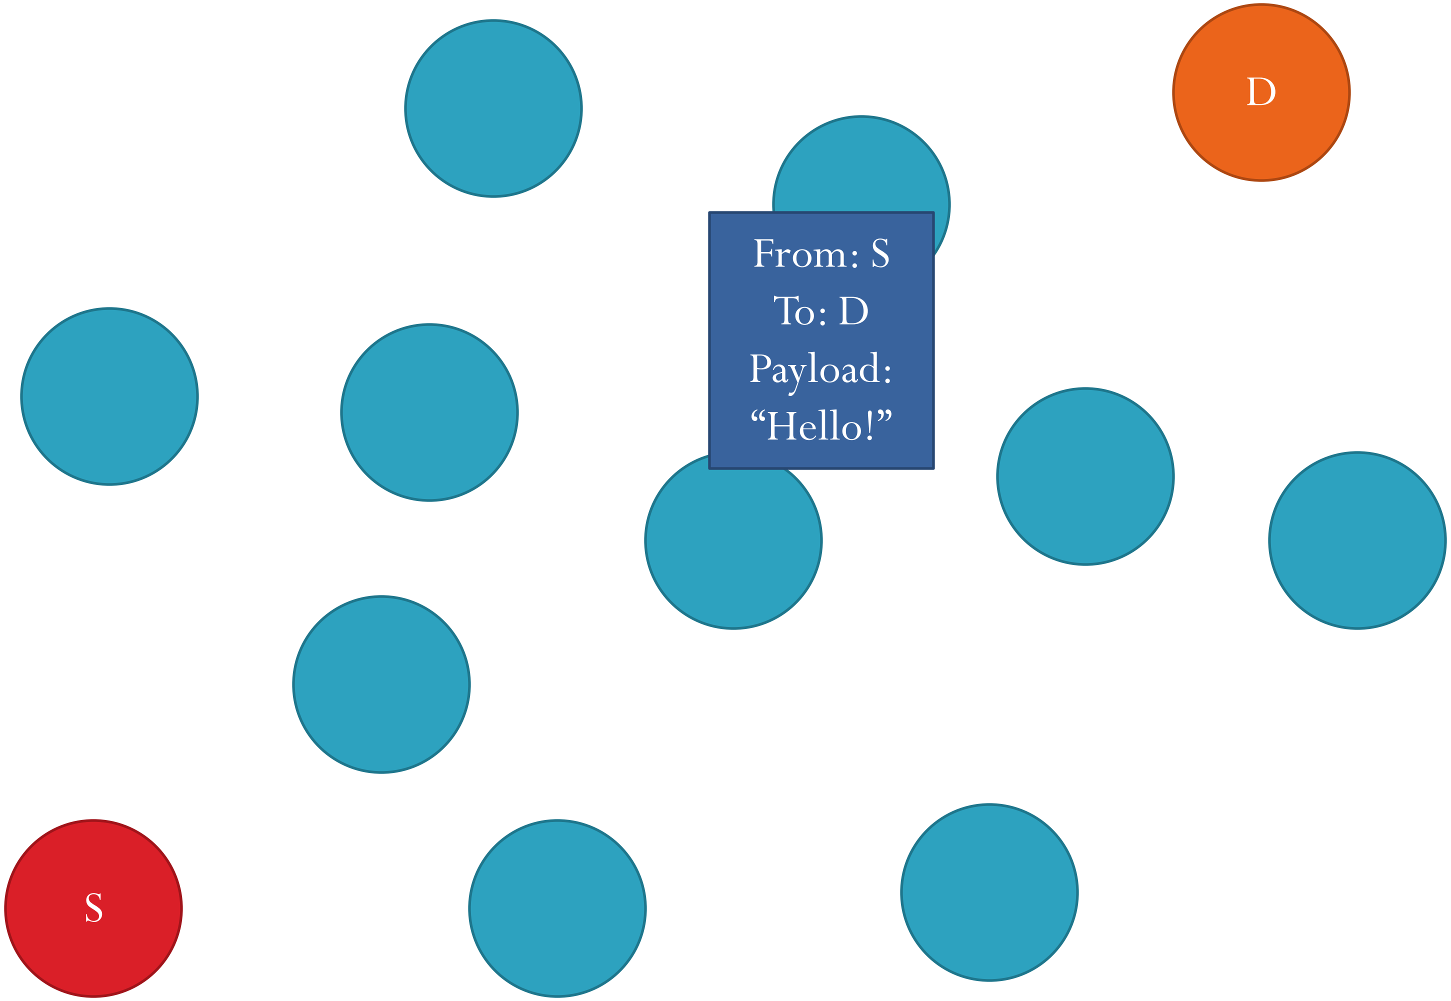

A Really Fast Intro to the Internet

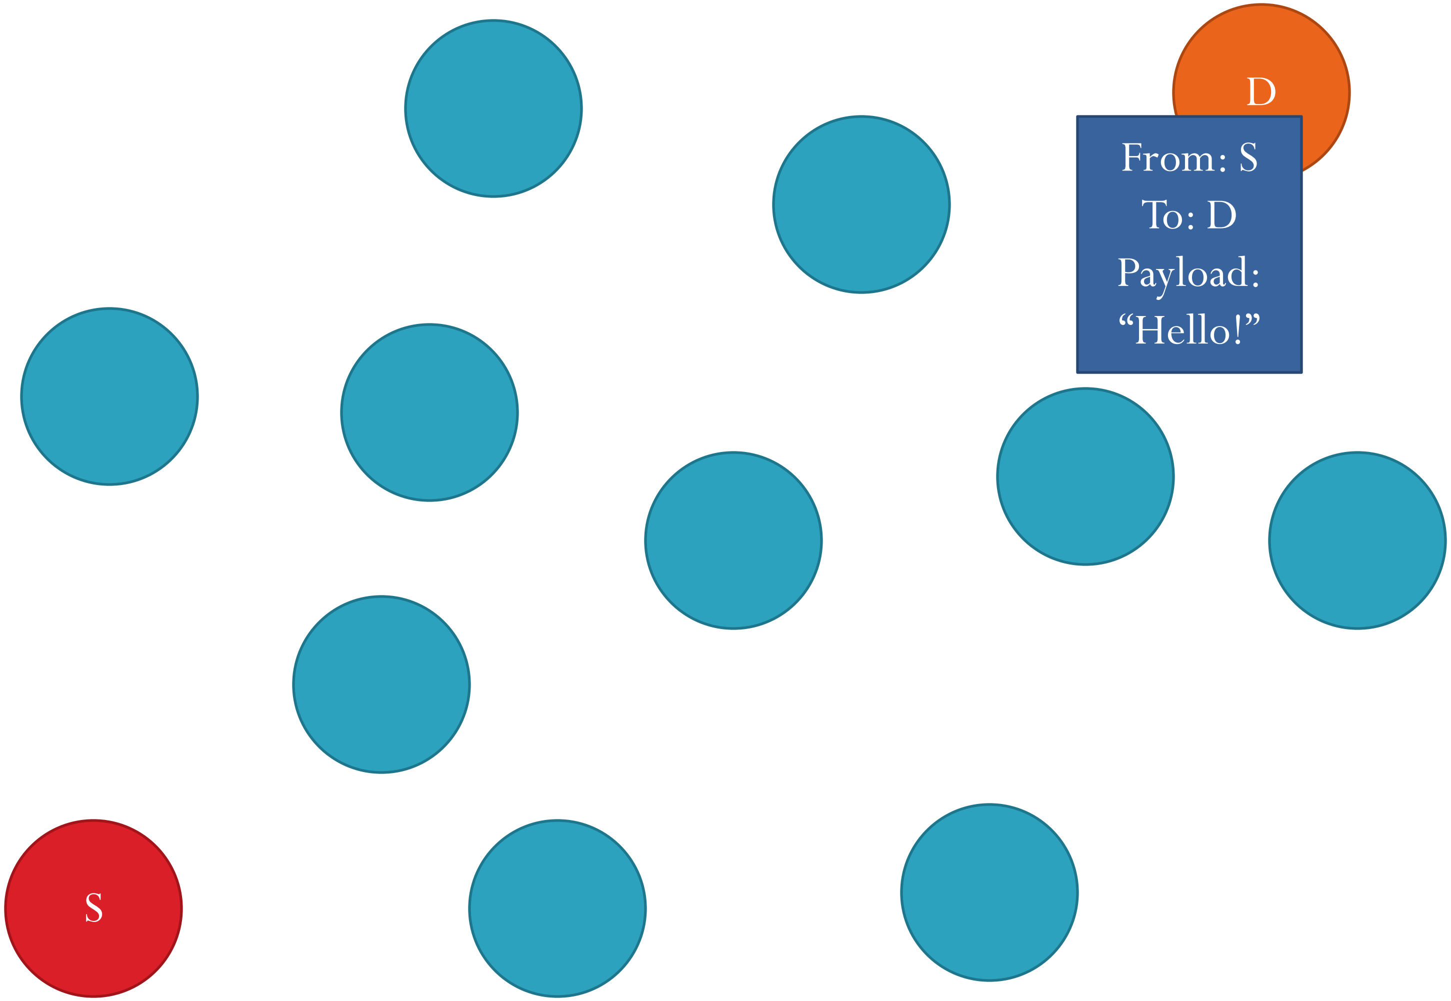

A Really Fast Intro to the Internet

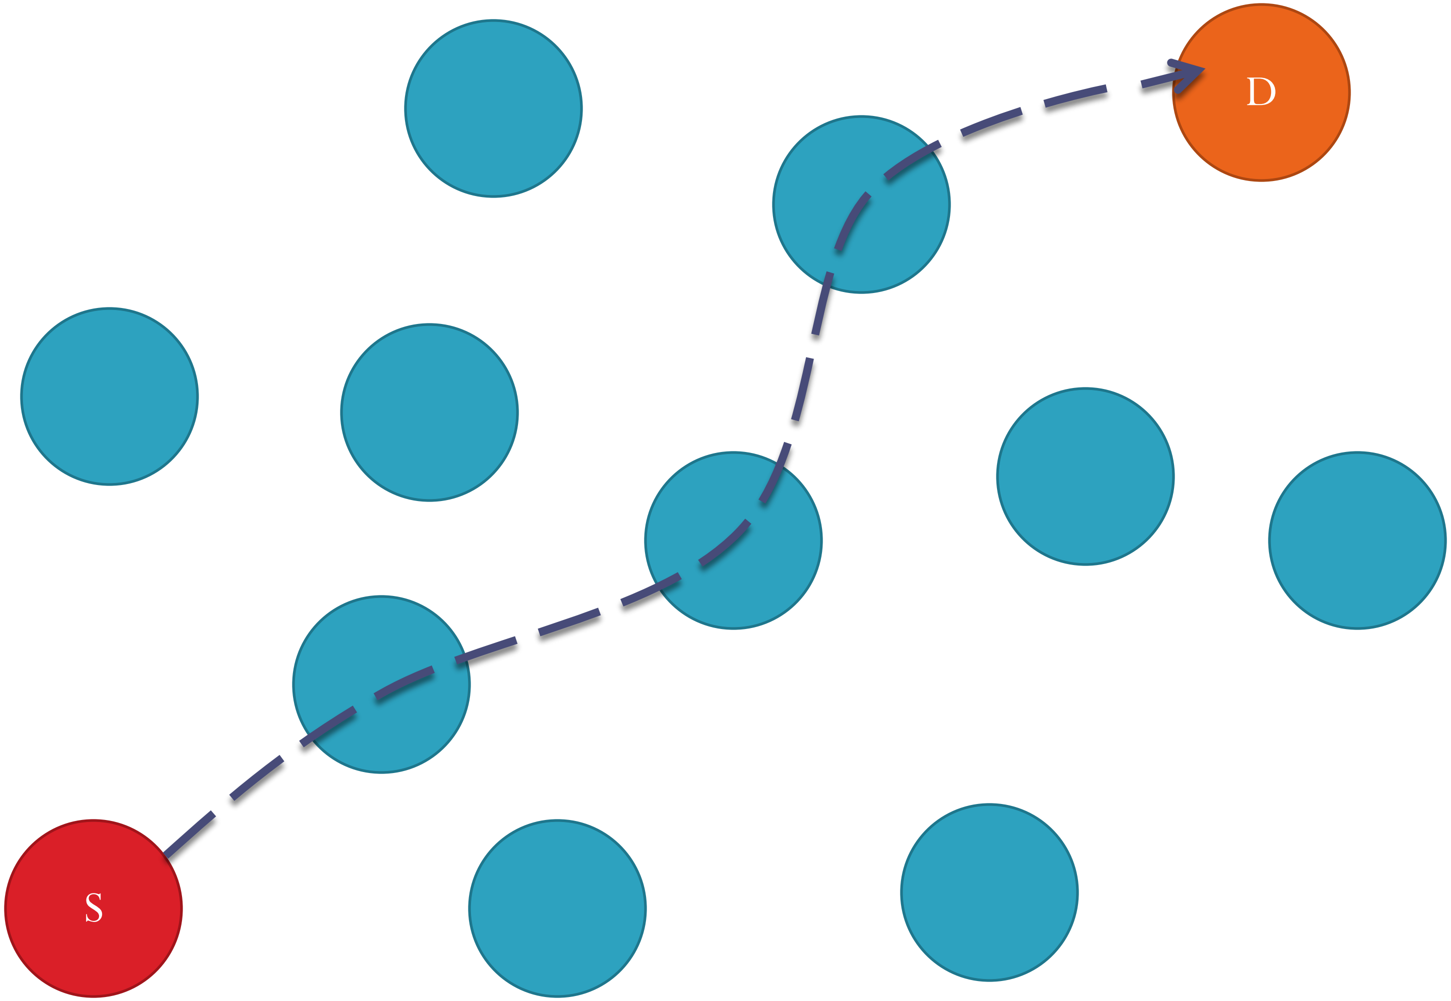

Agenda

- Understanding the Problem: Why Measure the Internet?
- Existing measurement tools
- A new basic tool: IP timestamp
- A large measurement tool: Reverse Traceroute
- Final thoughts

Agenda

- Understanding the Problem: Why Measure the Internet?
- Existing measurement tools
- A new basic tool: IP timestamp
- A large measurement tool: Reverse Traceroute
- Final thoughts

What does the Internet look like?

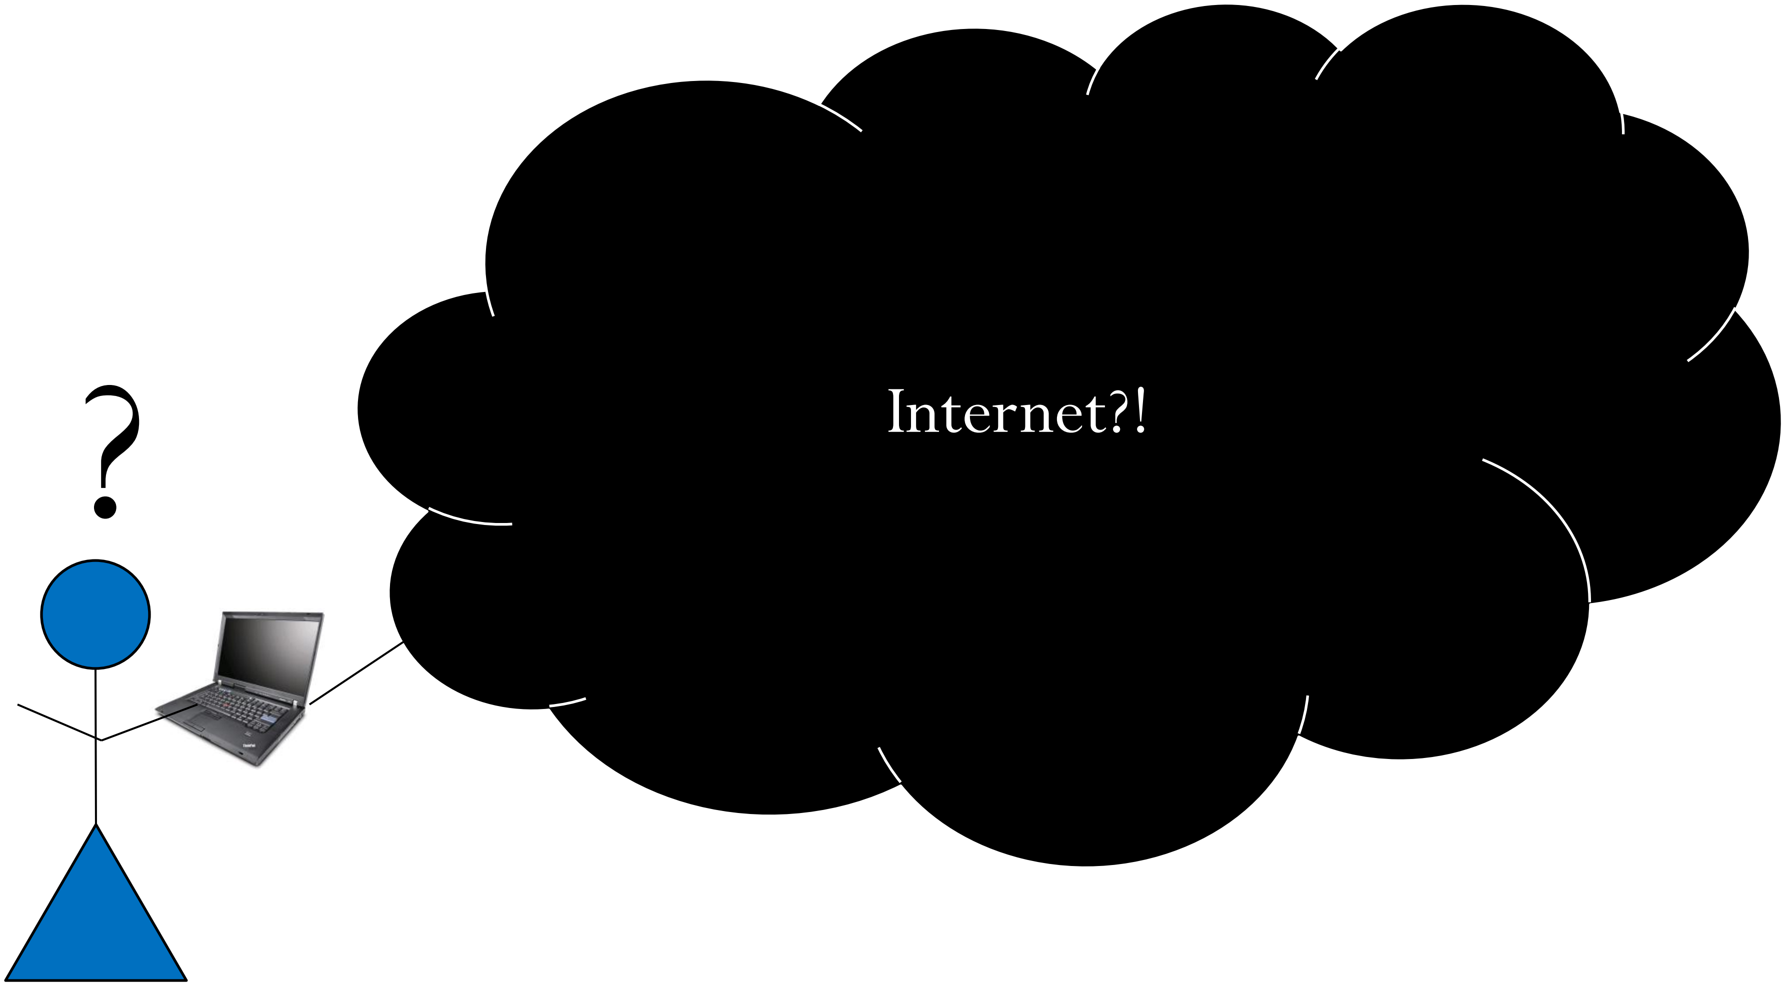

What does the Internet look like?

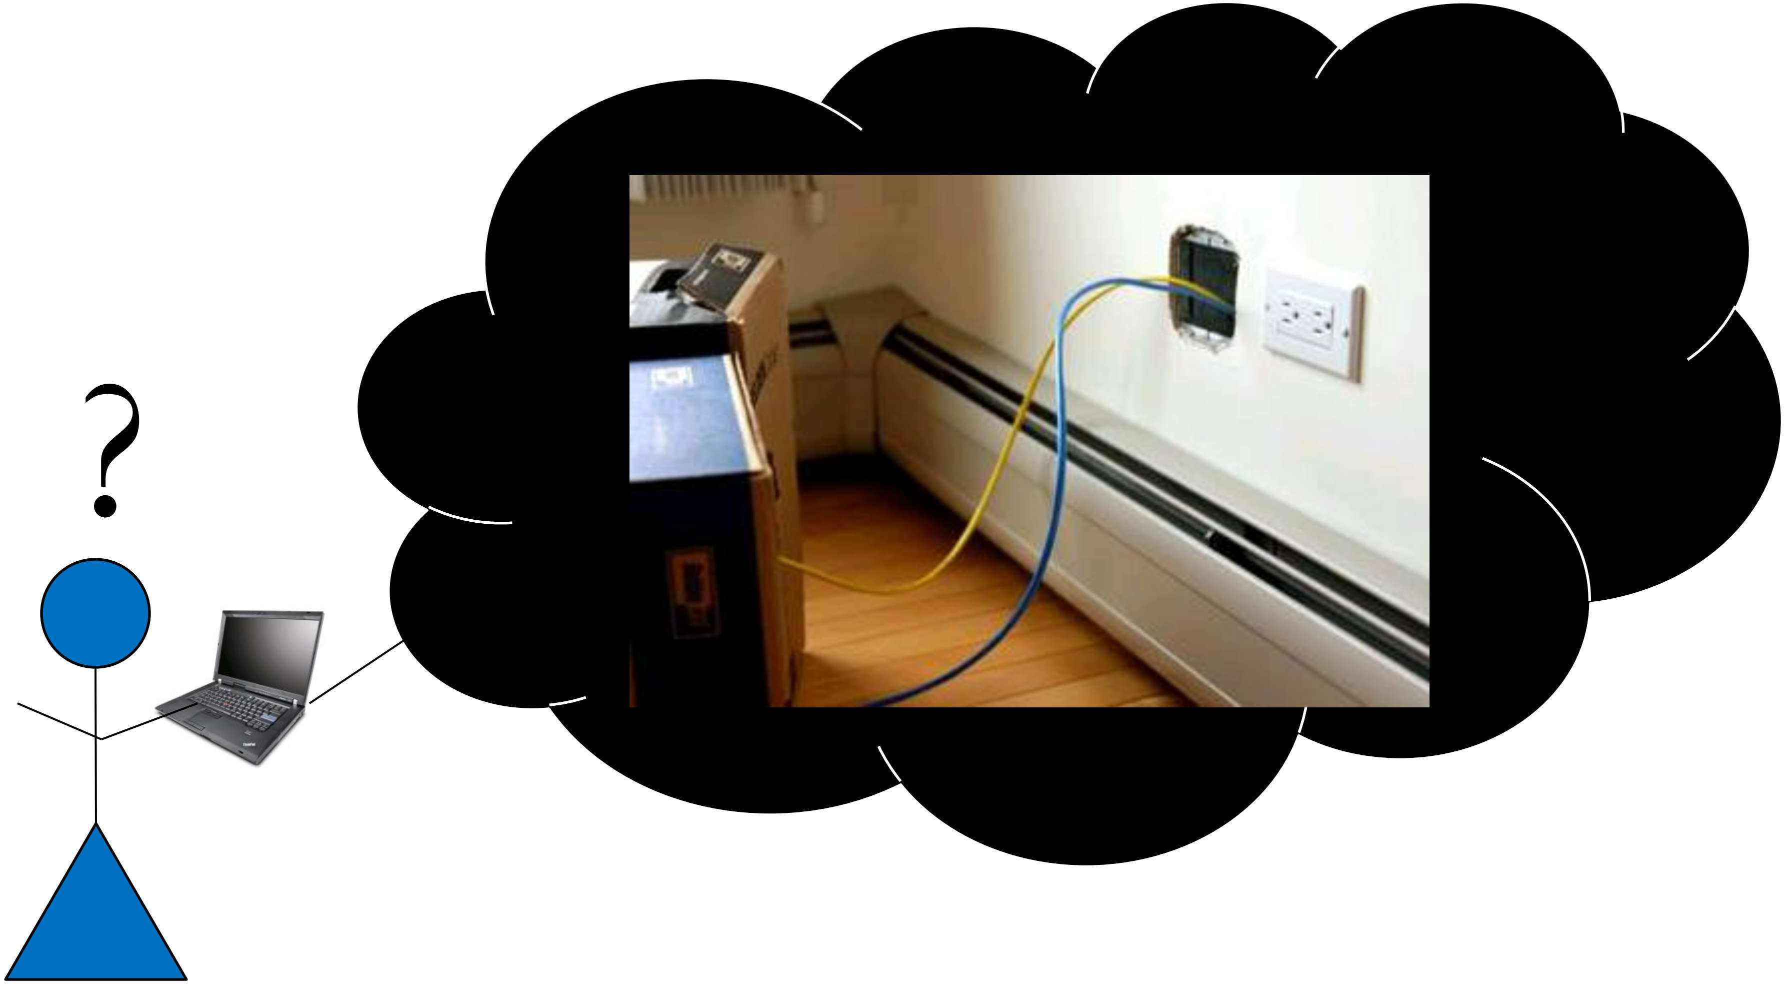

What does the Internet look like?

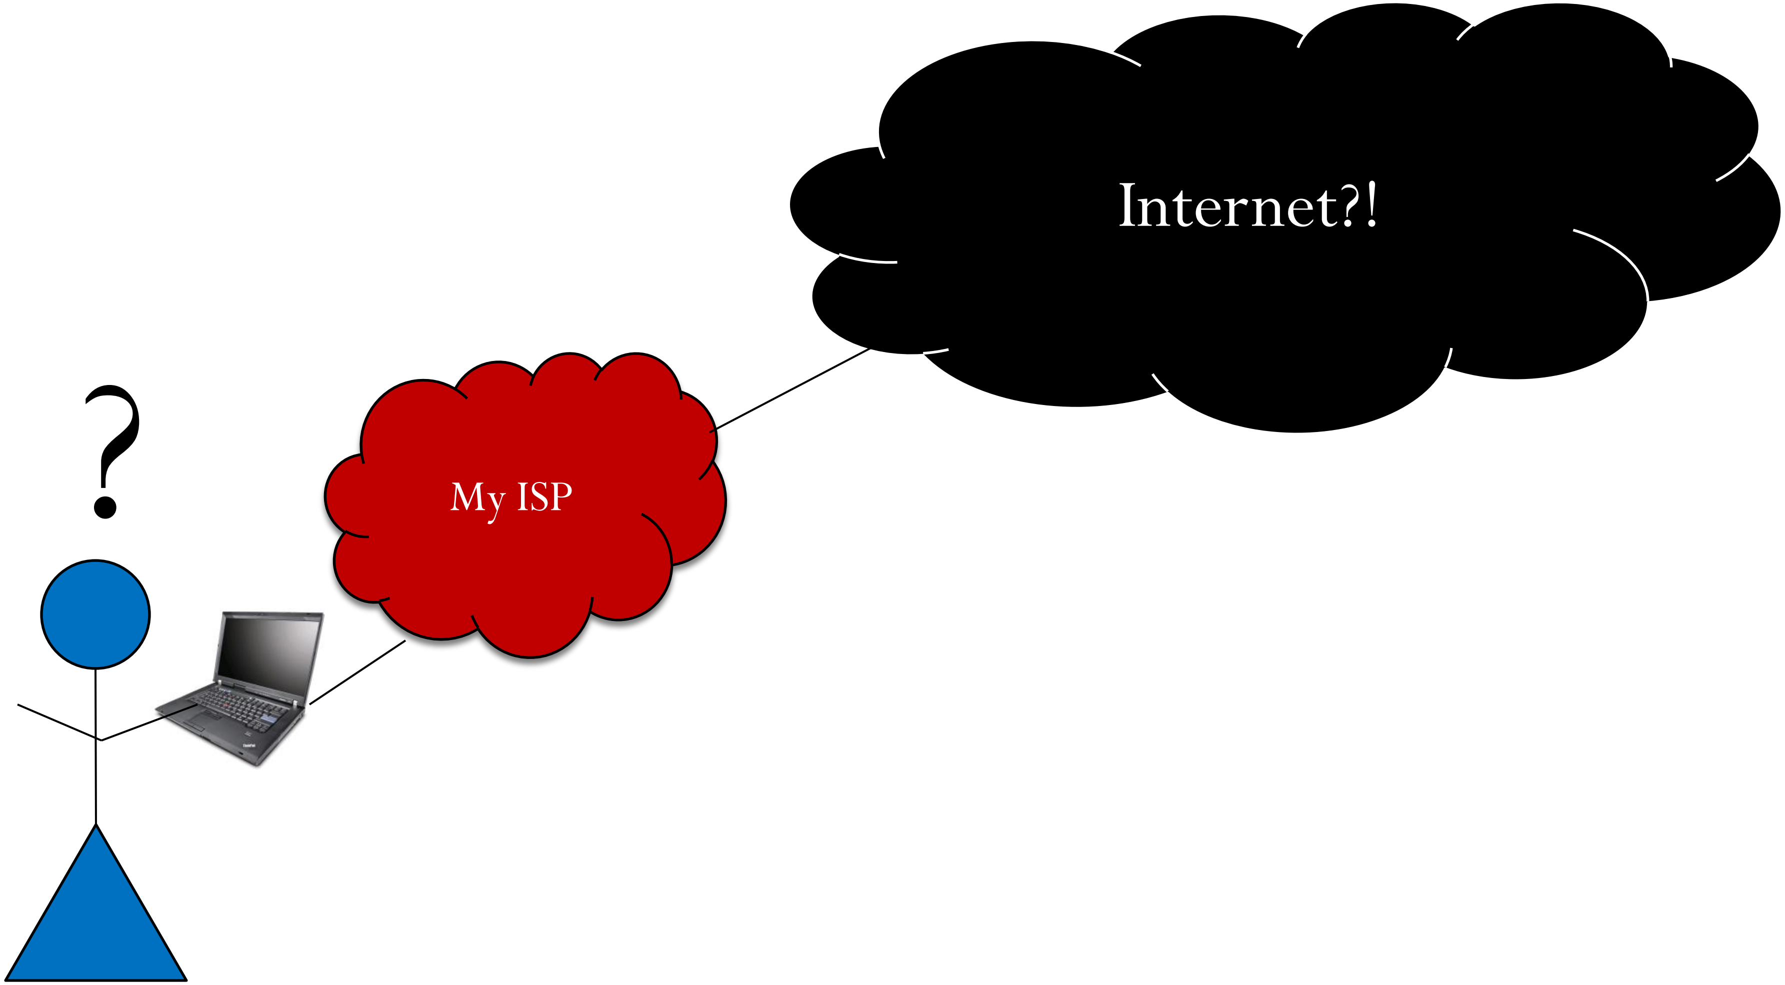

What does the Internet look like?

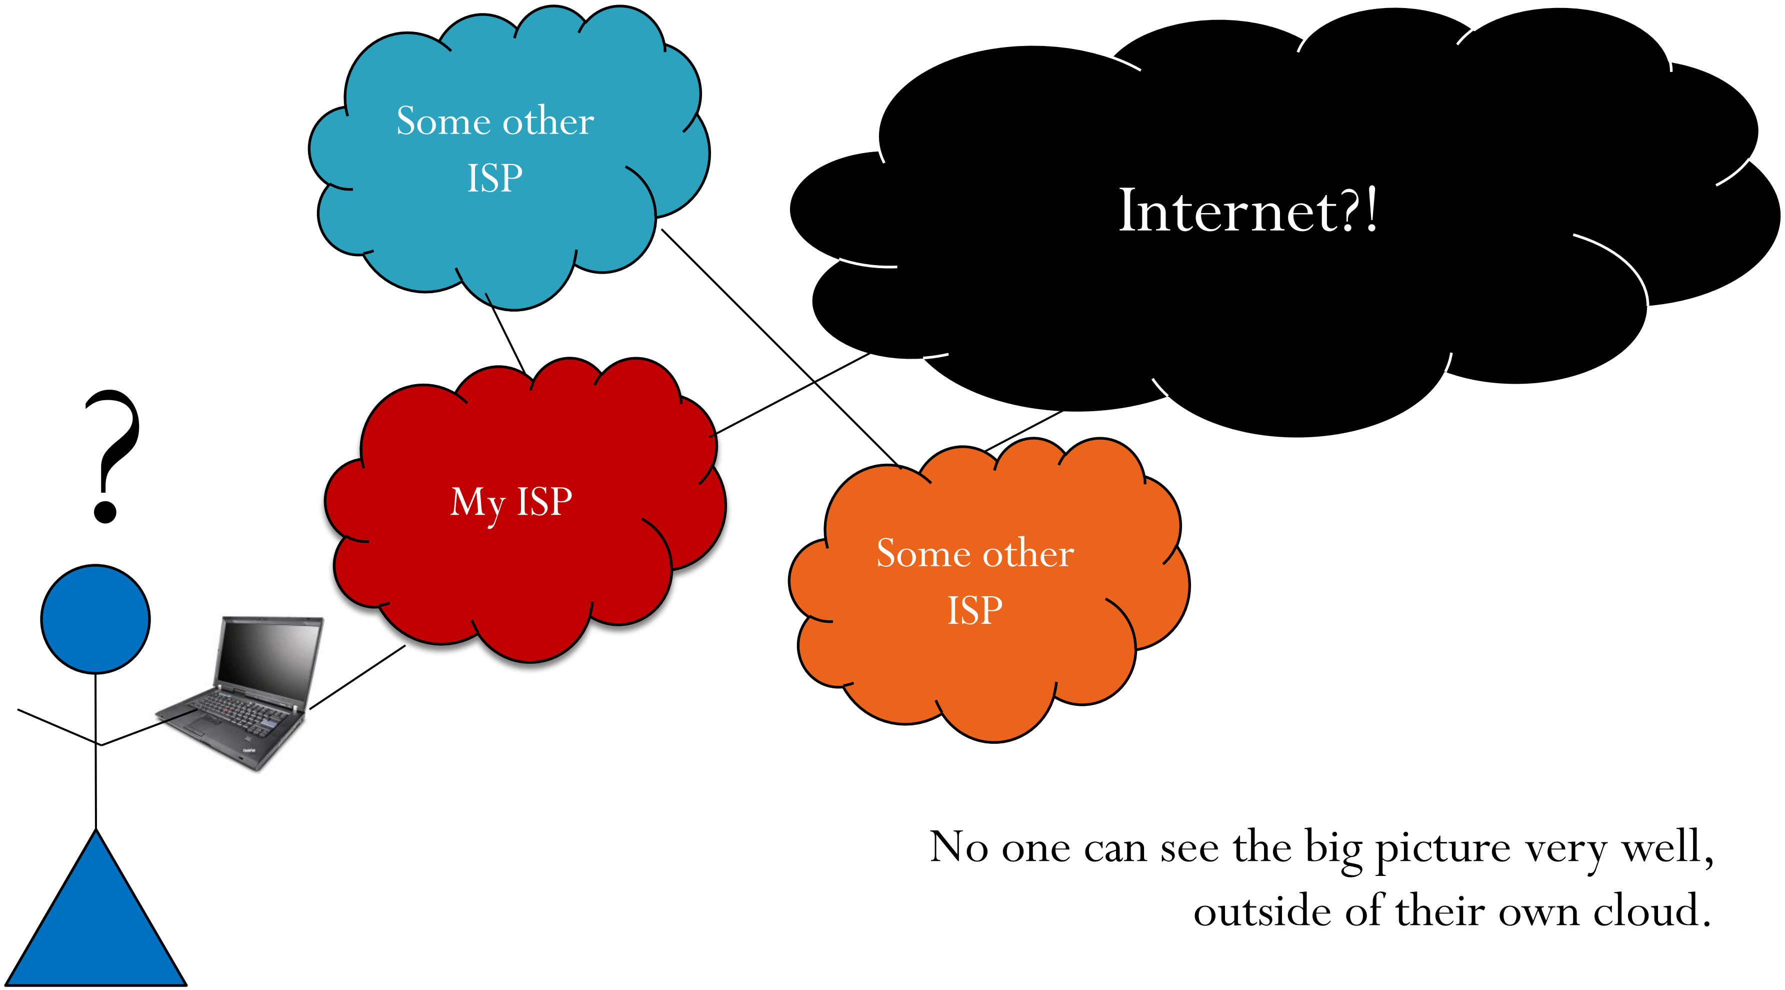

No one can see the big picture very well, outside of their own cloud.

Who cares?

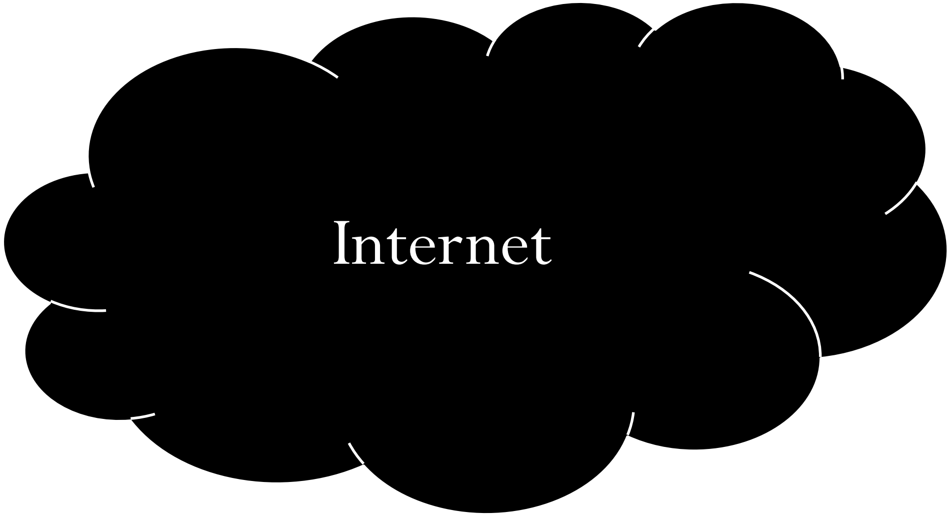

Who cares?

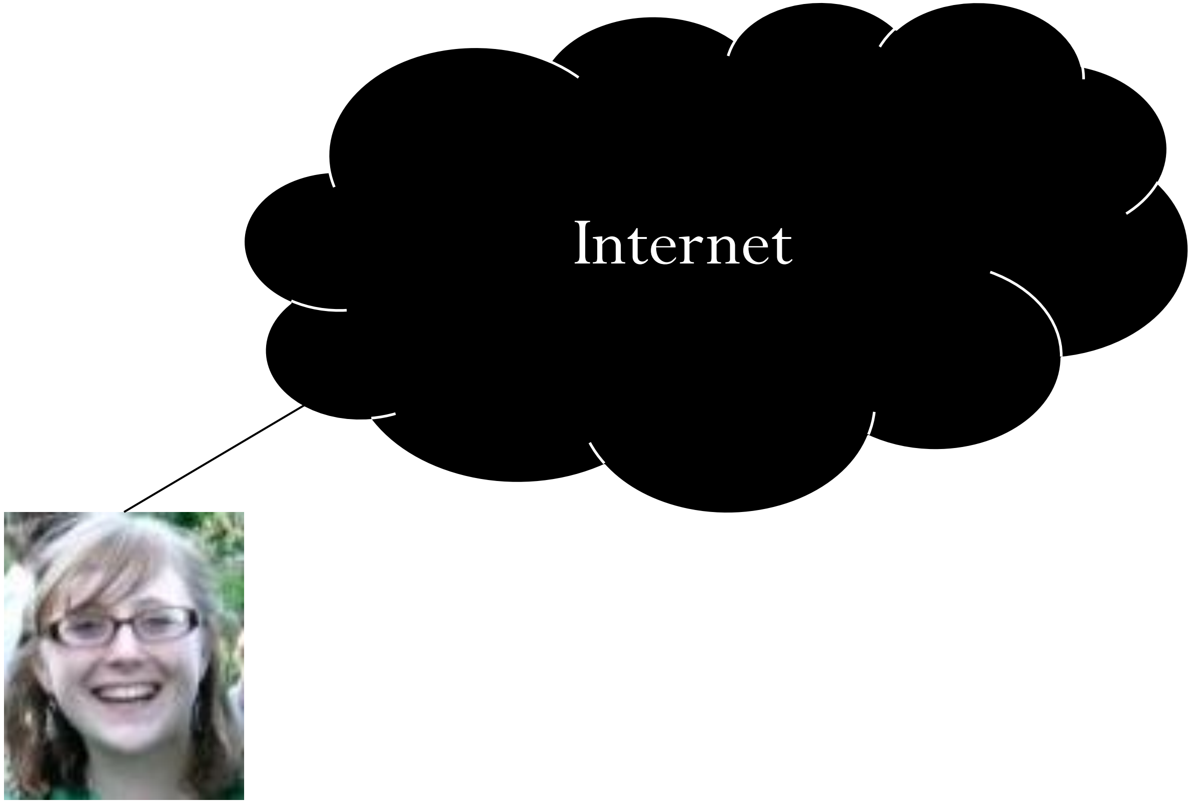

Internet

Who cares?

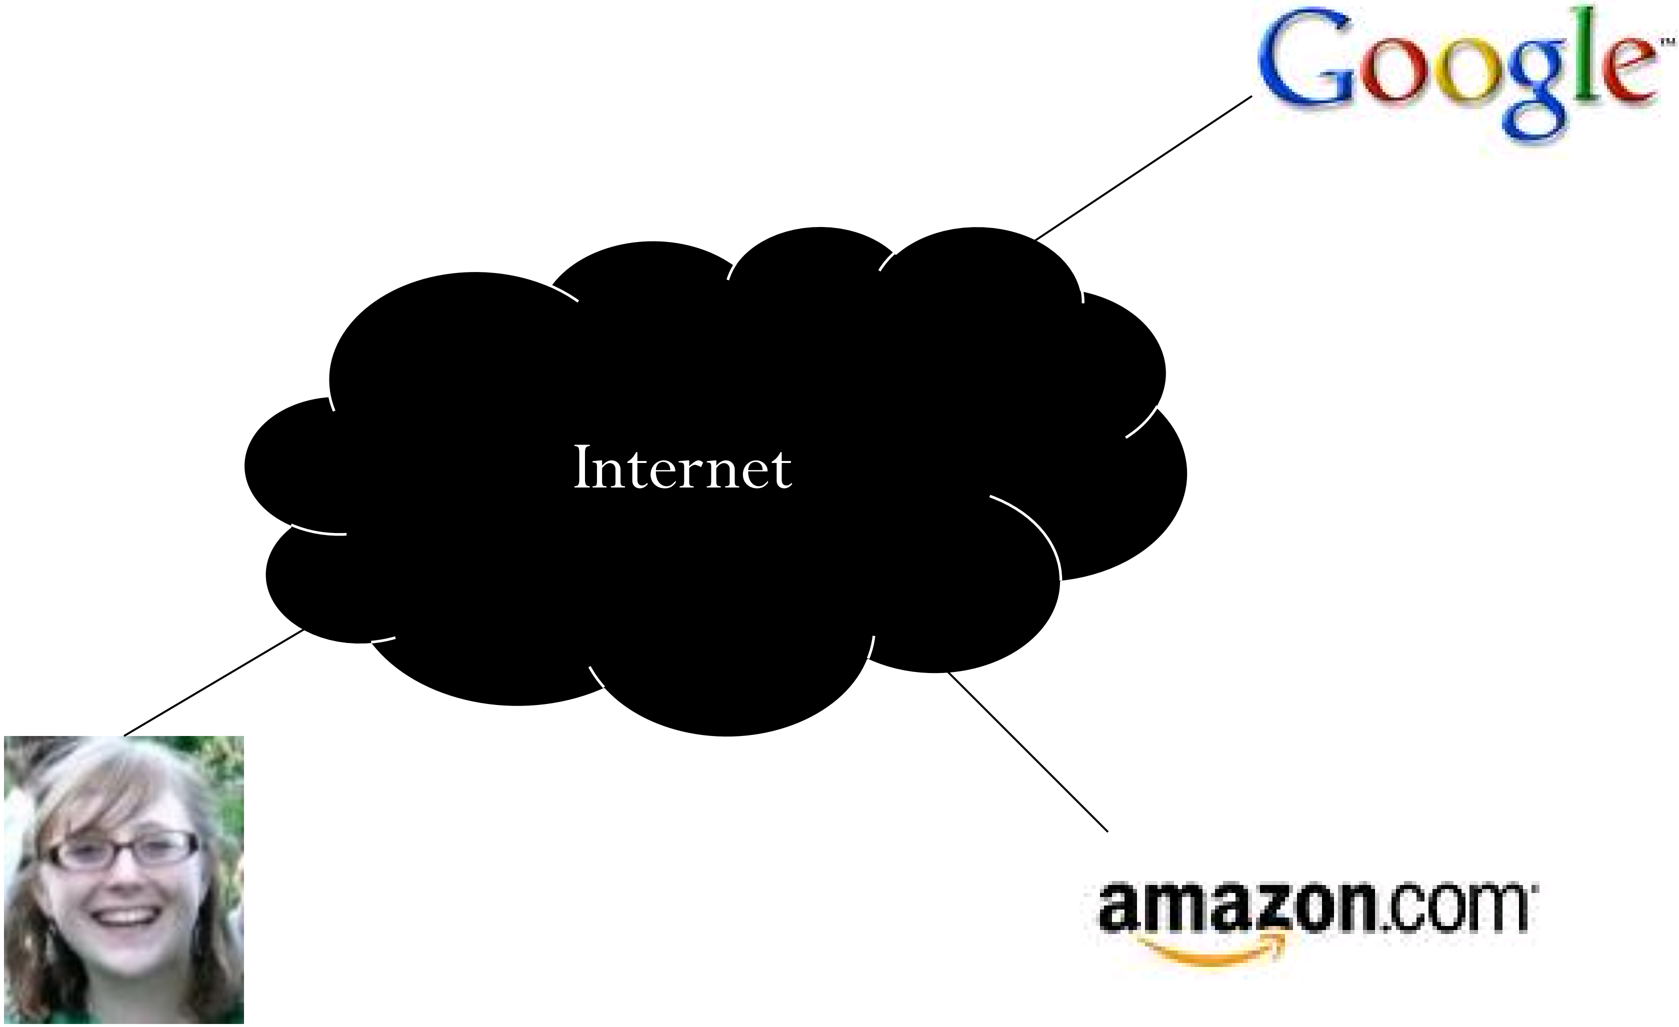

Internet

Who cares?

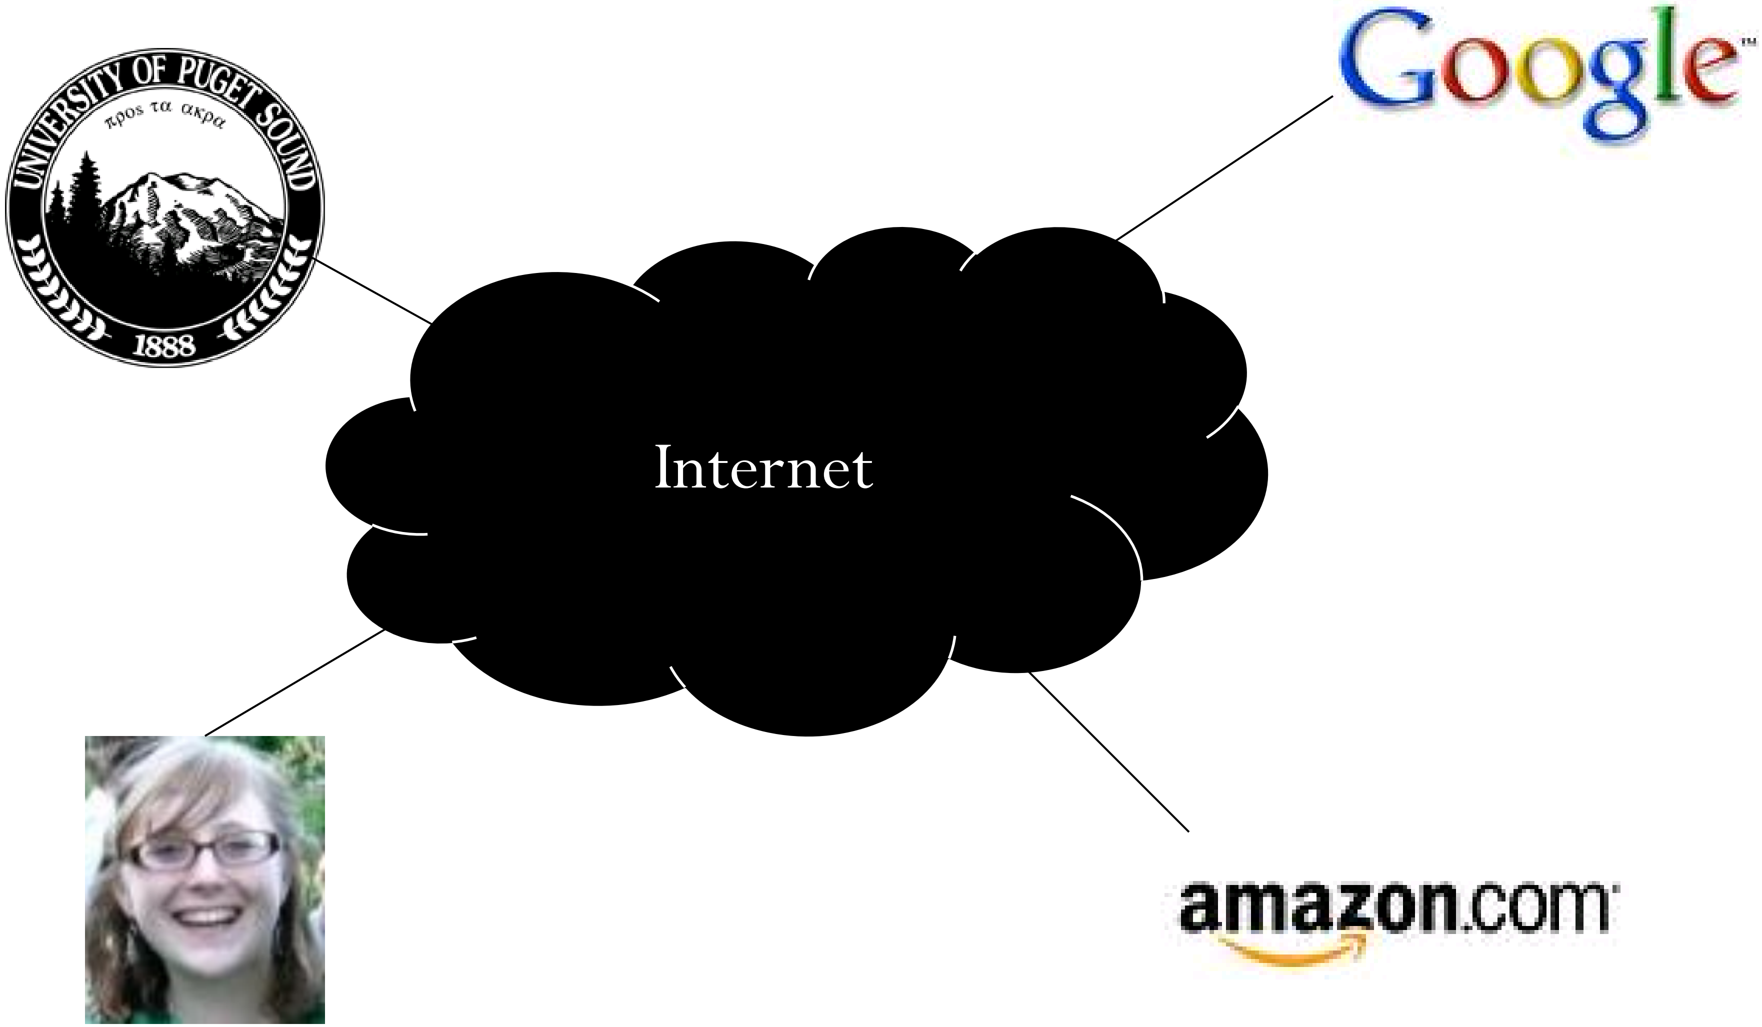
UNIVERSITY OF PUGET SOUND
προς τα ακρα
1888
Google
Internet
amazon.com

Who cares?

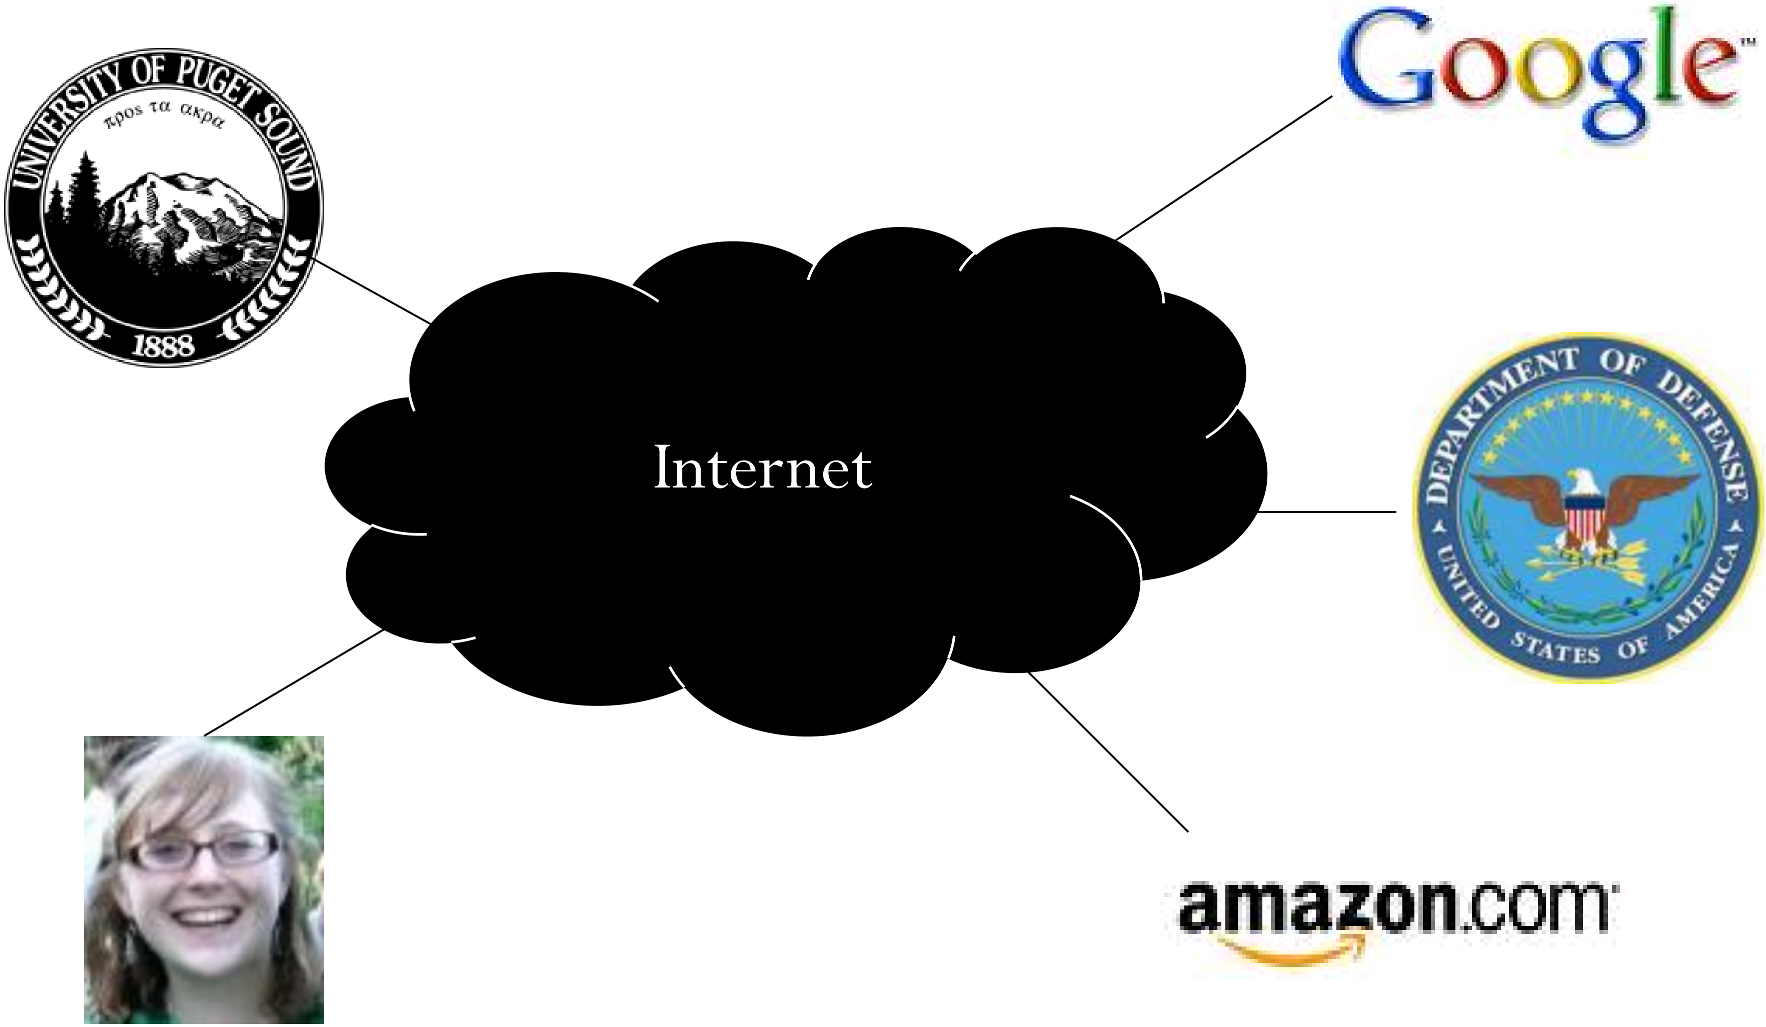

Internet

Agenda

- Understanding the Problem: Why Measure the Internet?
- Existing measurement tools
- A new basic tool: IP timestamp
- A large measurement tool: Reverse Traceroute
- Final thoughts

The Toolbox

- Small tools: small tricks that allow you to discover bits and pieces of information about the network.
- Large systems: huge applications that make use of many measurements from smaller tools to understand large-scale problems.

Ping

- Simple question: is this machine able to receive a message from me, and send me a response?

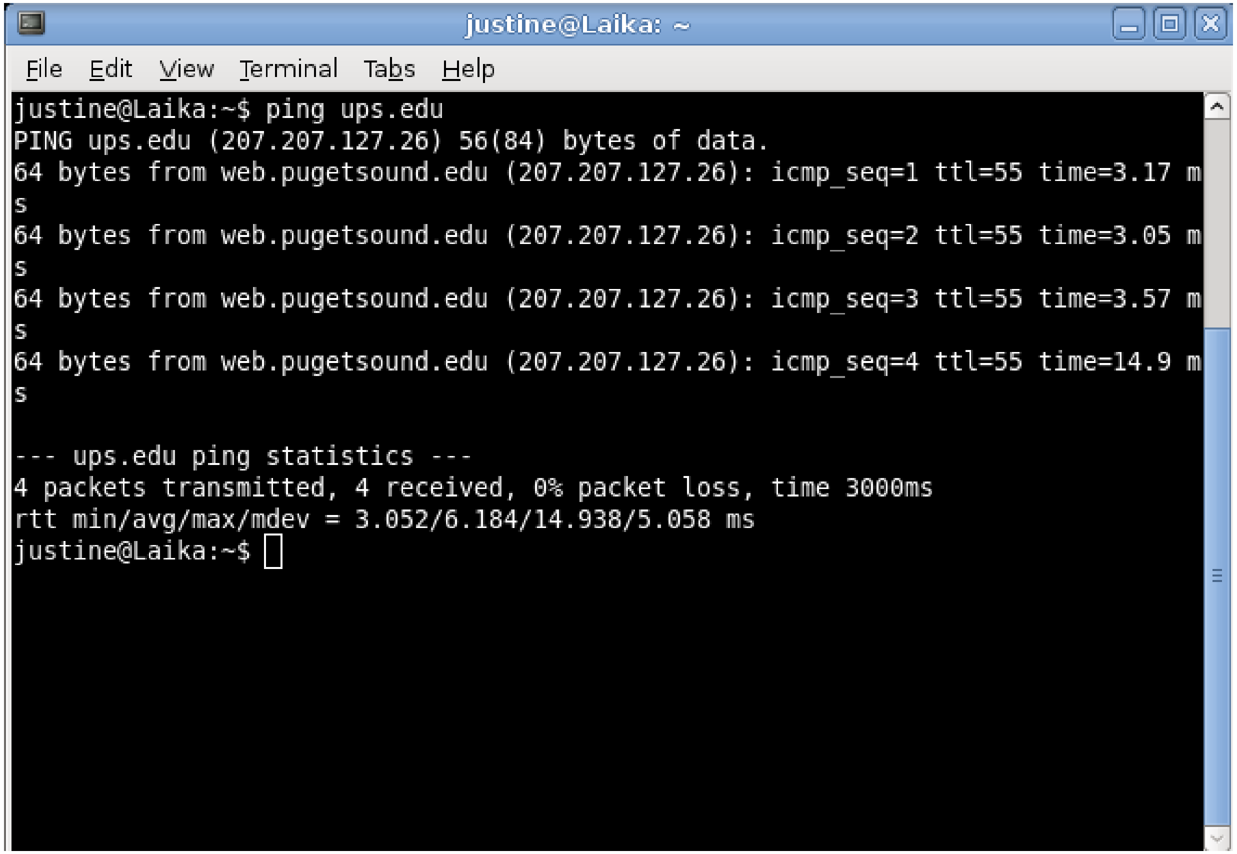

```
justine@Laika:~$ ping ups.edu
PING ups.edu (207.207.127.26) 56(84) bytes of data.
64 bytes from web.pugetsound.edu (207.207.127.26): icmp_seq=1 ttl=55 time=3.17 m
s
64 bytes from web.pugetsound.edu (207.207.127.26): icmp_seq=2 ttl=55 time=3.05 m
s
64 bytes from web.pugetsound.edu (207.207.127.26): icmp_seq=3 ttl=55 time=3.57 m
s
64 bytes from web.pugetsound.edu (207.207.127.26): icmp_seq=4 ttl=55 time=14.9 m
s

--- ups.edu ping statistics ---
4 packets transmitted, 4 received, 0% packet loss, time 3000ms
rtt min/avg/max/mdev = 3.052/6.184/14.938/5.058 ms
justine@Laika:~$ 
```

Traceroute

- Answers the question: what path does a packet take *from* the source *to* the destination?
- Provides an ordered list of IP addresses, one for each router that forwards the packet between the source and destination.
- How it works [extra credit]: Makes use of a neat trick with the Time-to-Live "expiration date" on the packet. By sending probes from the source to the destination with a low TTL value, the packet will expire along the way, generating an error message for the sender from a router in the middle.

Traceroute

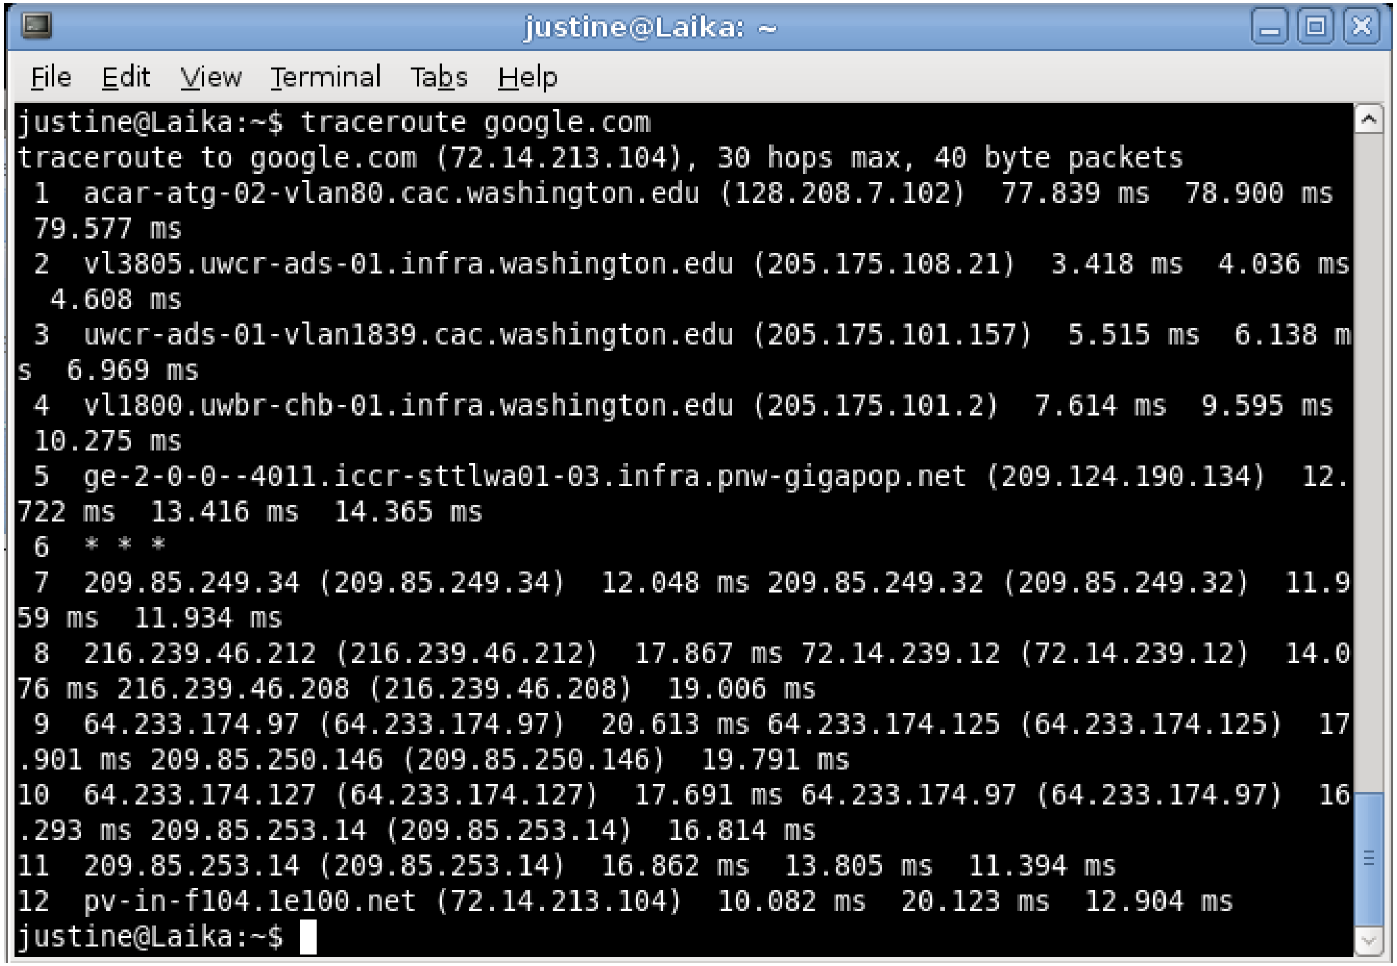

```
justine@Laika:~$ traceroute google.com
traceroute to google.com (72.14.213.104), 30 hops max, 40 byte packets
 1  acar-atg-02-vlan80.cac.washington.edu (128.208.7.102)  77.839 ms  78.900 ms
 79.577 ms
 2  vl3805.uwcr-ads-01.infra.washington.edu (205.175.108.21)  3.418 ms  4.036 ms
  4.608 ms
 3  uwcr-ads-01-vlan1839.cac.washington.edu (205.175.101.157)  5.515 ms  6.138 m
s  6.969 ms
 4  vl1800.uwbr-chb-01.infra.washington.edu (205.175.101.2)  7.614 ms  9.595 ms
 10.275 ms
 5  ge-2-0-0--4011.iccr-sttlwa01-03.infra.pnw-gigapop.net (209.124.190.134)  12.
722 ms  13.416 ms  14.365 ms
 6  * * *
 7  209.85.249.34 (209.85.249.34)  12.048 ms 209.85.249.32 (209.85.249.32)  11.9
59 ms  11.934 ms
 8  216.239.46.212 (216.239.46.212)  17.867 ms 72.14.239.12 (72.14.239.12)  14.0
76 ms 216.239.46.208 (216.239.46.208)  19.006 ms
 9  64.233.174.97 (64.233.174.97)  20.613 ms 64.233.174.125 (64.233.174.125)  17
.901 ms 209.85.250.146 (209.85.250.146)  19.791 ms
10  64.233.174.127 (64.233.174.127)  17.691 ms 64.233.174.97 (64.233.174.97)  16
.293 ms 209.85.253.14 (209.85.253.14)  16.814 ms
11  209.85.253.14 (209.85.253.14)  16.862 ms  13.805 ms  11.394 ms
12  pv-in-f104.1e100.net (72.14.213.104)  10.082 ms  20.123 ms  12.904 ms
justine@Laika:~$ 
```

…& friends

- There are a lot of tricks people use to gather data:
  - Reading BGP tables (RouteViews from U of Oregon)
  - IGMP "ASK_NEIGHBORS" requests
  - ICMP timestamp requests
  - UDP Ping
  - Etc, etc, etc.

Large Measurement Systems from UW

- Rocketfuel [Neil Spring & co., ‘02]: charting ISP topologies to help us better understand how different providers structure their networks.
- iPlane [Harsha Madhyastha & co., ‘06]: building basic map of the Internet to make accurate predictions about latencies and paths
- Hubble [Ethan K-B, & co., ‘08]: monitoring Internet ‘black holes’ to understand where and how Internet traffic disappears

Agenda

- Understanding the Problem: Why Measure the Internet?
- Existing measurement tools
- A new basic tool: IP timestamp
- A large measurement system: Reverse Traceroute
- Final thoughts

Introducing IP Timestamp

- IP Timestamp is an optional extension to the IP header. It allows the sender to request timestamp values from any machine which handles the packet by specifying its IP address.
- The specification allows you to request timestamps from up to four IP addresses at a time.
- IP Timestamp can help us answer some simple measurement questions.

Making a simple tool do neat tricks

- Push the limits: start asking what kinds of results you might get if you start sending out weird and abnormal measurements.
- We asked a weird question: what if we asked for the same address in all four prespecified requests?

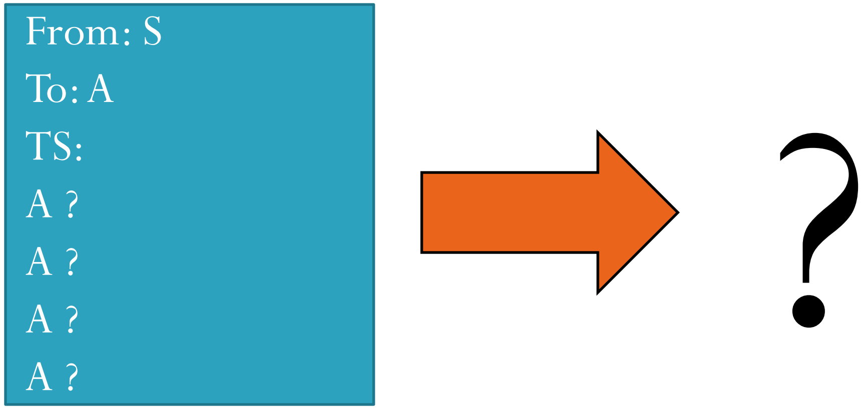

From: S
To: A
TS:
A ?
A ?
A ?
A ?

Making a simple tool do neat tricks

- Push the limits: start asking what kinds of results you might get if you start sending out weird and abnormal measurements.
- We asked a weird question: what if we asked for the same address in all four prespecified requests?

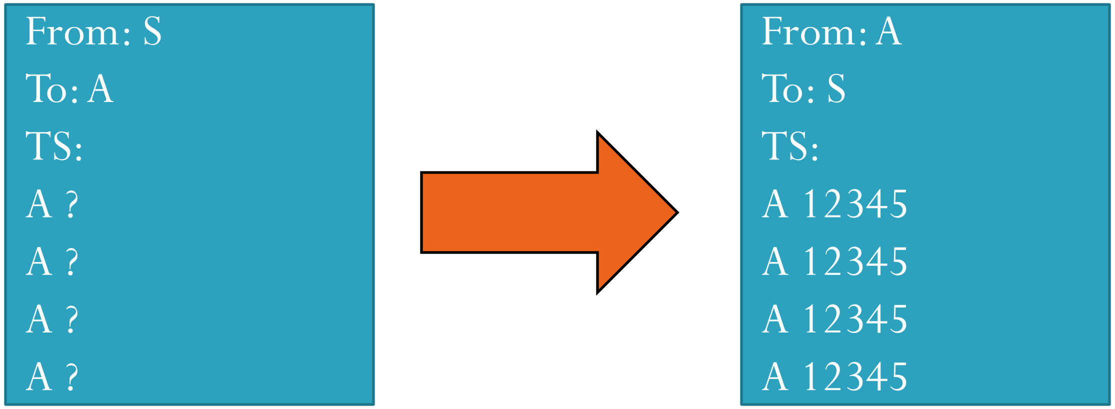

From: S
To: A
TS:
A ?
A ?
A ?
A ?

From: A
To: S
TS:
A 12345
A 12345
A 12345
A 12345

The IP Alias Resolution Problem

- Routers in the Internet have *many* IP addresses.
- Different measurements may interact with different IP addresses, but we want to know whether or not they interacted with the same machine.

Why Alias Resolution is Important

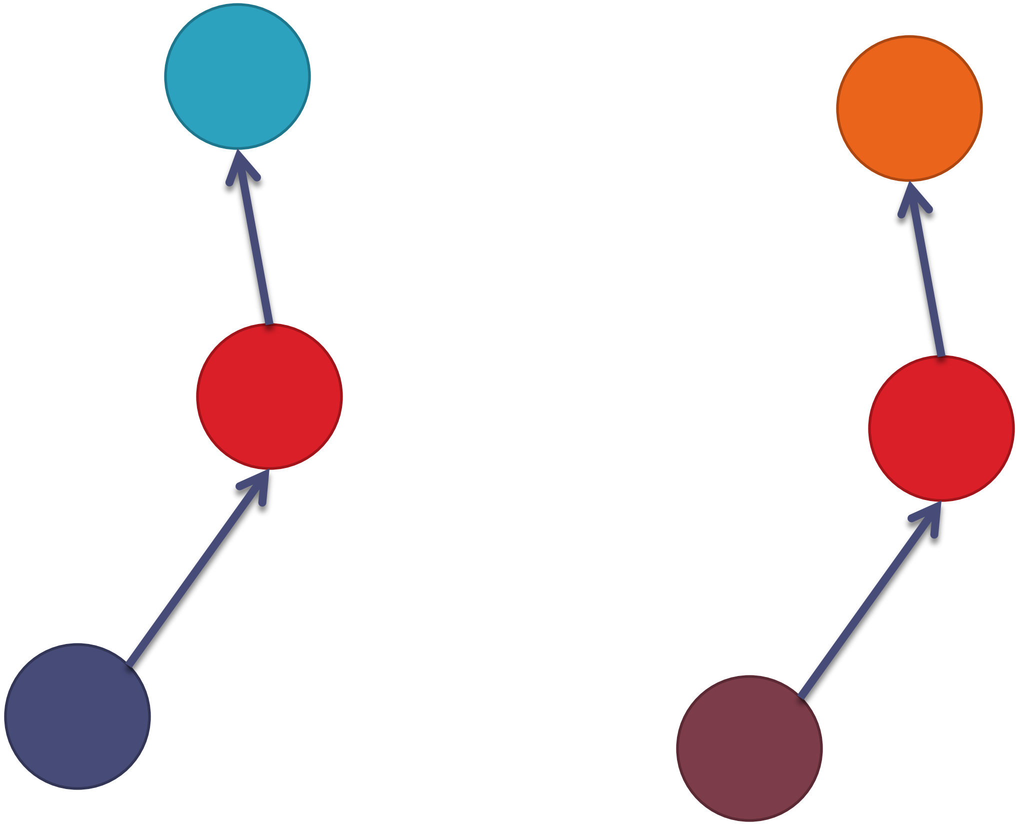

Oftentimes researchers try to intersect traceroutes to understand the network topology – but how to tell they intersect?

Why Alias Resolution is Important

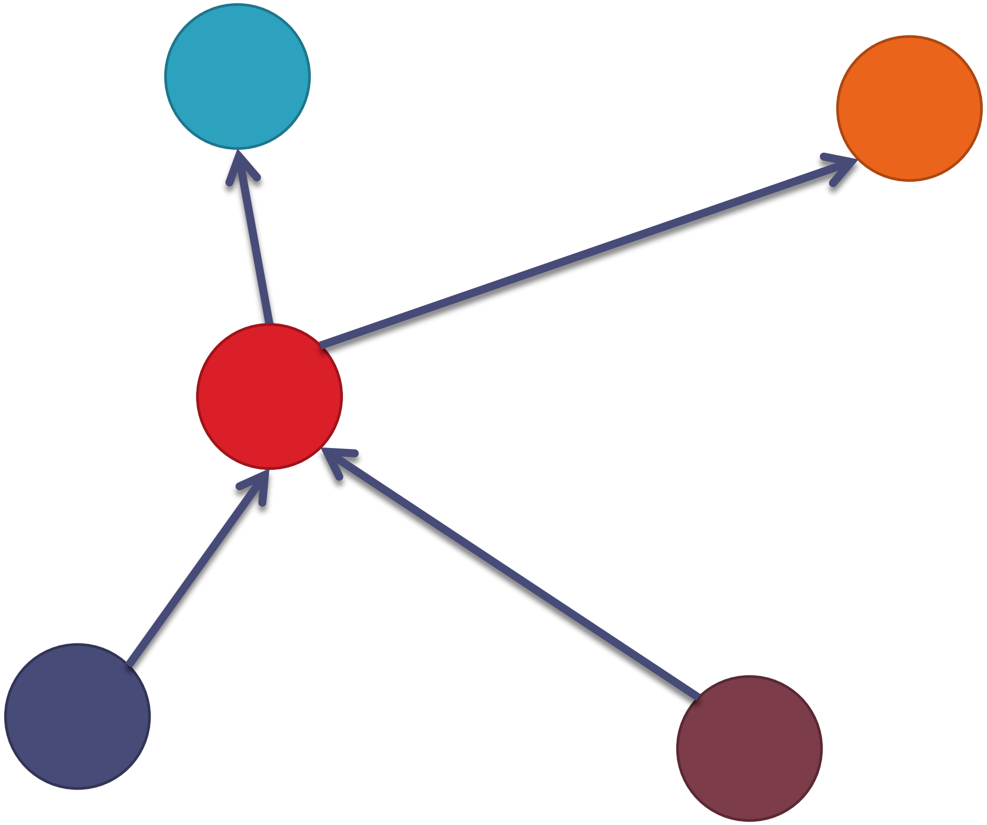

Oftentimes researchers try to intersect traceroutes to understand the network topology – but how to tell they intersect?

Identifying IP Aliases with Timestamps

- Our trick: place different IP addresses that might belong to the same machine, nested together in the same probe.

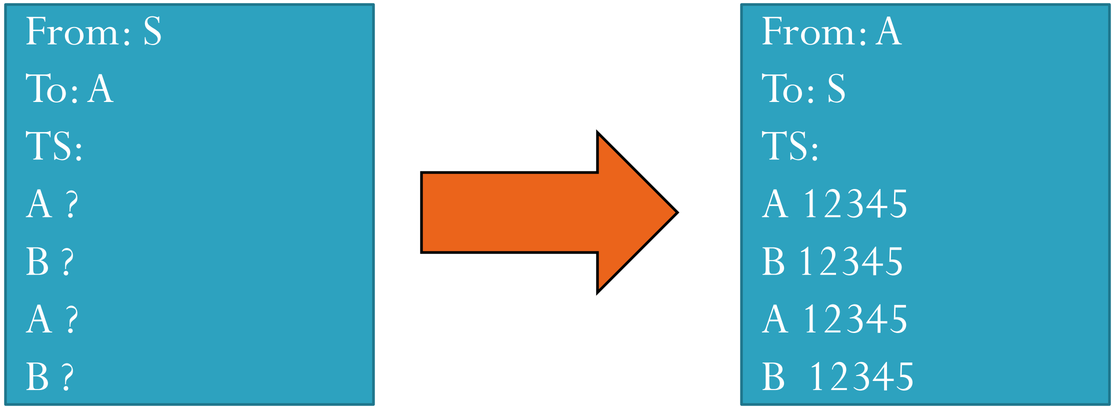

From: S
To: A
TS:
A ?
B ?
A ?
B ?

From: A
To: S
TS:
A 12345
B 12345
A 12345
B 12345

Agenda

- Understanding the Problem: Why Measure the Internet?
- Existing measurement tools
- A new basic tool: IP timestamp
- A large measurement tool: Reverse Traceroute
- More measurement research & final thoughts

Reverse traceroute is a lot harder than forward traceroute

- We’ve known since 1987 how to do forward traceroute.
- That technique works great to see the path from a local source to an uncontrolled destination, but it doesn’t work to find the path from an uncontrolled host back to your local source.
- Previous techniques often assumed that the forward and reverse paths were symmetric, but that is usually not accurate.
- Basically, people thought it was impossible.

Forward vs. Reverse Path

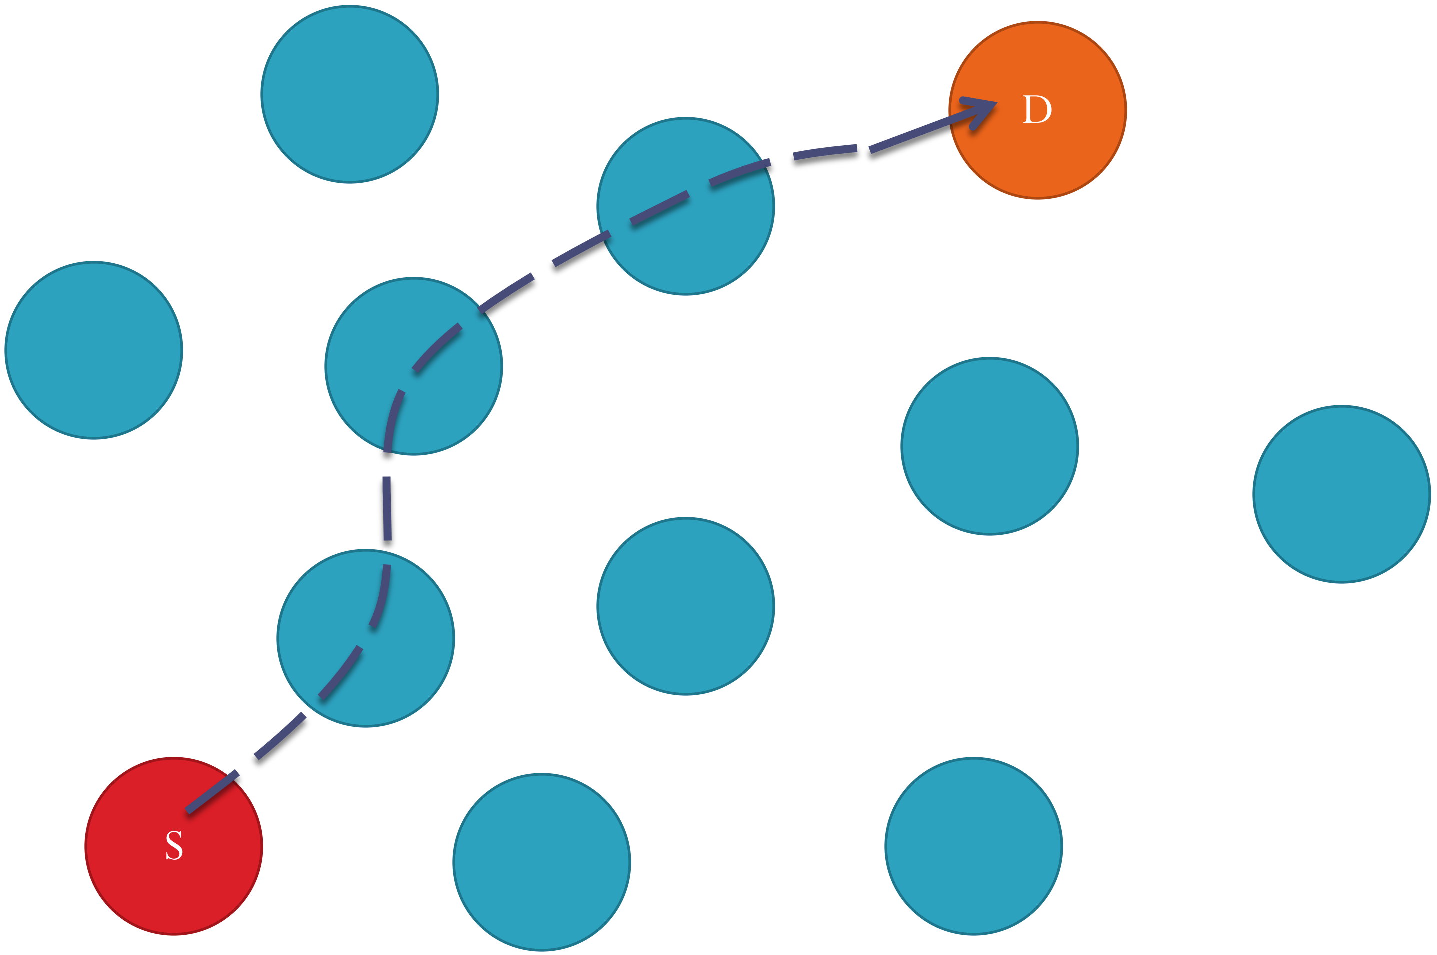

Forward vs. Reverse Path

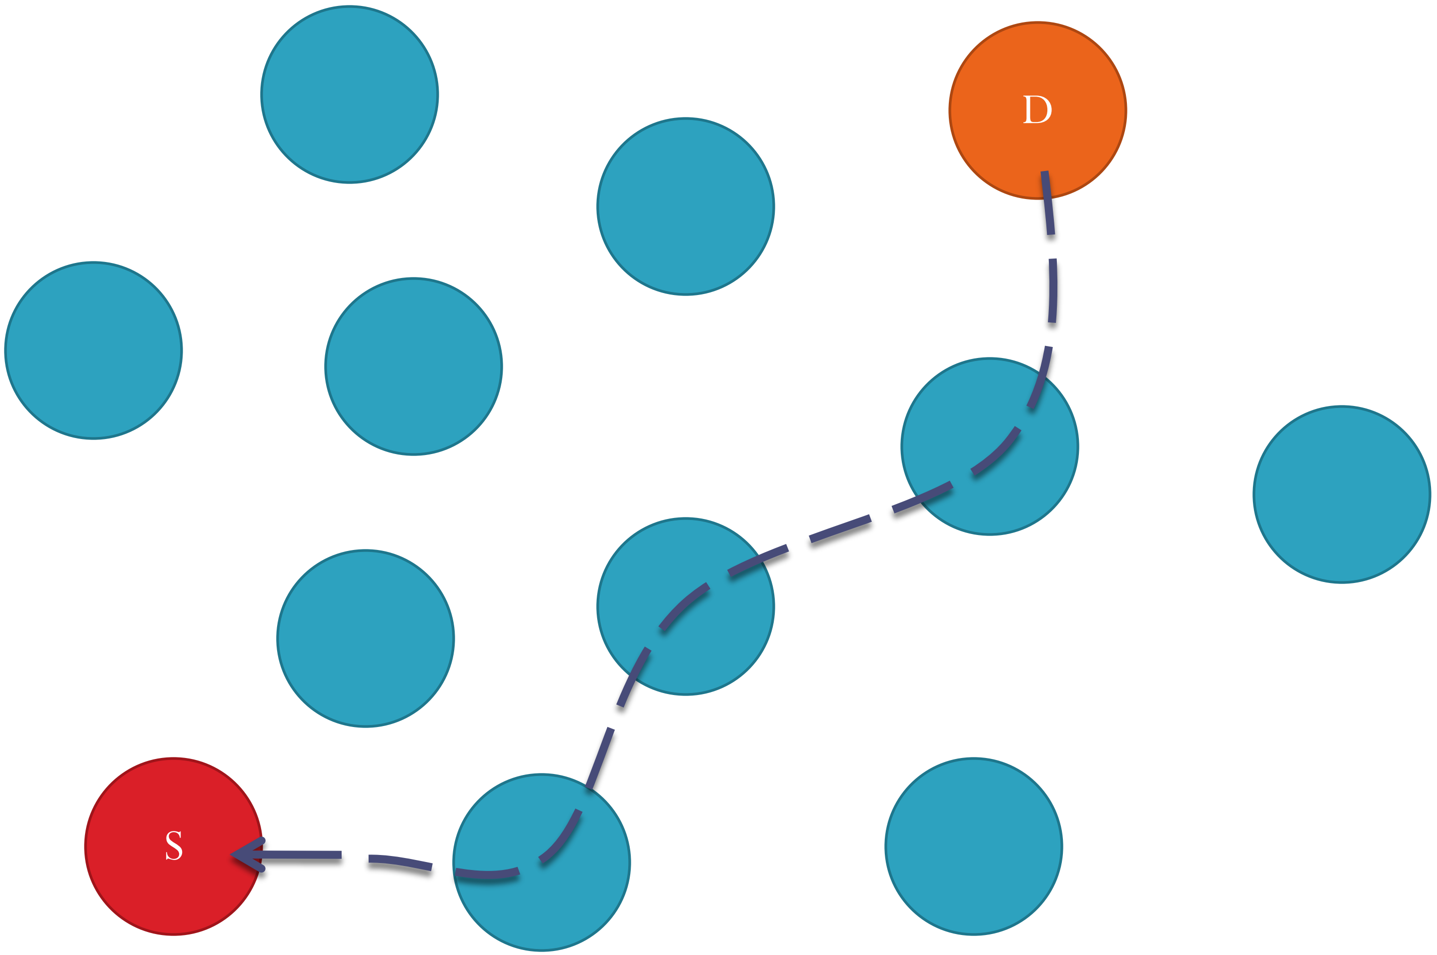

Record Route

- Similar to traceroute! Answers the question: what routers handled this packet while it was traveling?
- Unlike traceroute, IPs are not learned from response packets, but are instead recorded in the packet header. Each machine forwarding the packet affixes its own IP address to the packet before forwarding, up to a maximum of 9 addresses.

Record Route

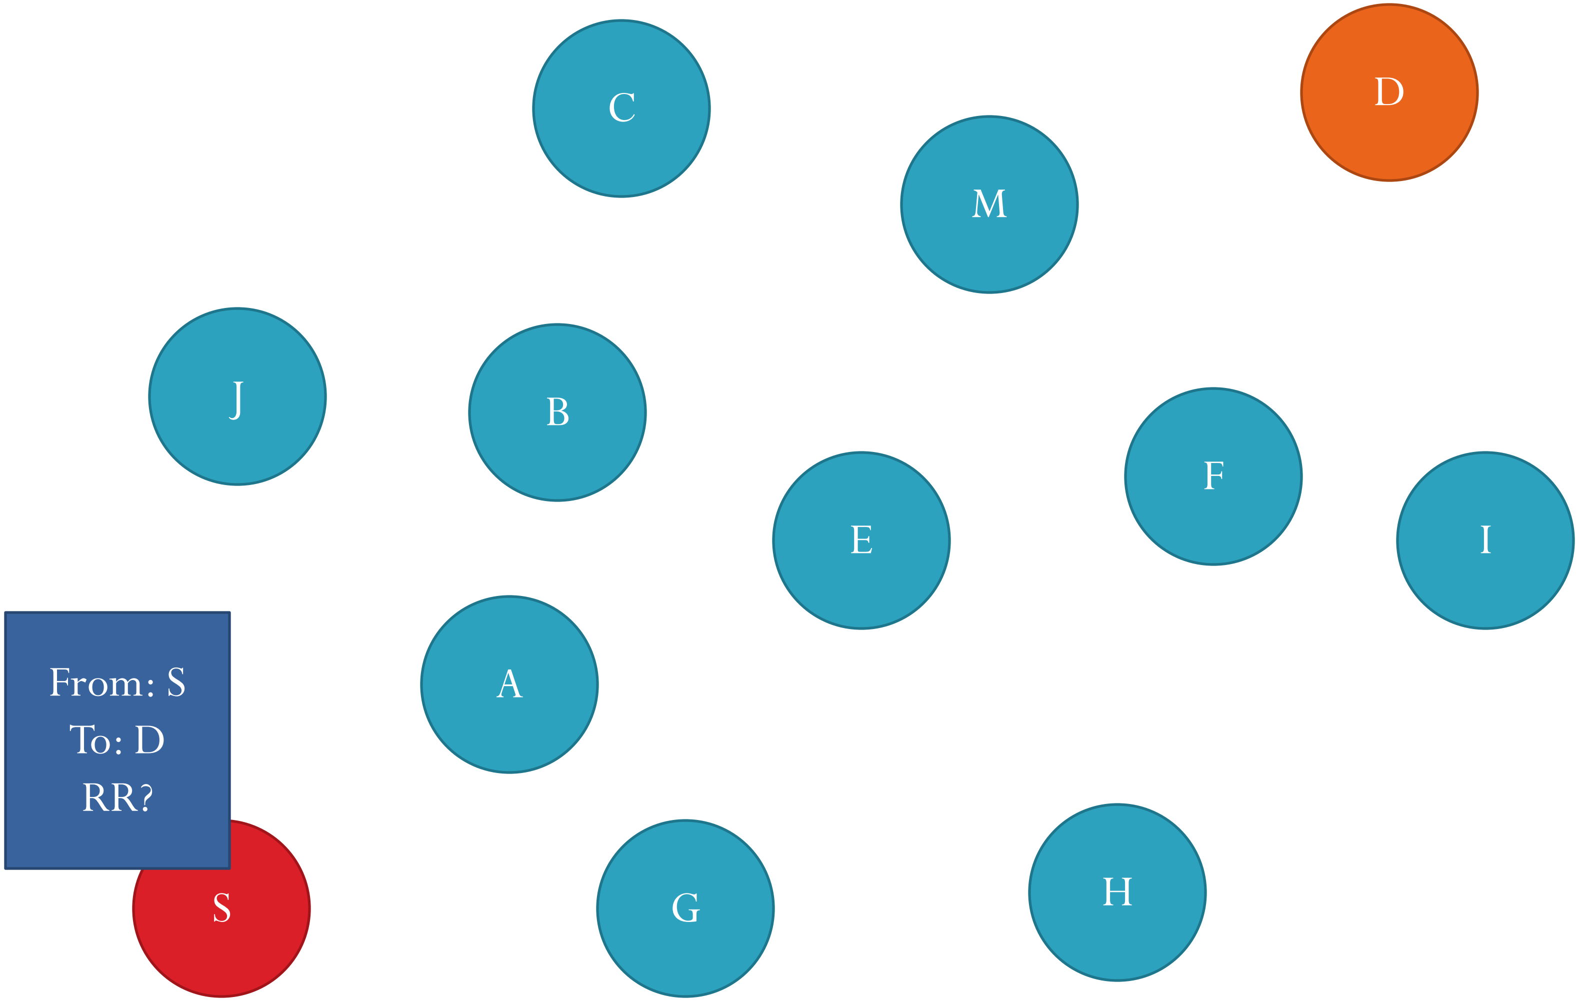

Record Route

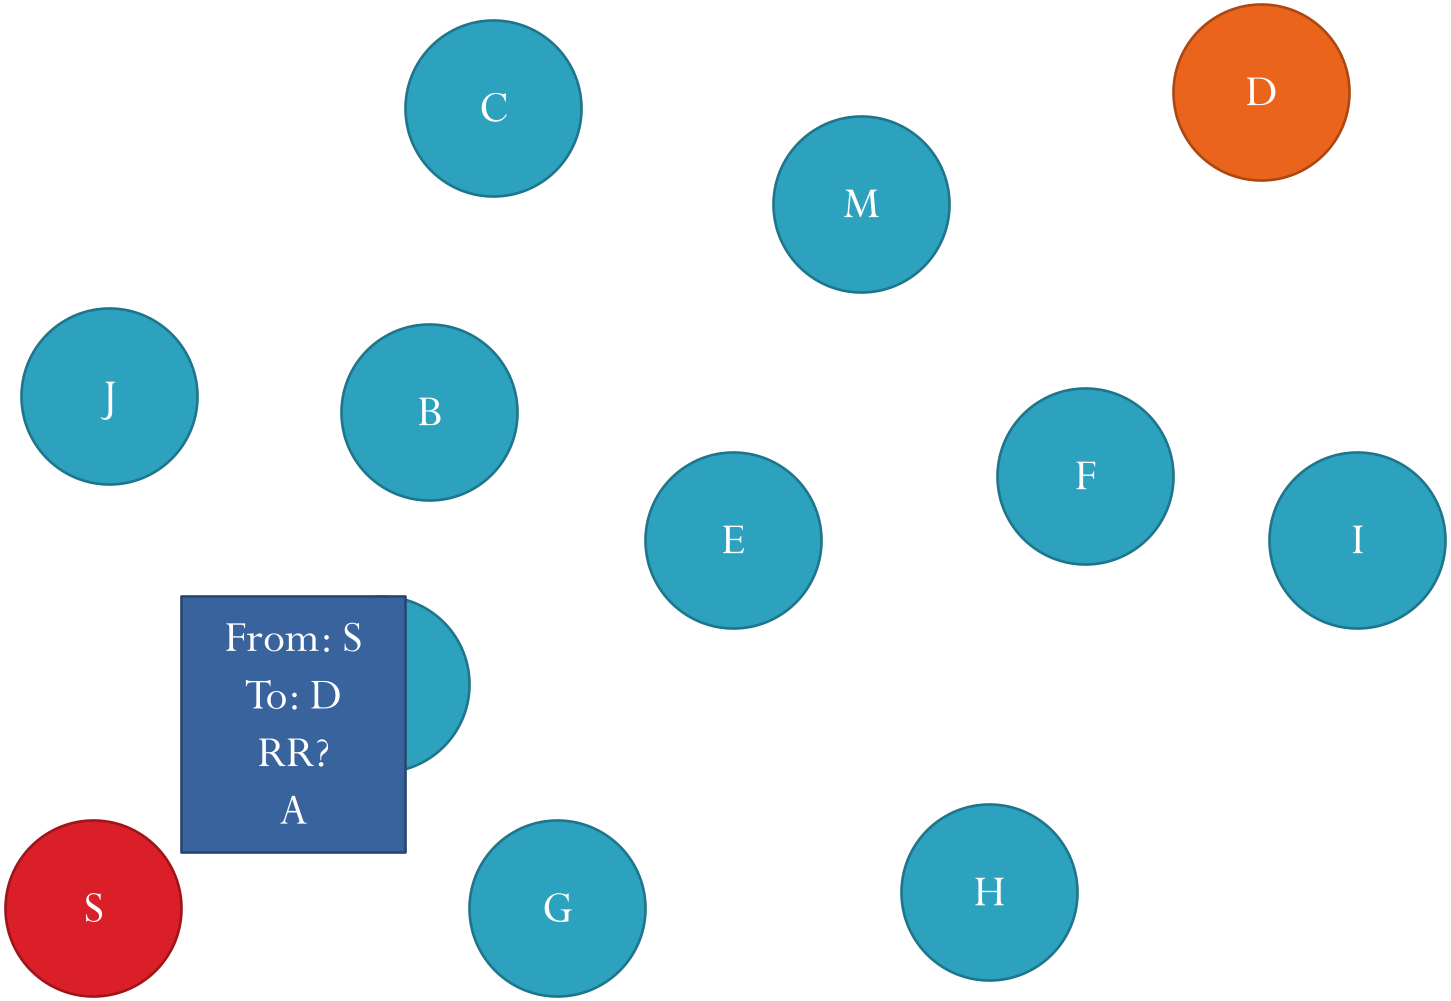

From: S

To: D

RR?

A

Record Route

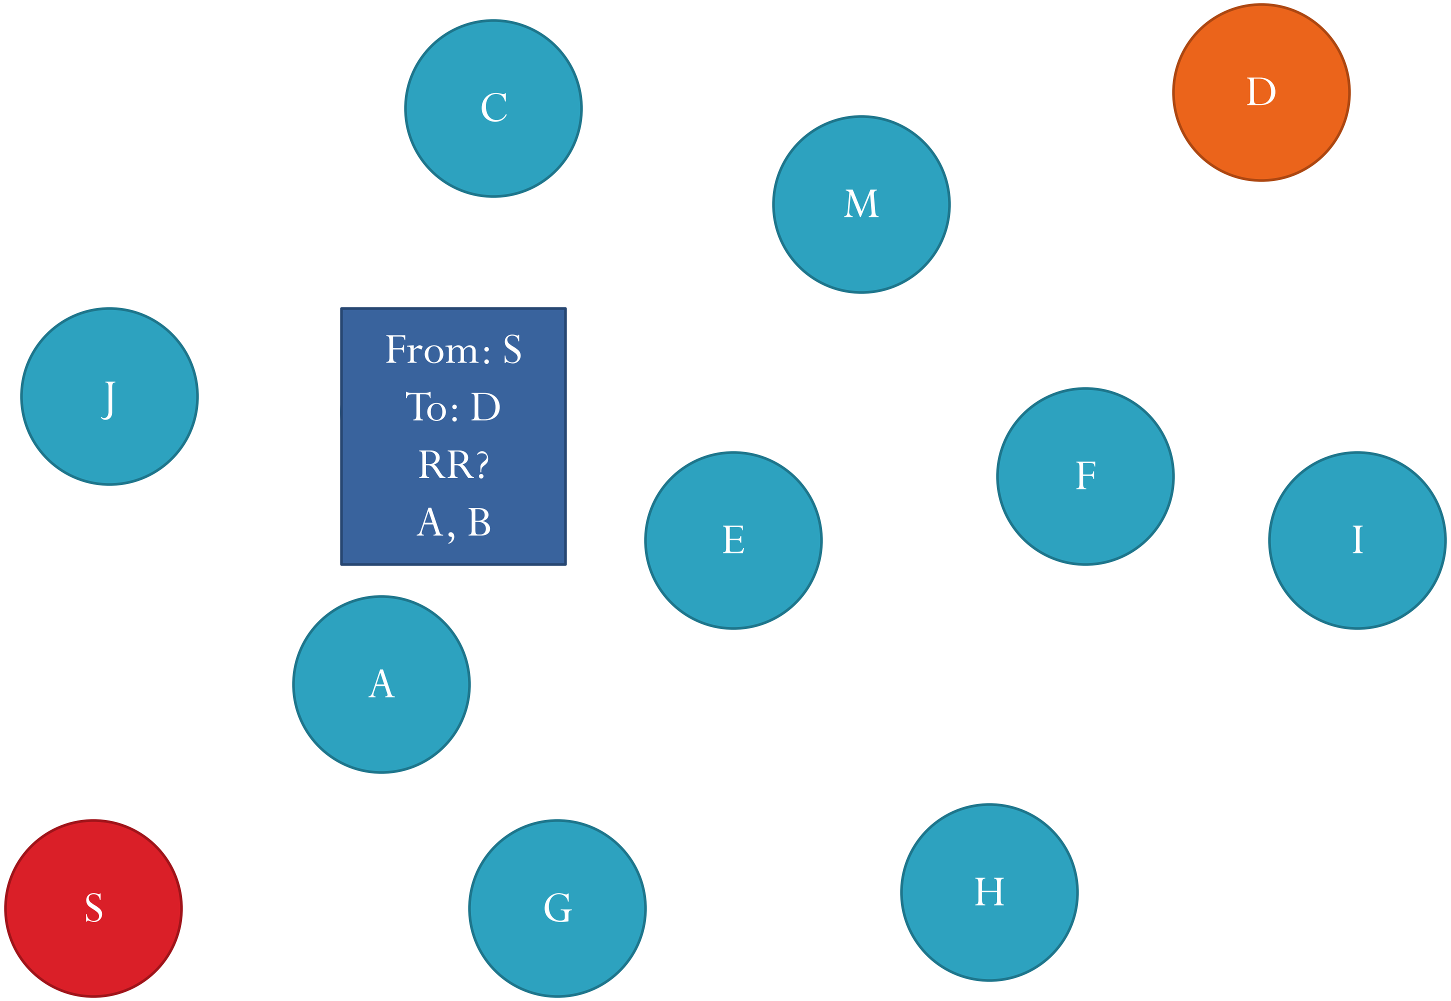

From: S

To: D

RR?

A, B

Record Route

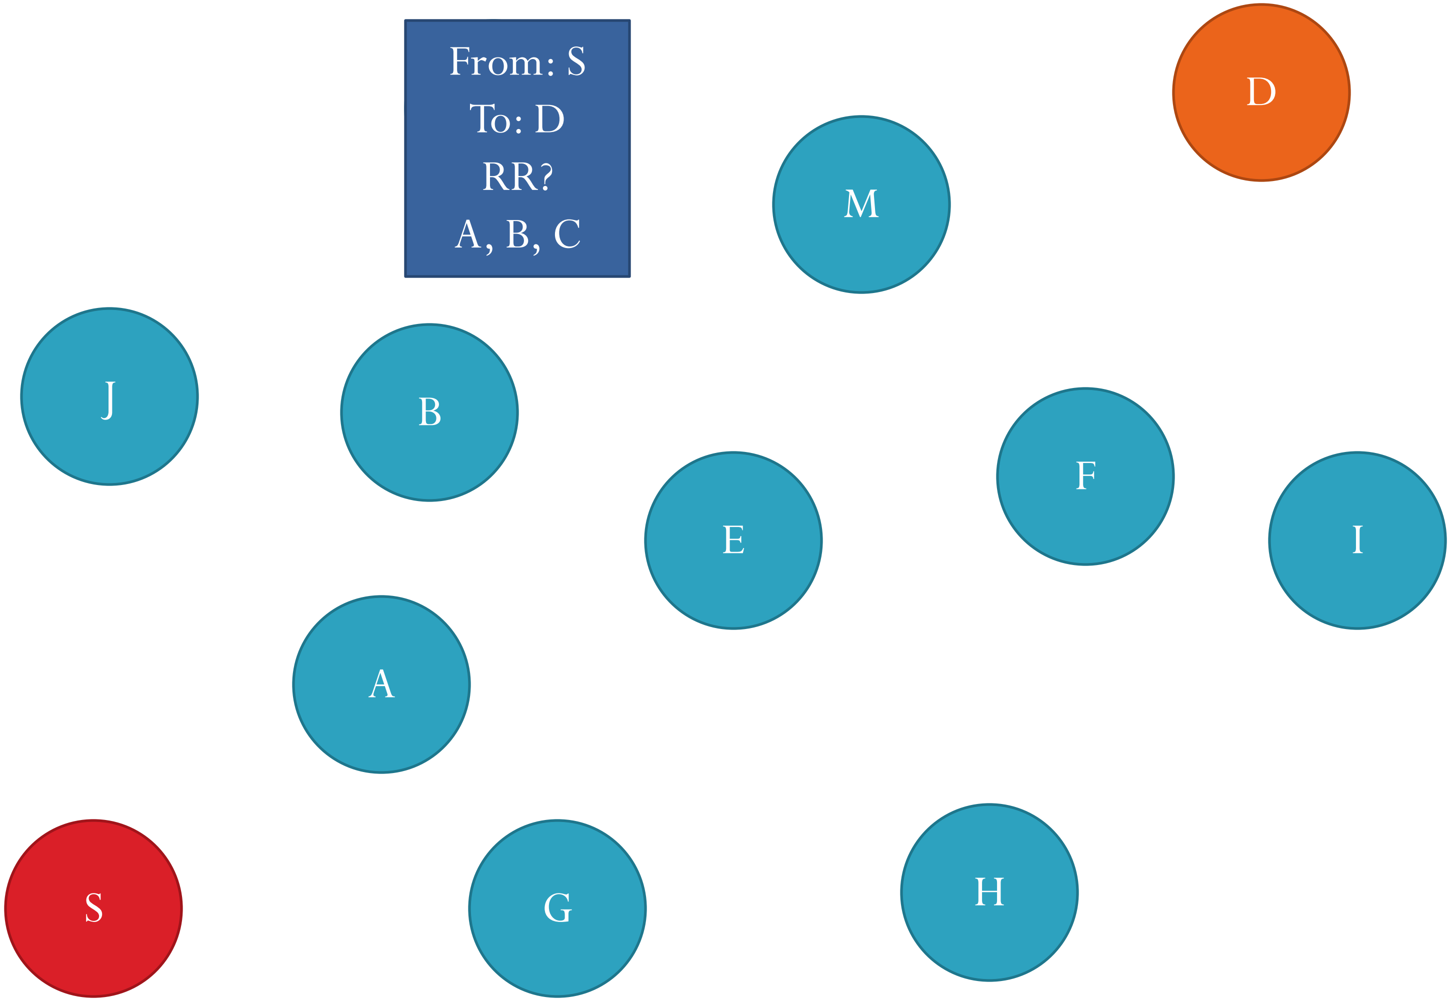

From: S

To: D

RR?

A, B, C

Record Route

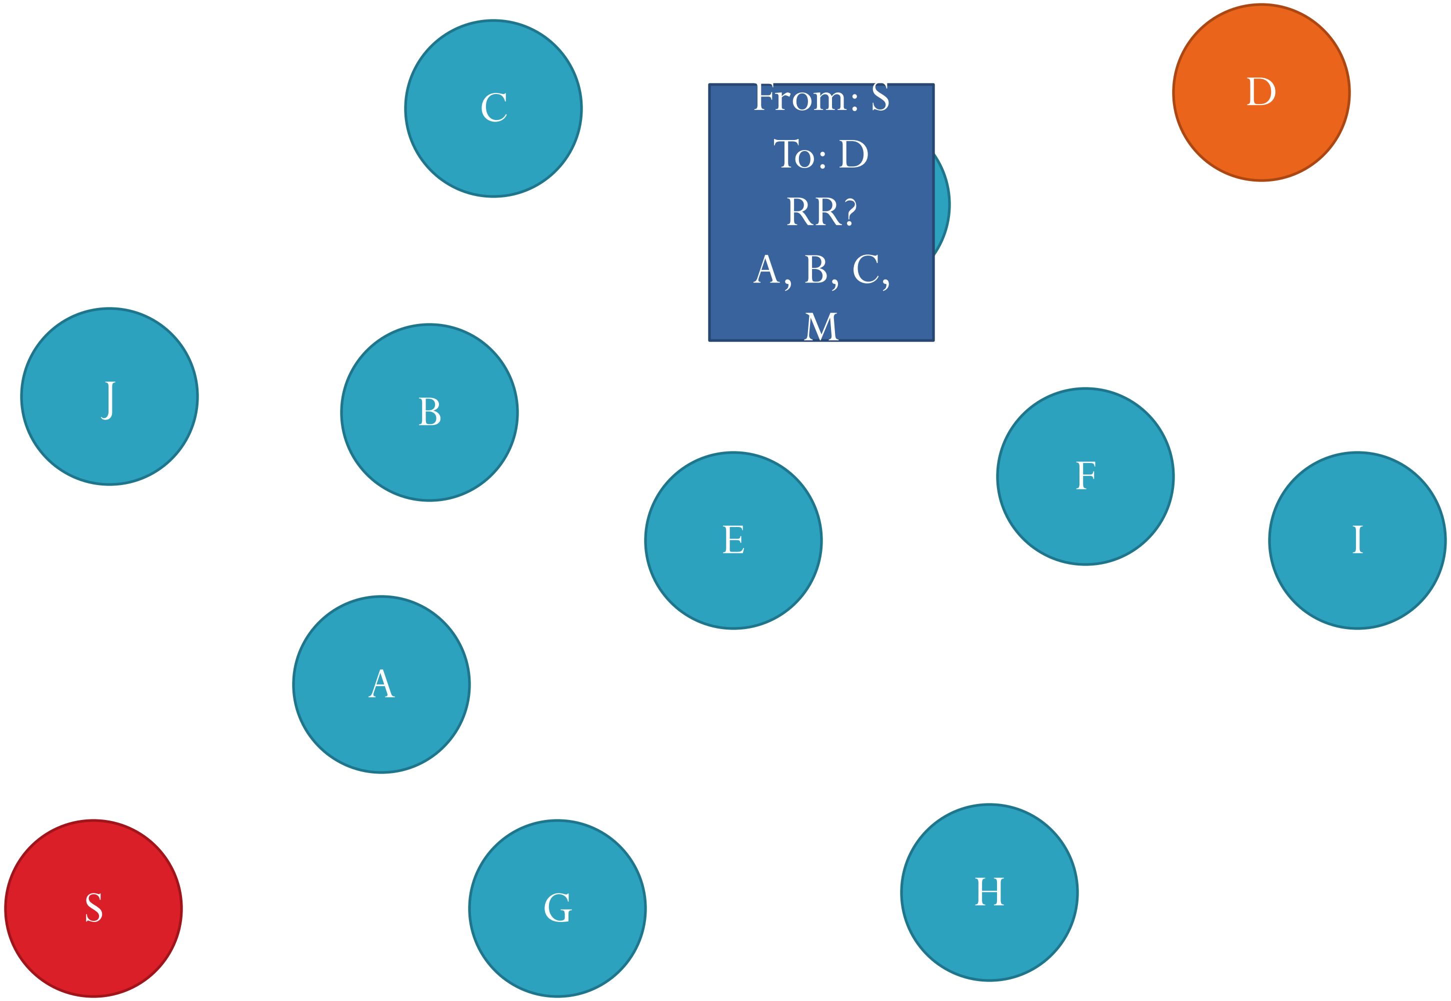

From: S
To: D
RR?
A, B, C,
M

Record Route

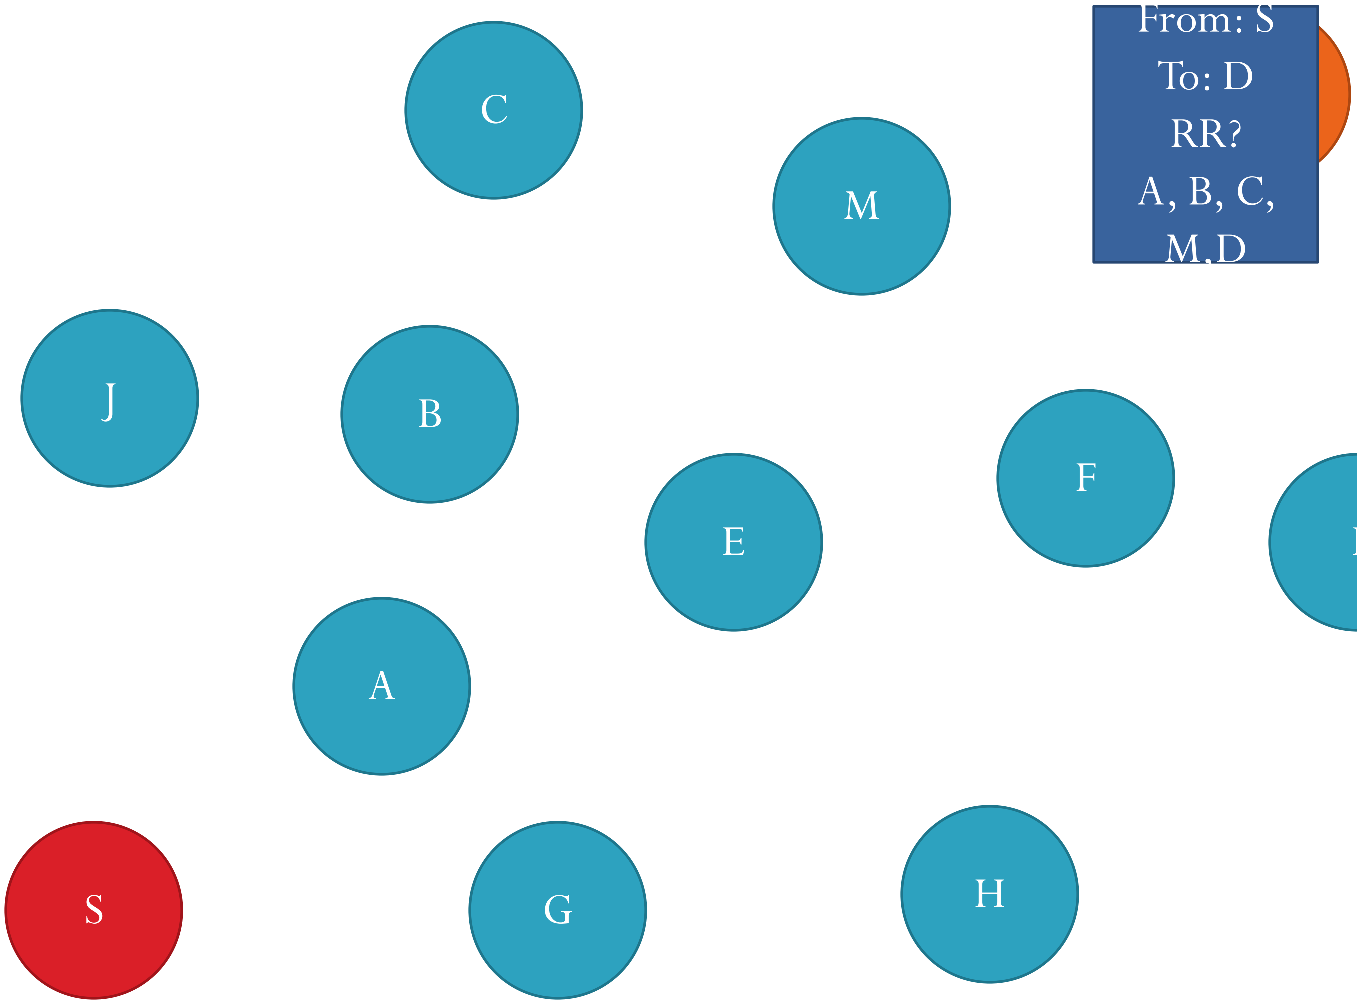

From: S
To: D
RR?
A, B, C,
M,D

Record Route

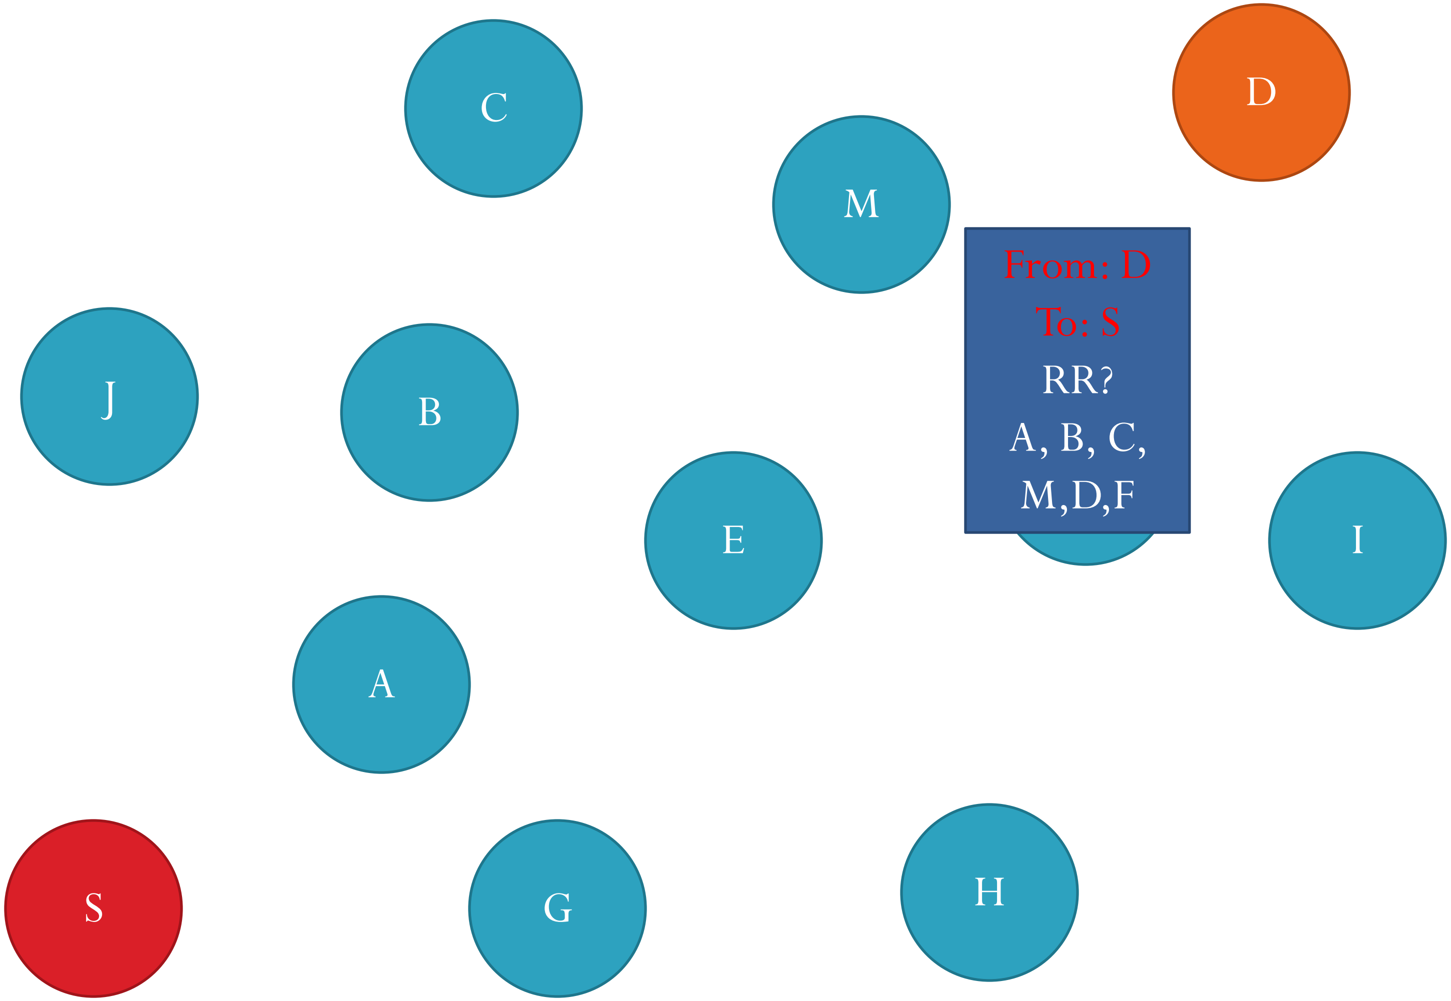

From: D

To: S

RR?

A, B, C,

M,D,F

Record Route

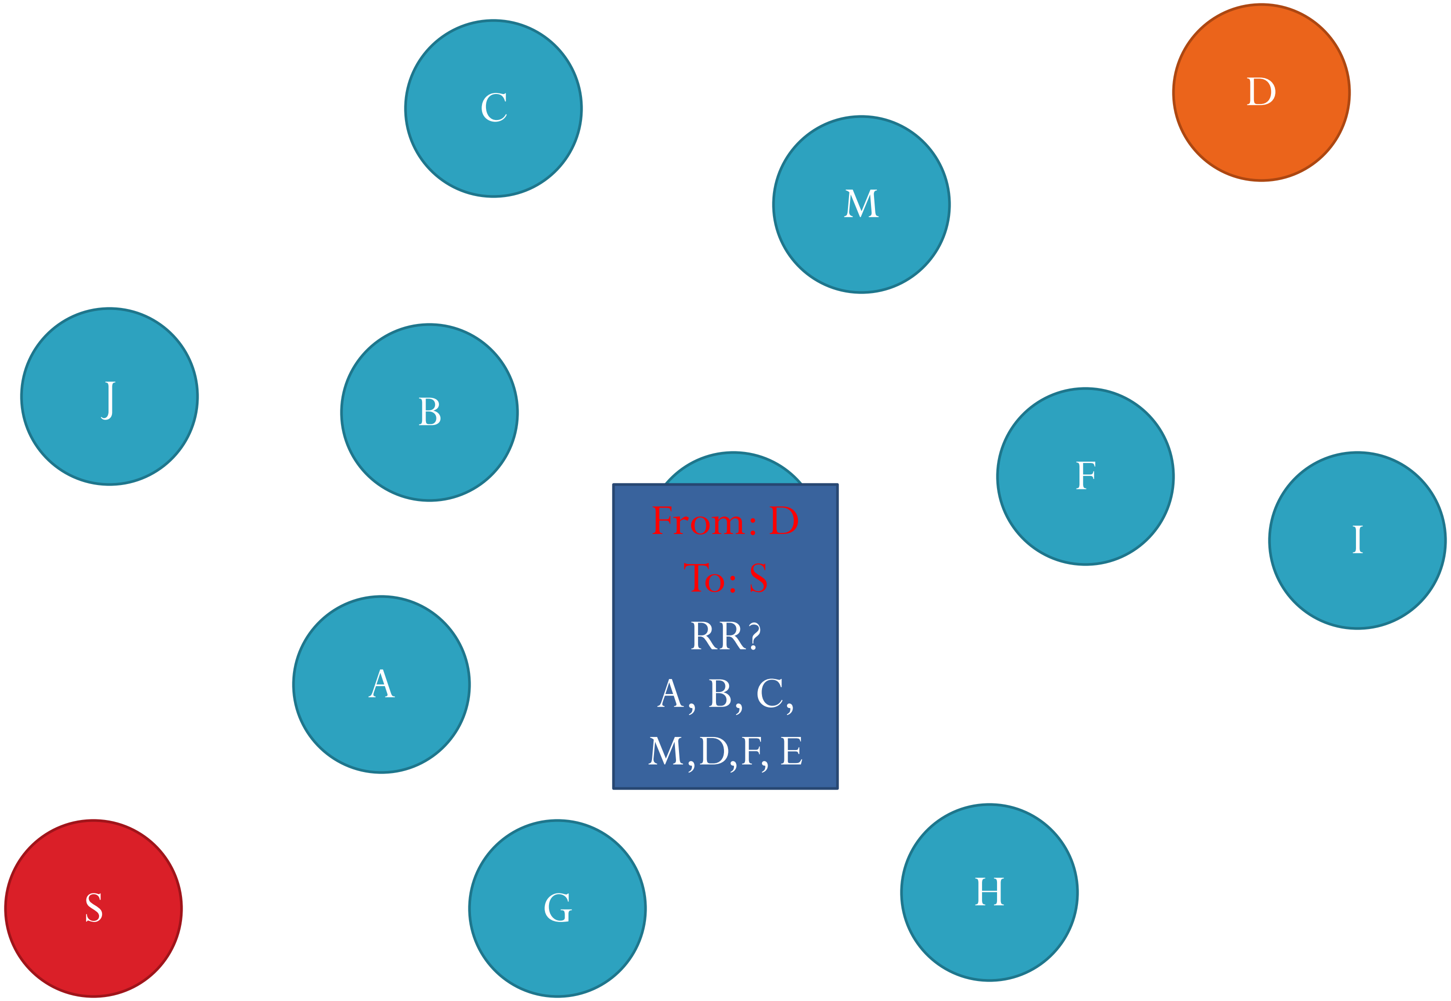

From: D

To: S

RR?

A, B, C,

M,D,F, E

Coordinate many many measurements

- We have several ways of getting a "peek" into the reverse path:
  - Record Route (previous slides)
  - We also issue lots of traceroutes
  - We even have a technique using timestamps
- We can use many computers (vantage points) located around the world to issue measurements together.
- Combining all of these resources – multiple vantage points, and multiple measurements, we can coordinate measurements from back at UW to piece together the reverse path in entirety!

Solving problems with RevTR

- *(Current practice)* Issue traceroute, check if indirect

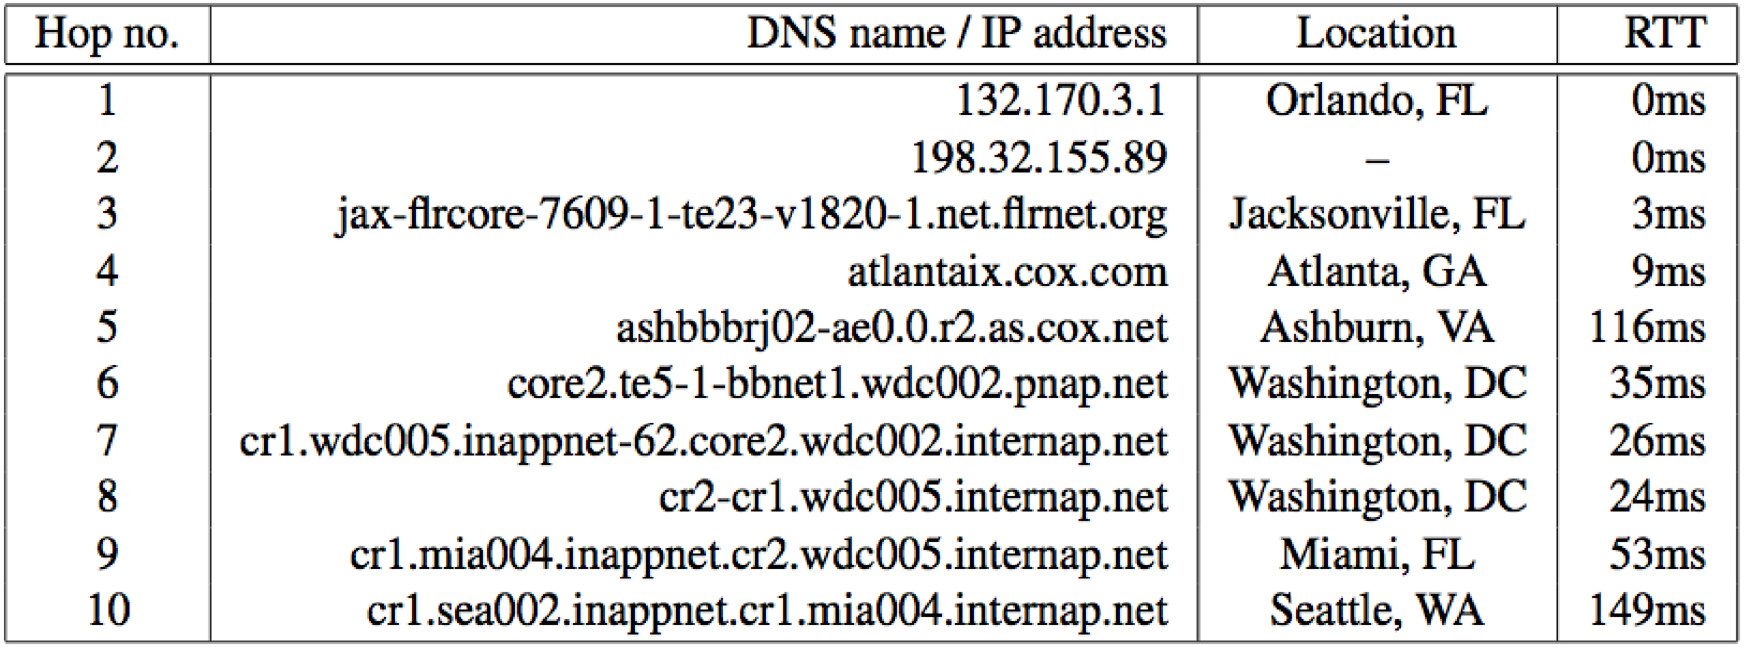

Hop no.	DNS name / IP address	Location	RTT
1	132.170.3.1	Orlando, FL	0ms
2	198.32.155.89	–	0ms
3	jax-flrcore-7609-1-te23-v1820-1.net.flrnet.org	Jacksonville, FL	3ms
4	atlantaix.cox.com	Atlanta, GA	9ms
5	ashbbbrj02-ae0.0.r2.as.cox.net	Ashburn, VA	116ms
6	core2.te5-1-bbnet1.wdc002.pnap.net	Washington, DC	35ms
7	cr1.wdc005.inappnet-62.core2.wdc002.internap.net	Washington, DC	26ms
8	cr2-cr1.wdc005.internap.net	Washington, DC	24ms
9	cr1.mia004.inappnet.cr2.wdc005.internap.net	Miami, FL	53ms
10	cr1.sea002.inappnet.cr1.mia004.internap.net	Seattle, WA	149ms

- Indirectness: FL→DC→FL, but does not explain huge latency jump from 9 to 10

Solving problems with RevTR

- *(With our tool)* Issue reverse traceroute, check rev path

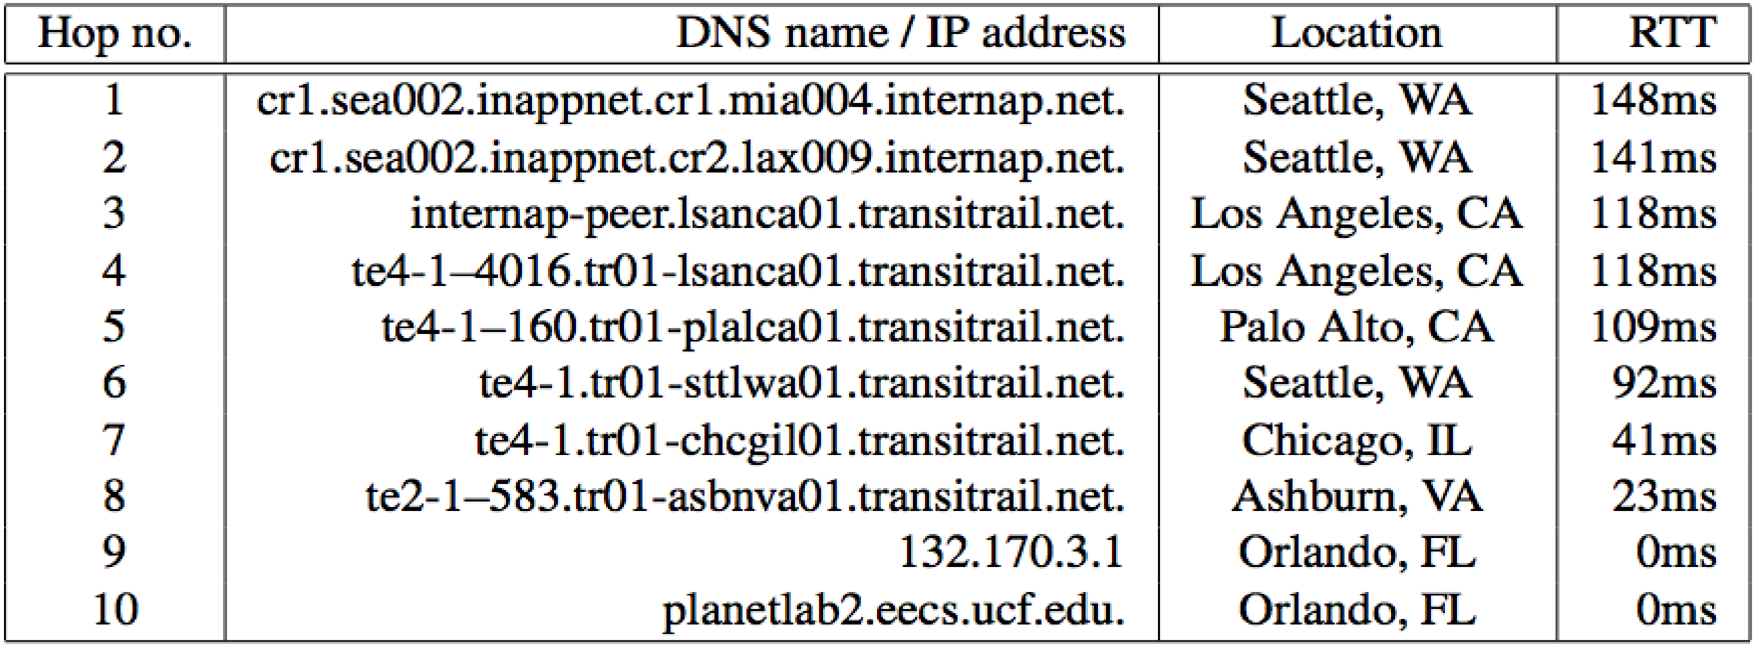

Hop no.	DNS name / IP address	Location	RTT
1	cr1.sea002.inappnet.cr1.mia004.internap.net.	Seattle, WA	148ms
2	cr1.sea002.inappnet.cr2.lax009.internap.net.	Seattle, WA	141ms
3	internap-peer.lsanca01.transitrail.net.	Los Angeles, CA	118ms
4	te4-1–4016.tr01-lsanca01.transitrail.net.	Los Angeles, CA	118ms
5	te4-1–160.tr01-plalca01.transitrail.net.	Palo Alto, CA	109ms
6	te4-1.tr01-sttlwa01.transitrail.net.	Seattle, WA	92ms
7	te4-1.tr01-chcgil01.transitrail.net.	Chicago, IL	41ms
8	te2-1–583.tr01-asbnva01.transitrail.net.	Ashburn, VA	23ms
9	132.170.3.1	Orlando, FL	0ms
10	planetlab2.eecs.ucf.edu.	Orlando, FL	0ms

- Indirectness: WA→LA→WA

  Bad rev path causes inflated round-trip delay

Agenda

- Understanding the Problem: Why Measure the Internet?
- Existing measurement tools
- A new basic tool: IP timestamp
- A large measurement tool: Reverse Traceroute
- Final Thoughts

On Internet Measurement

- The Internet does not have to remain a "black box" – we can make measurements to better understand important properties of the network we use every day.
- There are a suite of tools we can use to make useful measurements. Many of them are quite simple – anyone can issue a traceroute from their home computer. Others are large, complex pieces of research.

On Internet Measurement

- Our measurement research has helped to build better tools.
  - With IP Timestamp, we found novel uses for a previously overlooked basic measurement tool.
  - With Reverse Traceroute, we combined many measurements and coordinated distributed machines to solve the long-troublesome problem of discovering reverse paths in the Internet.

On Doing Undergrad Research

- There are a lot of opportunities out there to do amazing Computer Science research – ask your professors, apply to REU's, look for research positions at a local lab.
- You will be surprised at how fast you can learn about any topic. When thinking about what to work on, don't worry too much about having tons of experience in that area.
- Do ask yourself what makes you excited and what kinds of things you think are cool!

On Doing Undergrad Research

Some things other undergrads have worked on:

- Colin: building a system to guess, given an IP address, which vantage points are closest to it
- Peter: building the Reverse Traceroute website
- Ashoat: administering the Reverse Traceroute system
- Steve: cutting down on the number of Record Route measurements we have to make to find one that works
- Mary: predicting whether or not addresses will respond to timestamp measurements

I leave you with a pretty Internet Map

IPv4 & IPv6
INTERNET TOPOLOGY MAP
JANUARY 2009

AS-level INTERNET GRAPH

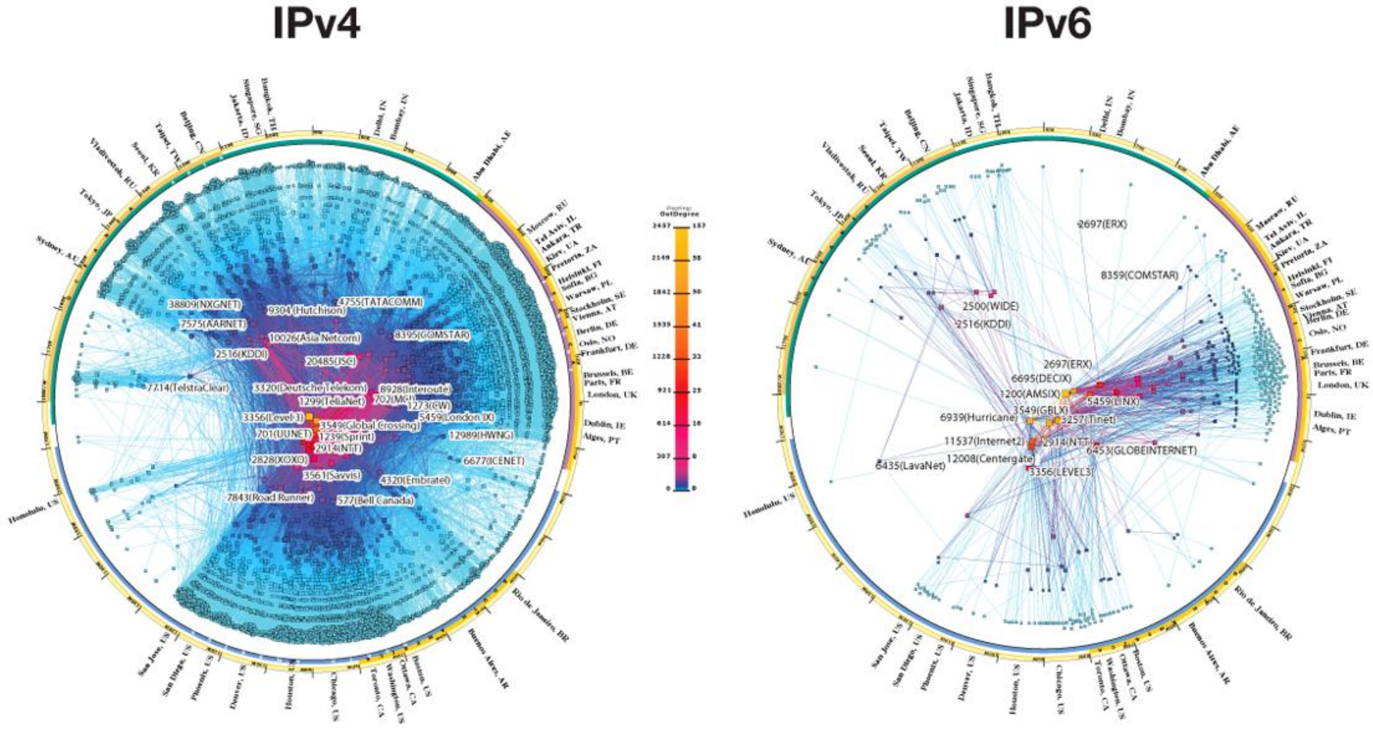

[This is from the awesome folks at CAIDA]