Applications of the IP Timestamp Option to Internet Measurement

Honors Thesis Presentation

Justine Sherry

March 2010

The Internet is a big black box

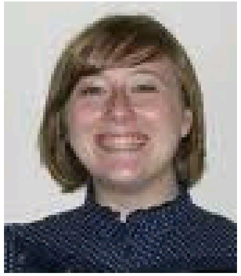

The Internet is a big black box

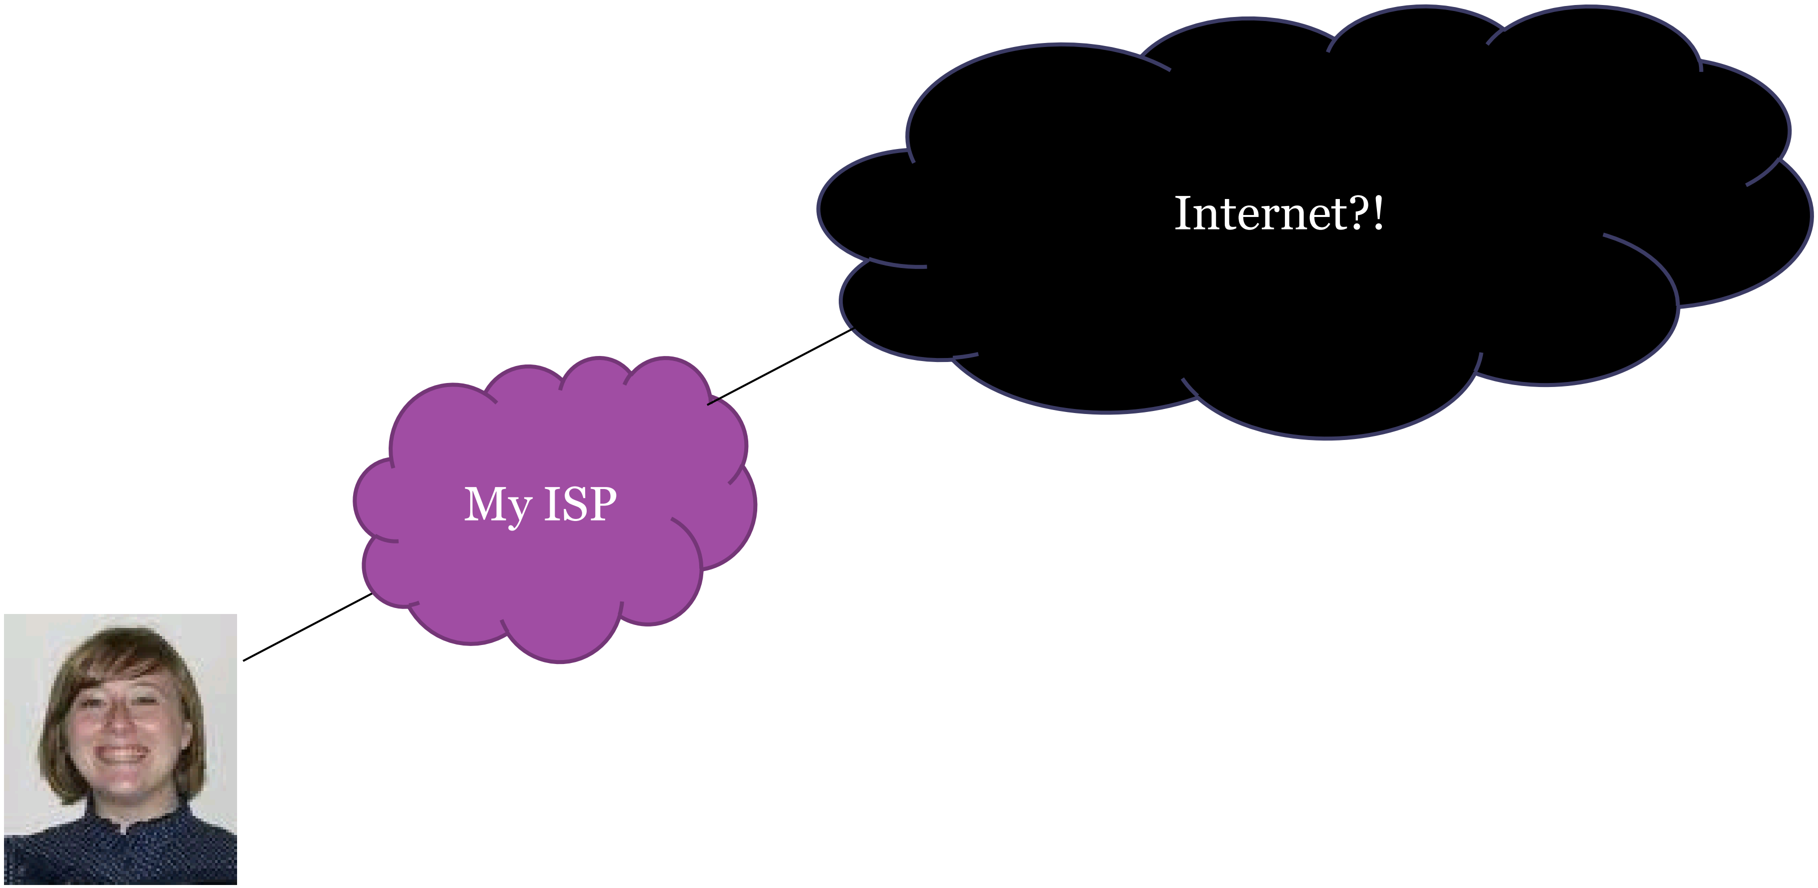

The Internet is a big black box

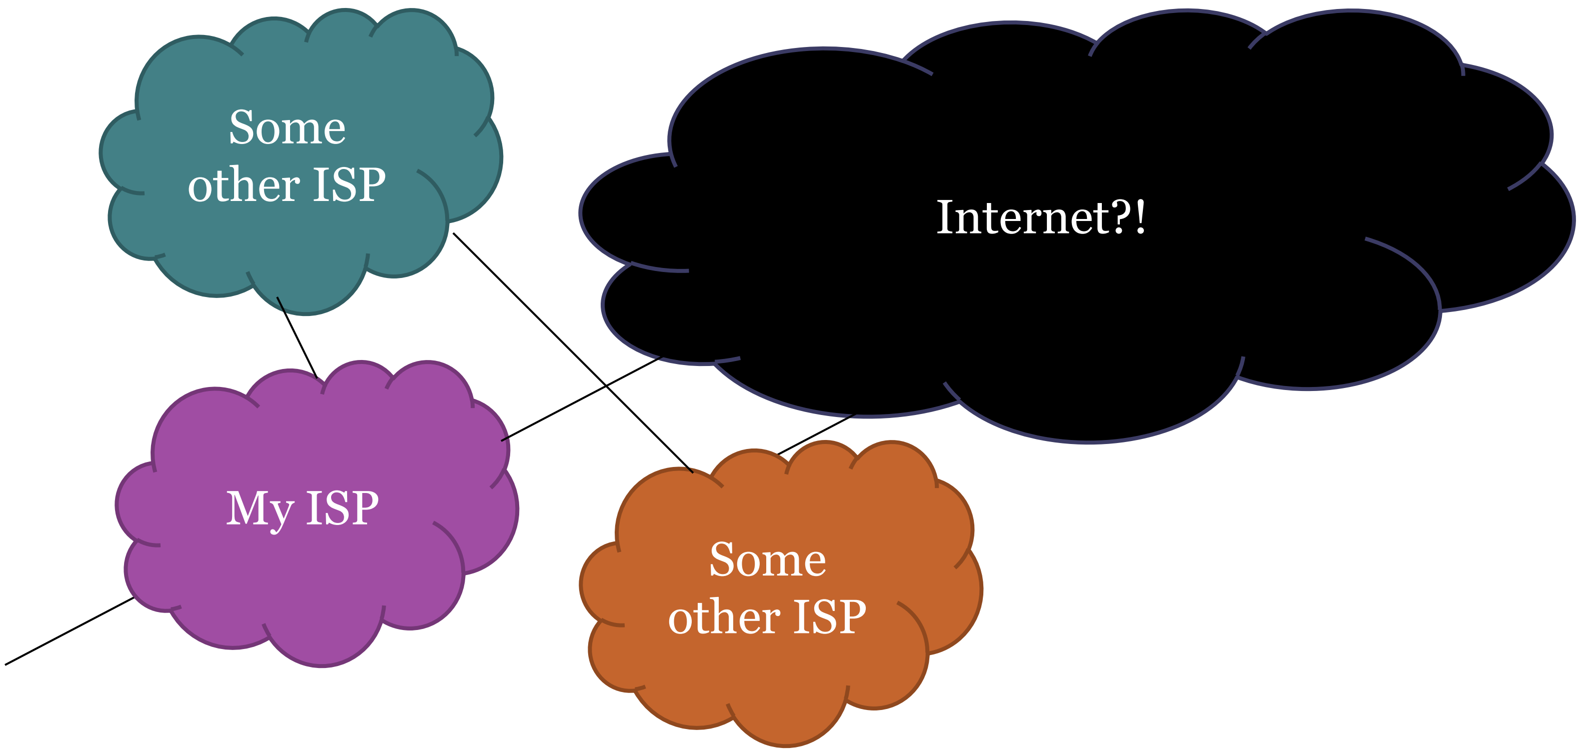

No one can see the big picture very well, outside of their own cloud.

Who cares?

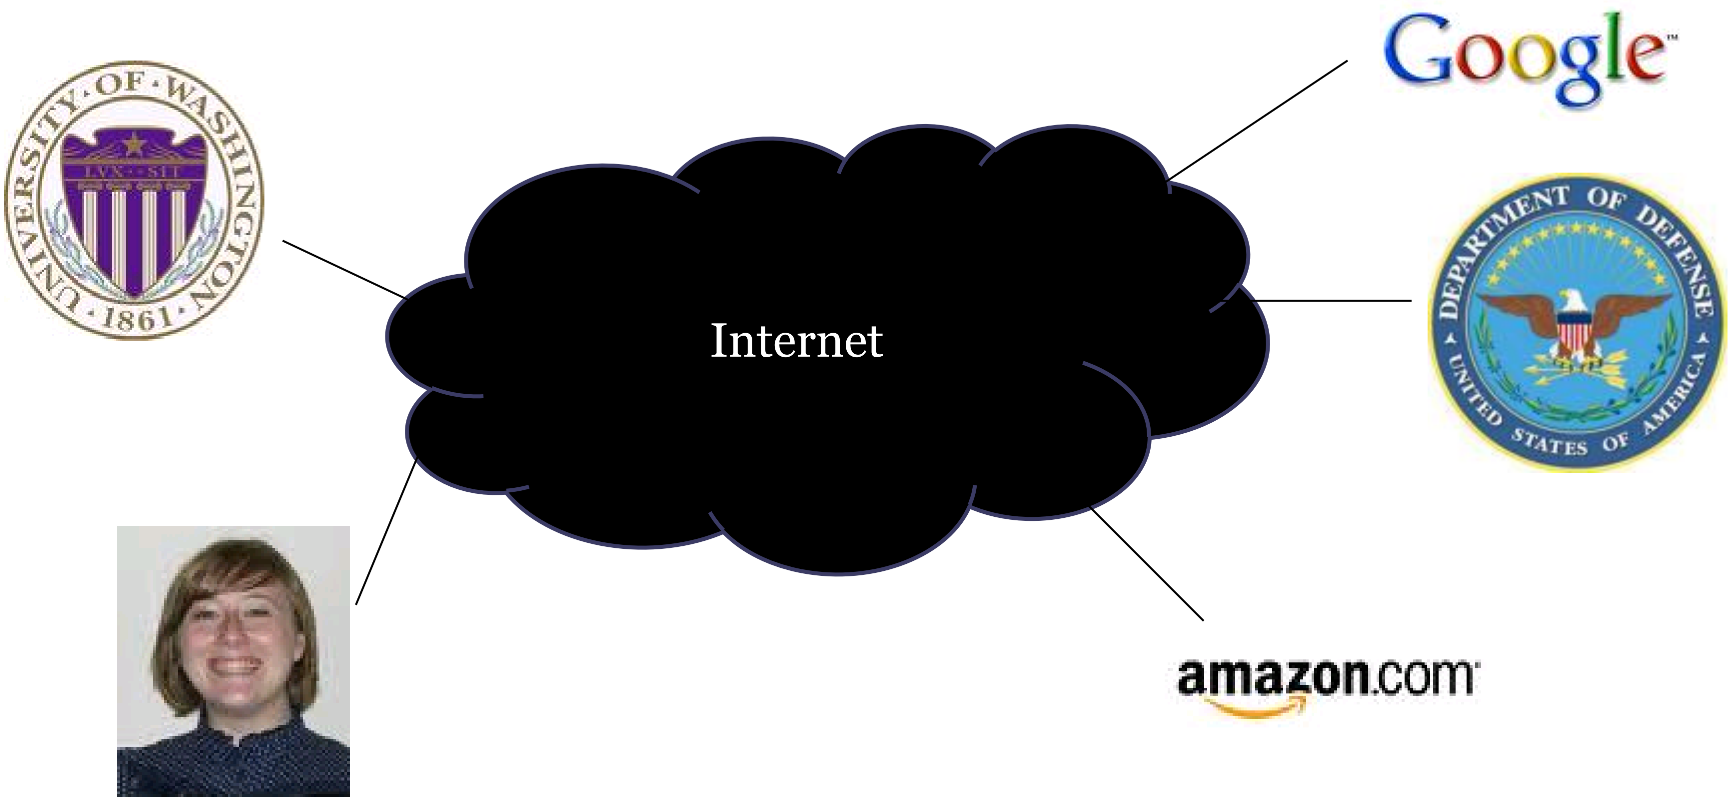
UNIVERSITY OF WASHINGTON
1861
Internet
Google
DEPARTMENT OF DEFENSE
UNITED STATES OF AMERICA
amazon.com

We have some tools to help us out

- Ping – classic! Is this machine connected to the network, responsive, can I reach it?
- Traceroute – show me the routers on the forward path between me and a distant target
- Record route – show me the routers within the first 9 hops the packet takes
- And others - most involve tricking routers into providing some information about themselves after receiving specially crafted packets

Tons of questions are still hard though!

Some Questions I Want to Answer

- Traceroute gives us the forward path a packet takes, how do I tell if a router is on the reverse path (which may be different)?
- Can I tell when two IP addresses belong to the same machine?
- How long does it take for a packet to travel from one router to the next?

Agenda

- Motivating Internet Measurements
- Understanding IP Timestamp
- Three Use Cases:
  - Reverse Path Visibility
  - IP Alias Resolution
  - One-Way Link Latency

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 it's IP address.
- IP Timestamp can help us answer some of these questions.

Timestamp Specification

- The sender lists up to four IP addresses in the packet header
- Each router along the way checks if it's own IP address is the first unstamped IP address
- If it does indeed own that IP address, then it provides a timestamp before forwarding

Example 1

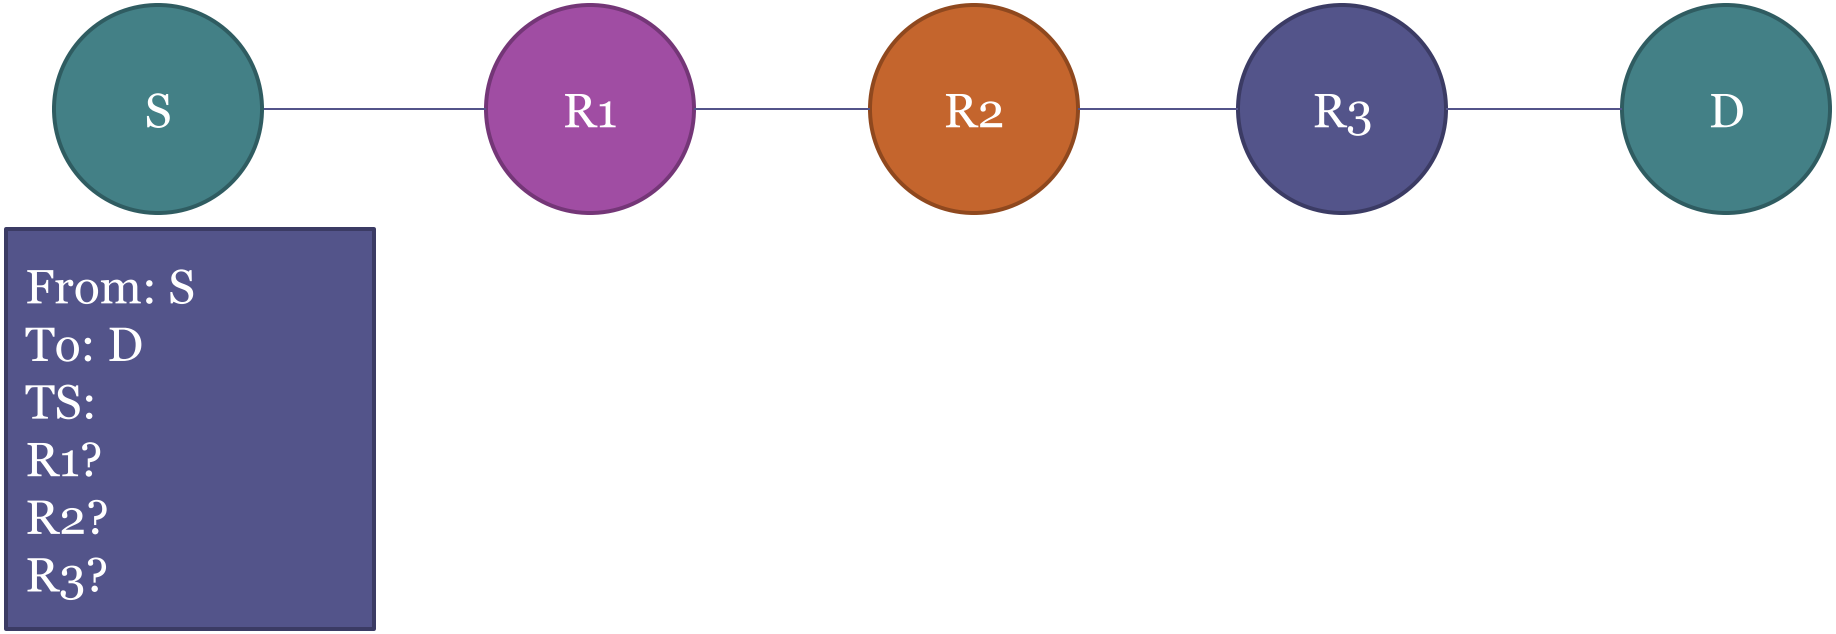

Example 1

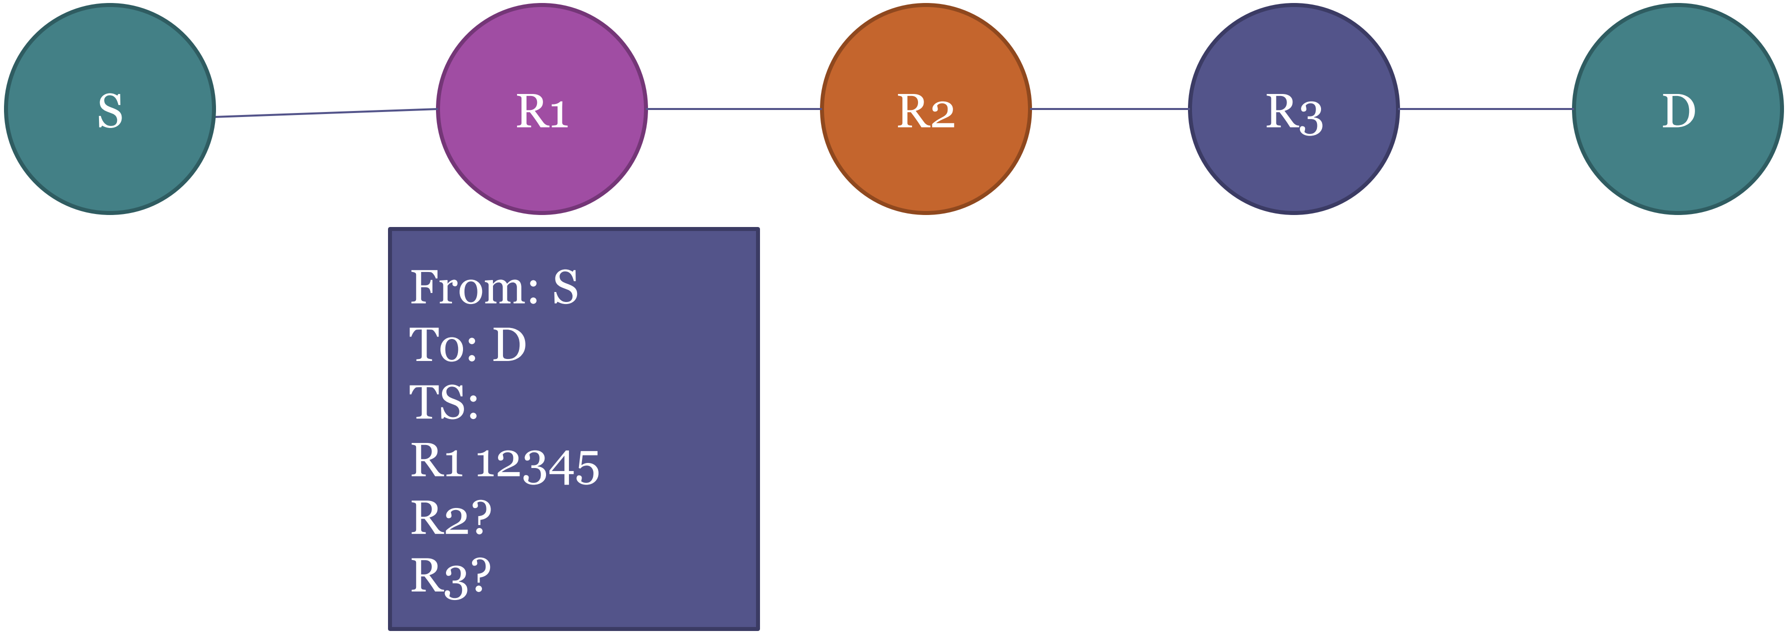

Example 1

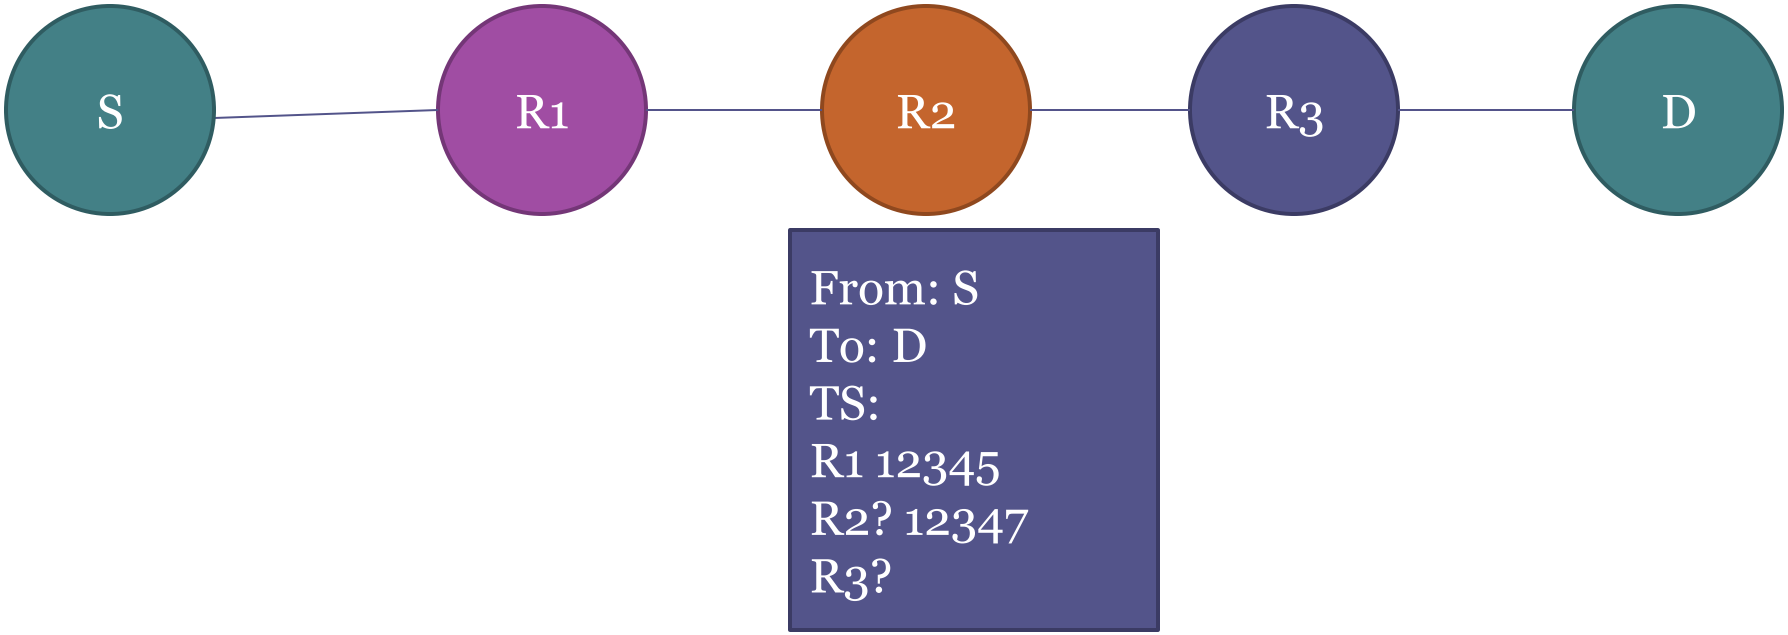

Example 1

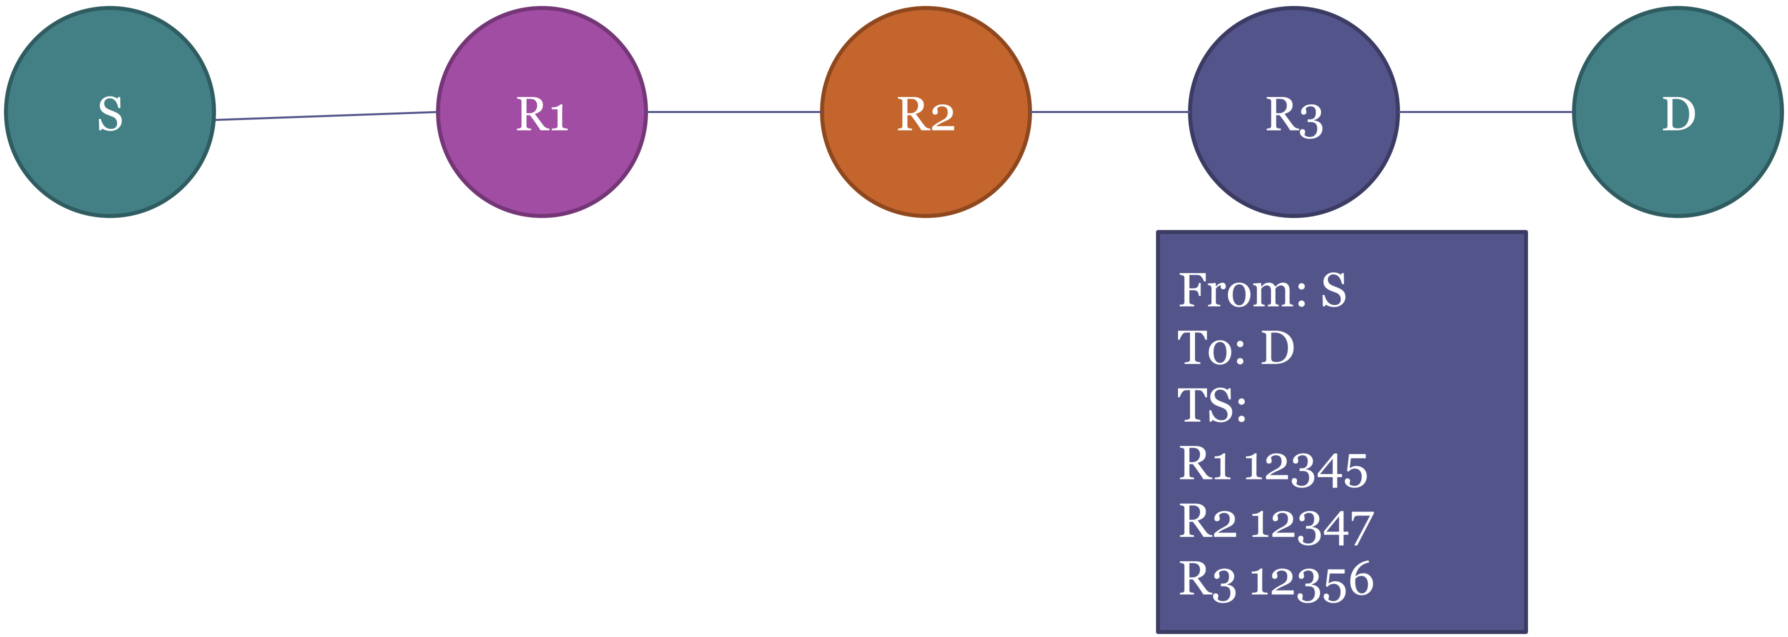

Example 1

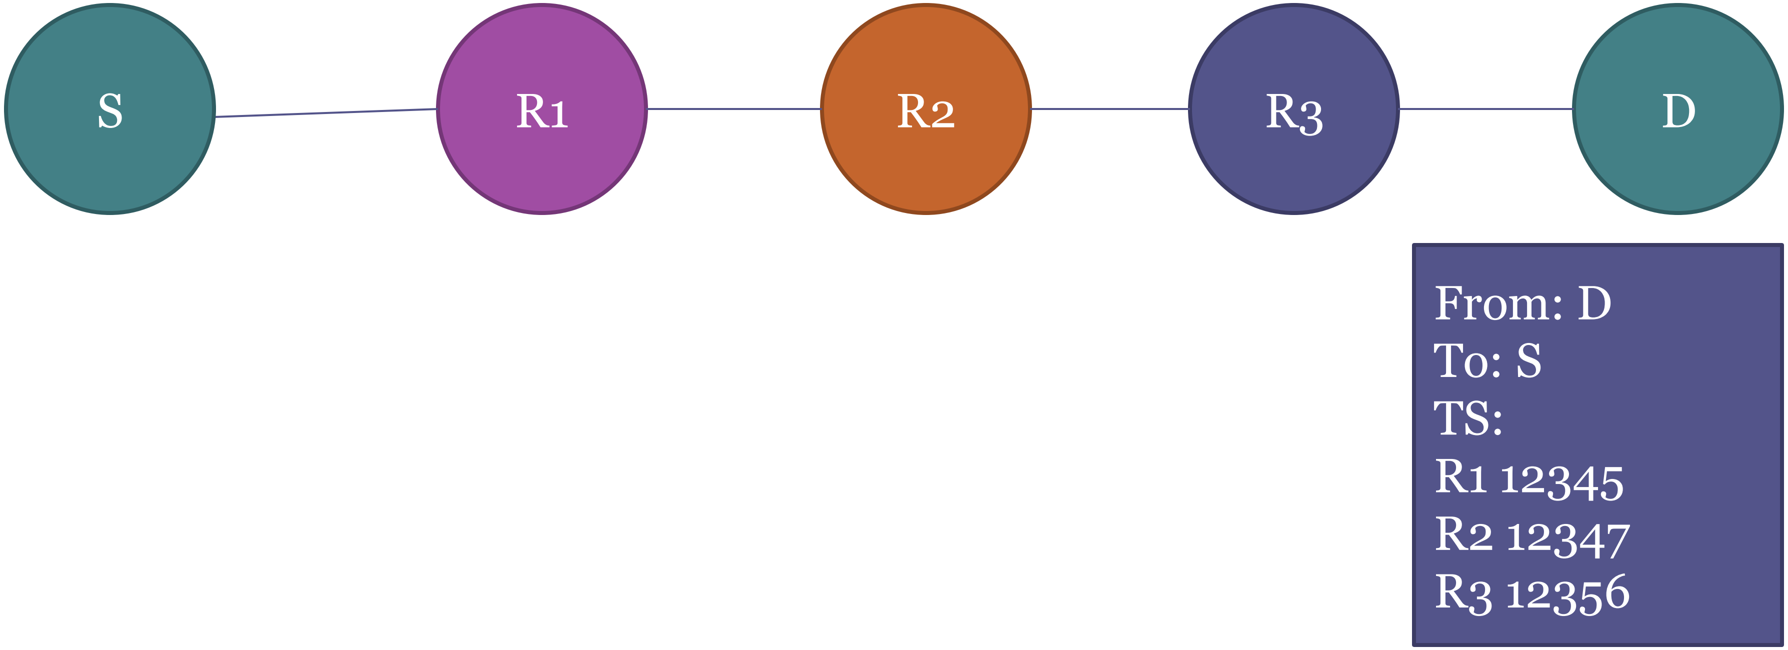

Example 2

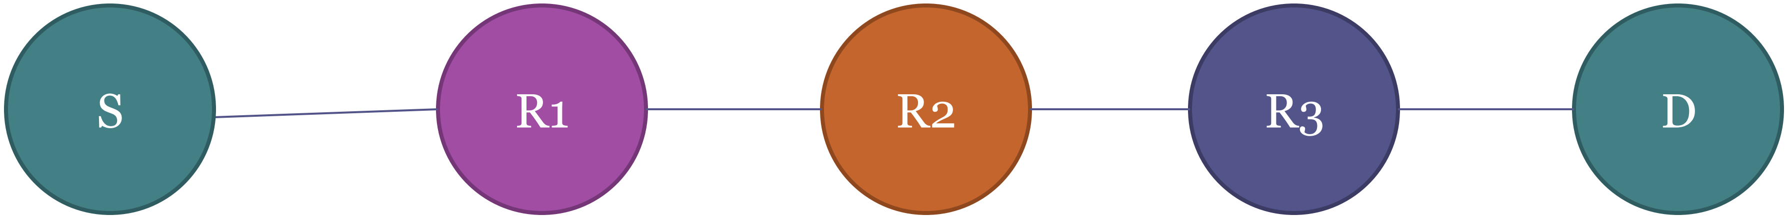

From: S
To: D
TS:
R1 ?
R3 ?
R2 ?

Example 2

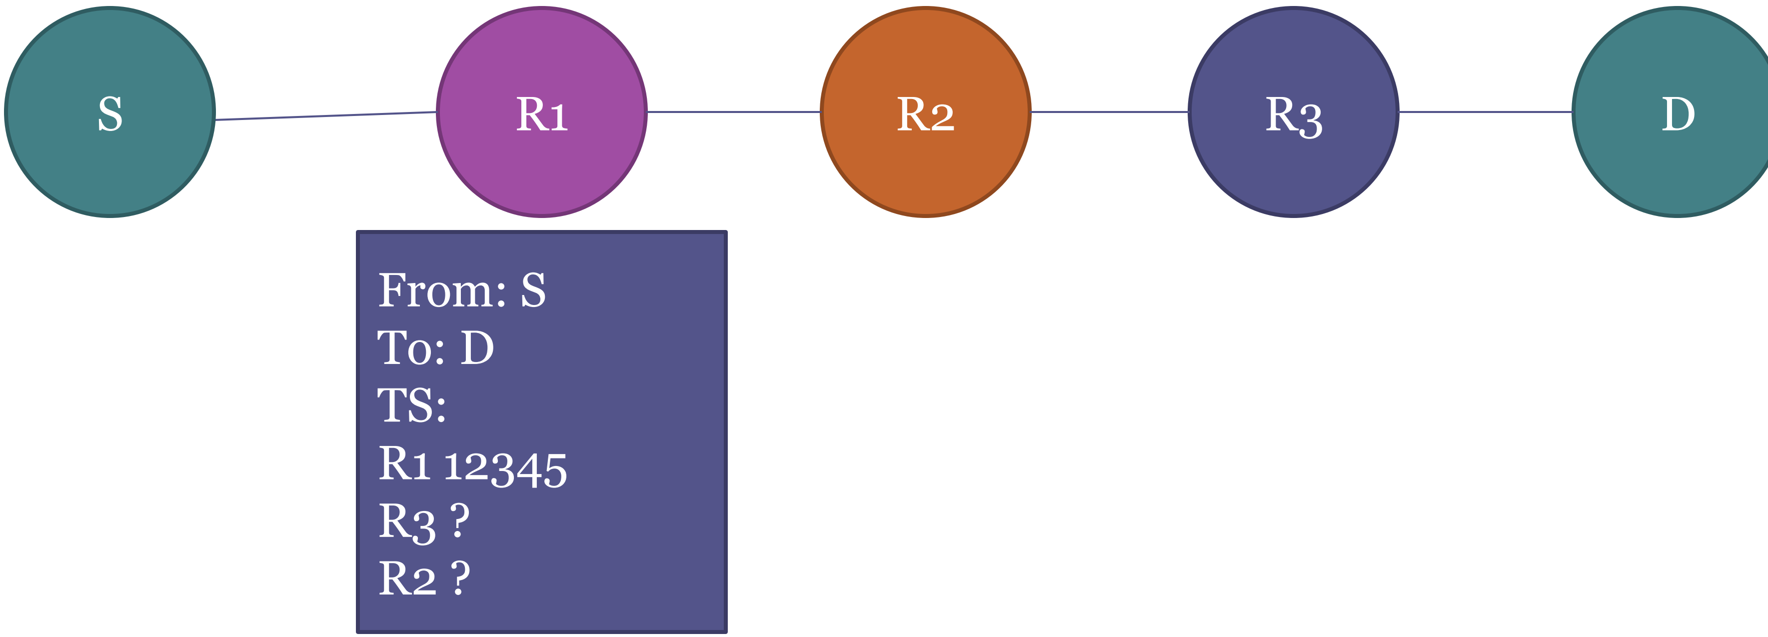

Example 2

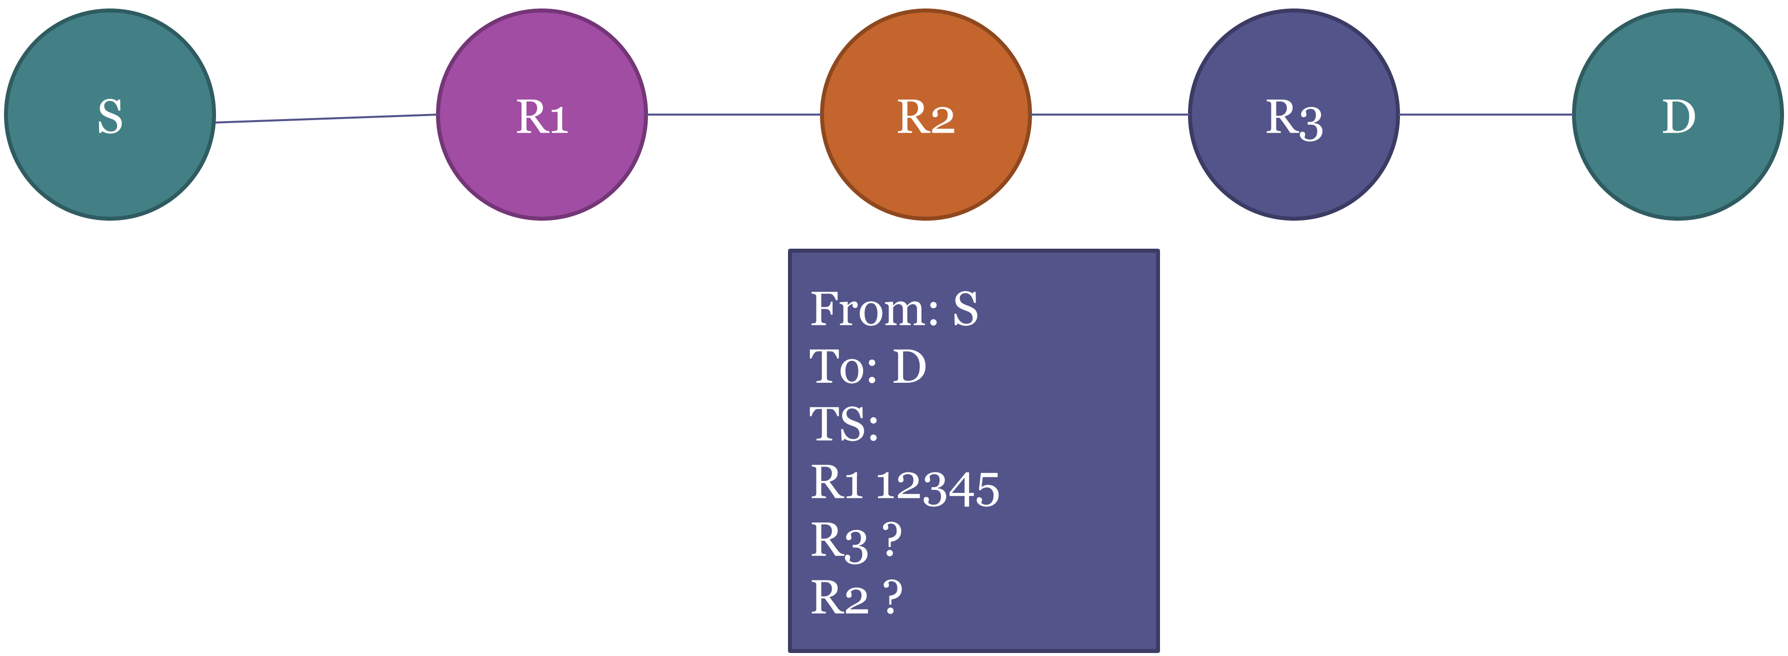

Example 2

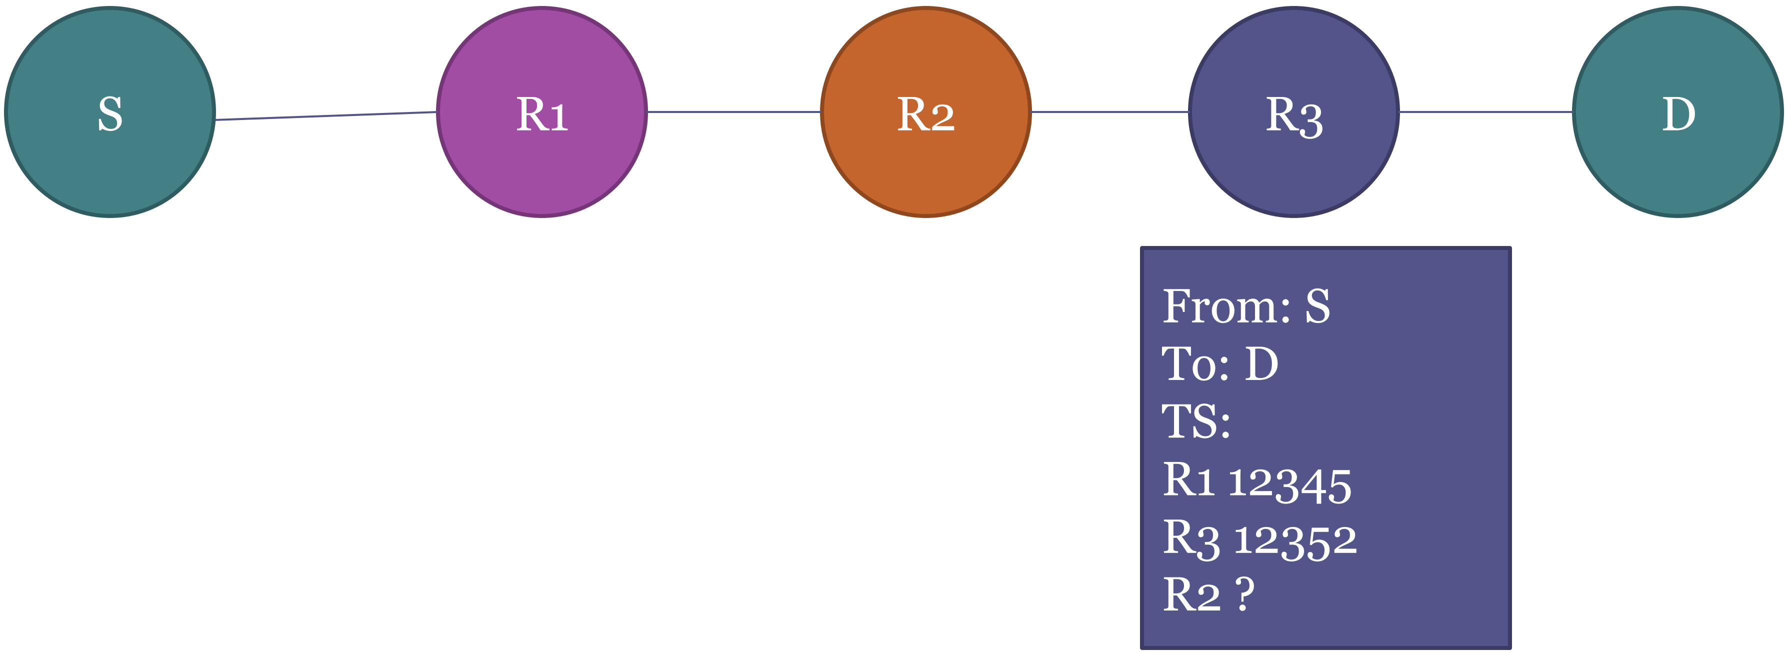

Example 2

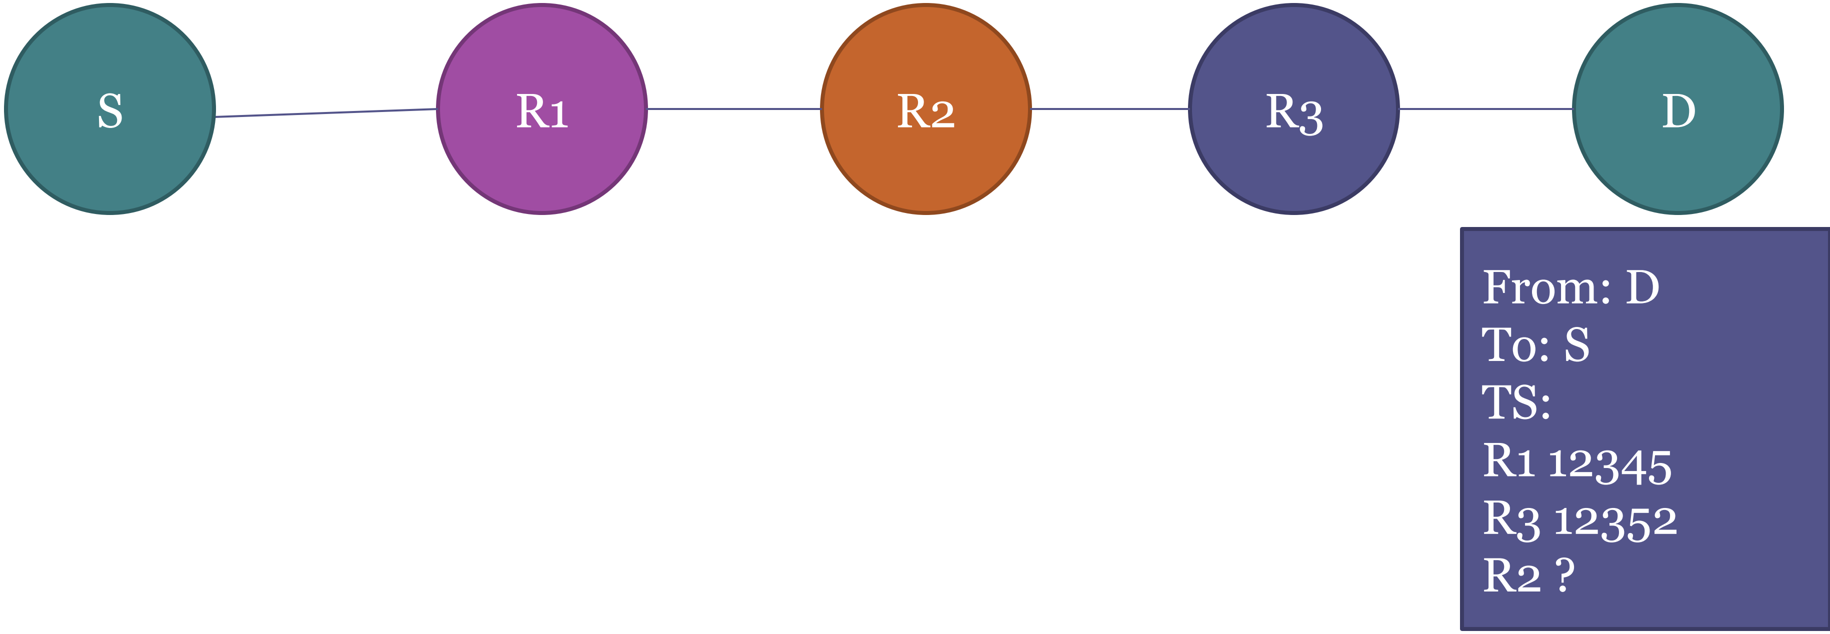

Let’s assume the path is symmetric for this example –
that the packet takes the same path back to S that it came from.

Example 2

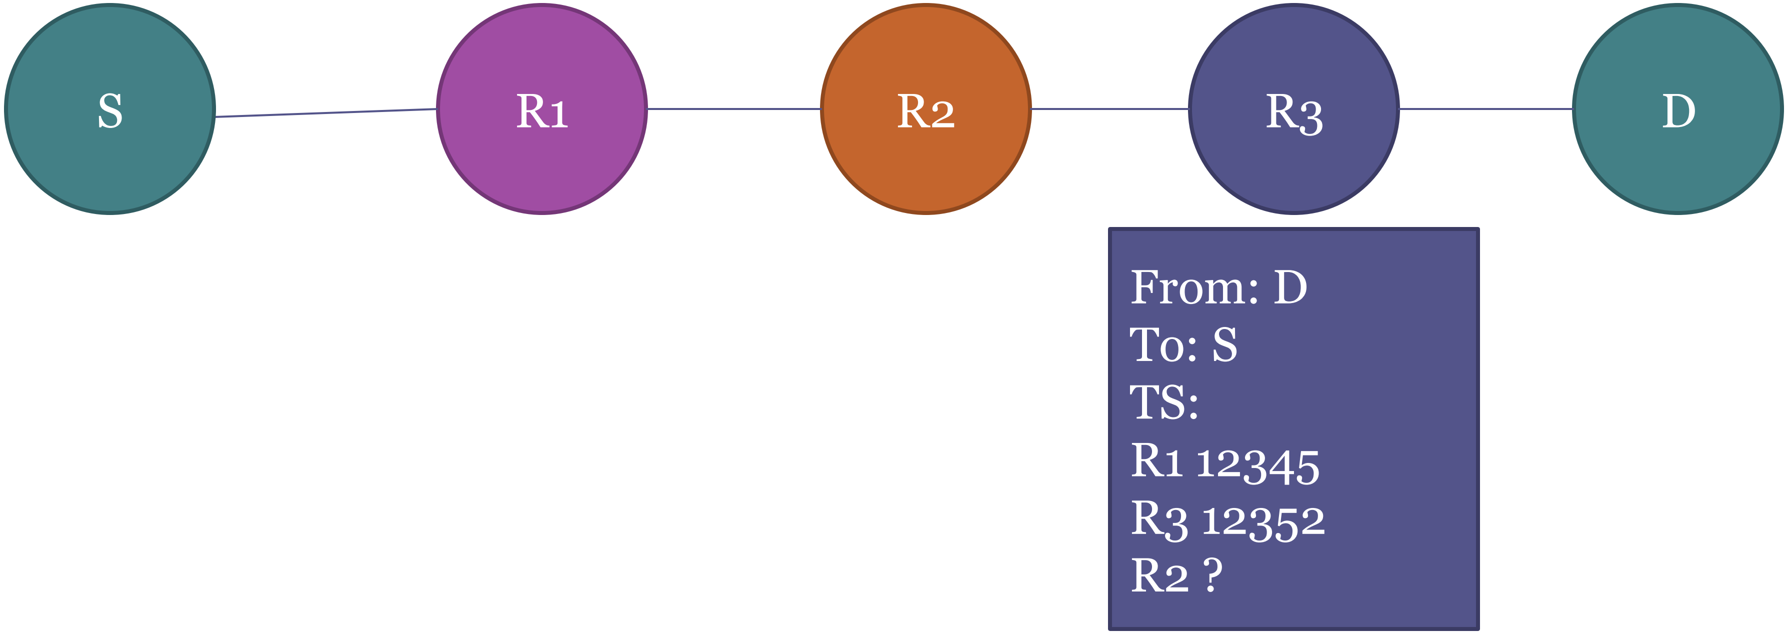

Example 2

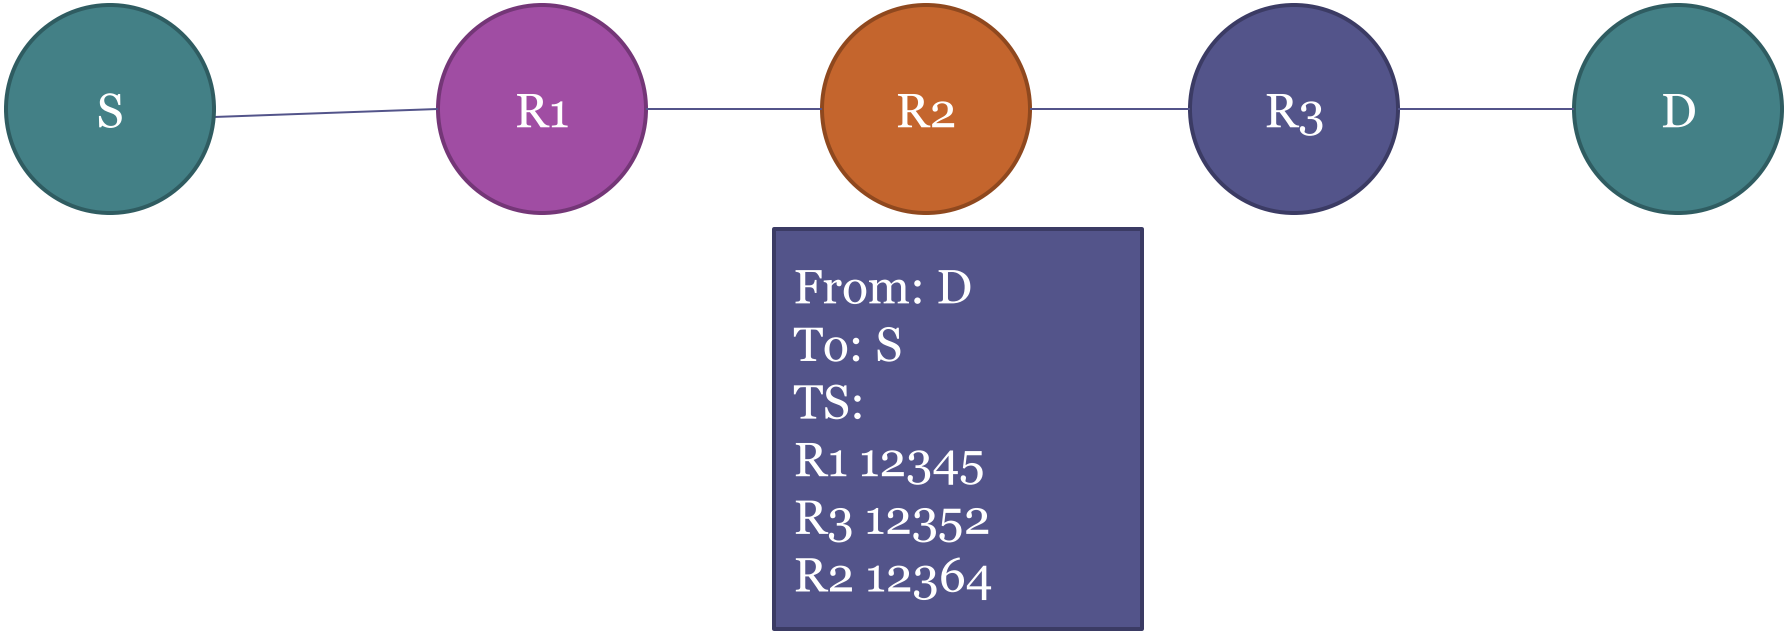

Example 2

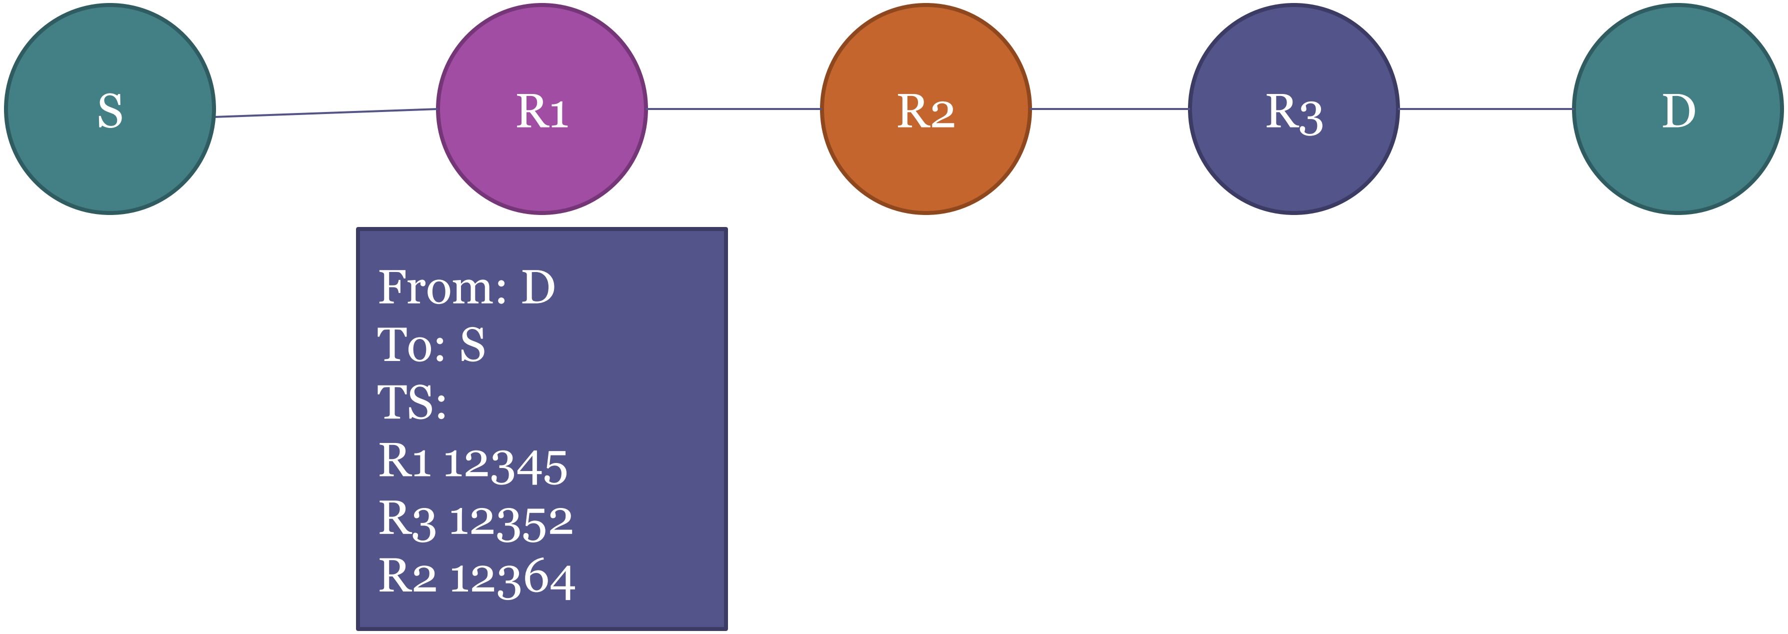

Example 2

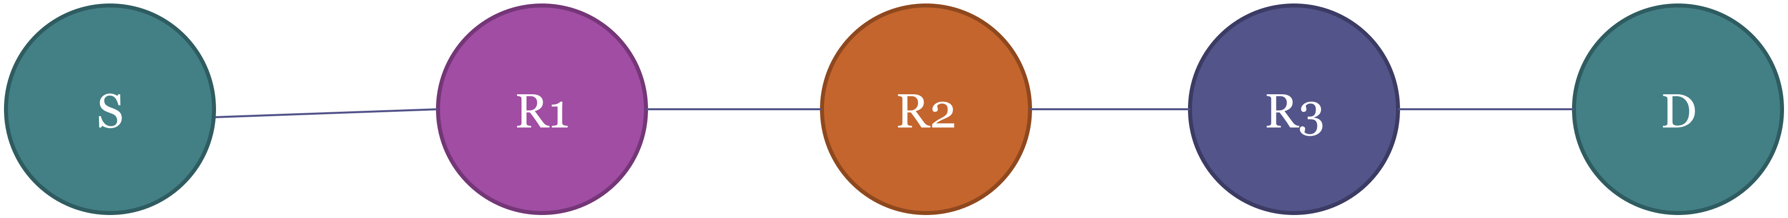

From: D
To: S
TS:
R1 12345
R3 12352
R2 12364

Unique Characteristics

- Probe can be stamped in-transit on the forward or the reverse path
- Can query multiple IP addresses in a single probe
- Timestamp sequence implies ordering between routers
- Literal timestamp values provide milliseconds since midnight UTC

How often are timestamps supported?

We performed studies of several different datasets. In our least successful run, we saw

- 55.5% dropped the packet
- 19.5 % do not provide timestamps
- 25% do provide timestamps

in measurements to all targets from over 20 PlanetLab vantage points.

Not all timestampers are the same

We targeted 153,565 routers with a direct request and asked for their own timestamp in all four prespecified requests.

From: S
To: D
TS:
D?
D?
D?
D?

Not all timestampers are the same

We targeted 153,565 routers with a direct request and asked for their own timestamps in all four prespecified requests. Of those that provided timestamps,

- 34.9% provided one stamp

From: D
To: S
TS:
D 12345
D?
D?
D?

Not all timestampers are the same

We targeted 153,565 routers with a direct request and asked for their own timestamp in all four prespecified requests. Of those that provided timestamps,

- 34.9% provided one stamp
- 43.9% provided two stamps

From: D
To: S
TS:
D 12345
D 12345
D ?
D ?

Not all timestampers are the same

We targeted 153,565 routers with a direct request and asked for their own timestamp in all four prespecified requests. Of those that provided timestamps,

- 34.9% provided one stamp
- 43.9% provided two stamps
- 13.5% provided four stamps

From: D
To: S
TS:
D 12345
D 12345
D 12345
D 12345

Agenda

- Motivating Internet Measurements
- Understanding IP Timestamp
- Three Use Cases:
  - Reverse Path Visibility
  - IP Alias Resolution
  - Single Link One-Way Delay

Traceroute shows the forward path

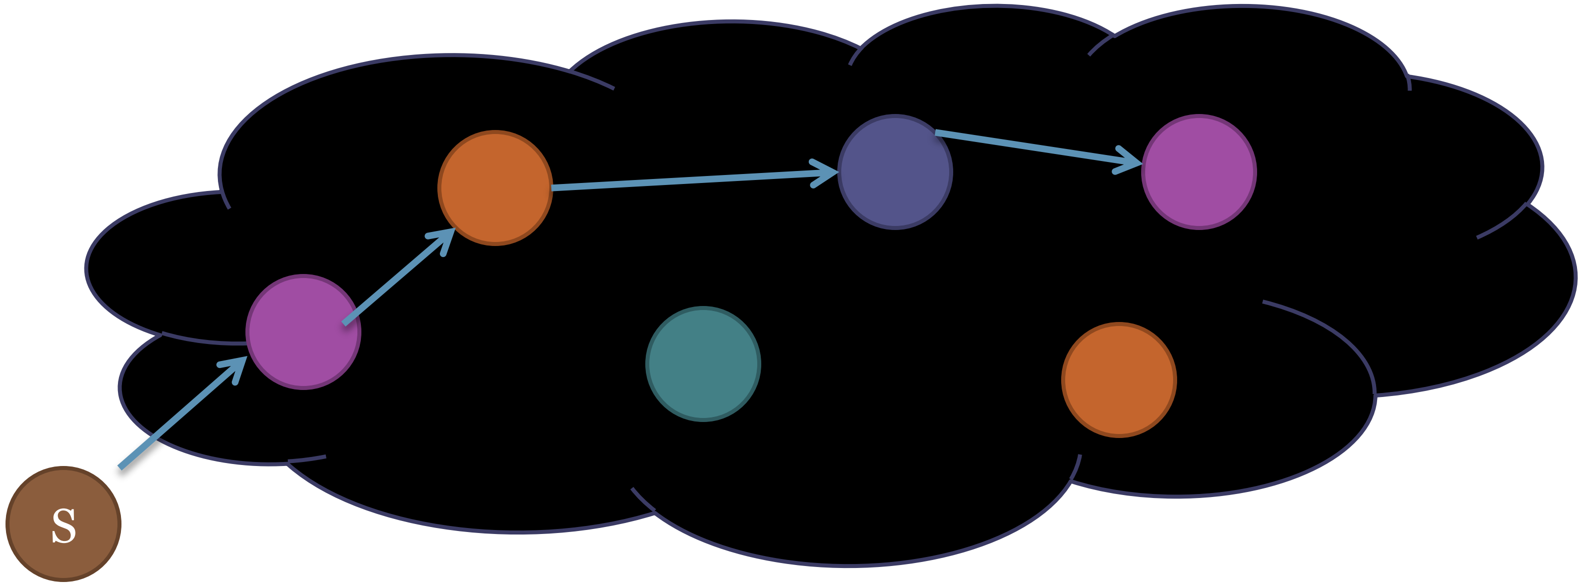

Reverse path is often very different

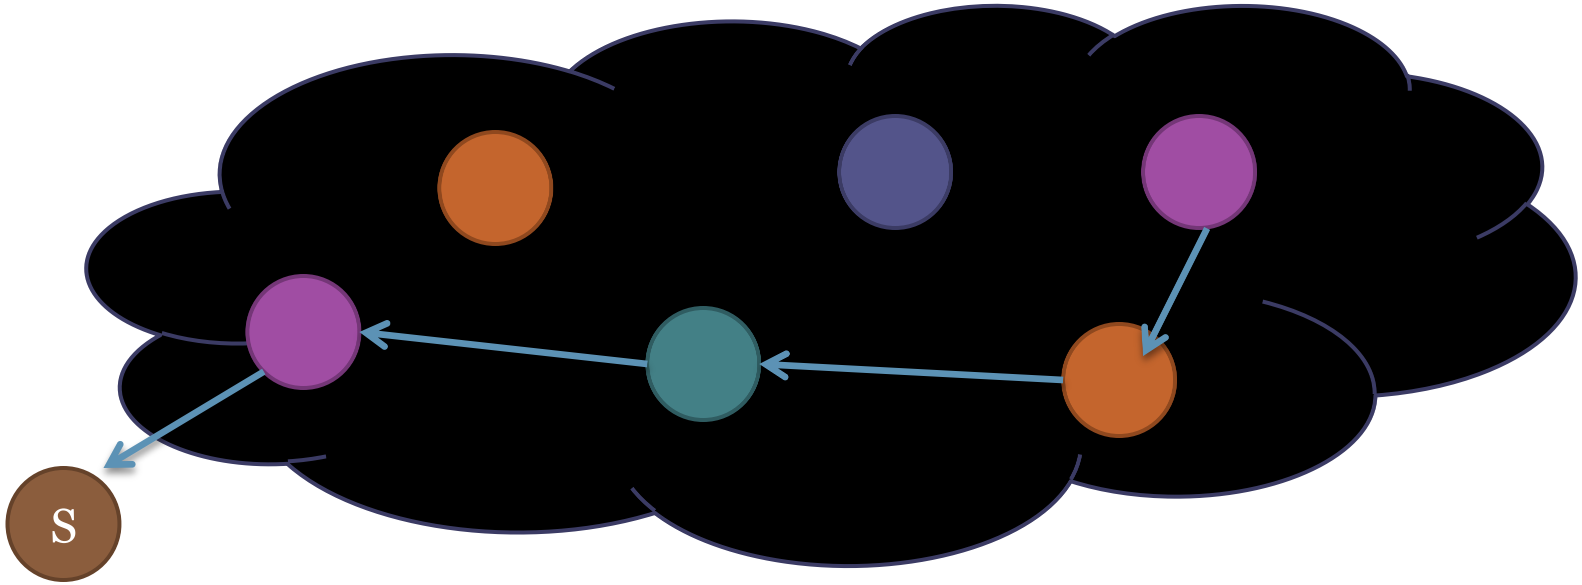

Making a guess: sometimes we have an idea that a router might be on the reverse path

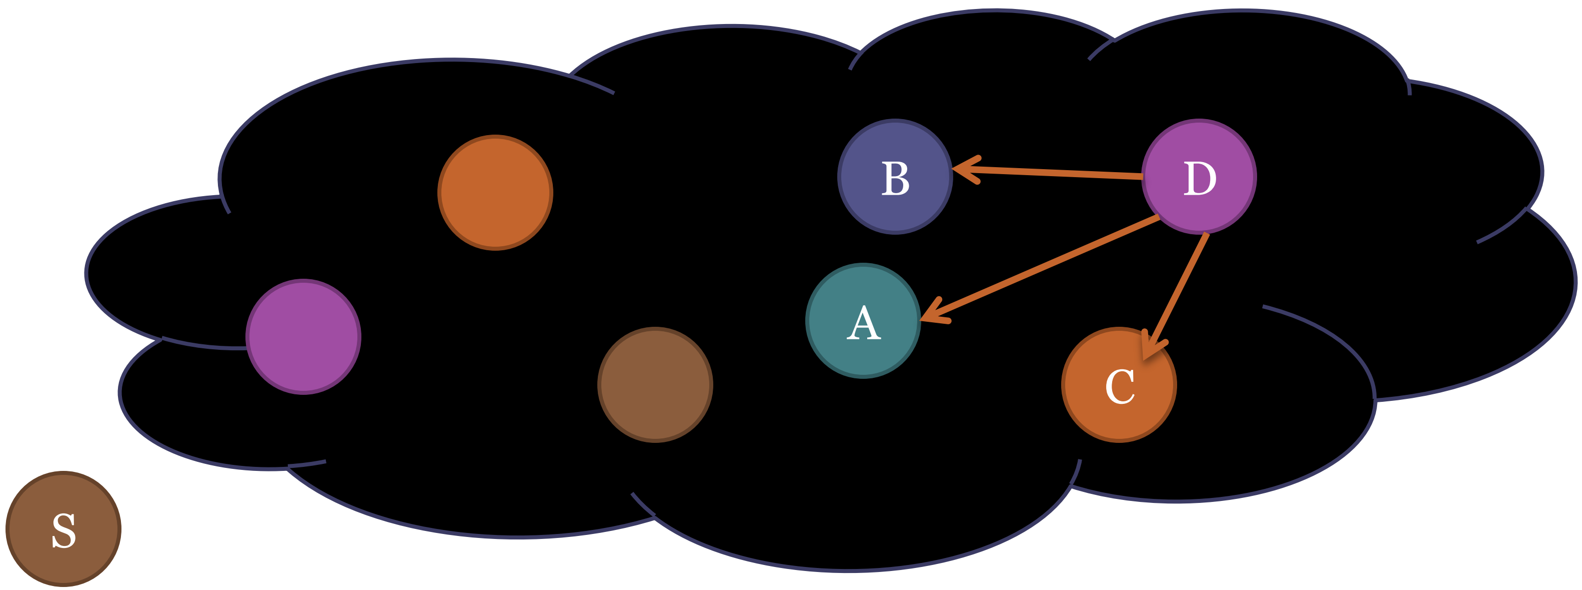

Send a timestamp probe and request stamps from destination, followed by our guess hop, to verify

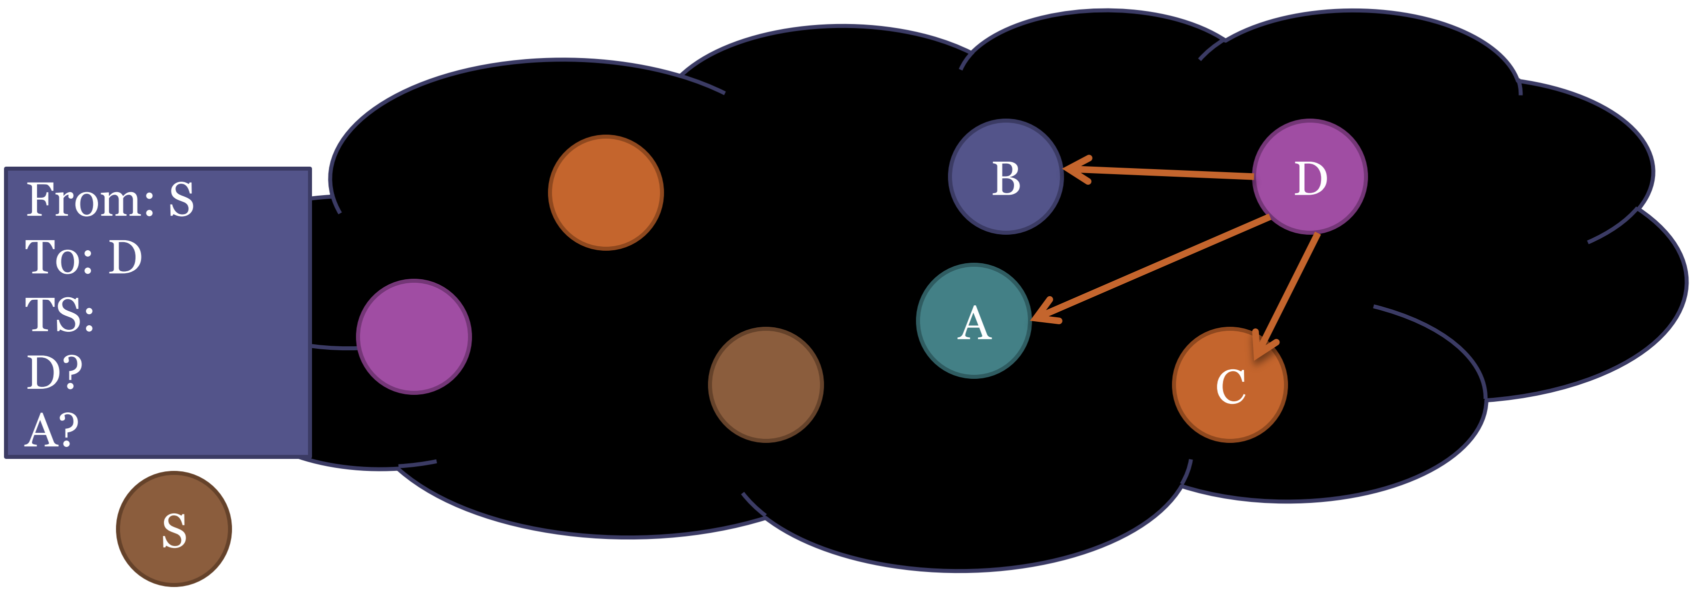

Send a timestamp probe and request stamps from destination, followed by our guess hop, to verify

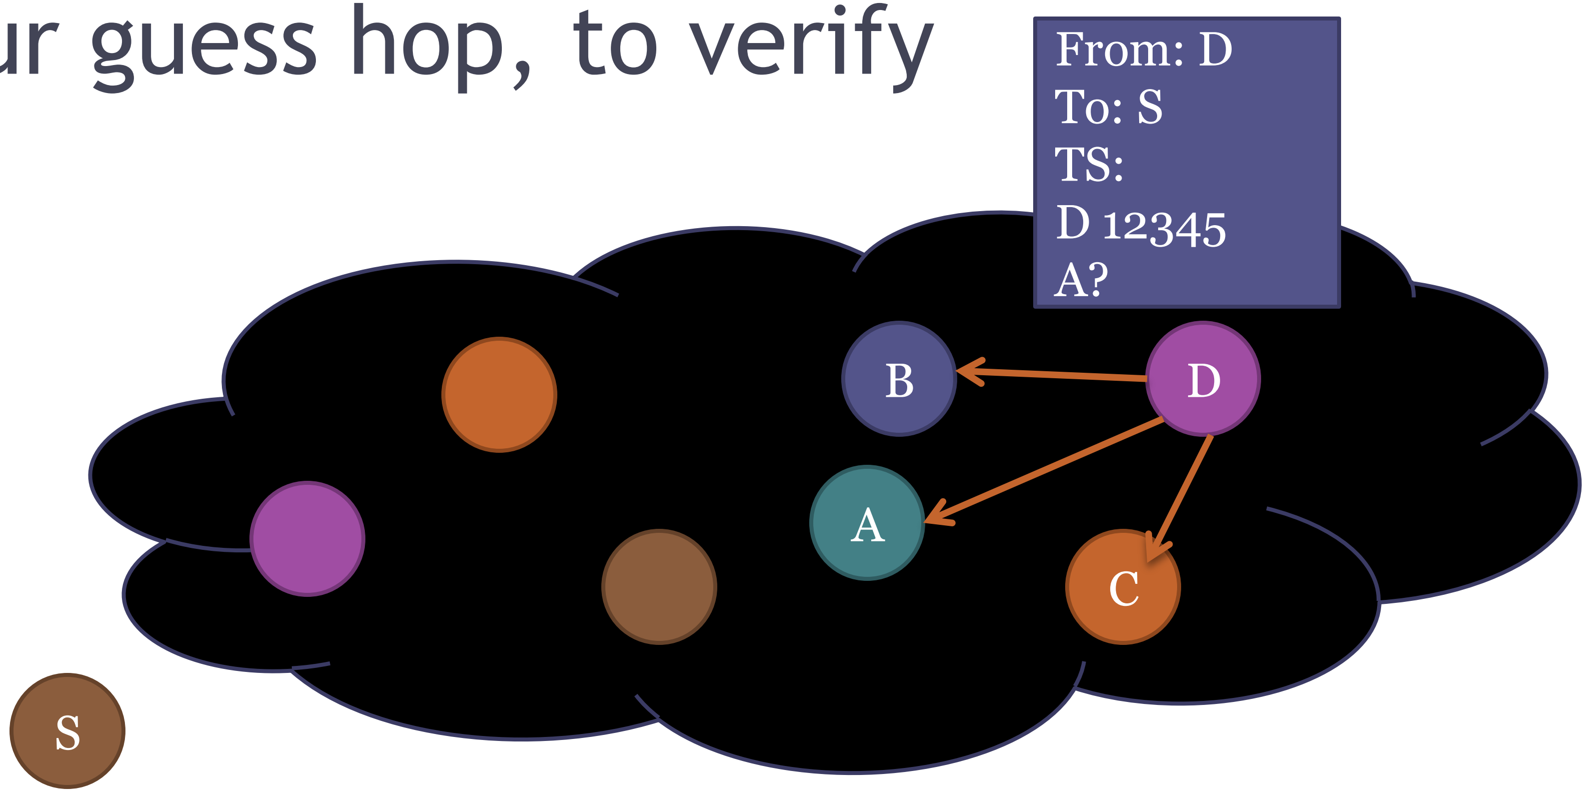

Send a timestamp probe and request stamps from destination, followed by our guess hop, to verify

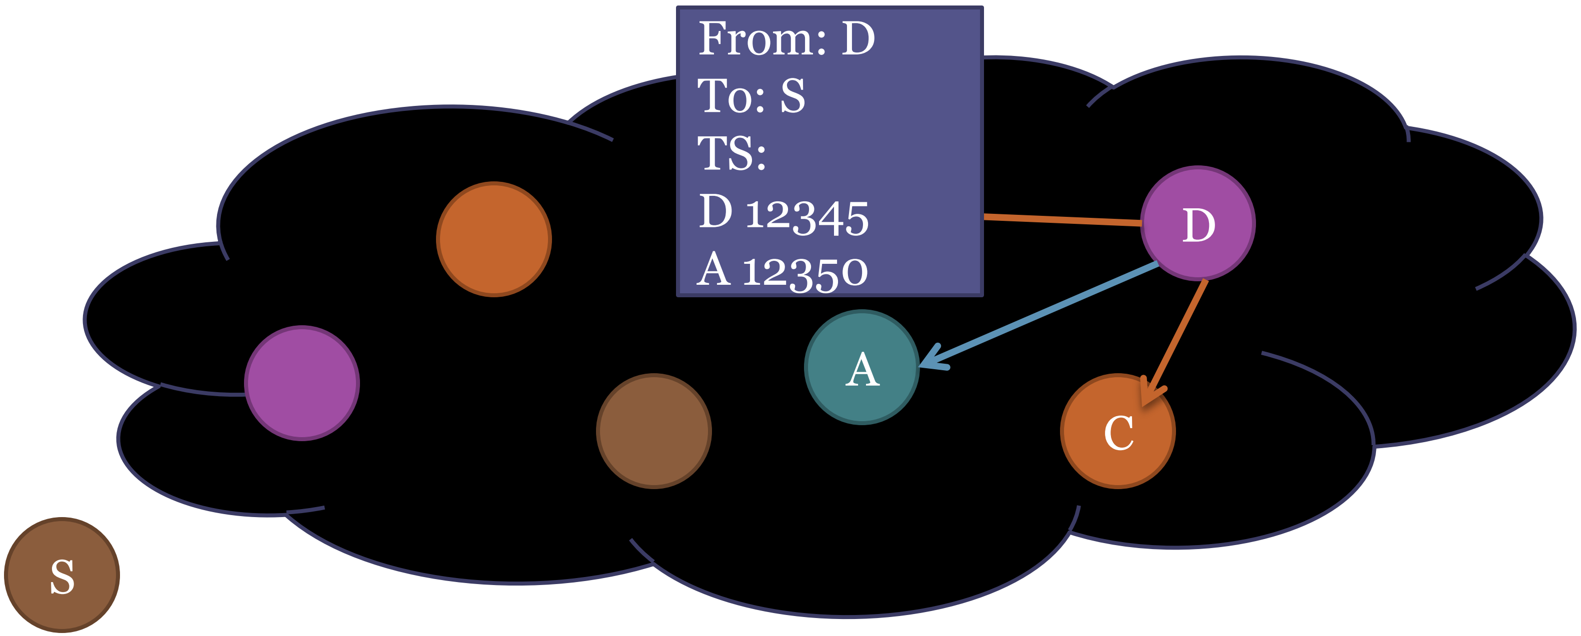

Some More Complicated Cases

- What if D doesn’t provide timestamps at all, but A still does?
- What about anomalous timestamp implentations?
  - Addressed in written thesis
- How can we combine timestamp with other techniques to view the whole path?
  - Talk to Ethan and read Reverse Traceroute ☺

Timestamp Gains for Reverse Traceroute

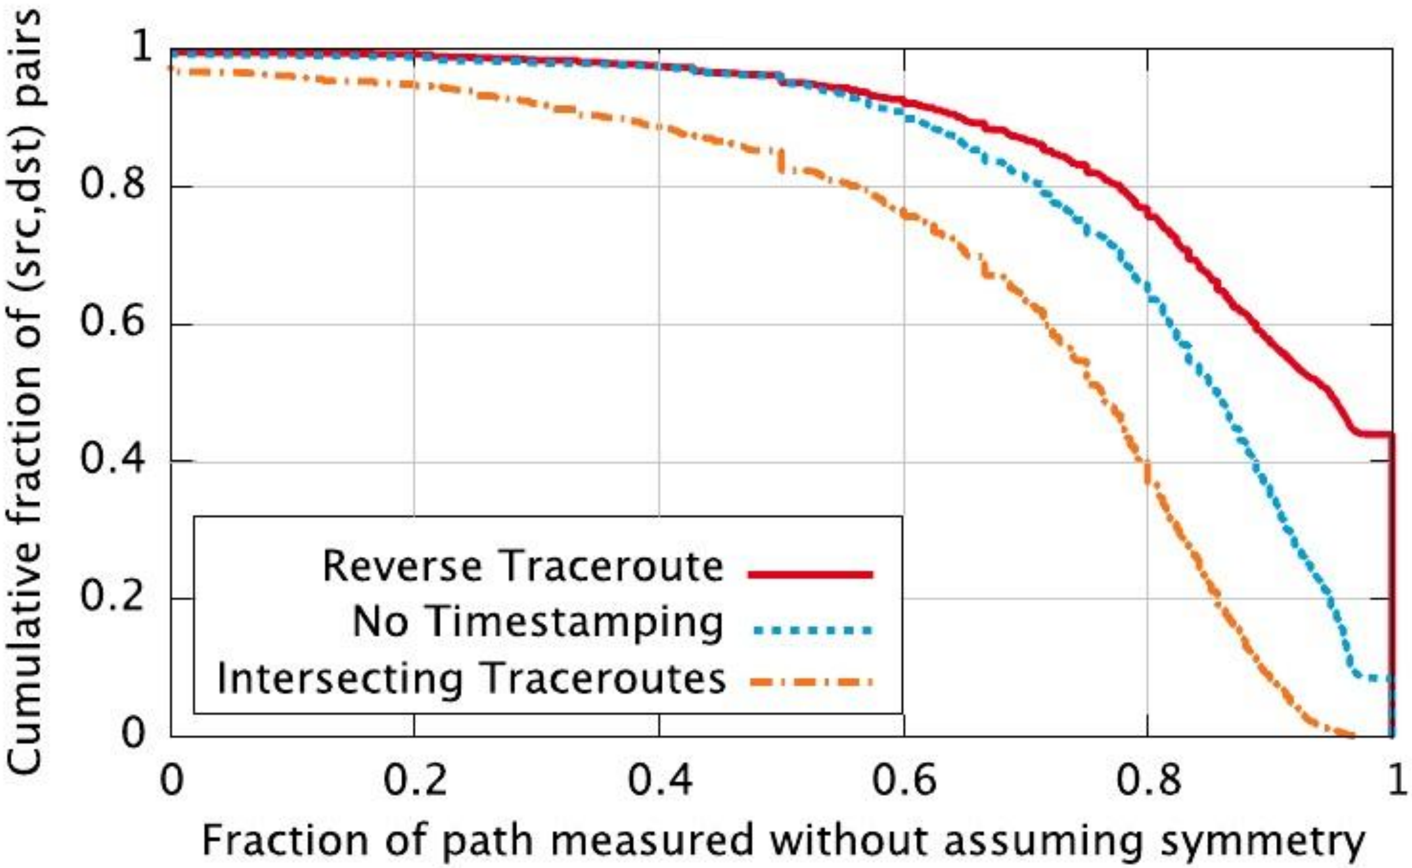

Agenda

- Motivating Internet Measurements
- Understanding IP Timestamp
- Three Use Cases:
  - Reverse Path Visibility
  - IP Alias Resolution
  - Single Link One-Way Delay

IP Aliasing Problem

- Routers may have dozens of IP addresses assigned to them
- We use IP addresses as identifiers for routers
- Different measurements of the same router may be associated with several IP addresses

Many Attempts to Resolve Aliases

- Ally, Radargun, DisCarte, Mercator…
- All rely on various tricks to make the router reveal the association between different IPs
- Timestamp can fill in some of the gaps left by these existing techniques

Timestamp Aliasing

- We can use timestamps to place constraints on the relationship between a candidate pair A and B:
  - Topological constraint: order of the timestamps implies the order that the packet traversed the routers
  - Shared clock constraint: timestamp values can inform us whether A and B may access a common clock

Timestamp-Based Alias Resolution

- Say we have a candidate alias pair A,B.
- Send a probe to A, and request timestamps from A and B interleaved
- Send a probe to B, and do the same

To: A
From: Justine
TS:
A?
B?
A?
B?

To: B
From: Justine
TS:
B?
A?
B?
A?

Understanding a Timestamp Reply

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

What configuration of A and B might have generated this response?

Understanding a Timestamp Reply

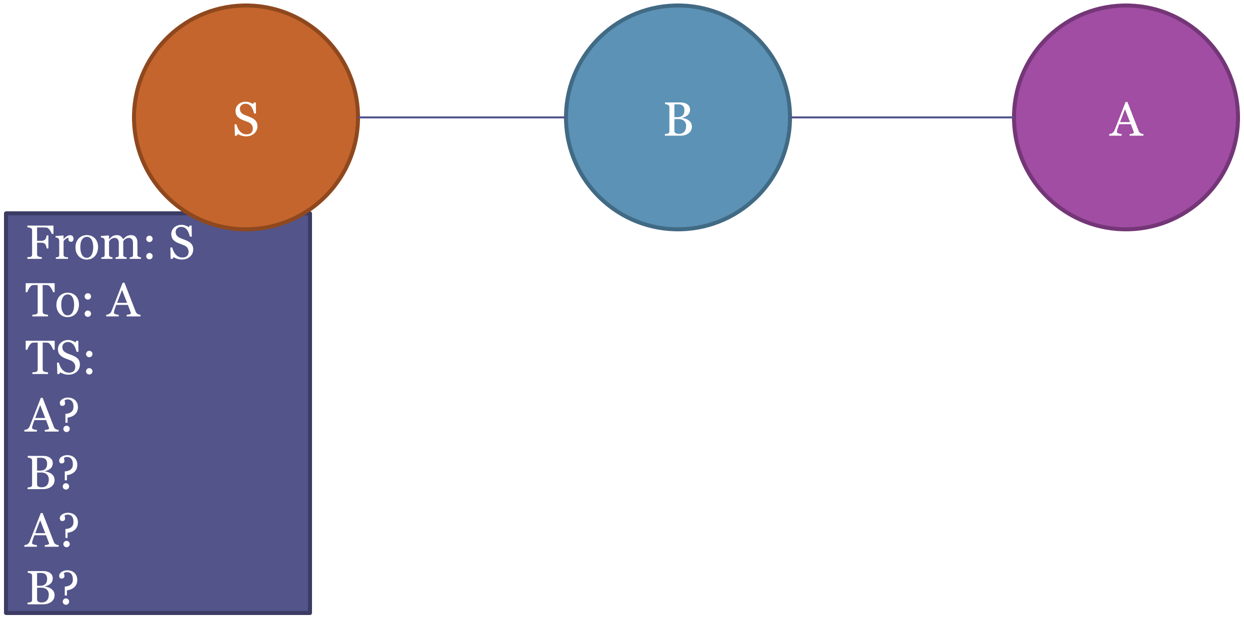

Understanding a Timestamp Reply

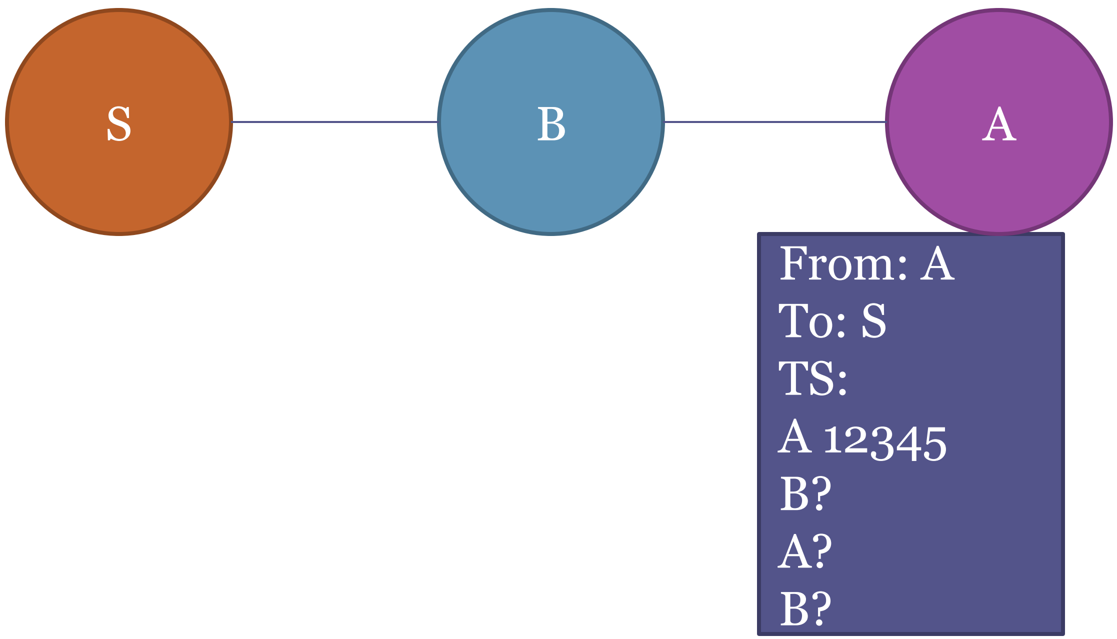

Understanding a Timestamp Reply

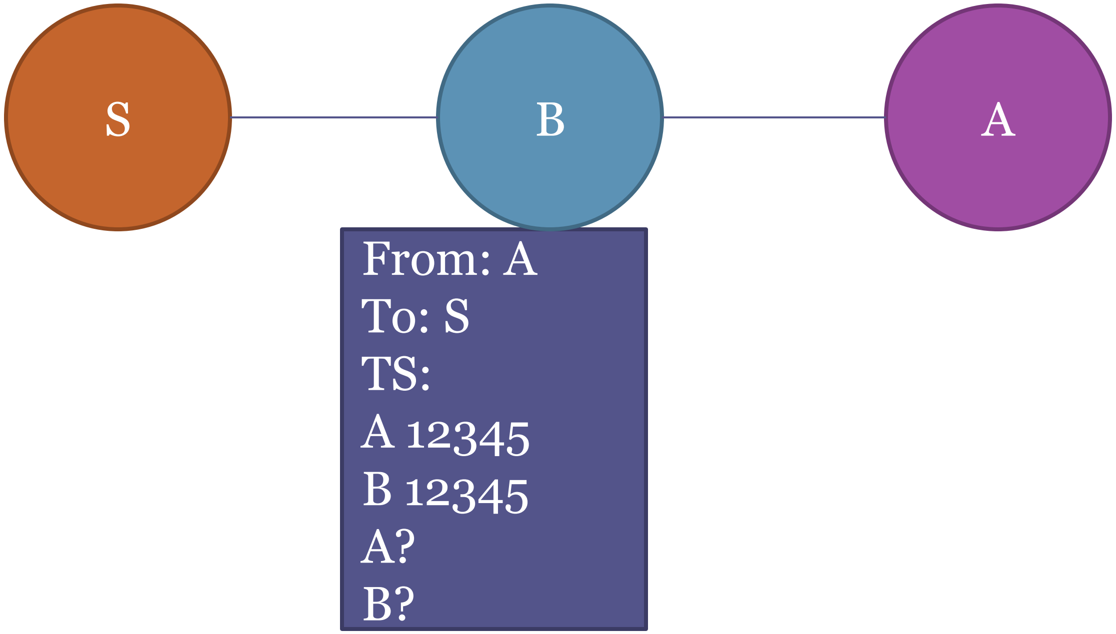

Understanding a Timestamp Reply

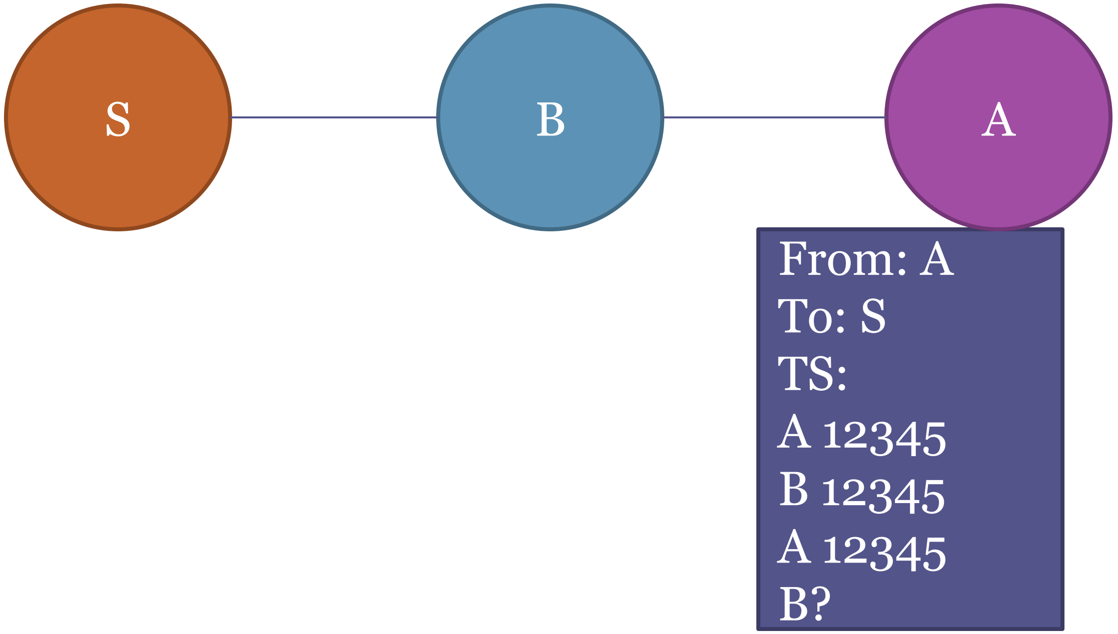

Understanding a Timestamp Reply

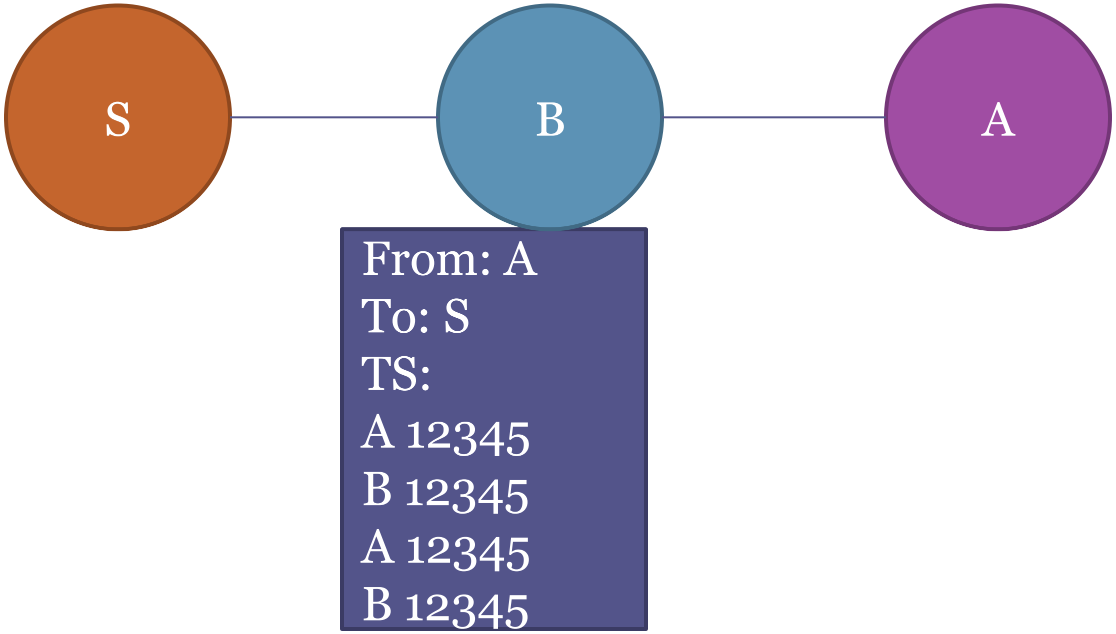

Understanding a Timestamp Reply

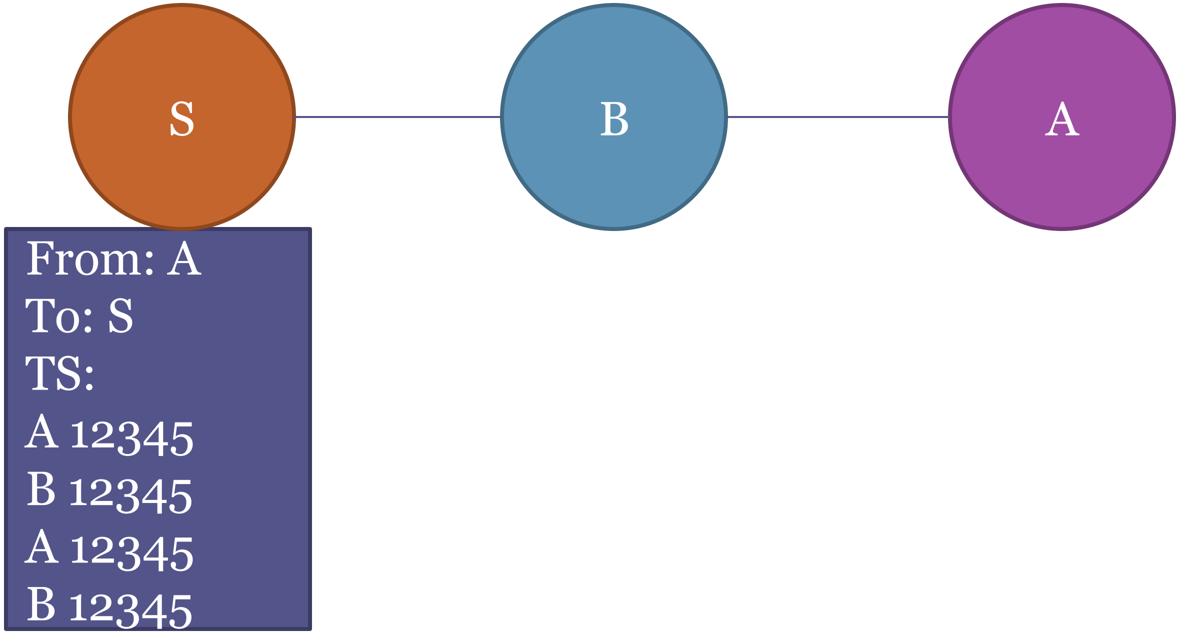

A and B are Aliases

- That forwarding pattern was wacky!
- Furthermore, the timestamp values are identical across all four stamps, despite A and B forwarding the packet back and forth four times.

It makes much more sense if A and B were just aliases, and the packet was stamped and forwarded once.

In Practice

- We can use similar arguments even if we only get stamps from A and B twice (rather than all four times).
- Measuring over a set of ground-truth alias pairs, we found that Radargun, an existing technique, was unable to address 79.3% of the targets it measured. Of those 79.3%, we were able to successfully confirm alias pairs for 19.3%.
- In September, we confirmed 43,847 alias pairs using the timestamp technique, generating 8,697 alias clusters.

Agenda

- Motivating Internet Measurements
- Understanding IP Timestamp
- Three Use Cases:
  - Reverse Path Visibility
  - IP Alias Resolution
  - One-Way Link Latency

Latency

- Typically measured with Round-Trip Times (RTTs)
- However, many applications require more precise, more accurate measurements.
  - Like geolocation

A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps

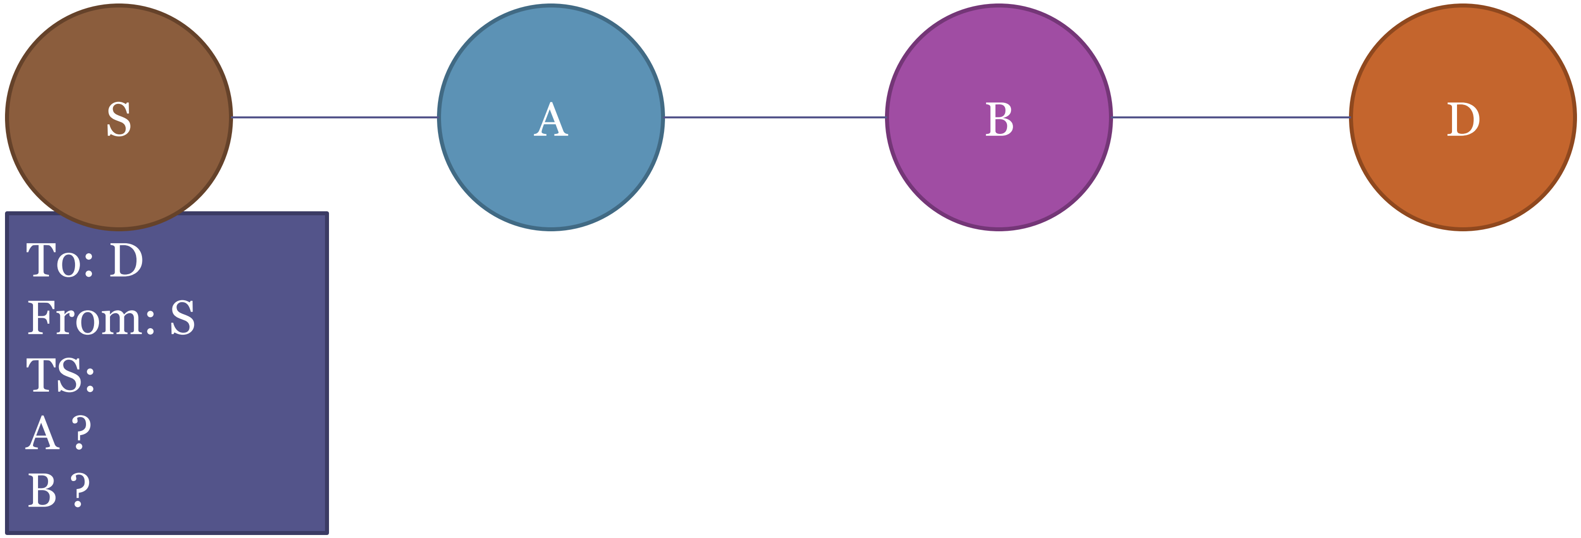

A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps

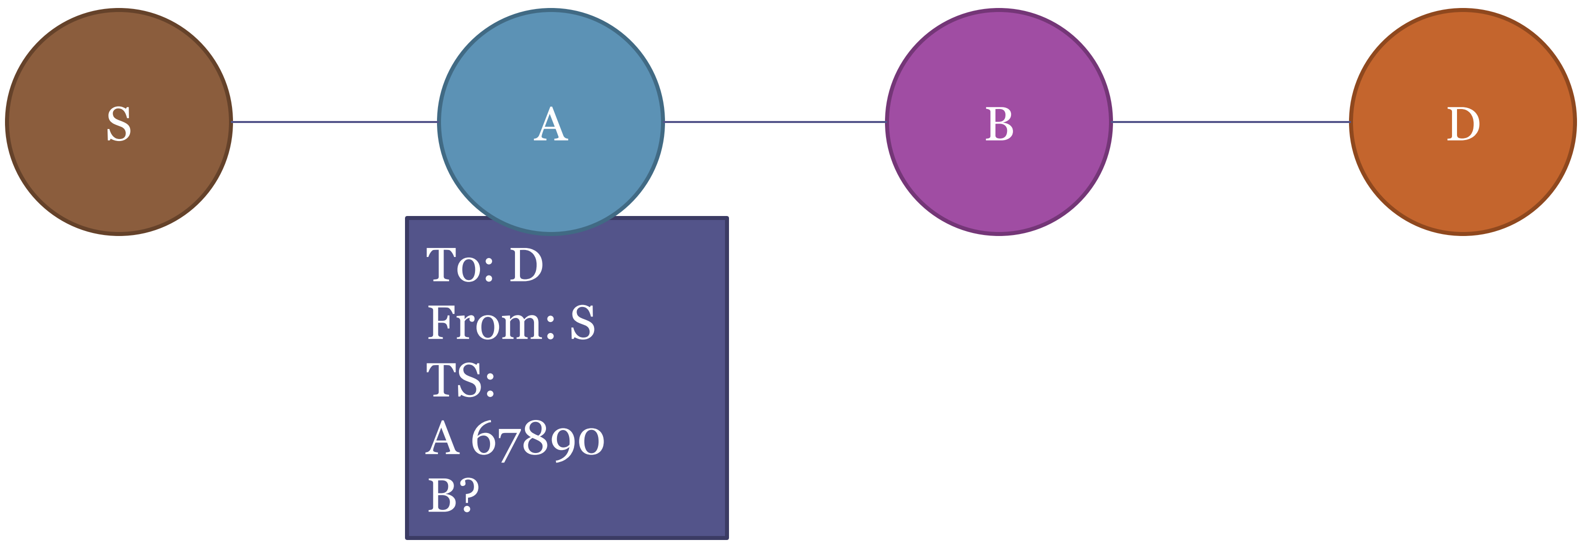

A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps

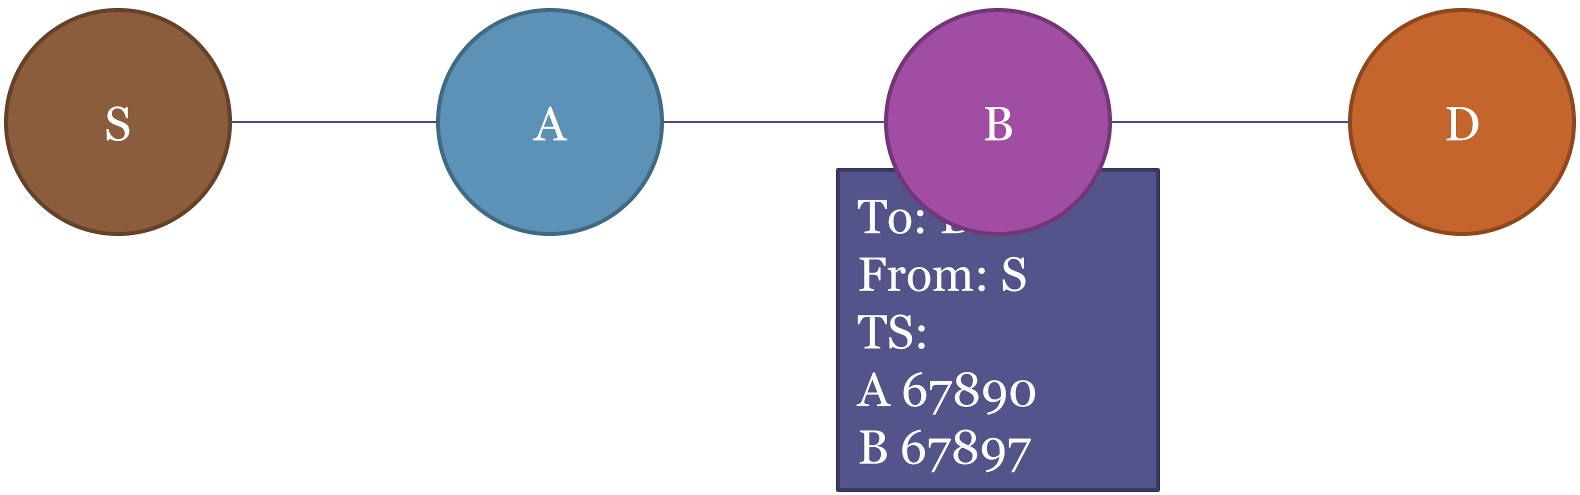

A Timestamp Measurement

- Send a probe that traverses an A-B link, and ask A and B each for timestamps

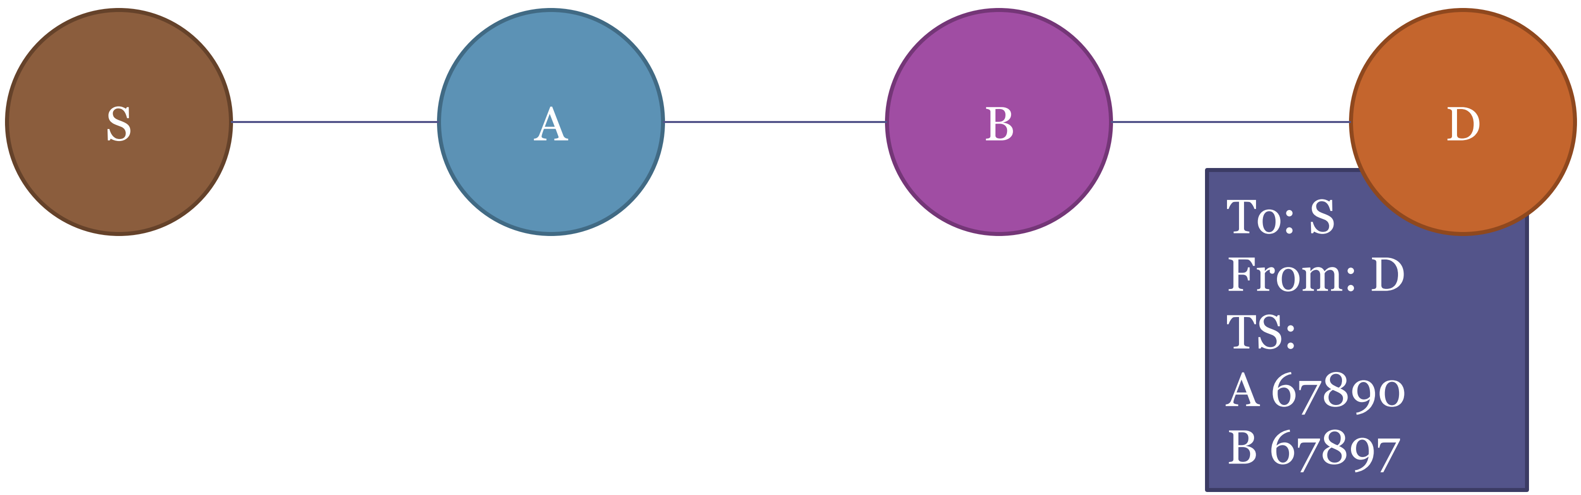

Components of Timestamp Values

- We can subtract A's timestamp from B
- The difference is 7 milliseconds
- But what does this difference comprise?

To: Steve
From: D
TS:
A 67890
B 67897

TS(A) – TS(B) = latency + skew(A,B) + queue

Components of Timestamp Values

- We can subtract A's timestamp from B
- The difference is 7 milliseconds
- But what does this difference comprise?

To: Steve
From: D
TS:
A 67890
B 67897

TS(A) – TS(B) = latency + skew(A,B) + queue

Can ignore by
taking the min
across several
measurements

Components of Timestamp Values

- We can subtract A's timestamp from B
- The difference is 7 milliseconds
- But what does this difference comprise?

To: Steve
From: D
TS:
A 67890
B 67897

TS(A) – TS(B) = latency + skew(A,B) + queue

Still need to get rid of this!

Canceling out Skew

- What if we could measure the B-A Link in the opposite direction?
- With many PlanetLab nodes, we can find a path that crosses the link in the opposite direction

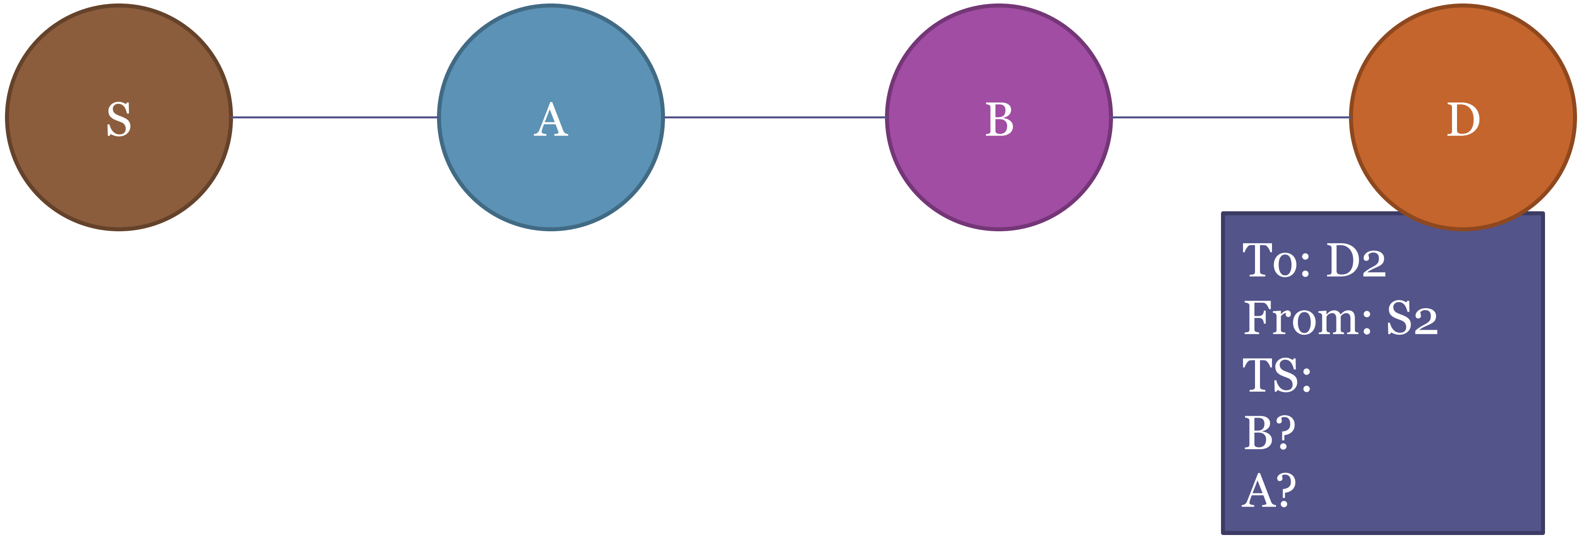

Canceling out the Skew

$\Delta 1 = TS(A) - TS(B) = \text{latency} + \text{skew}(A,B) + \text{queue}$

To: Steve
From: D
TS:
A 67890
B 67897

$\Delta 2 = TS(A) - TS(B) = \text{latency} - \text{skew}(A,B) + \text{queue}$

To: Ethan
From: D2
TS:
B 67900
A 67912

So… $\text{latency} = \Delta 1 + \Delta 2 / 2$

Work in Progress

We tested over links in the Internet2 backbone

For 11 out of 13 links, we were within a millisecond of the estimates provided by measurements at the routers themselves!

Further experiments required to see if technique will be successful in more diverse networks.

Conclusion

- IP Timestamps are a practical tool for Internet Measurement.
- Timestamps are s*upported by over 25% of routers*.
- Measurement techniques can take advantage of *unique characteristics* of timestamp to confirm if a router lies on the reverse path a packet takes, declare IP aliases, and measure the delay of a single link.

Acknowledgements

Ethan Katz-Bassett, Arvind Krishnamurthy, Tom Anderson, Colin Scott, Mary Pimenova, Harsha Madhyastha, and probably like half the people in this room.