# Measuring Reverse Paths

Ethan Katz-Bassett, Harsha V. Madhyastha,
Arvind Krishnamurthy, Thomas Anderson

University of Washington
August 2008

This work partially supported by Cisco, Google, NSF

# Reverse Paths Would Be Useful

Many distributed systems would benefit from reverse path information

- **Hubble** to isolate failures and group problems
- **iPlane**, Path-Stitching to provide more accurate path and property predictions
- Ark, etc., for more complete topologies
- Google to find inflated paths back from clients
- ISPs to find inflated paths back to customers

# Current Tools Don't Provide That Info

- ping, traceroute
  - Simple tools proven useful for many systems
  - Only provide forward path or round-trip info
- Existing one-way tools require control of both ends
  - RIPE's TTM infrastructure
  - owping
- Vantage points could solve problems
  - Prober in every home?

# Goals

- Techniques for **reverse traceroute** and **one-way ping** when we do not control destination
- Evaluate how often they work
- Demonstrate how they help us understand Internet
  - Systems from earlier slide: iPlane, topology, Google
  - Asymmetry
  - Daily reverse map from world back to PlanetLab

Preliminary/ongoing for now

Talk will focus on **reverse traceroute**

# Reverse Traceroute Approach

- Exploit destination-based routing
- IP options carried over to response packets
  - Timestamp option (TS): time-query 4 ordered IPs
  - Record route option (RR): first 9 routers recorded
- Spoofing to overcome:
  - Lack of vantage points in most prefixes
  - Max 9 hops recorded with RR
  - Limited support/ filtering of options

# Spoofing?? Isn't that bad?

- We use only a restricted version
  - Only spoofing as nodes we control
  - Rate limit, restrict destinations (no broadcast IPs)
- Millions of spoofed probes sent to 10s of thousands of IPs, no complaints
- **Hubble** and this work show utility
- Lets us approximate:
  - Having control of destinations
  - One-hop detouring/ loose source routing
  - One VP sending to another, bouncing through dst

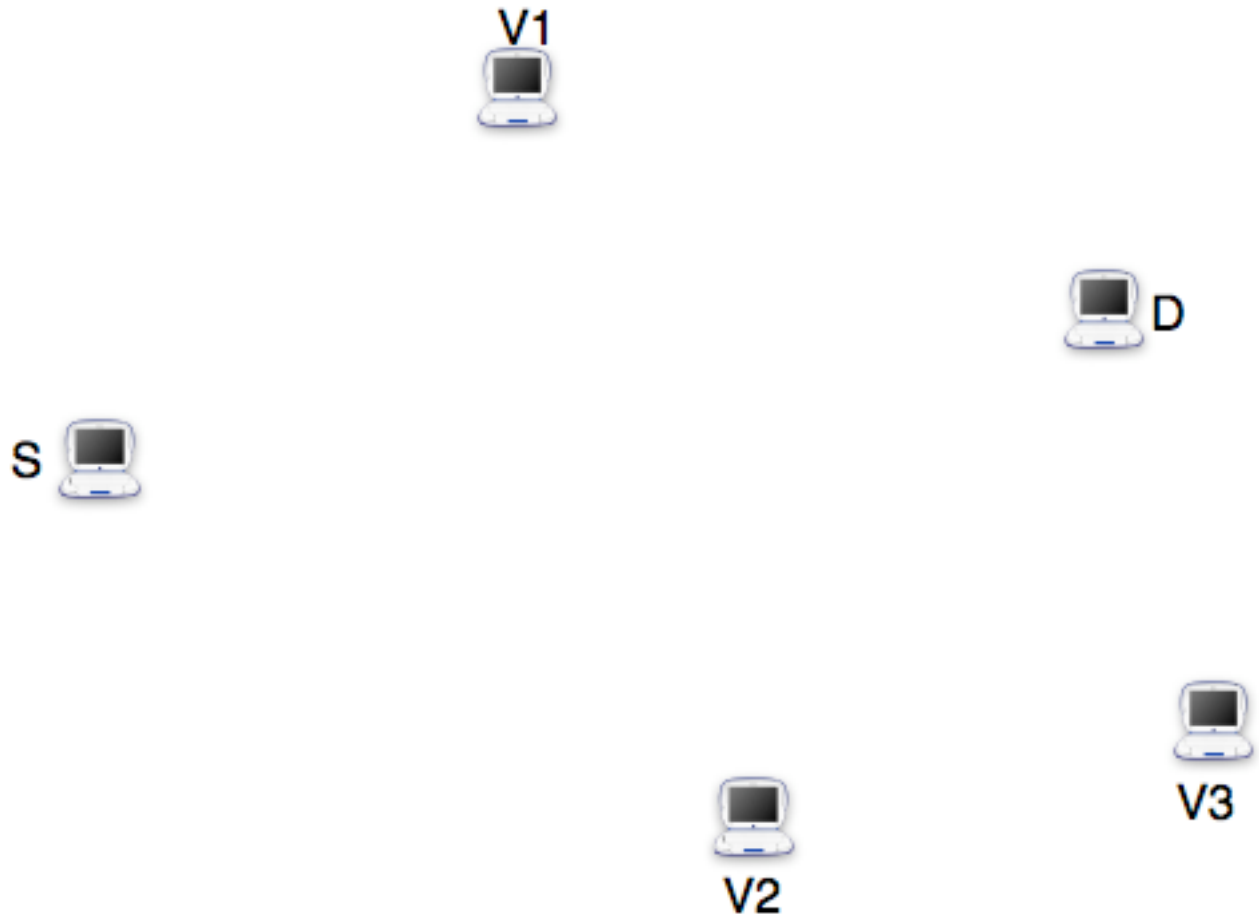


- Want reverse path from **D** back to **S**, but don't control **D**
- Set of vantage points, some of which can spoof

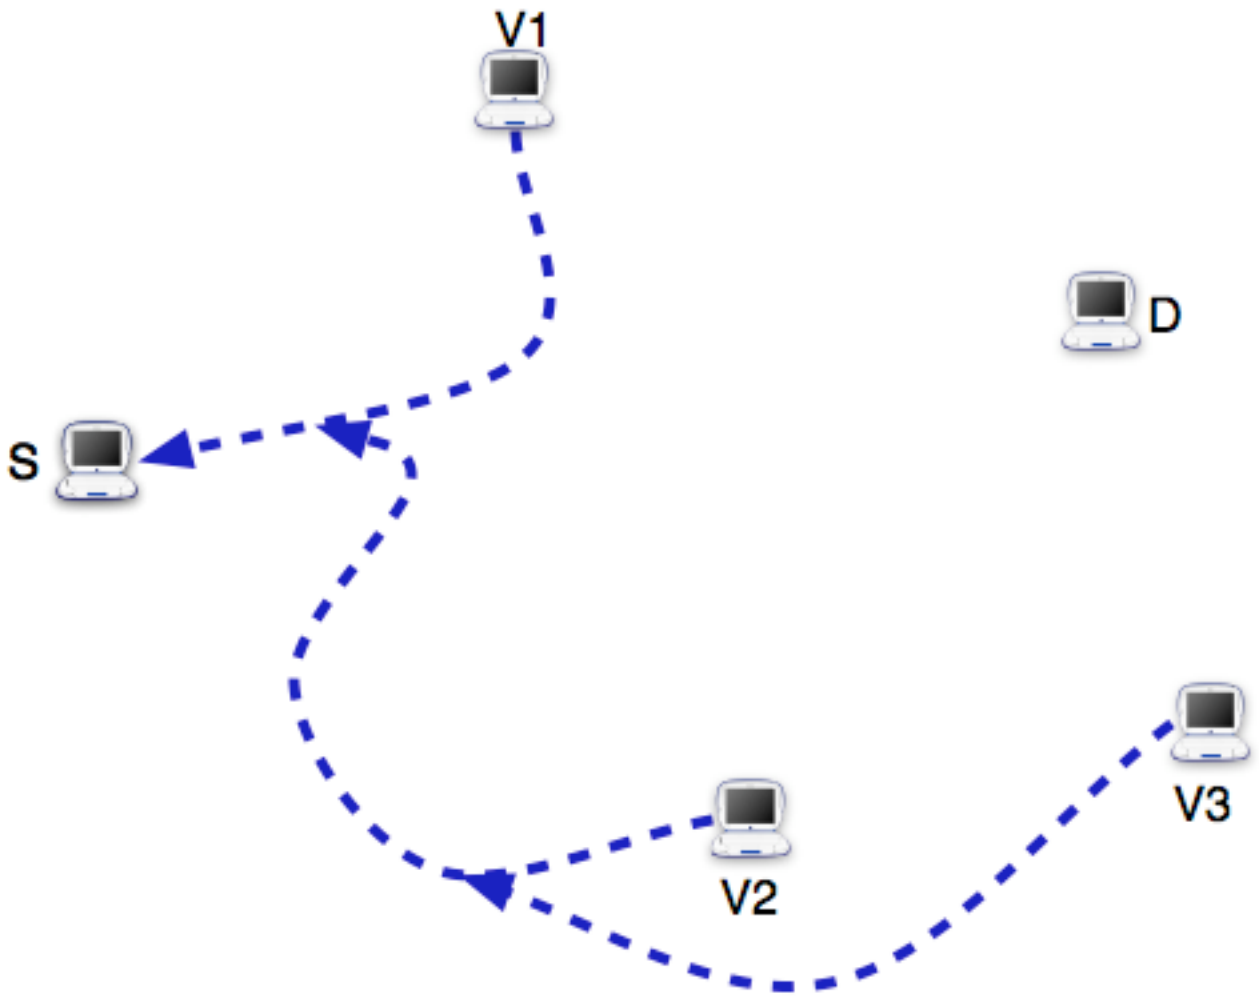


- Traceroute from all vantage points to **S**
- Gives atlas of paths to **S**; if we hit one, we know rest of path

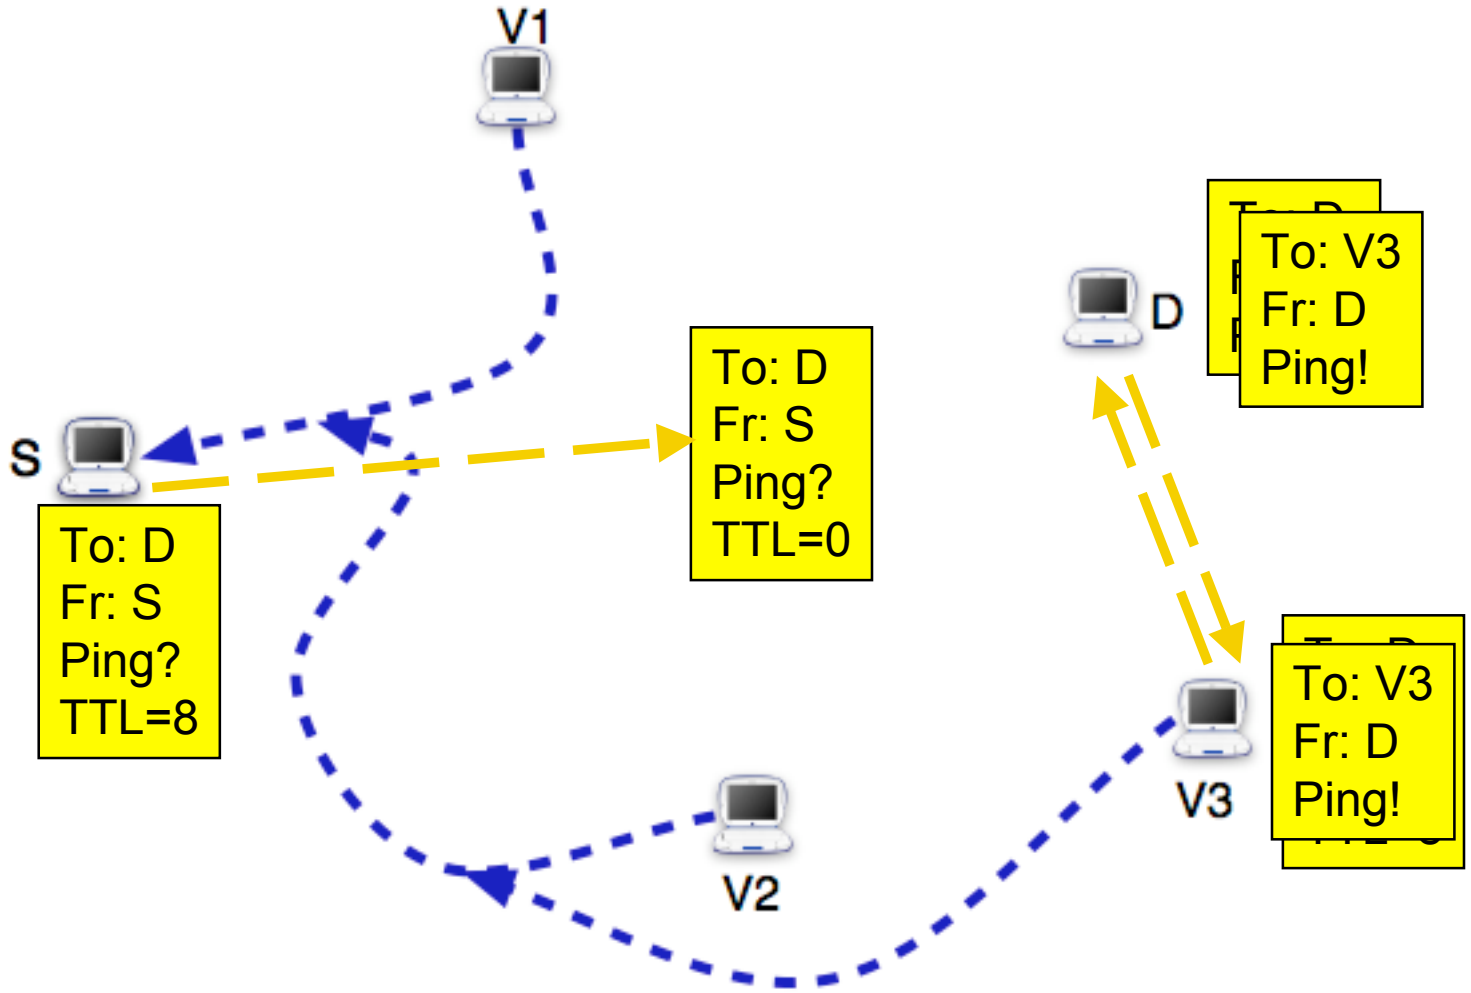


- From all vantage points, ping **D** with TTL=8 to find those within 8 hops
- Record route does 9 hops, so these will give us return hop(s)

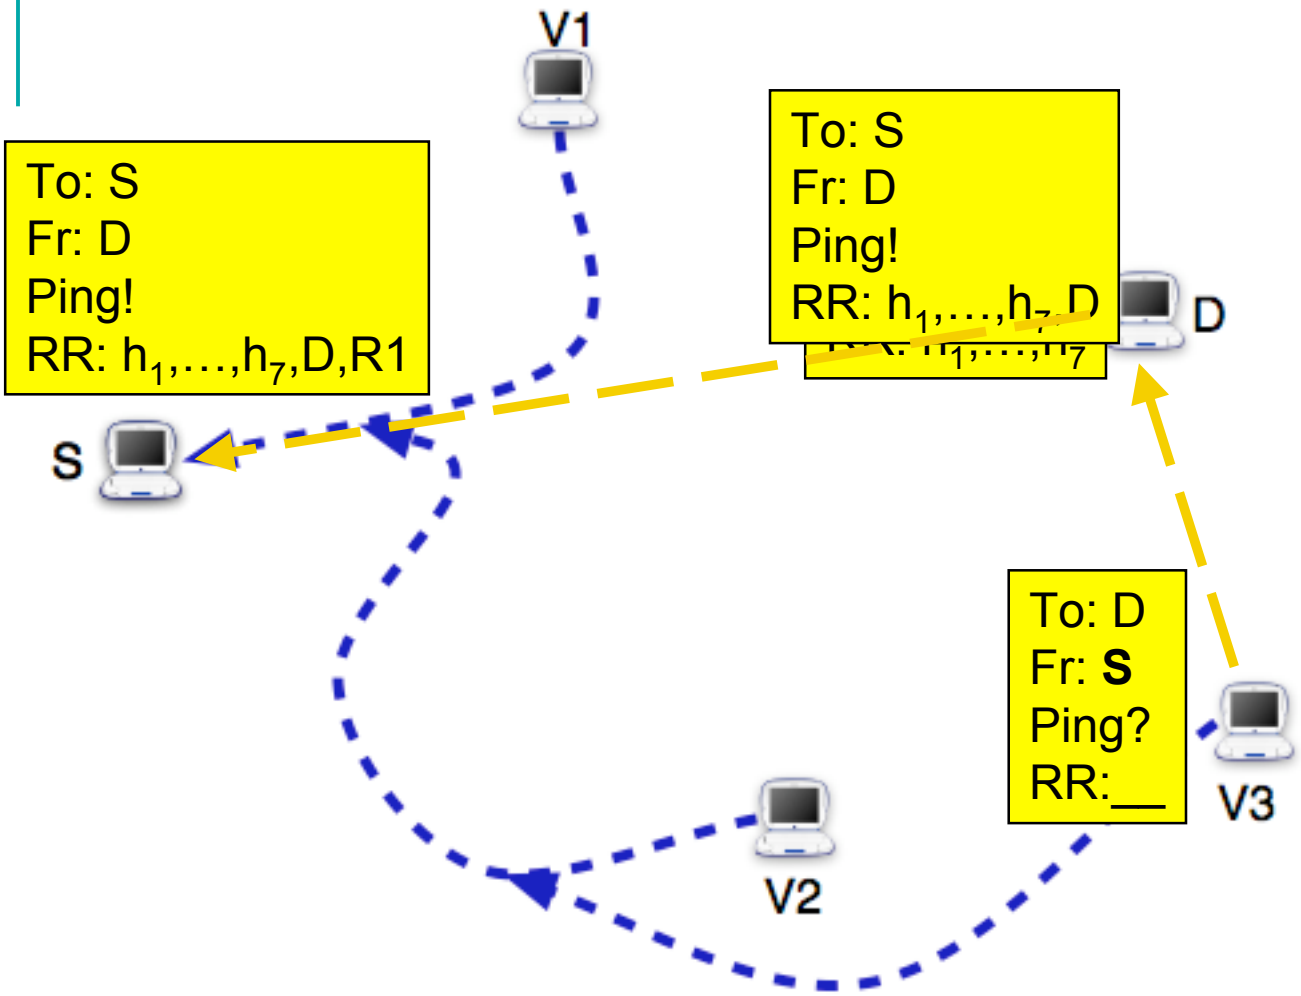


- From vantage point within 8 hops of **D**, ping **D** spoofing as **S** with record route option
- **D**'s response will contain recorded hop(s) on return path

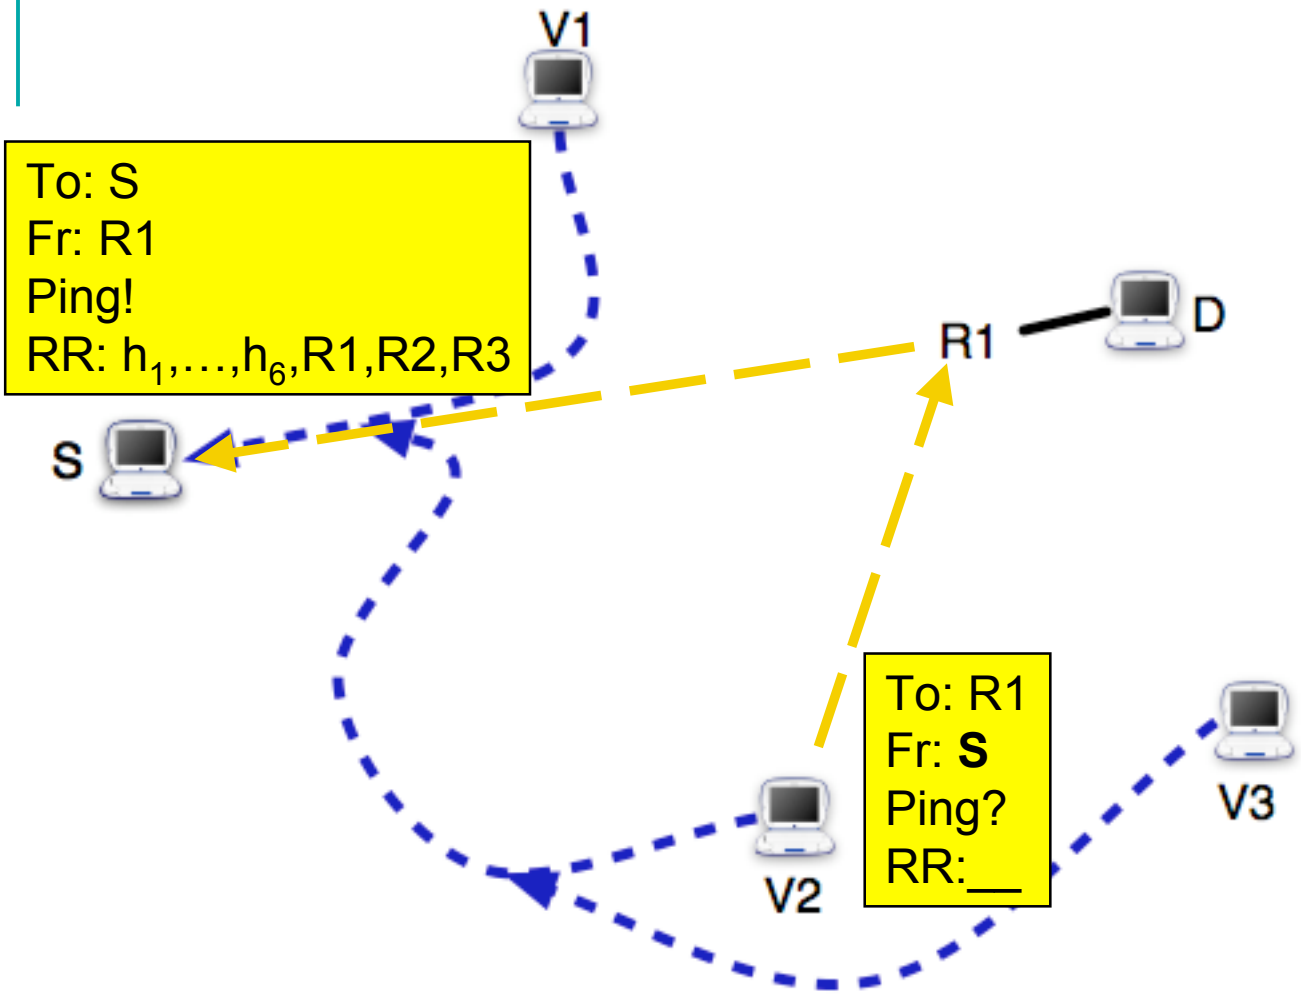


- Iterate, performing TTL=8 pings and spoofed RR pings for each router we discover on return path

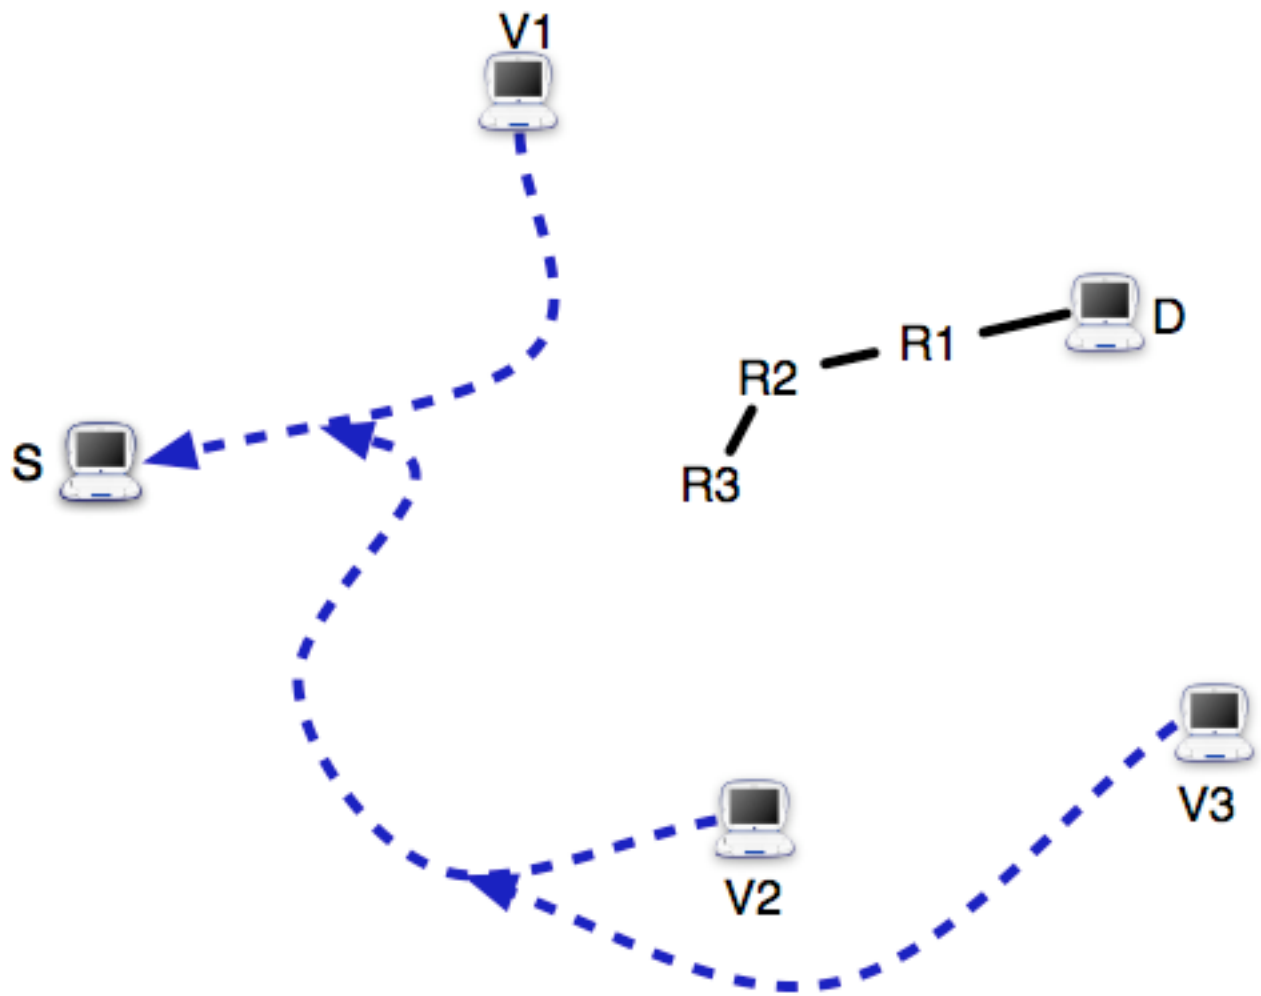
V1
S
R2
R1
D
R3
V2
V3

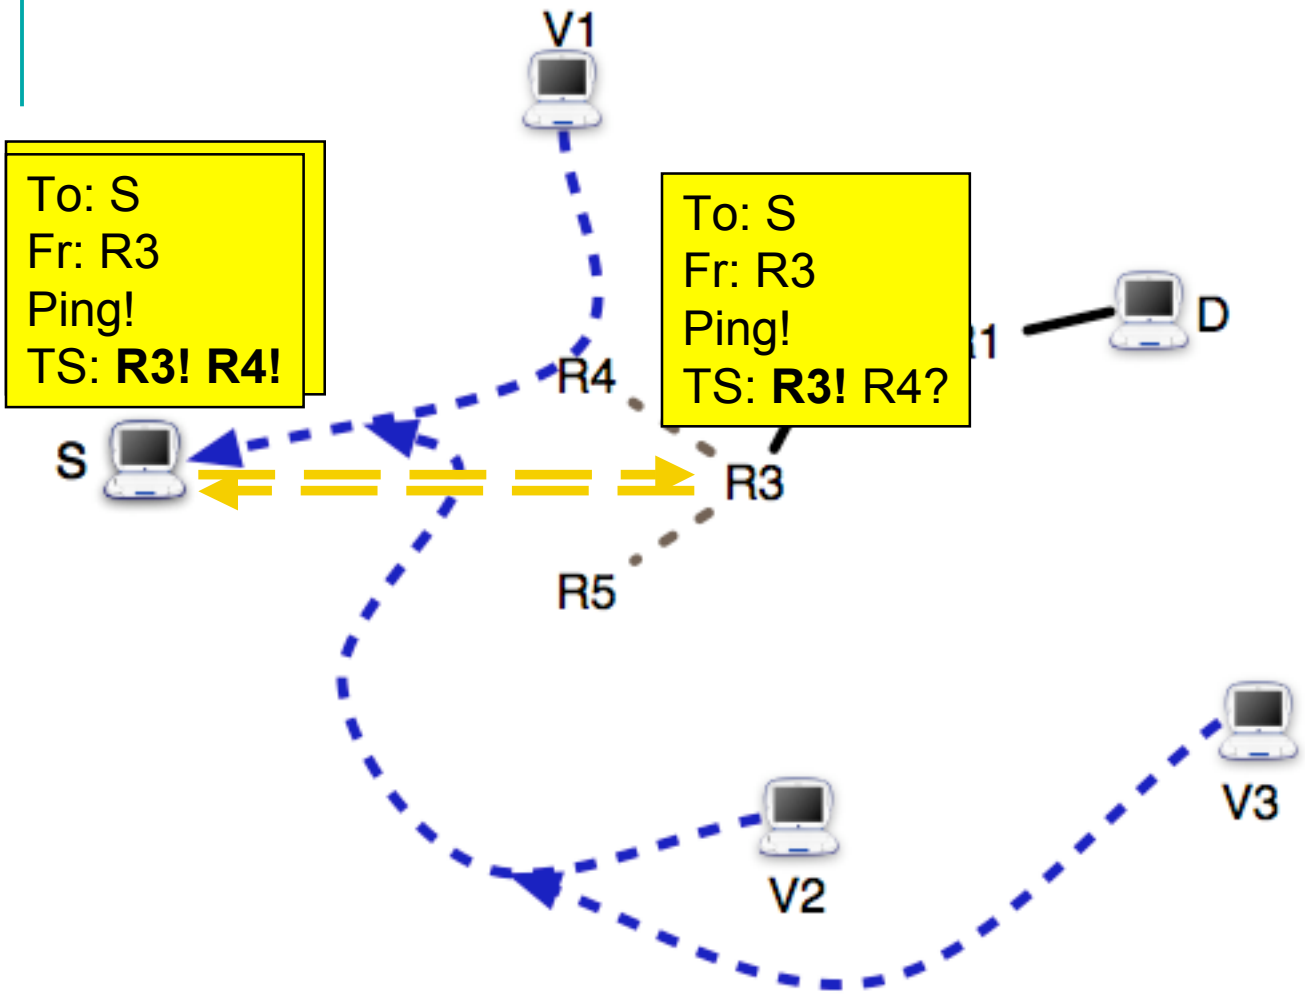


- If no spoofing vantage points within 8 hops, consider set of routers directly connected to **R3** (in pre-measured topology)
- Use timestamp option to try to verify which is on return path

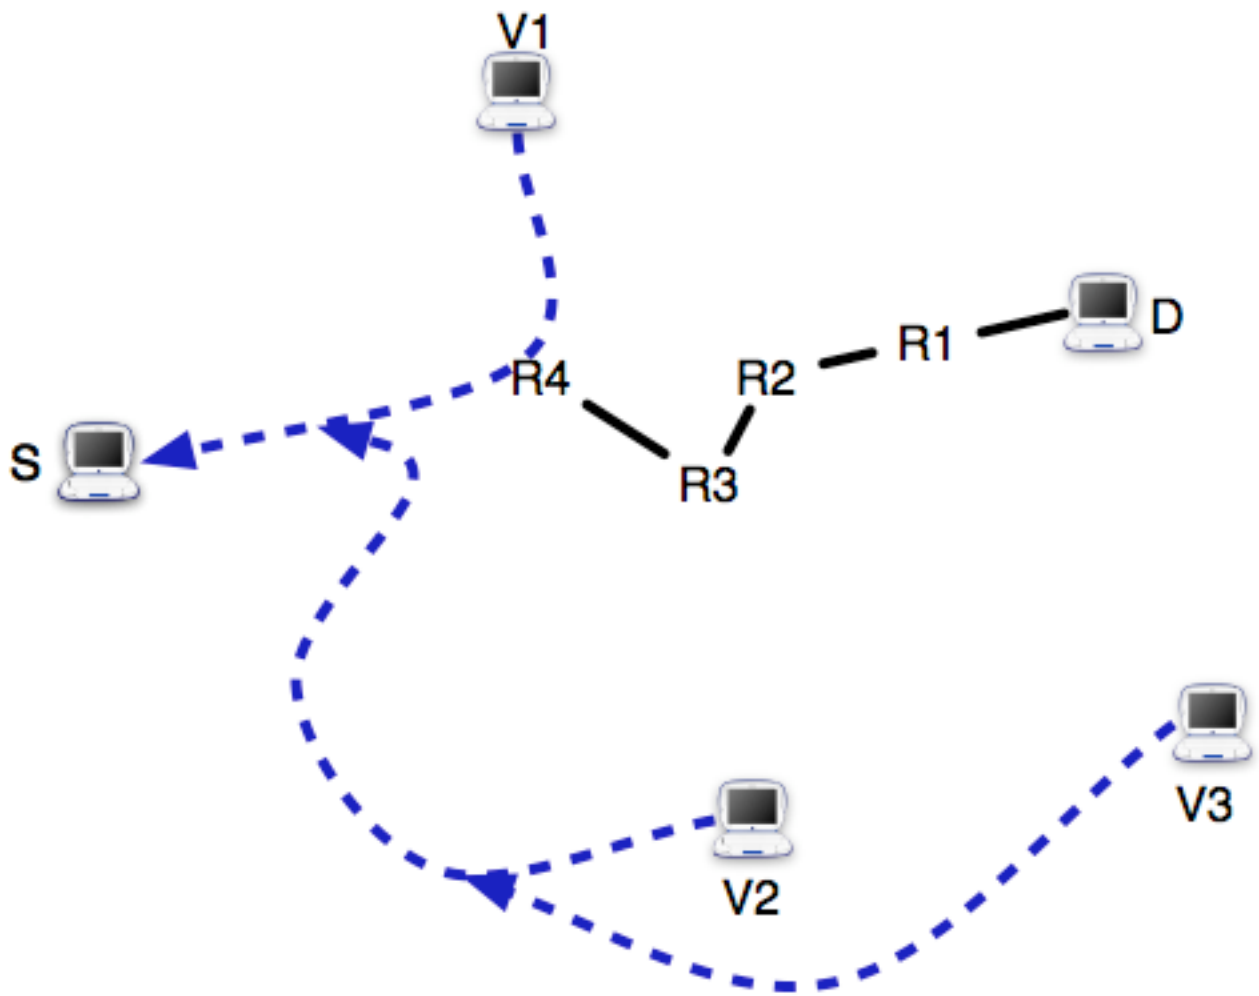
V1
D
R1
R2
R4
S
R3
V3
V2

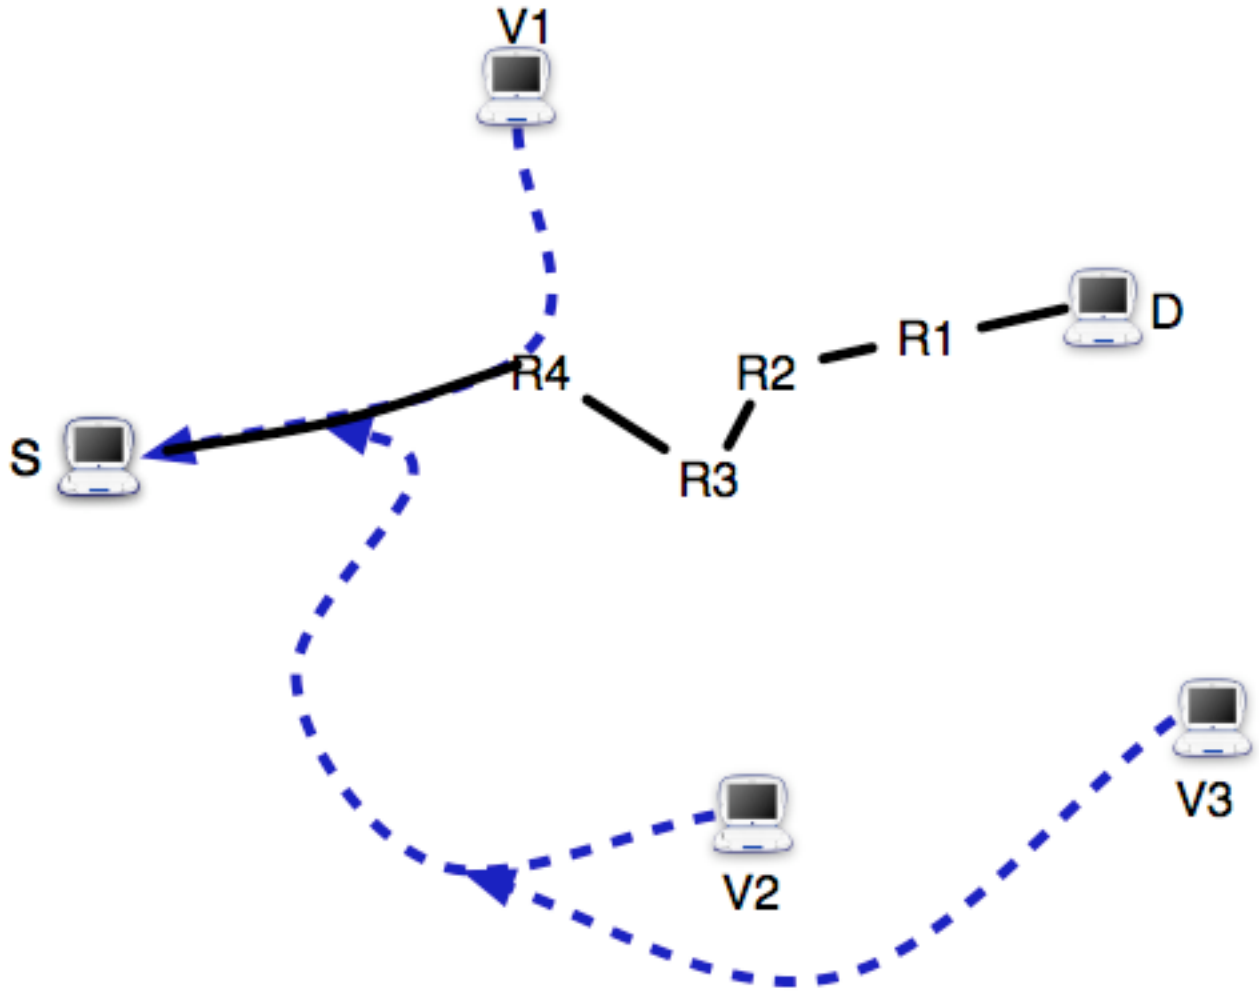


- Once we see a router on a known path, we know remainder

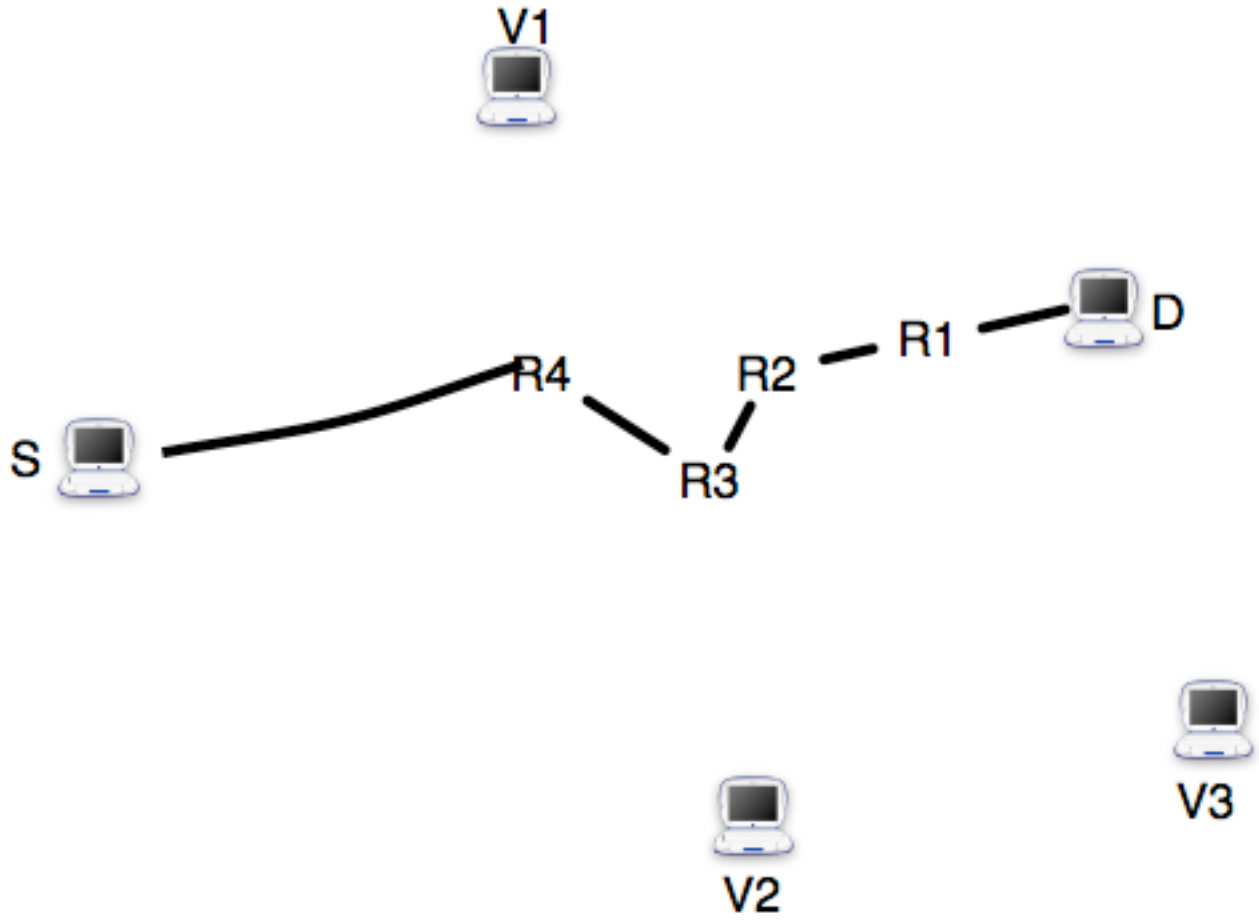


- Techniques combine to give us complete path
- We have additional techniques for inferring reverse hops

# Preliminary results

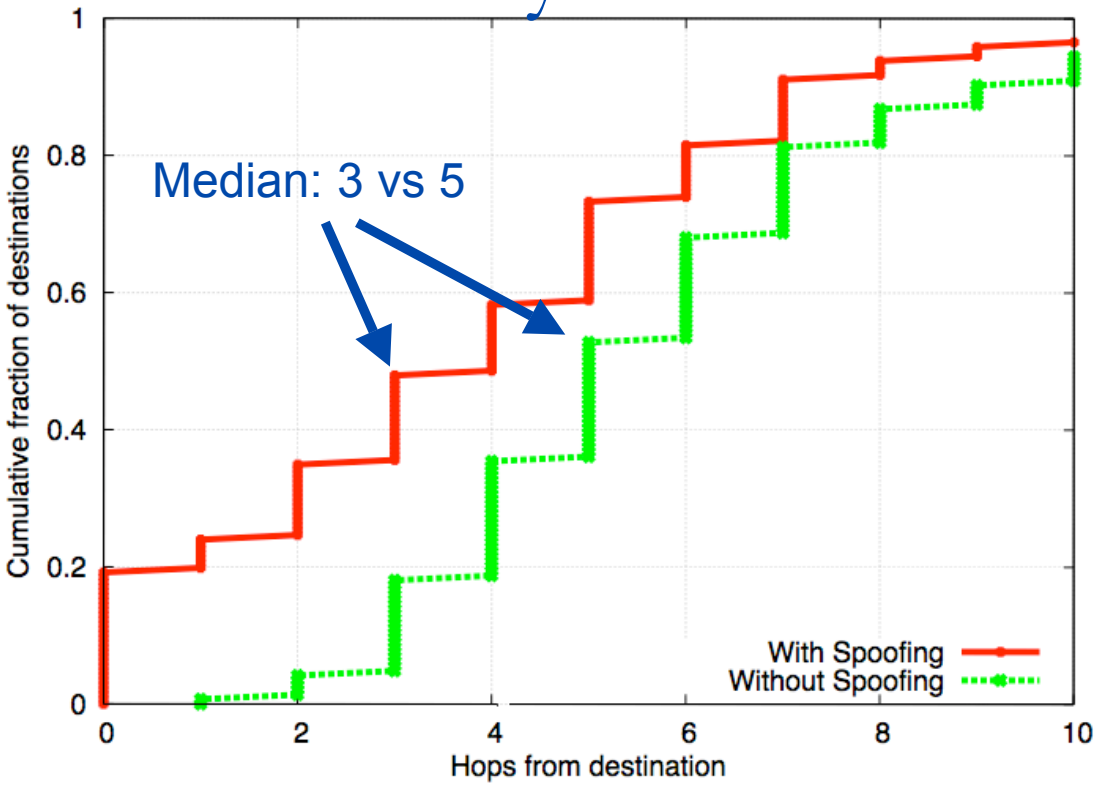


- Reverse paths from PL sites back to UW
- Measurements:
  - TR PL to UW
  - RR PL to UW
  - Spoofed RR as UW
  - Pick dst, exclude site
- How many hops back from dst need to be given before we can construct a complete path for rest of reverse TR?

- Spoofing gives a few extra hops to connect to measured paths
- End hosts like PL are a few extra hops from routers
- PL-PL measurements more likely to share paths (GREN)

# Reverse Path Summary

- Reverse path info can be very useful to systems
- Ongoing work on **reverse traceroute** and **one-way ping** for when we don't control destination
- Preliminary results here and in **Hubble** show techniques can work

- Limiting factors:
  - Restricted support for options
  - Current prober deployment
    - Need diverse paths back to our test sources
    - Need spoofing vantage points in diverse network locations
    - Any we can use?

# Measurement Work at UW

- Real Internet-scale measurement-based systems
  - Hubble - Monitoring black holes on the Internet
  - iPlane - Providing Internet path and path property predictions
- Ongoing work
  - Reverse path techniques
  - ***Massive software prober deployment***
  - Evaluating prober deployments

# Massive Prober Deployment

- Goal: on-demand probes from any prefix
- Talking with RIPE Science Group about 3 tier brain/ controller/ prober architecture
- Different classes of probers operating under standard controllers
  - Super probers - TTM, PlanetLab
  - Hardware probers - simple USB dongles
  - Software probers - next slide

# Software Probers: Incenting End-users

- Plan to develop software prober plugin
- Deploy in different vehicles that incent users to contribute measurements by providing benefit of measurements
  - BitTorrent client
  - Reliability-focused detouring - Firefox plugin
  - Apps built on iPlane predictions

# Measurement Work at UW

- Real Internet-scale measurement-based systems
  - Hubble - Monitoring black holes on the Internet
  - iPlane - Providing Internet path and path property predictions
- Ongoing work
  - Reverse path techniques
  - Massive software prober deployment
  - ***Evaluating prober deployments***

# Reviewers (properly) suspicious of PL

Actual (paraphrased) comments from reviews:

- "Needs evaluation of likely coverage of all paths in Internet given small size of PlanetLab"
- "Let me know how much of Internet is observable and suggest vantage points to improve coverage"
- "Oddities of Abilene are hard to reason about"
- "Including more text on limitations of PlanetLab"
- "Include discussion on how well you see this technique working in the global Internet."

# Assessing prober deployment

Previous work either focuses on:

- Measurements between vantage points
- Cumulative topology

Our focus:

- Paths to prefixes

# Goals

Techniques to help with:

- node selection for a system: # and which
- node deployment: where to place new nodes
- assessing how set of vantage points represents overall diversity of paths and how results of a study would vary with a different deployment

# Questions to answer

- Ideal is every end host.  How close is our data to that?  How much does spoofing help?
- Is PlanetLab limited primarily by # of sites or also by network locations of sites?
- How many vantage points do we "need?"
- How much does it help to select vantage points per target vs one set for all targets?
- How can we characterize which nodes are most useful to add?

# Measurement Work at UW

- Real Internet-scale measurement-based systems
  - Hubble - Monitoring black holes on the Internet
  - iPlane - Providing Internet path and path property predictions
- Ongoing work
  - Reverse path techniques
  - Massive software prober deployment
  - Evaluating prober deployments

Would love to talk about or collaborate on any of this.