Reverse Traceroute

Ethan Katz-Bassett, Harsha V. Madhyastha,
Vijay K. Adhikari, Colin Scott,
Justine Sherry, Peter van Wesep,
Arvind Krishnamurthy, Thomas Anderson

AIMS, February 2010

This work partially supported by Cisco, Google, NSF

Motivation: Google Wants Reverse Paths

"*The number one go-to tool is traceroute.*
Asymmetric paths are the number one plague.
The reverse path itself is completely invisible."

Richard Steenbergen, CTO, nLayer Communications
NANOG Network operators troubleshooting tutorial, 2009

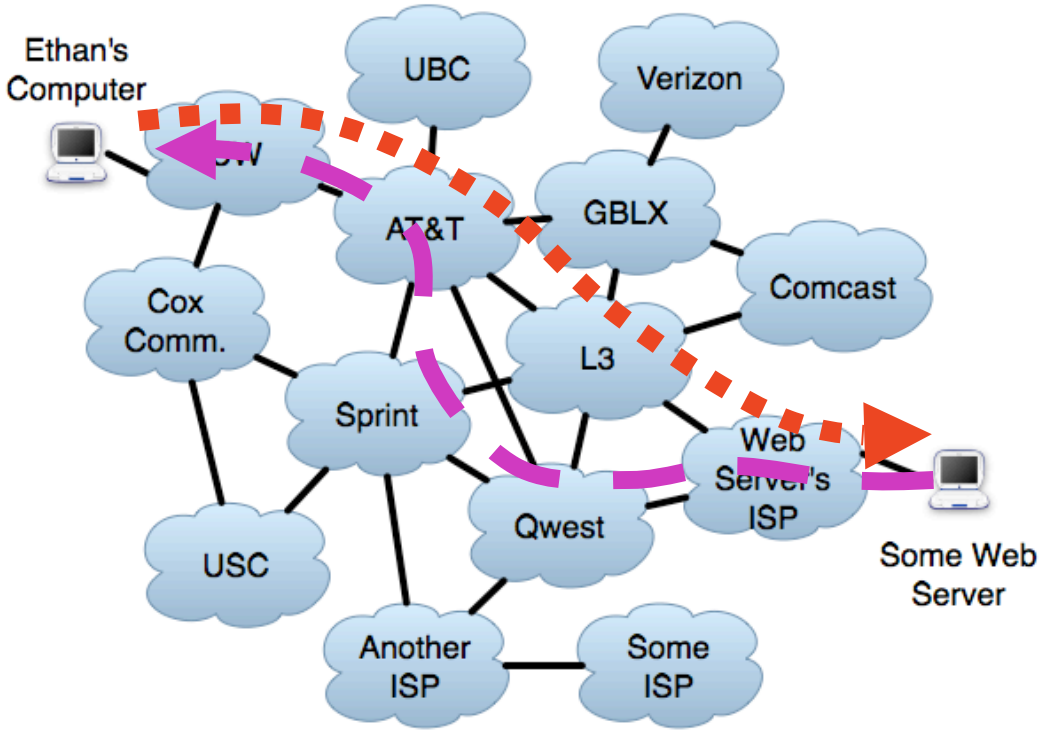

"*To more precisely troubleshoot problems, [Google] needs the ability to gather information about the reverse path back from clients to Google.*"

Google IMC paper, 2009

Goal: **Reverse traceroute**, without control of destination

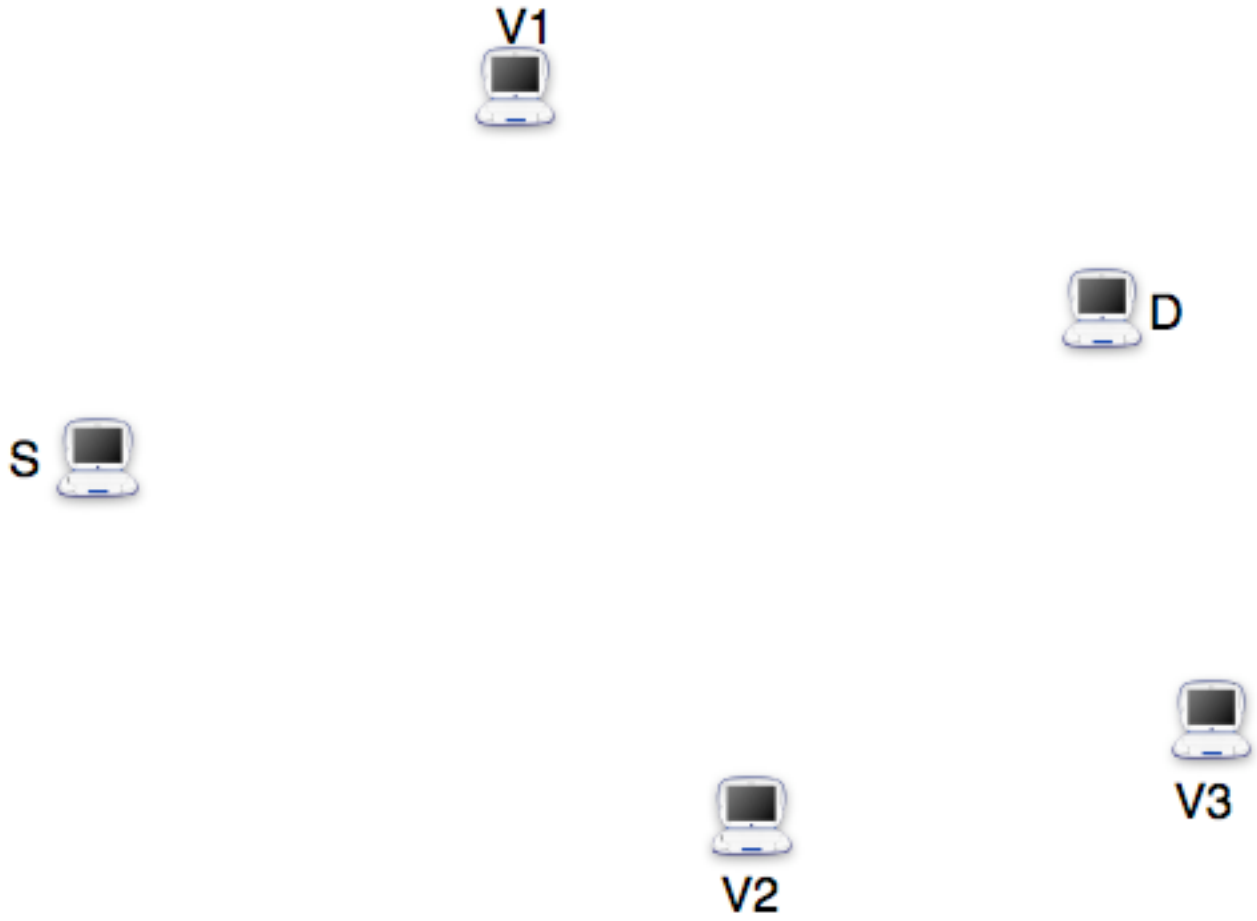

- Want reverse path from **D** back to **S**, but don't control **D**
- Set of vantage points around the world

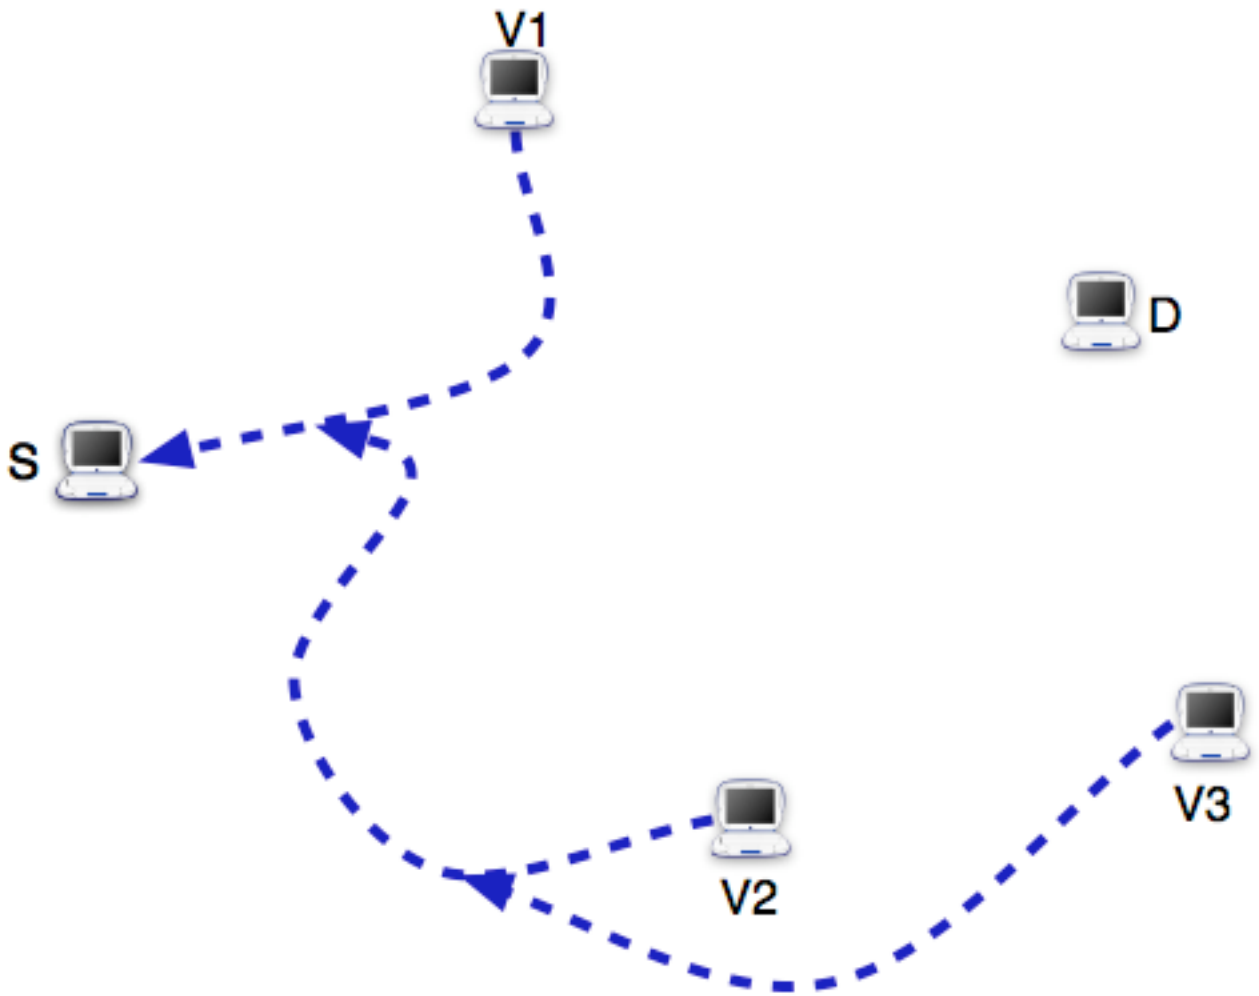

- Traceroute from all vantage points to **S**
- Gives atlas of paths to **S**; if we hit one, we know rest of path

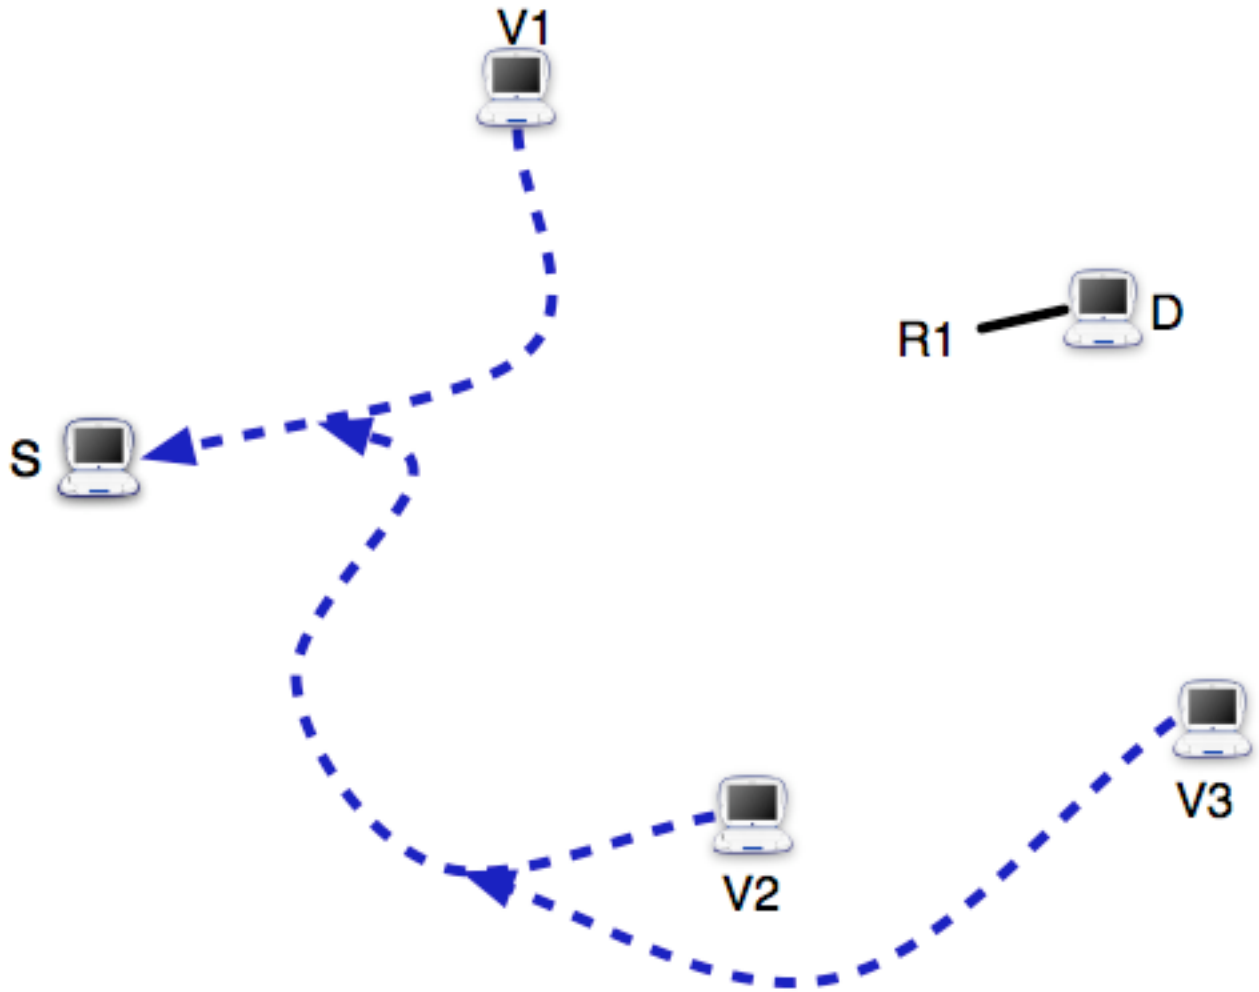

- Build back hop-by-hop to atlas (assumes destination-based routing)
- Set of techniques to measure hops using IP options

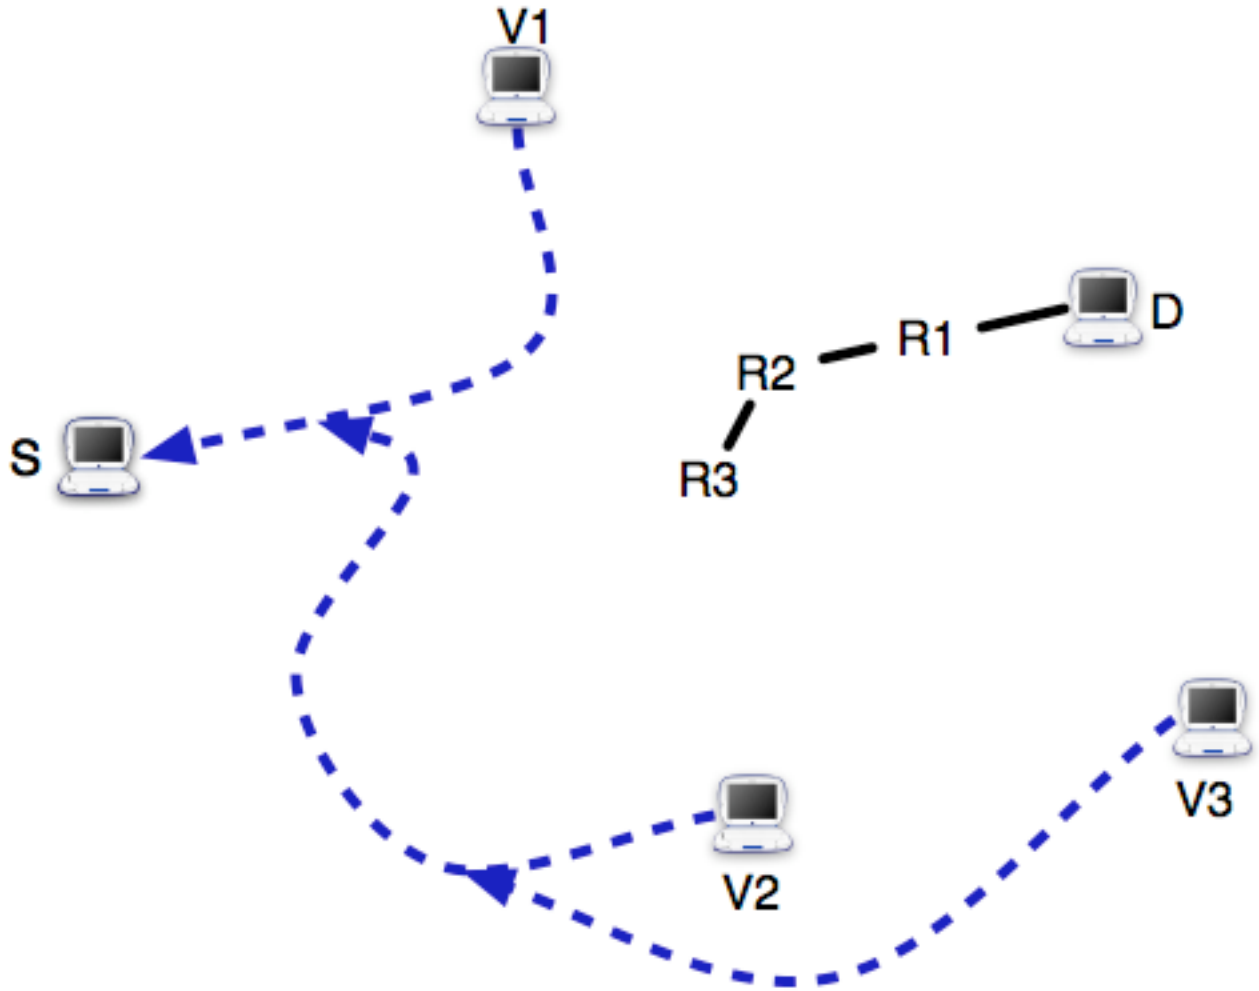

- Build back hop-by-hop to atlas (assumes destination-based routing)
- Set of techniques to measure hops using IP options

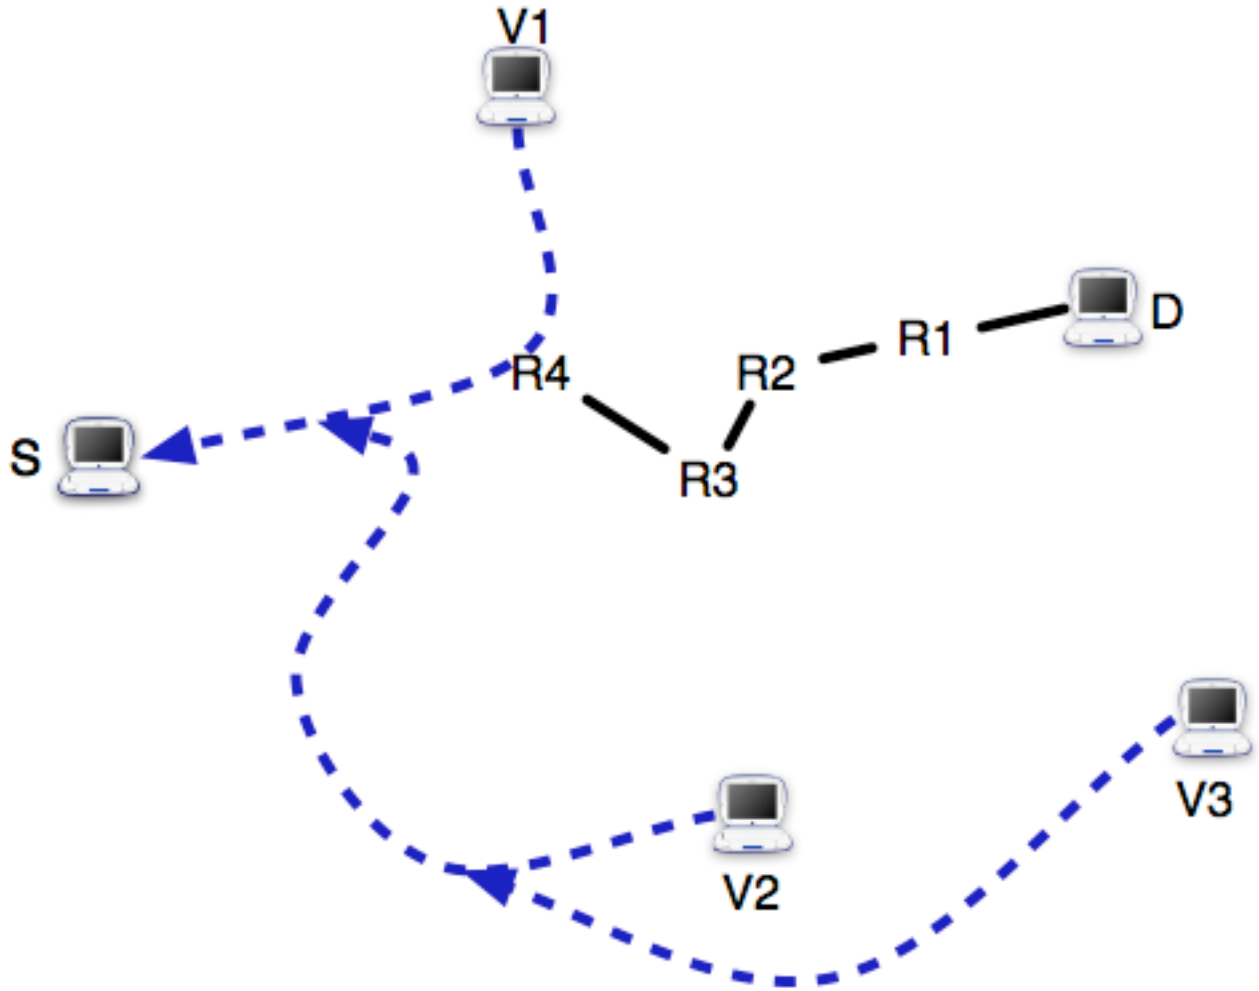

- Build back hop-by-hop to atlas (assumes destination-based routing)
- Set of techniques to measure hops using IP options

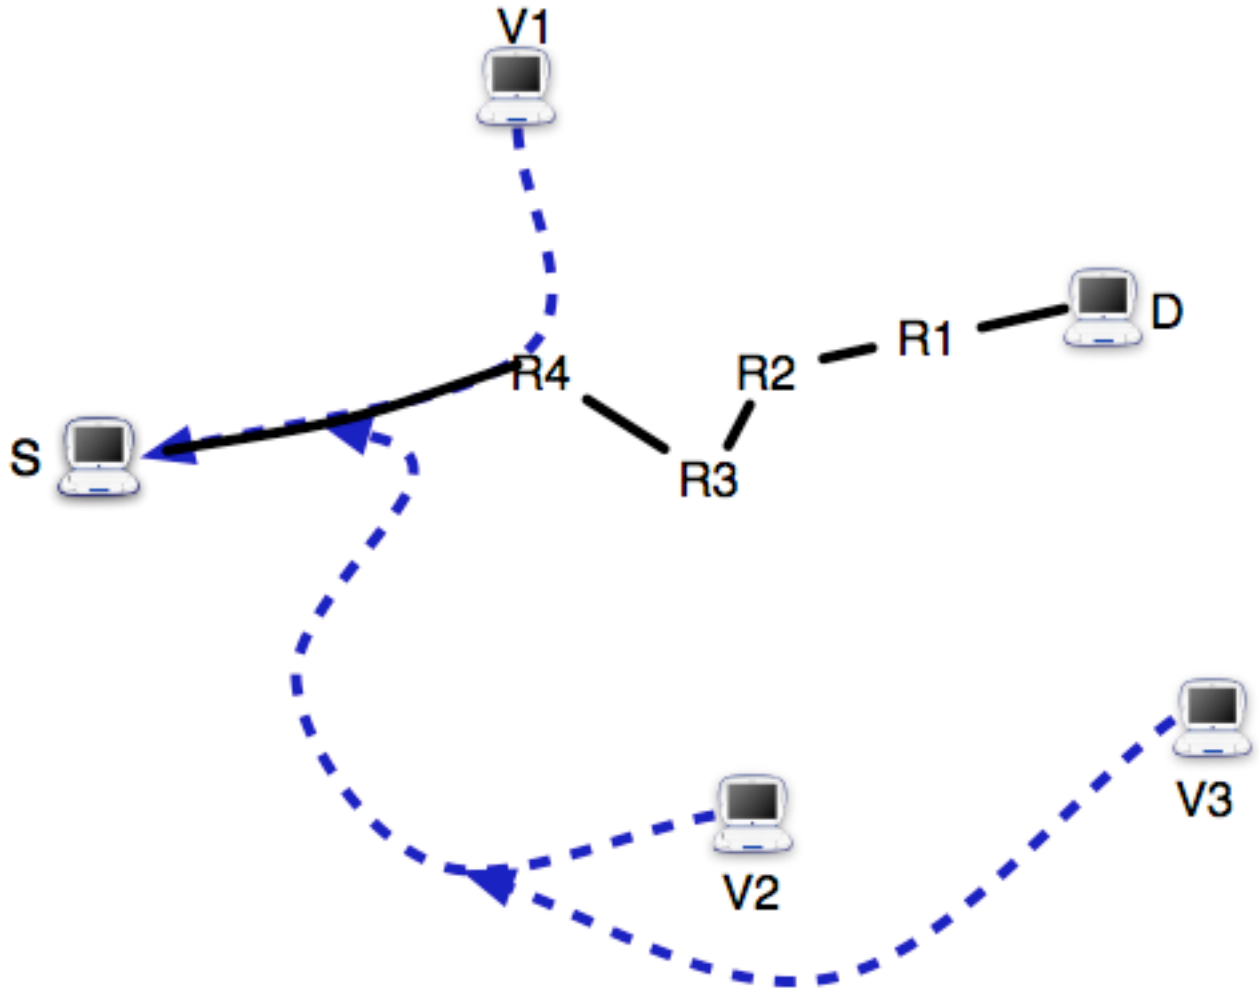

- Once we see a router on a known path, we know remainder

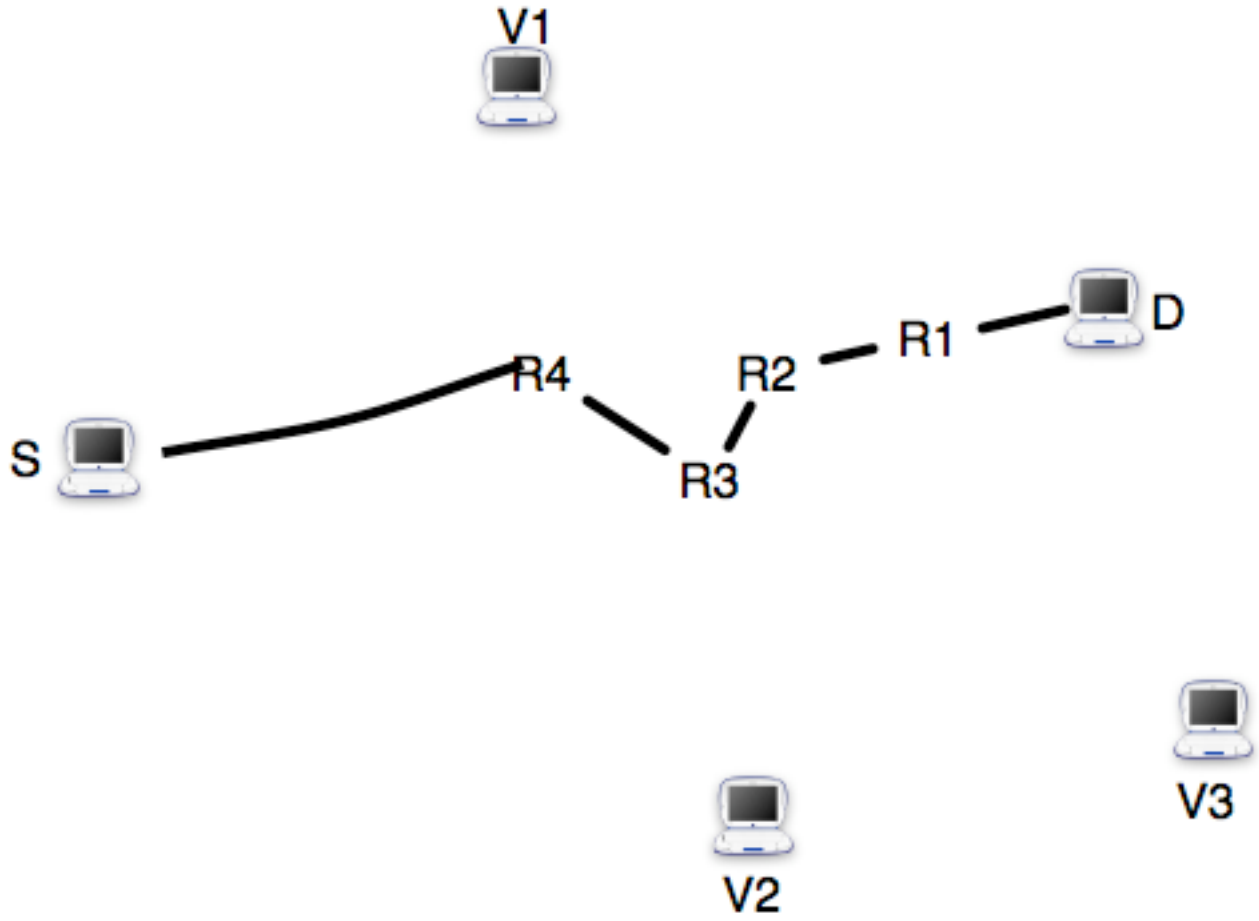

Techniques combine to give us complete path

Status of Project and This Talk

- Appearing in *NSDI 2010*
- http://revtr.cs.washington.edu
  - PlanetLab and MeasurementLab nodes
  - Measure paths from arbitrary IPs to PL nodes
  - Revising system to improve scalability, overhead
  - Plan to use Scamper (thanks Matthew!)
  - Then open system to let users measure to themselves
- This talk: applications to link latency and topology mapping
- NSDI paper: technique, accuracy, coverage

Motivation: Apps Want Link Latencies

- Traceroute/ping give round-trip time (RTT)
- … but many apps want one-way link latency
  - Peter's and Noa's geolocation talks yesterday
  - Path performance estimation (iPlane)
  - ISP comparison (Netdiff)
  - Troubleshooting poor performance

Measuring Link Latency

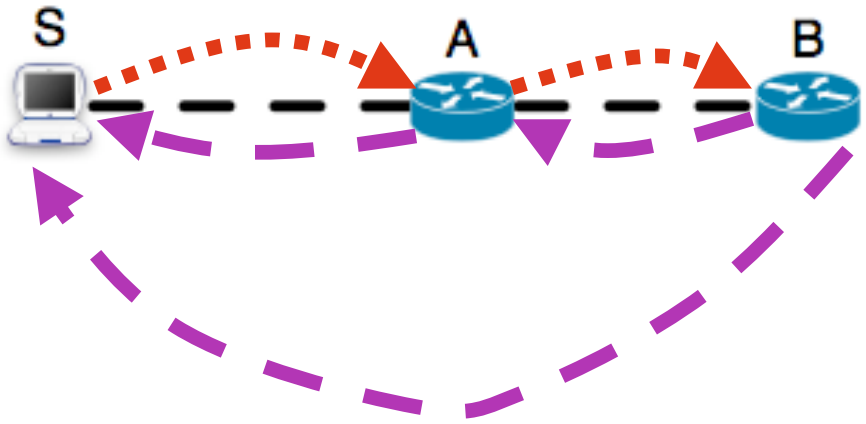

- Traditional approach:

  Delay(A,B) = (RTT(S,B) – RTT(S,A)) / 2
- Asymmetry skews link latency inferred from traceroutes

Reverse Traceroute Detects Symmetry

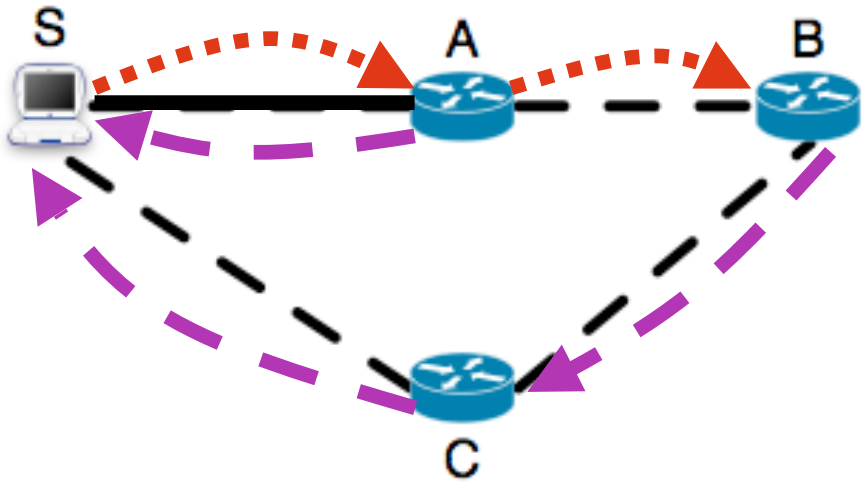

- Reverse traceroute identifies symmetric traversal
  - Identify cases when RTT difference is accurate
  - Many links traversed symmetrically
    from some vantage points, not others

Reverse TR Constrains Link Latencies

RTT(A,B)=Delay(S,A) + Delay(A,B)
+ Delay(B,C) + Delay(C,S)

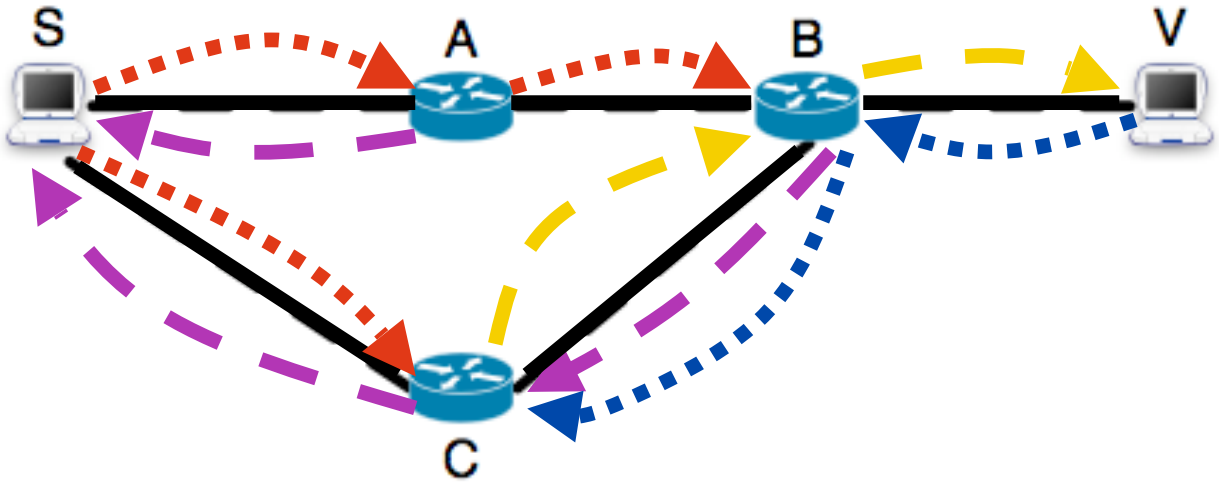

- Build up system of constraints on link latencies of all intermediate hops
  - Traceroute and reverse traceroute to all hops
  - RTT = Forward links + Reverse links
- Open issues: Treat unbound links as segment? MPLS?

Measuring Sprint's Link Latencies

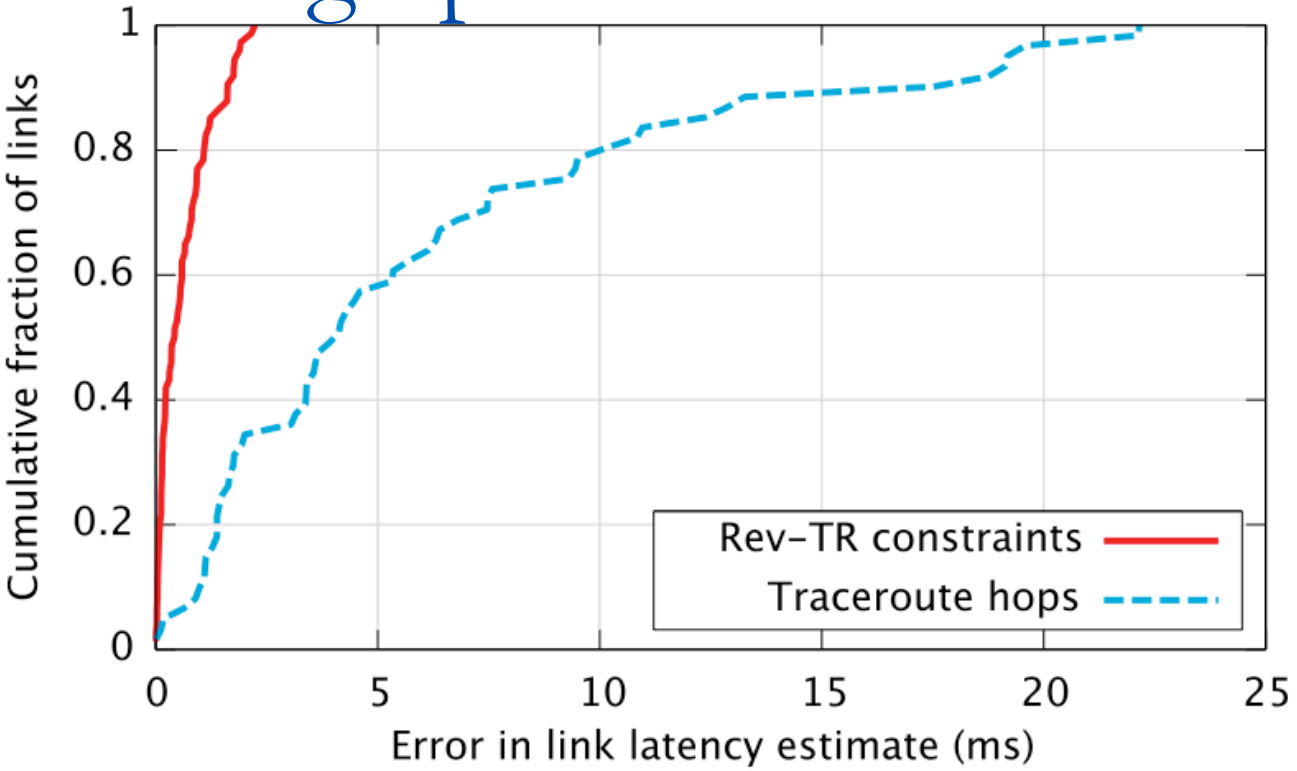

- We see 79 of Sprint's 89 inter-PoP links, whereas traceroute only sees 61
- Median (0.4ms), mean (0.6ms), worst case (2.2ms) error all 10*x* better than with traditional approach

Motivation: Ricardo Wants Peering Links

"New inference techniques are needed to capture or estimate peer links" Ricardo Oliveira, SIGMETRICS '08

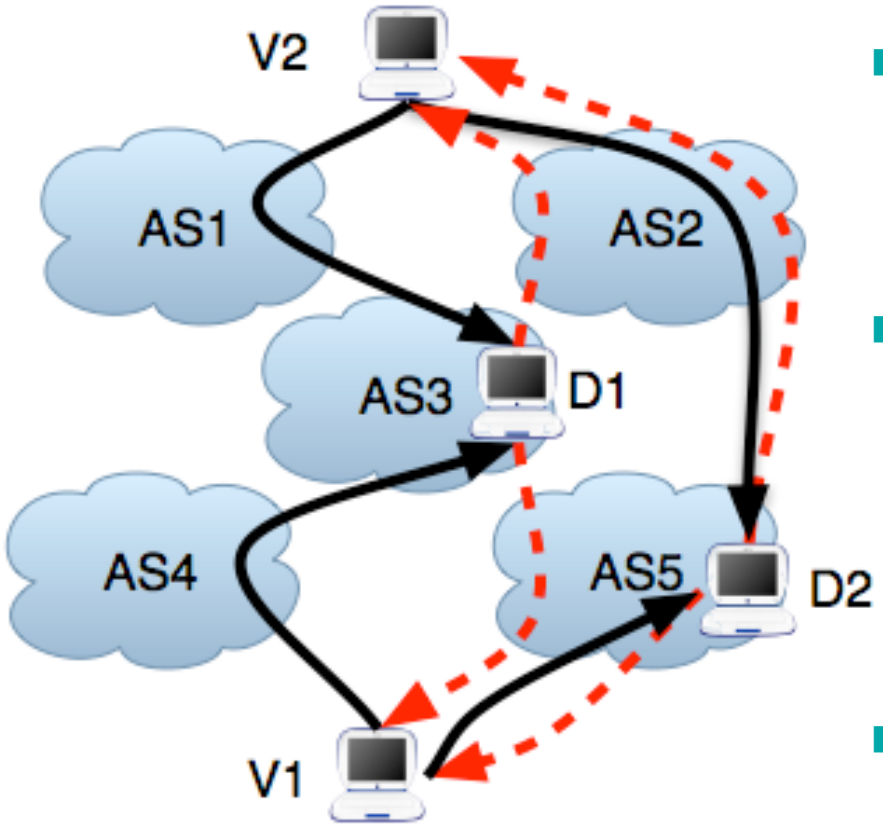

AS3 peers w/ other ASes shown

- Only AS and its customers see/use its peer links
  - No path will traverse > 1
- Trad. methods miss links
  - V1 and V2 can't traceroute AS3-AS2, AS3-AS5
  - Most peer links invisible to RouteViews, RIS
- Reverse traceroute sees AS3-AS2 and AS3-AS5

How many extra links do we see?

- Considered just peering links at IXPs
- Baseline:
  58,534 IXP links on 51,832 AS pairs
  - ***IXPs: Mapped?*** *[B. Augustin, B. Krishnamurthy, and W. Willinger. IMC '09]*
  - Most exhaustive study of IXPs yet
  - Traceroutes from 1000s of hosts, source routing
- Reverse traceroute enriches the study:
  9096 additional IXP links (16%)
  5057 additional distinct AS pairs (10%)
  1910 of those also not in iPlane or UCLA data

Reverse Traceroute Vs Ono

Complementary approaches to measuring more routes

Reverse traceroute

- Use existing VPs to measure any destination
- Relies on IP options, spoofing
- (Future) On-demand measurements for all
- Paths from arbitrary locations (used in apps)
- Scalable? (I built it)

Ono

- Use P2P (need / have peers everywhere)
- Relies on standard traceroute
- On-demand? Arbitrary targets? For all?
- Paths reflect actual end-user traffic, edge
- Scalable (Dave built it)

Conclusion and Questions for You

- Traceroute is very useful, but can't provide reverse path
- Our reverse traceroute system addresses limitation, providing complementary information
- Gives most hops as if you issued traceroute from remote site
- Useful in wide range of situations, including:
  - Accurately measuring link latencies
  - Exposing "hidden" topology
- What should we measure?
- Ideas on more vantage points?