

The Record Route Option Is an Option!

Brian J Goodchild,

Yi-Ching Chiu, Rob Hansen, Haonan Lu, Matt Calder,
Matthew Luckie, Wyatt Lloyd,
David Choffnes, Ethan Katz-Bassett

•Rutgers--Camden • USC • Columbia • Northeastern • Microsoft • Waikato •

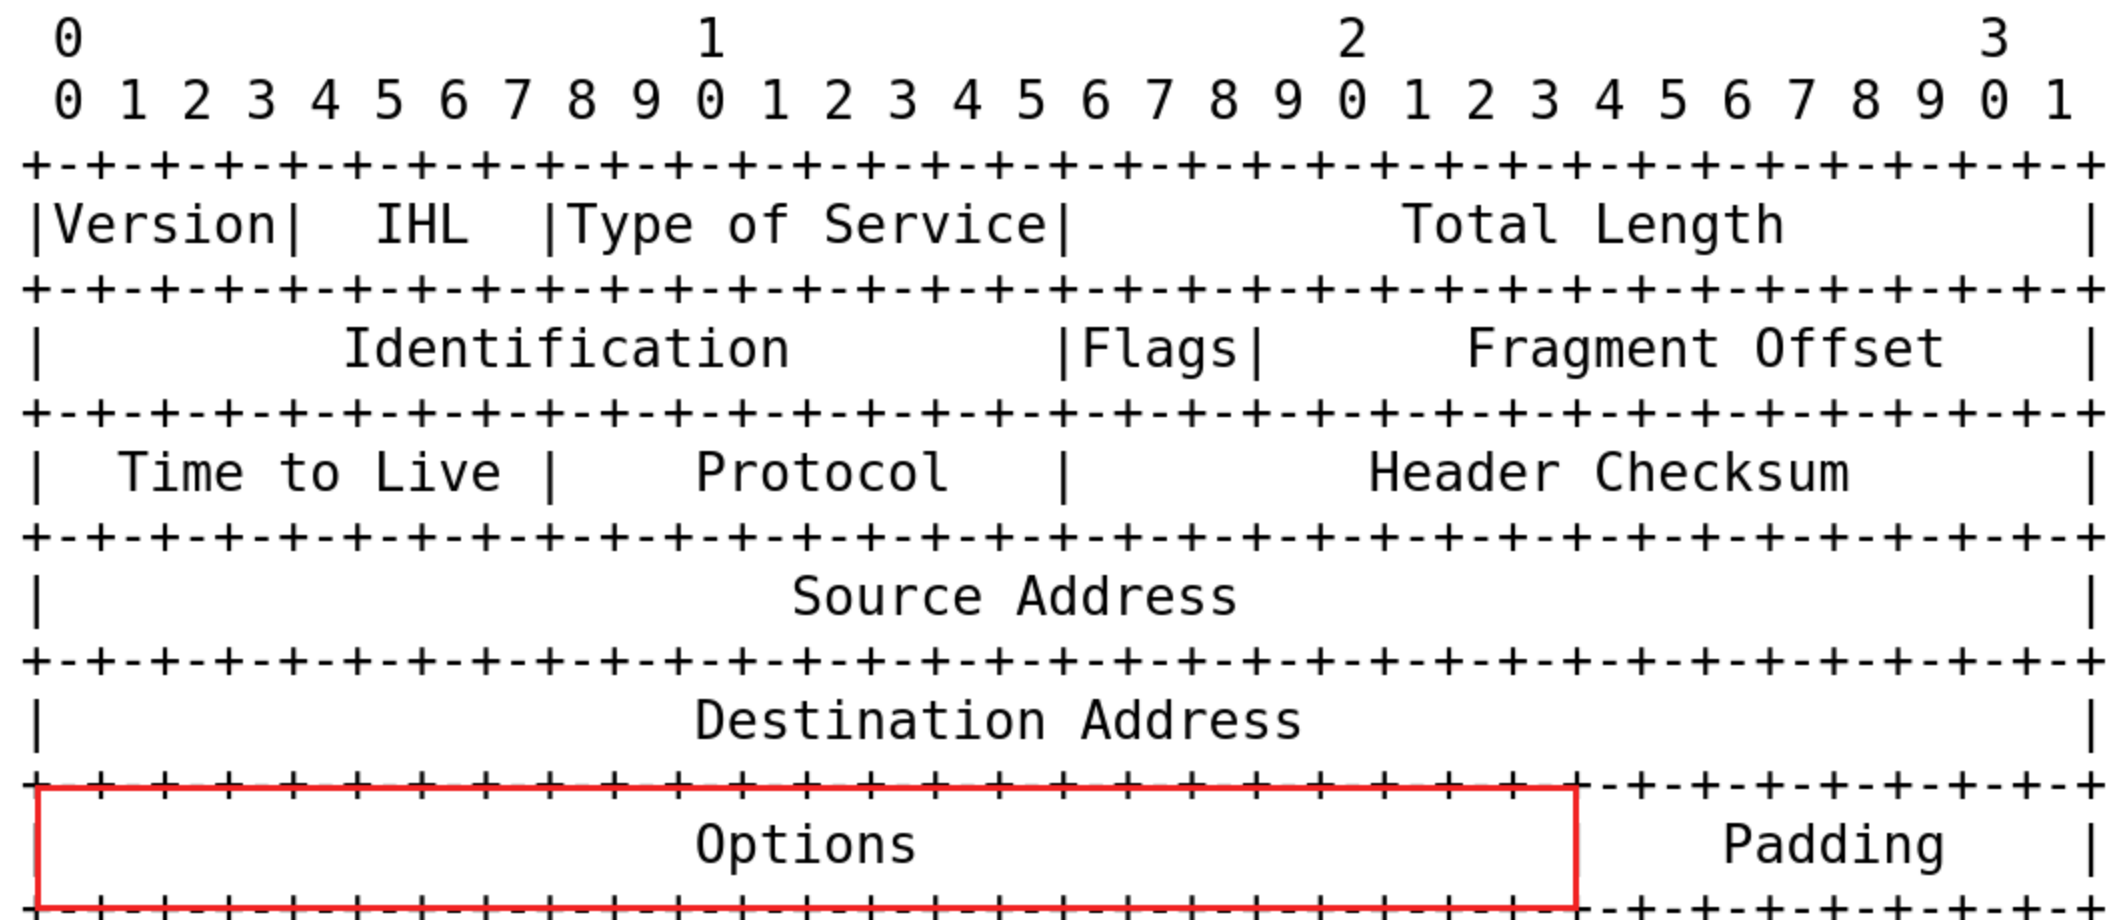
IMC November 2017

Observation & Motivation

- Measuring Internet routing is difficult.
- Limited set of tools for measuring routes.
 - Traceroute: measures $s \rightarrow d$ paths.
- Many aspects of routing remain opaque.
- Any new tools could help.

Record Route: IP Option

- Standardized IP Option.
- Allocates extra space in packet header.
- Instruct routers to record IP address in header.

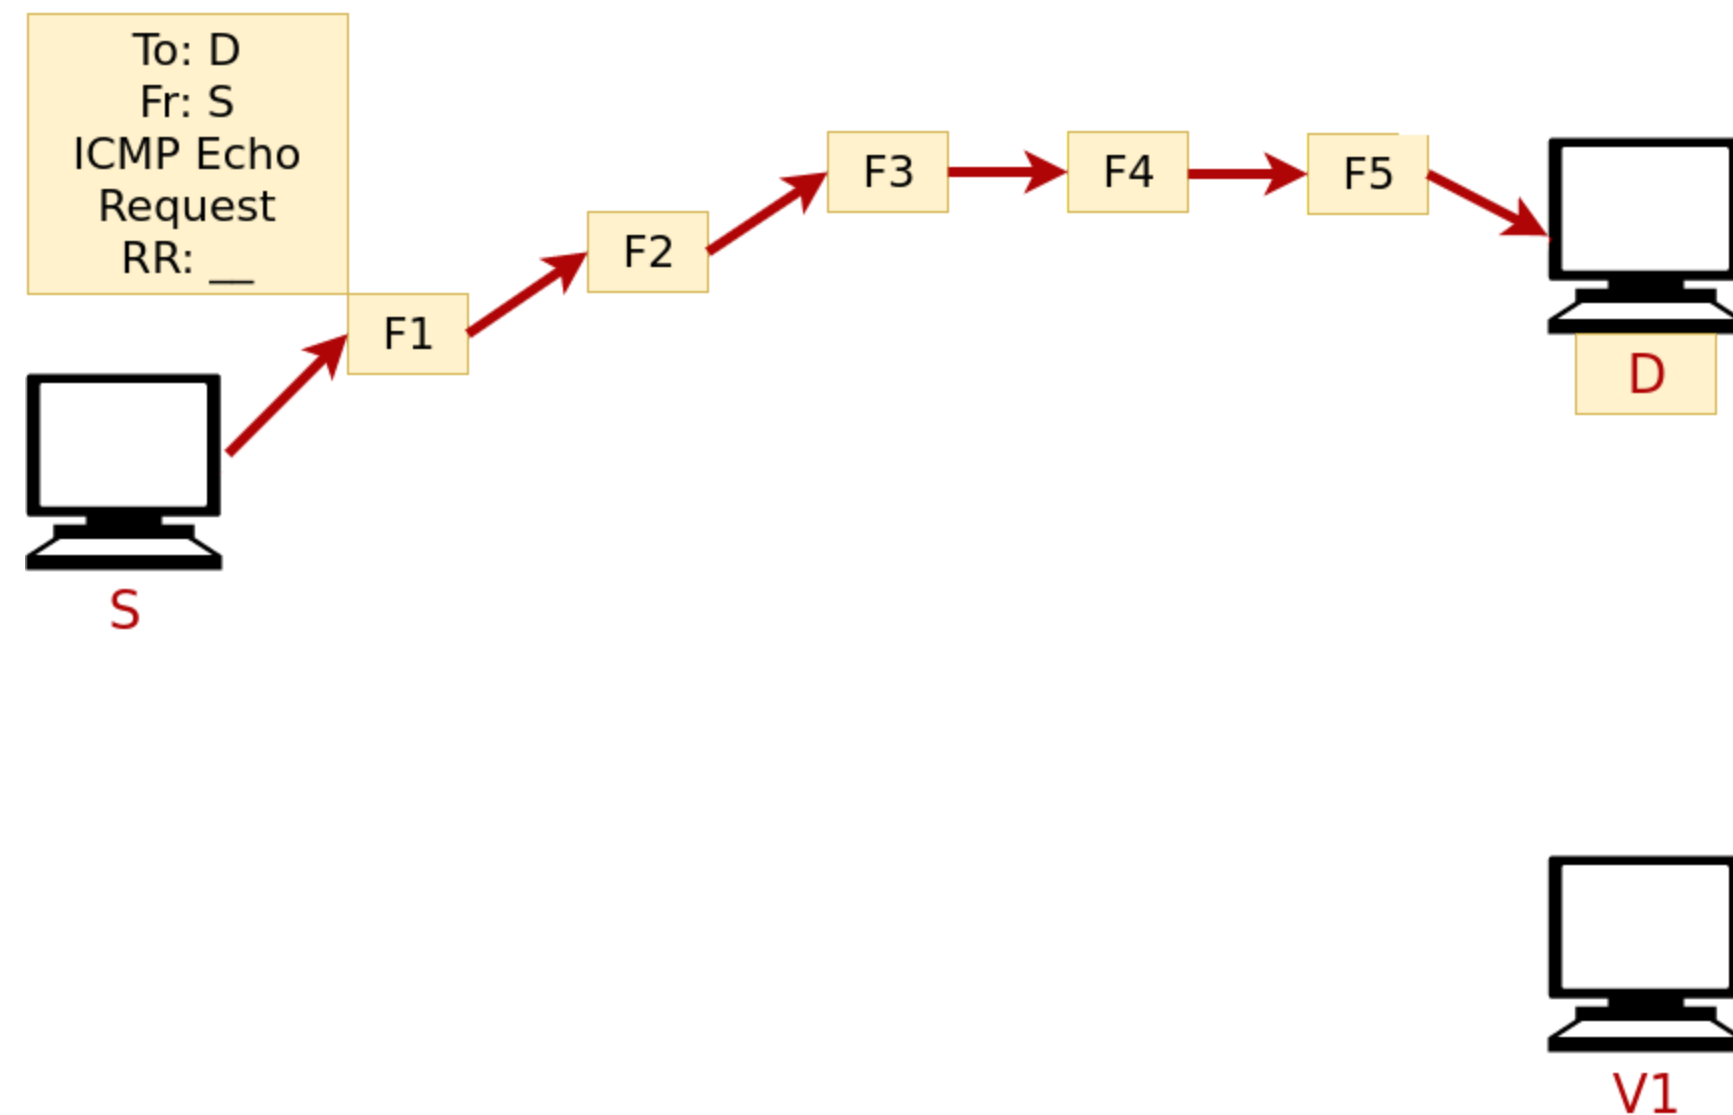


Record Route: advantages

- Returns traceroute-like path information.
- Information complementary to traceroute.
- Can record up to **9 hops** in single packet.
- Demonstrated Uses:
 - *Discarte (topology discovery)*
(Sherwood, et al; SIGCOMM 2008)
 - *Reverse Traceroute*
(Katz-Bassett, et al; NSDI 2010)
 - *Measuring Networks Using IP Options*
(Marchetta, et al; IEEE Network 2017)

Record Route: example

- Suppose path is 6 hops long.
- \Rightarrow Record Route will capture entire path in a single packet:
F1 \rightarrow F2 \rightarrow F3 \rightarrow F4 \rightarrow F5 \rightarrow D
- Room for 3 more in header.
- Question:
What happens to the packet when D responds?



Are IP Options supported on the Internet?



Ripe Atlas

to me ▾

Sep 12



> I am working on a study involving IP Options. I was wondering if it is
> currently possible to send pings with IPv4 options enabled?

The Atlas probe code current does not support IPv4 options.

IPv4 options are believed to be widely filtered and have not seen much use in the past few decades. For this reason there are currently no plans to support them.

"*IP Options Are Not An Option*"

- 2005 Technical Report from Berkeley; Fonseca, et al.
- **Results:**
 - ~50% of paths between PlanetLab sites drop options packets.
- **Their conclusion:**
 - Not universally supported.
 - Not an option for protocol extensibility
- **Popular interpretation:**
 - IP Options are generally useless.

"*IP Options Are Not An Option*"

- 2005 Technical Report from Berkeley; Fonseca, et al.
- **Results:**
 - ~90% of drops occurred in a handful of source and destination ASes.
 - ⇒ Most ASes do not drop options packets.
- **Our Interpretation:**
 - Options could be useful for measurement.
 - No need for universal support.

Goals

*To investigate the level of support for
Record Route in the wild.*

*To reassess the suitability of the Option
for use in **Internet measurement.***

Our Questions

- Do destinations *respond* to RR?
- Are destinations *reachable* within the 9 hop limit?
- Has reachability changed over time?

Results: Do destinations respond to RR?

Vantage Points: 55 PlanetLab and 86 M-Lab sites.

	Number	Percent	
One destination per BGP prefix	510,305	100%	• <i>Spread across the globe.</i>

Results: Do destinations respond to RR?

Vantage Points: 55 PlanetLab and 86 M-Lab sites.

	Number	Percent	
One destination per BGP prefix	510,305	100%	
Ping-Responsive	394,644	74%	<i>• Destinations that respond to regular ping (no RR).</i>

Results: Do destinations respond to RR?

Vantage Points: 55 PlanetLab and 86 M-Lab sites.

	Number	Percent	
One destination per BGP prefix	510,305	100%	
Ping-Responsive	394,644	74%	
RR-Responsive	296,734	58%	<i>• Destinations that respond to RR ping sent from at least one VP.</i>

75% of Ping-Responsive destinations are RR-Responsive!

Our Questions

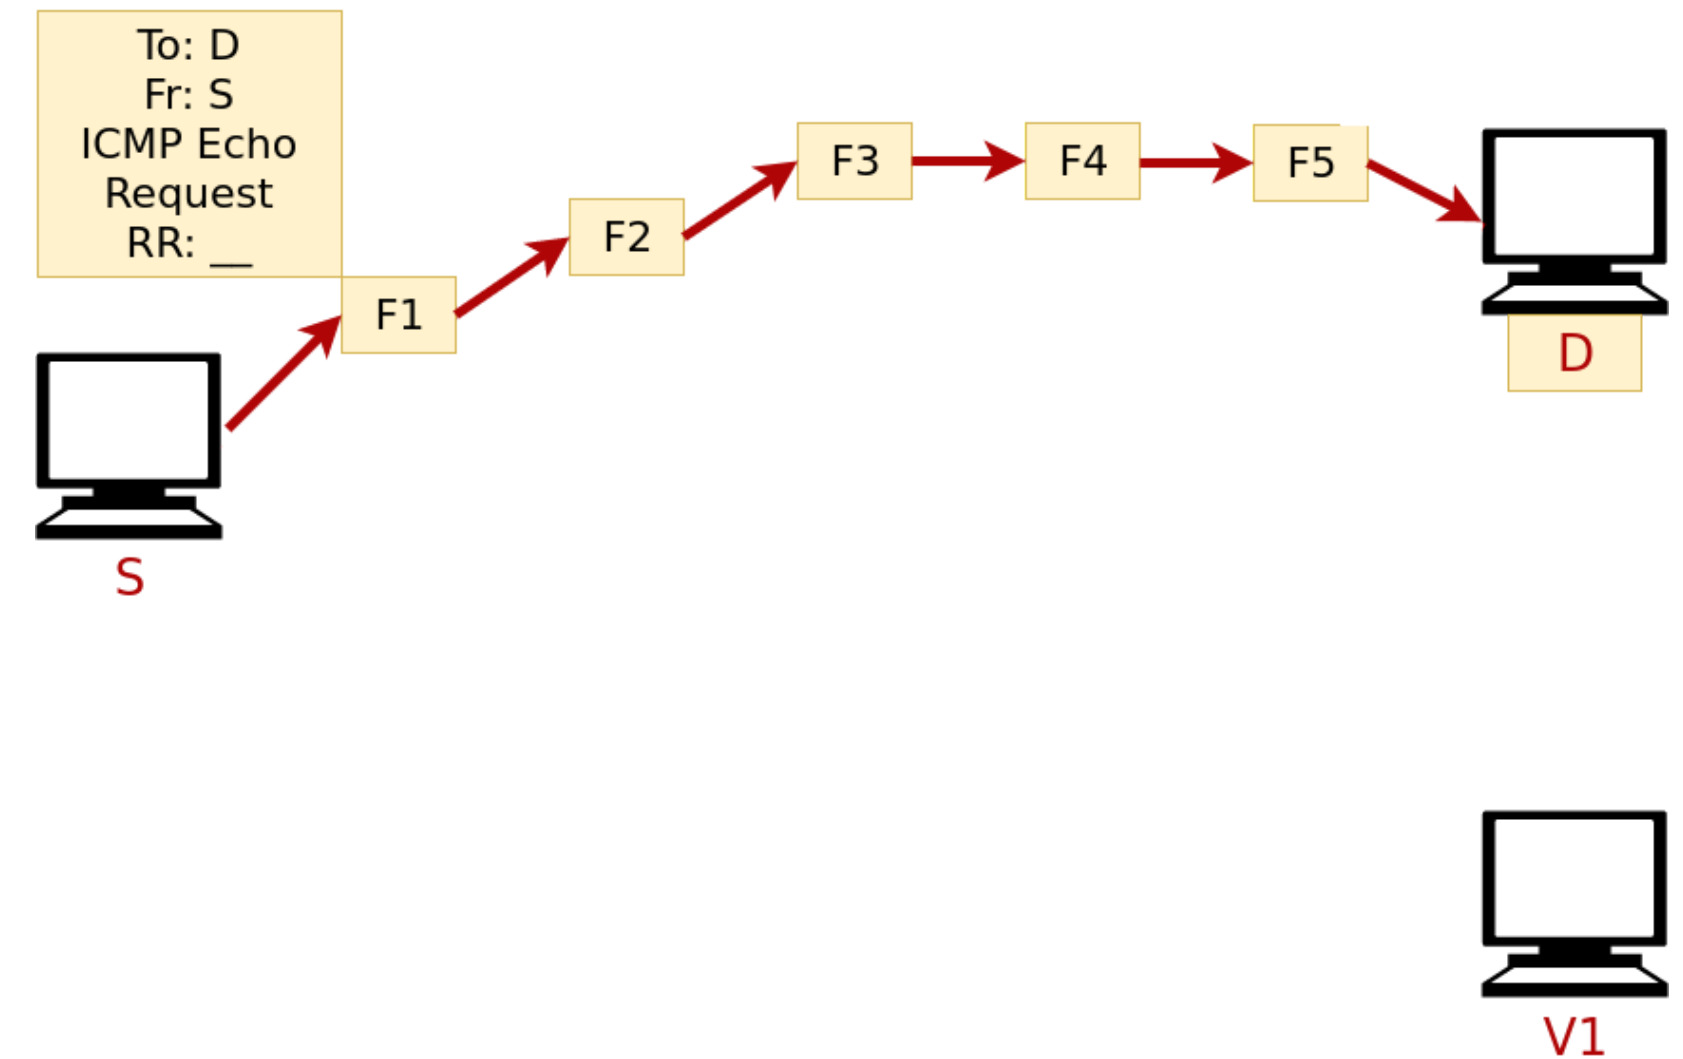
- Do destinations *respond* to RR? **yes!**
- Are destinations *reachable* within the 9 hop limit?
- Has reachability changed over time?

Are destinations reachable within 9 hop limit?

- Largely, yes.
 - 66% of RR-Responsive destinations within 9 hops of closest VP.
- Even better: Are destinations reachable within **8 hops**?

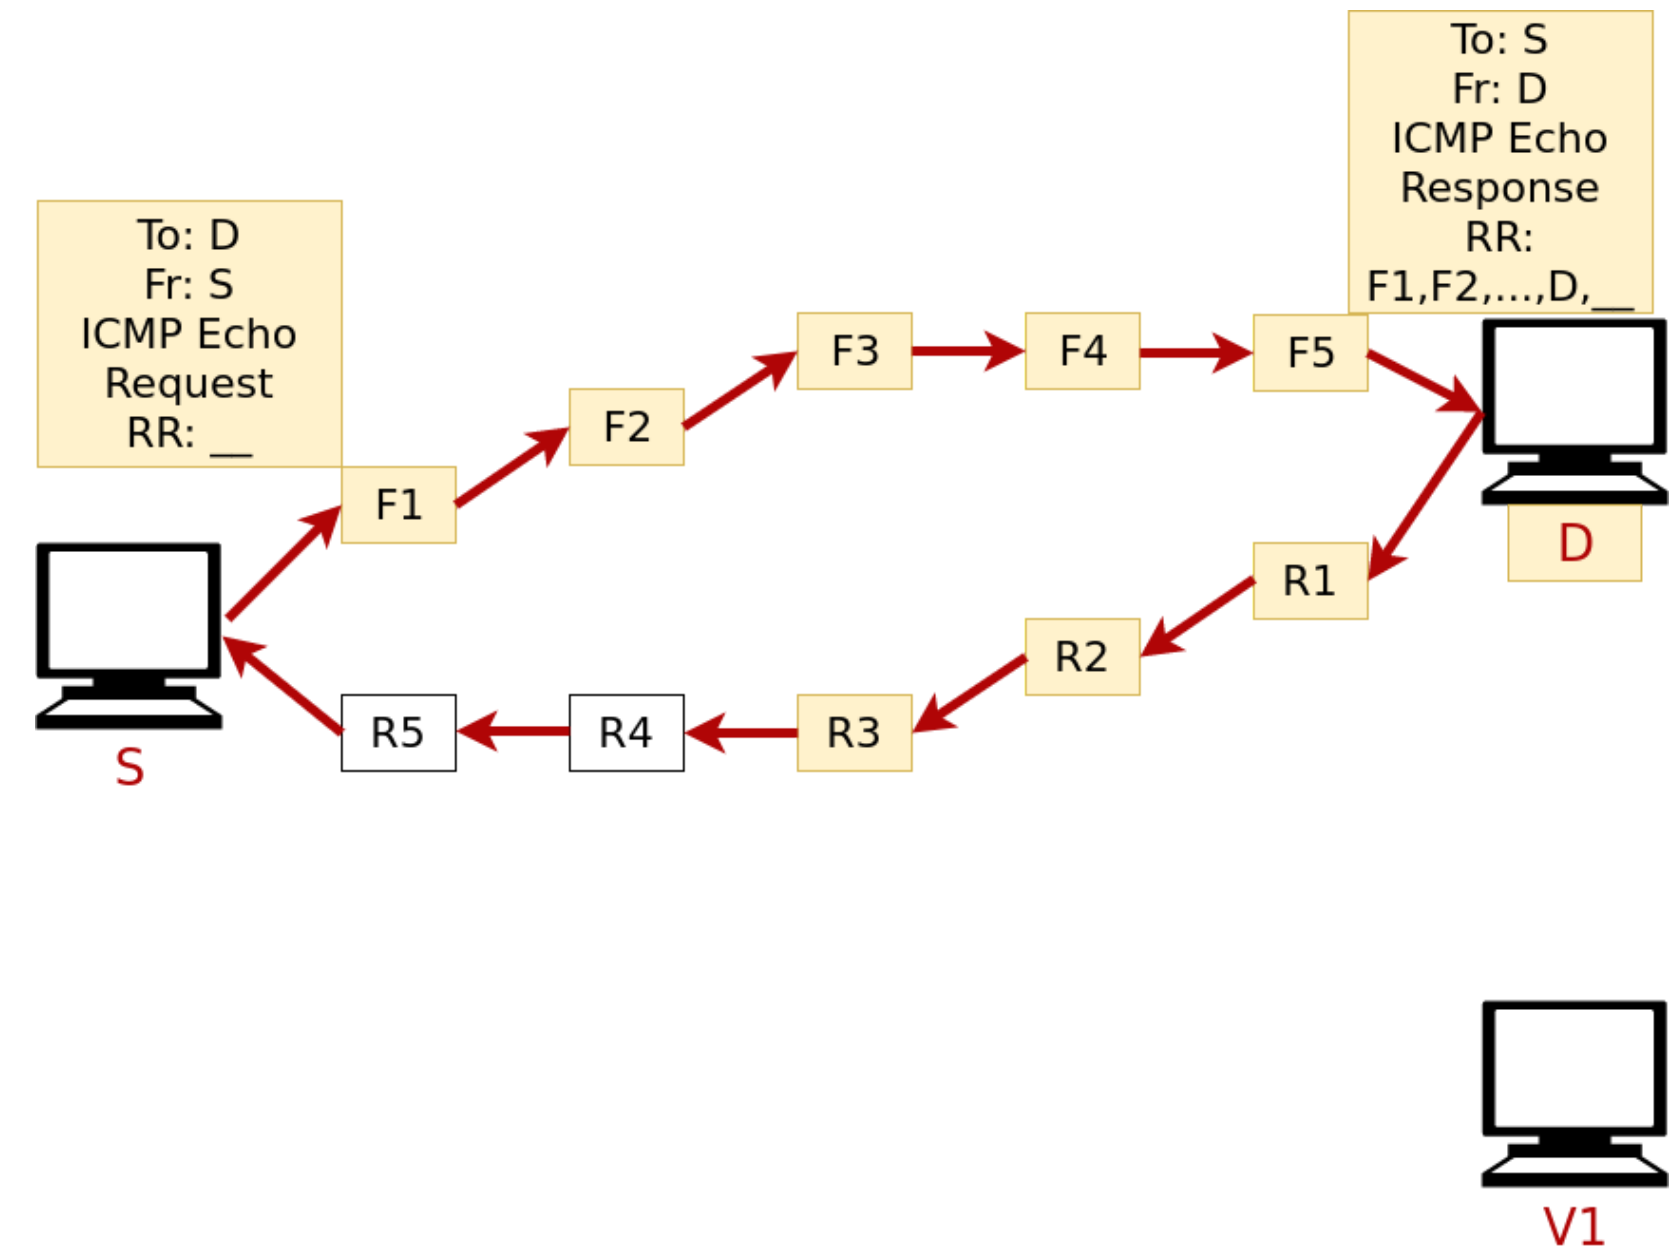
Reverse Traceroute: Source within 8 hops

- What happens to the packet when D responds?



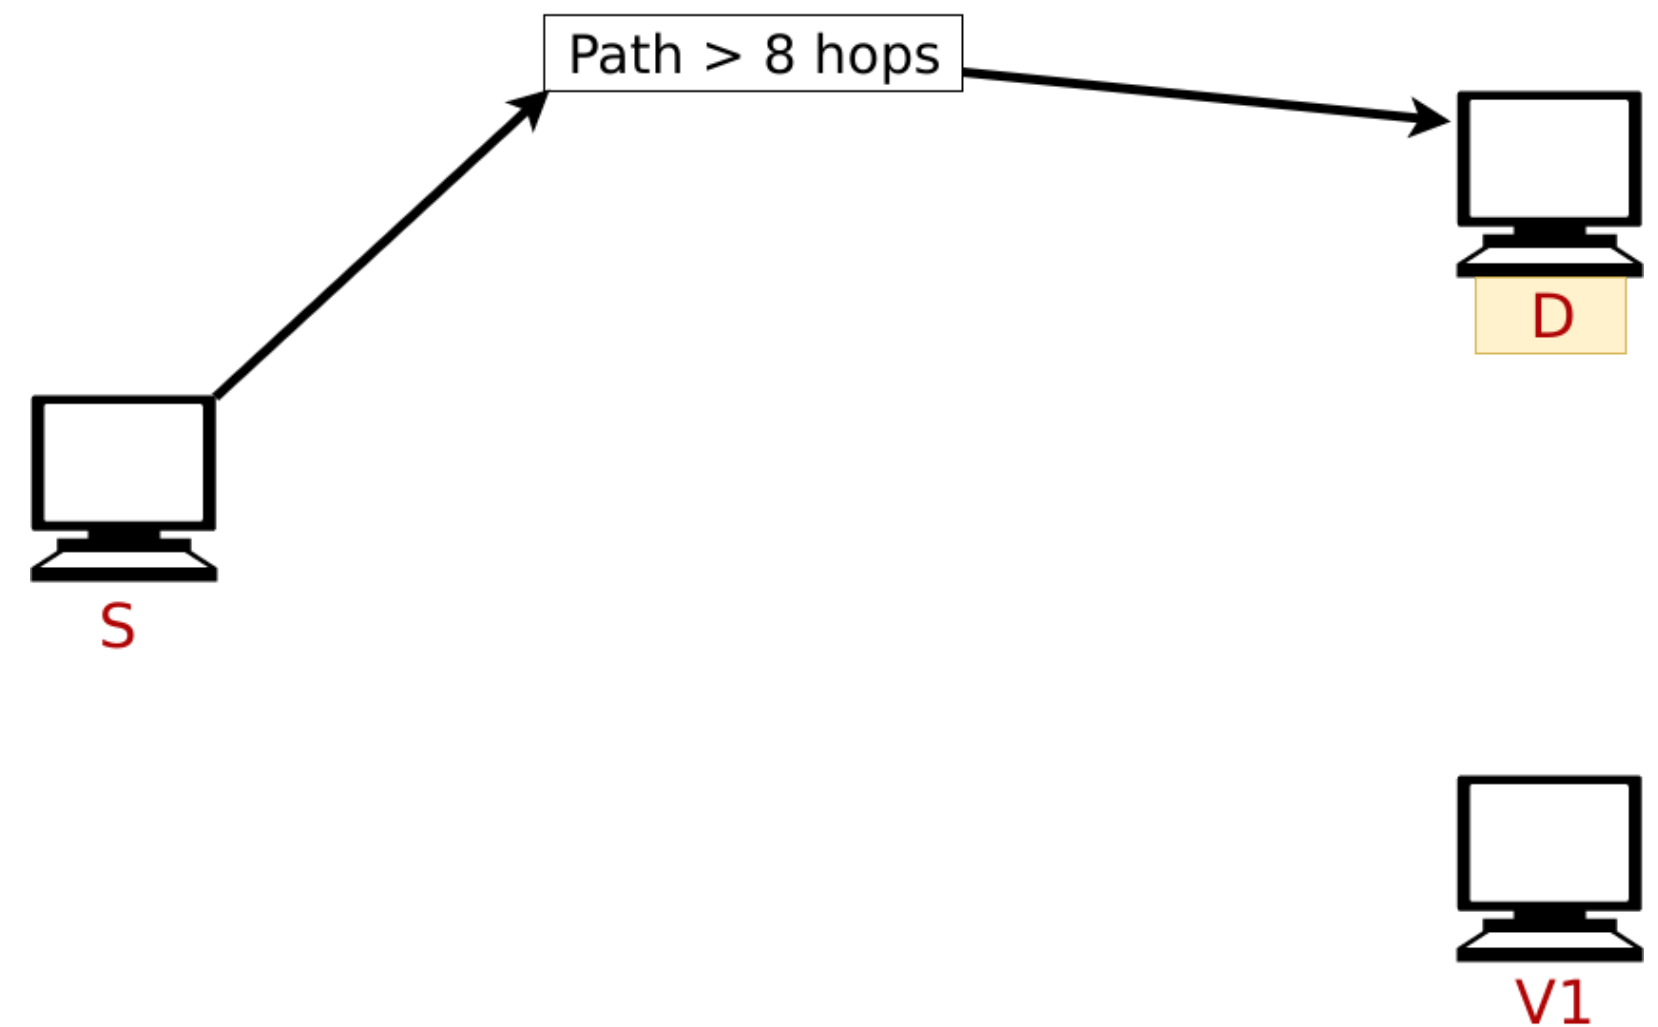
Reverse Traceroute: Source within 8 hops

- D's response still contains Record Route!
- Keeps recording hops on reverse path.
- **R1→R2→R3**
- Can stitch path together.



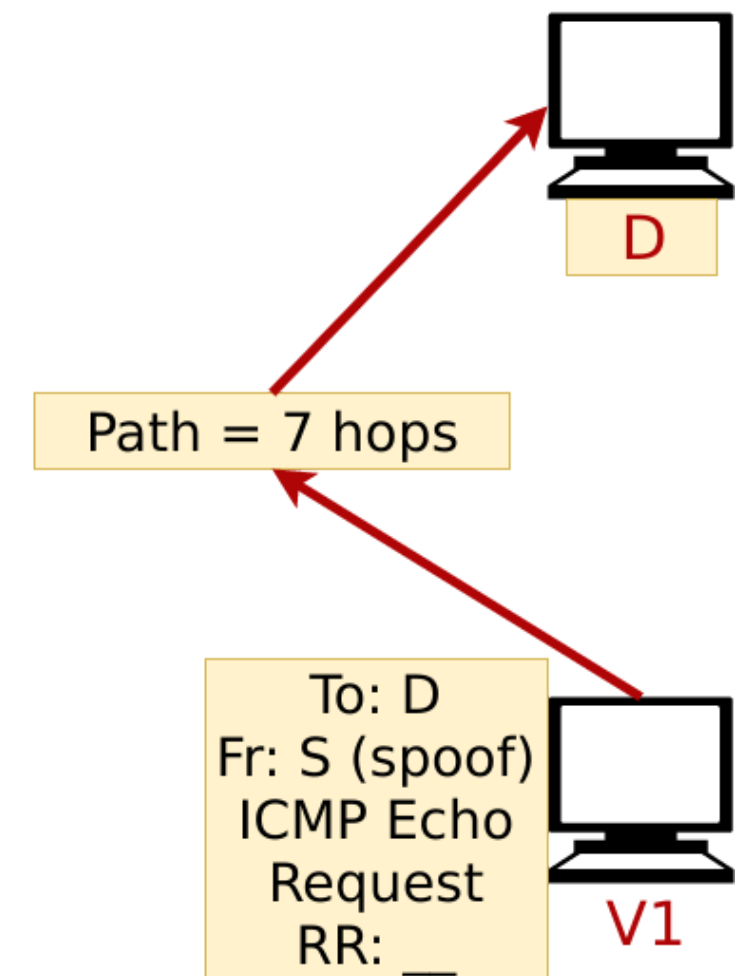
Reverse Traceroute: Source out of range

- What if D is more than 8 hops away from S?



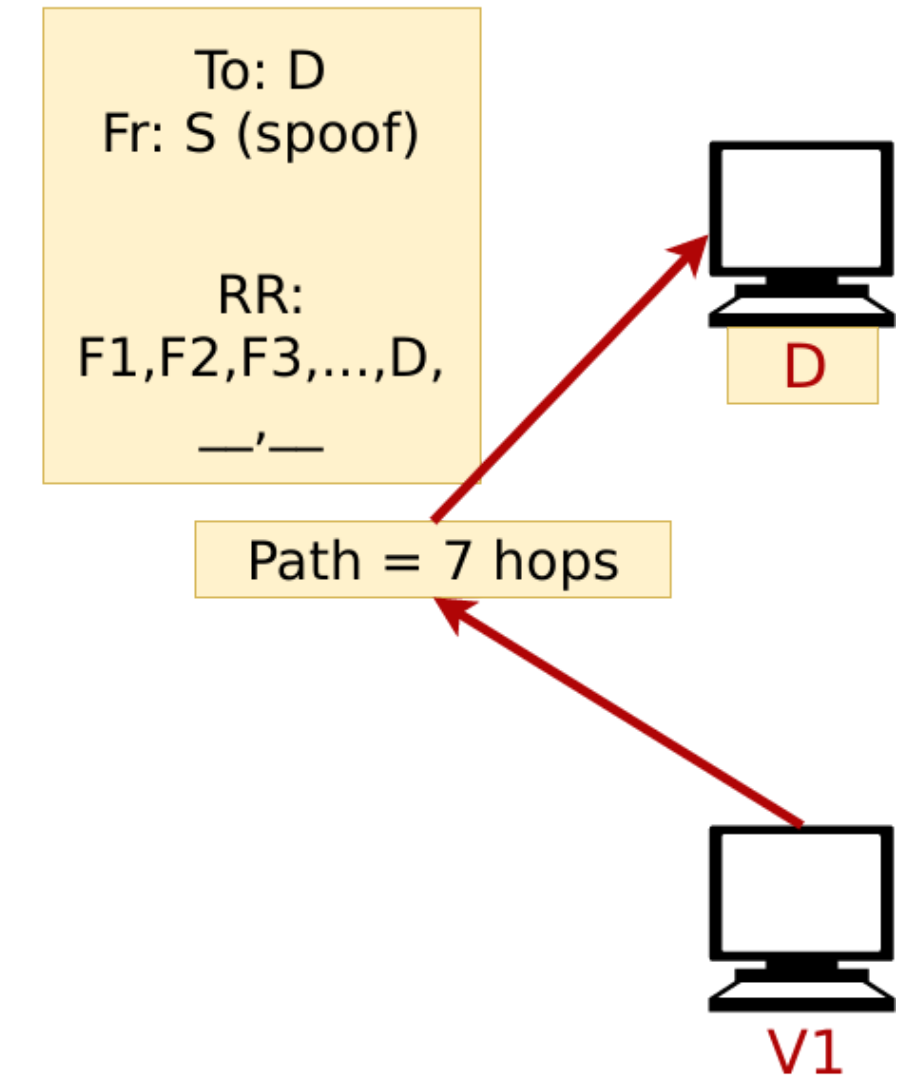
Reverse Traceroute: Spoofing

- What if D is more than 8 hops away from S?
- **Step 1:** Find a VP ≤ 8 hops away (V1).



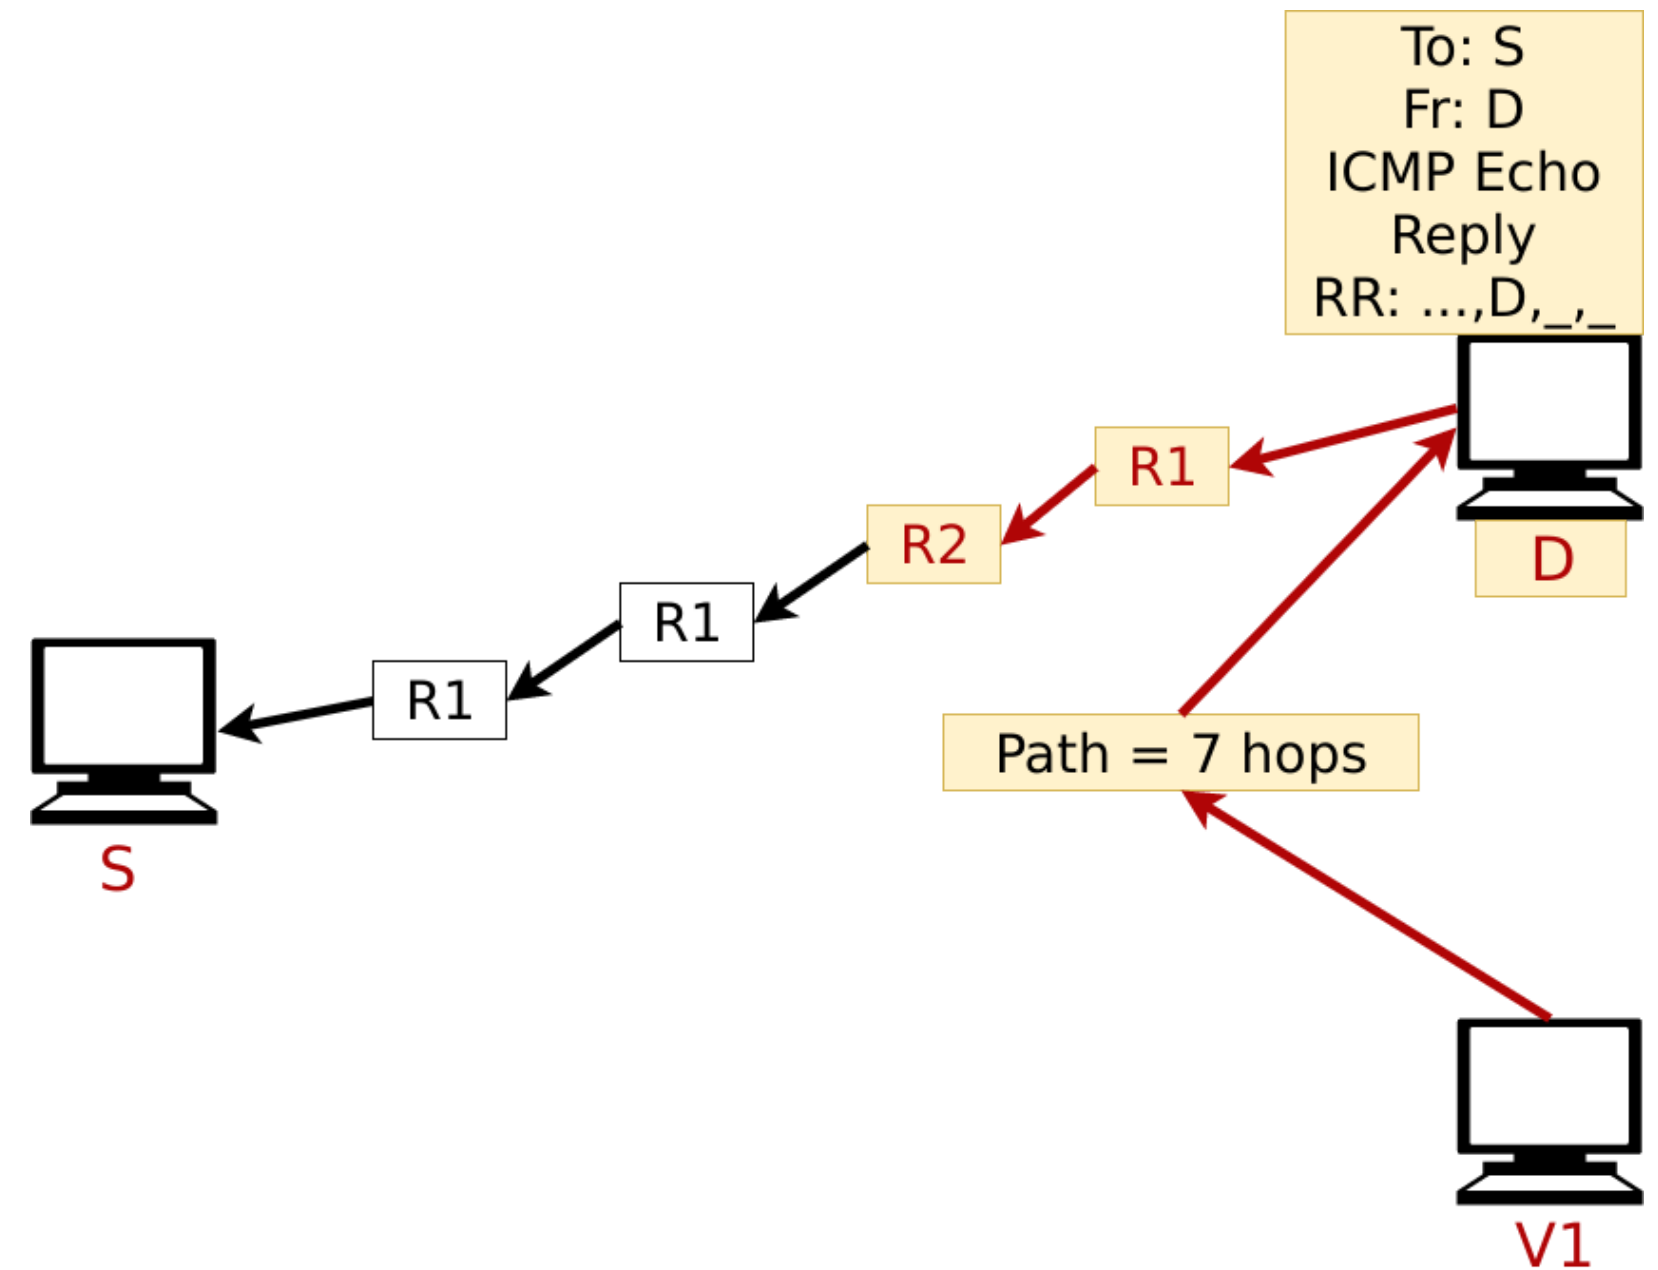
Reverse Traceroute: Spoofing

- What if D is more than 8 hops away from S?
- **Step 1:** Find a VP ≤ 8 hops away (V1).
- **Step 2:** Spoof S's address from V1.



Reverse Traceroute: Spoofing


- What if D is more than 8 hops away from S?
- **Step 1:** Find a VP ≤ 8 hops away (V1).
- **Step 2:** Spoof S's address from V1.
- **Step 3:** D sends response to S instead.
- **R1**→**R2** recorded.



We care about *closest* VP.

Results: Are destinations reachable in 9 hops?

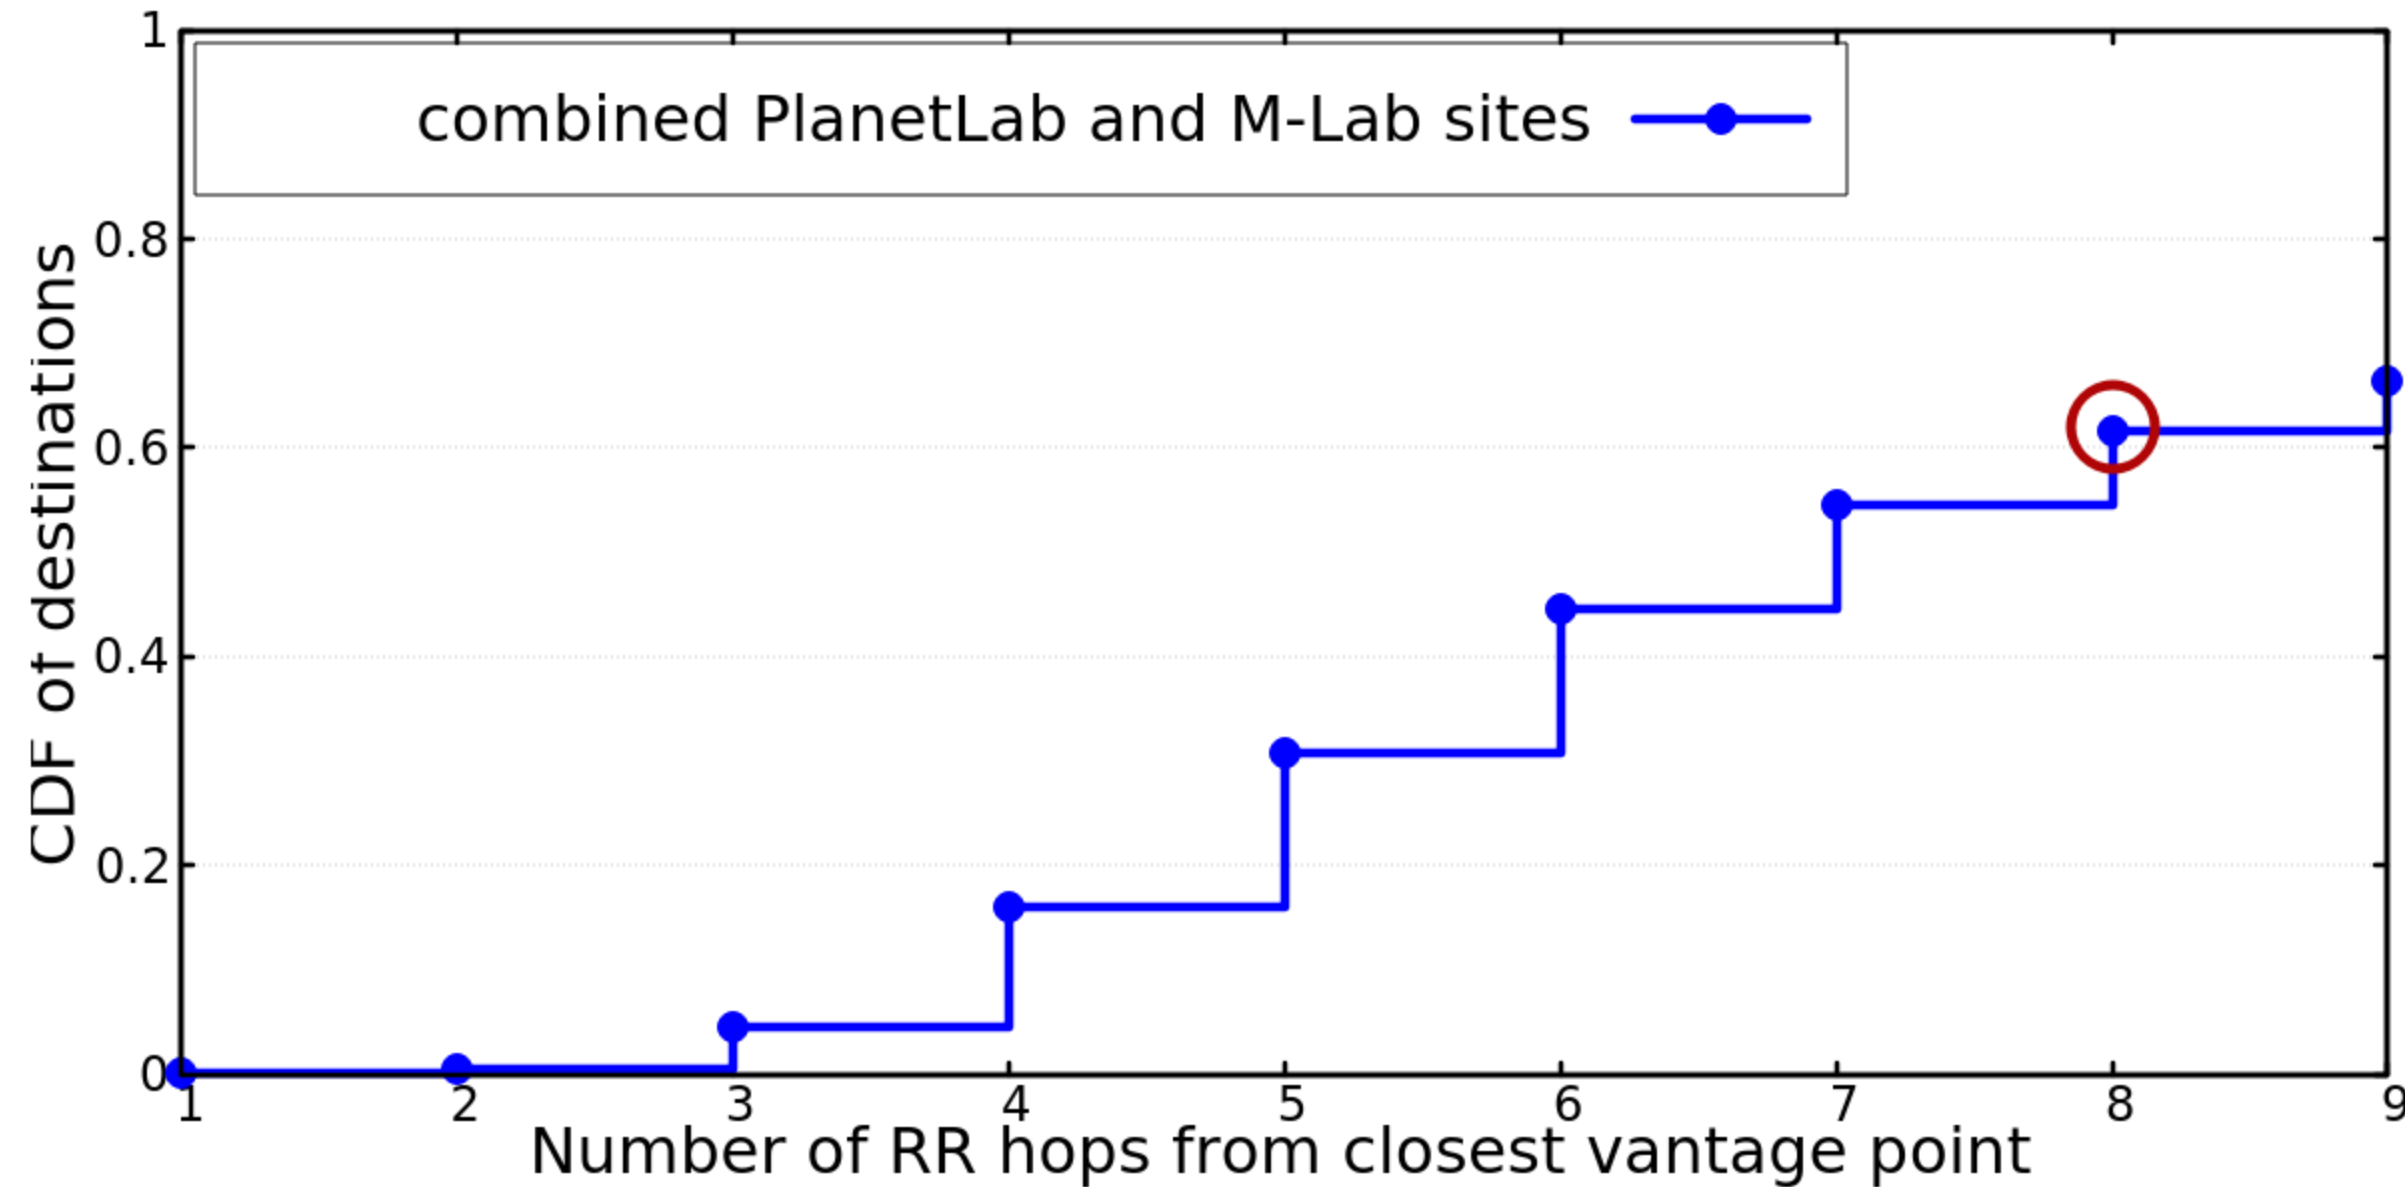
296,734 destinations responding to RR



- 66% of destinations are within 9 hops of closest VP.

Results: Are destinations reachable in 8 hops?

296,734 destinations responding to RR



- 62% of destinations are within 8 hops of closest VP.


Our Questions

- Do destinations *respond* to RR? **yes!**
- Are destinations *reachable* within the 9 hop limit?
yes!
- **Has reachability changed over time?**


Has reachability changed over time?

- Compared our results to measurements we made in 2011.
- Nearly identical methodology.
- 2016 destinations: all BGP-advertised prefixes.
- 2011 destinations: all /24 prefixes.
- Compared only VPs used in both years.

Results: Has reachability changed over time?



Results: Has reachability changed over time?



- 50% of destinations within 8 hops of closest VP in 2016.

- 6% of destinations within 8 hops of closest VP in 2011.

Why? Increased peering \Rightarrow better M-Lab reachability, MPLS.

Questions Asked in the Paper

- Do destinations *respond* to RR? **yes!**
- Are destinations *reachable* within the 9 hop limit? **yes!**
- Has reachability changed over time? **yes, gotten better**

- Do any ASes refuse to *stamp* RR packets? **no**
- Could RR be useful to cloud providers? **yes**
- How can we use RR responsibly? **read the paper!**


Conclusion:

The Record Route Option is an Option!*

- 75% of ping responsive destinations respond to RR. **for measurement*
- Destinations closer to our VPs now than in the past.
- Majority close enough to **measure reverse paths**.
- **What new uses for Record Route can we discover?**


Results: PlanetLab vs. M-Lab

296,734 destinations responding to RR



Backup: PlanetLab vs. M-Lab

296,734 destinations responding to RR



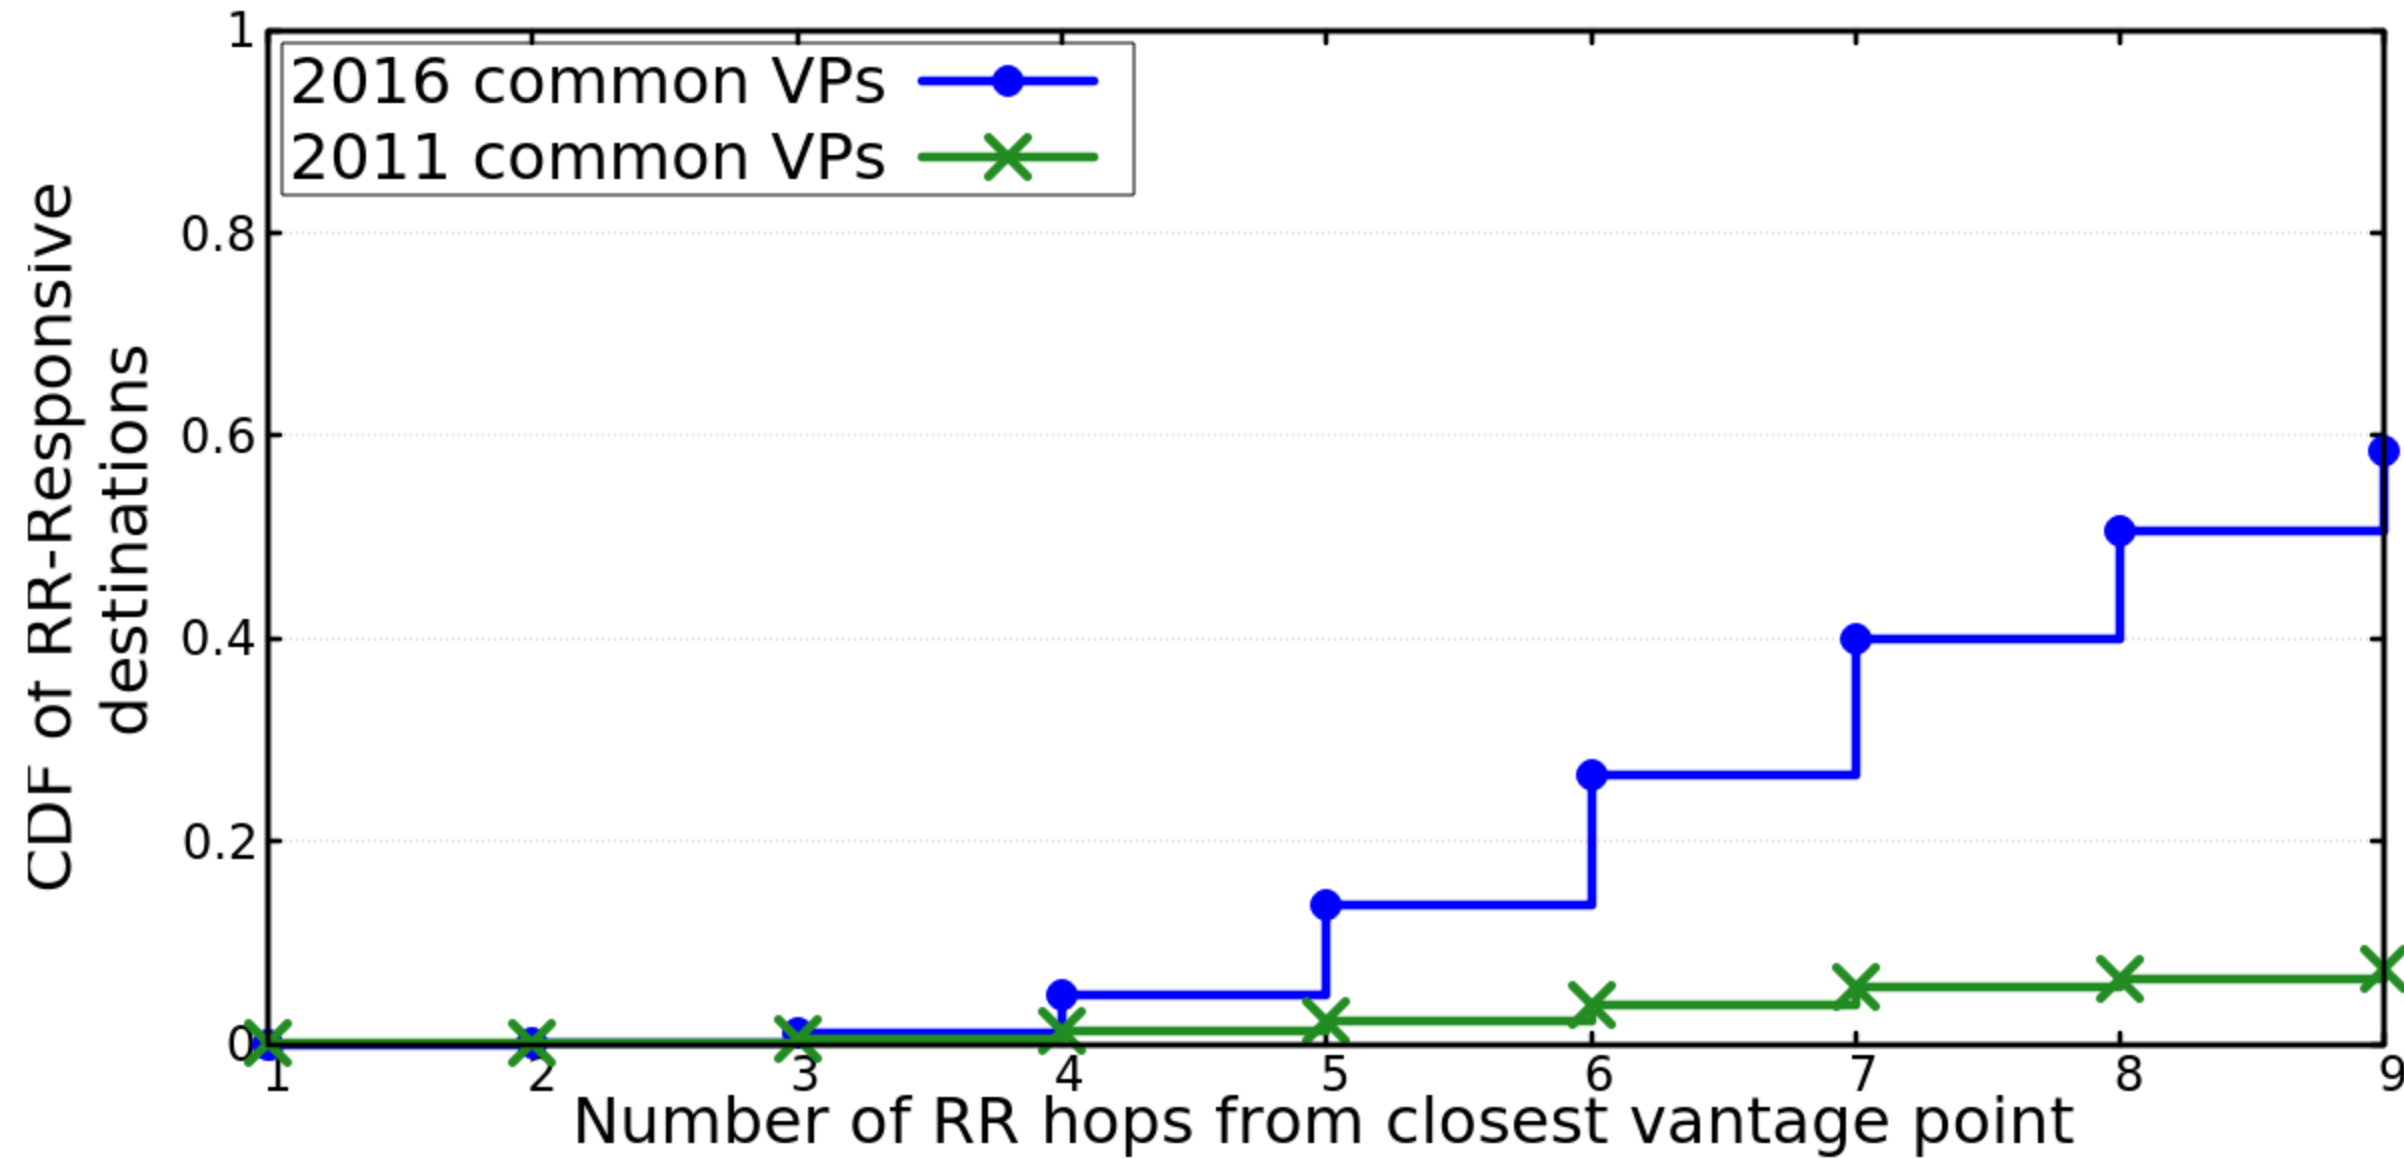
- 61% of destinations within 8 hops of closest M-Lab VP.
- 35% of destinations within 8 hops of closest Planetlab VP.

Evidence of the decline of Planetlab, the rise of M-Lab

Backup: Responsiveness by AS Type

	Total	Transit/Access	Enterprise	Content	Unknown
	510,305 (100%)	388,959 (100%)	61,204 (100%)	44,295 (100%)	15,847 (100%)
By IP					
All Probed	510,305 (100%)	388,959 (100%)	61,204 (100%)	44,295 (100%)	15,847 (100%)
Ping Responsive	394,644 (77%)	296,011 (76%)	51,579 (84%)	37,299 (84%)	9,755 (62%)
RR-Responsive	296,734 (58%)	225,000 (58%)	34,917 (57%)	28,786 (65%)	8,031 (51%)

Backup: Reachability over time



- Changes in set of VPs alone cannot account for difference

Backup: Limiting network impact

