A second look at "*Detecting third-party addresses in traceroute traces with the IP timestamp option*"

Matthew Luckie, k claffy
mjl@caida.org

CAIDA - University of California, San Diego

Passive and Active Measurement Conference (PAM) 2014

RELATED WORK ON TRACEROUTE SHORT-COMINGS

- Significant volume of literature reporting the short-comings of traceroute

  - **Oliveira** *et al.*: Observing the Evolution of Internet AS Topology. SIGCOMM 2007

  - **Willinger** *et al.*: Mathematics and the Internet: a source of enormous confusion and great potential. AMS 2009

  - **Zhang** *et al.*: Quantifying the pitfalls of traceroute in AS connectivity inference. PAM 2010

RELATED WORK ON TRACEROUTE SHORT-COMINGS

- Significant volume of literature reporting the short-comings of traceroute

  - **Oliveira** *et al.*: Observing the Evolution of Internet AS Topology. SIGCOMM 2007

  - **Willinger** *et al.*: Mathematics and the Internet: a source of enormous confusion and great potential. AMS 2009

  - **Zhang** *et al.*: Quantifying the pitfalls of traceroute in AS connectivity inference. PAM 2010

  - P. **Marchetta**, W. de Donato, A. Pescape: Detecting third-party addresses in traceroute traces with IP timestamp option. PAM 2013

ASSUMED TRACEROUTE BEHAVIOR

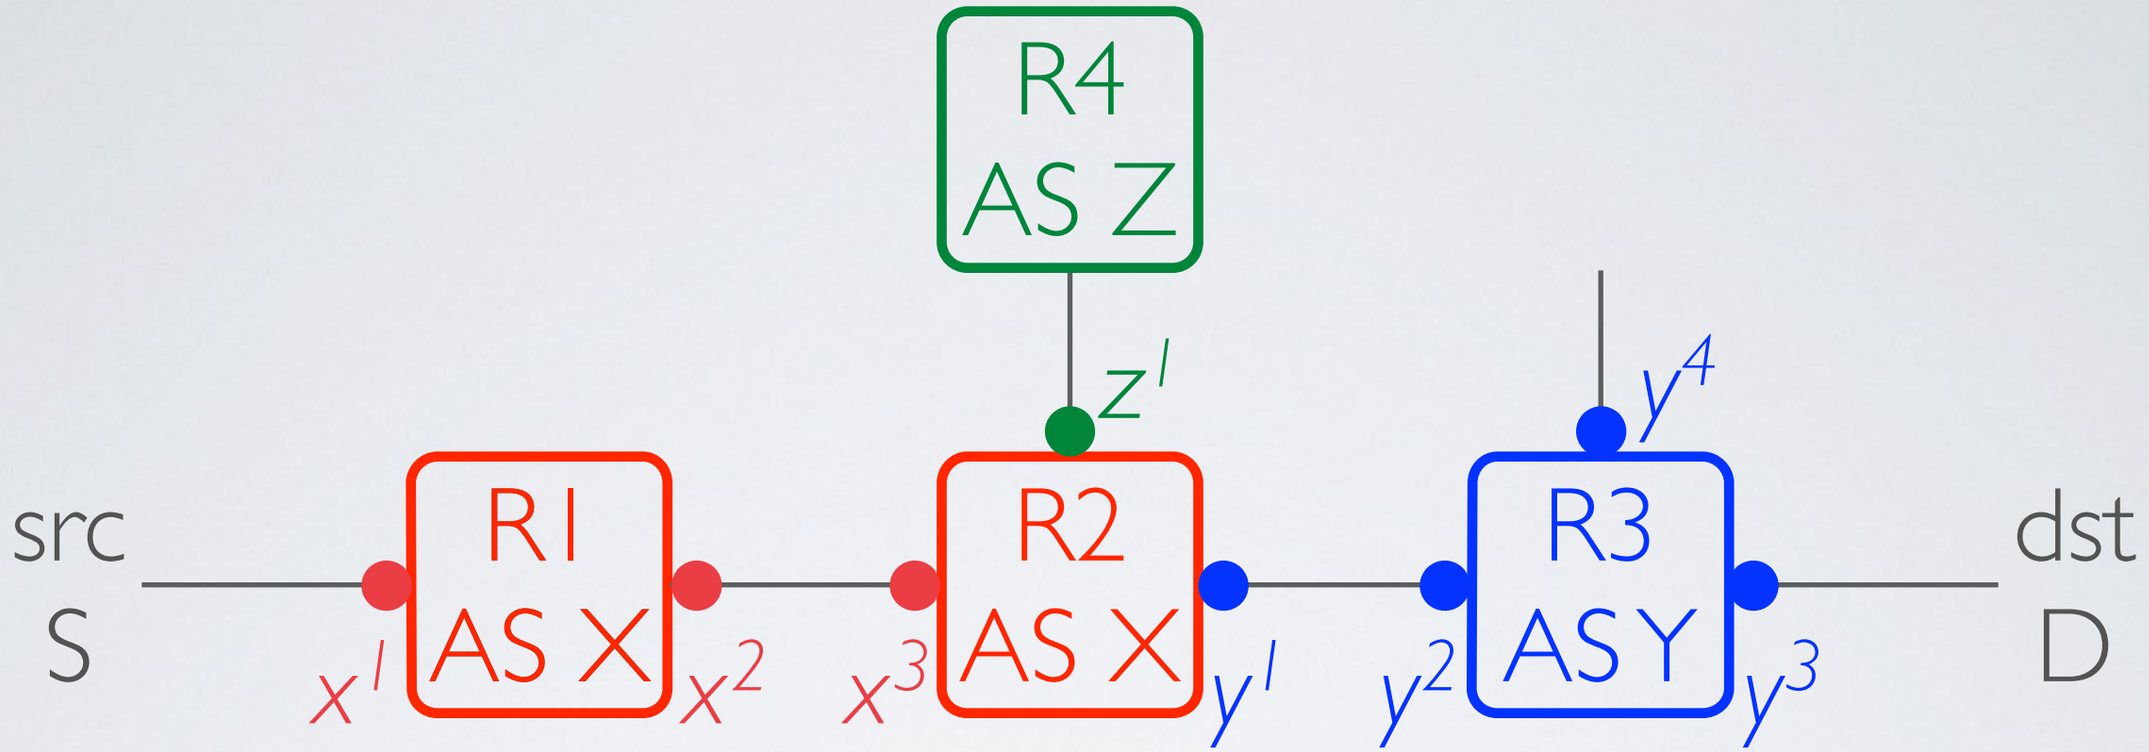

ASSUMED TRACEROUTE BEHAVIOR

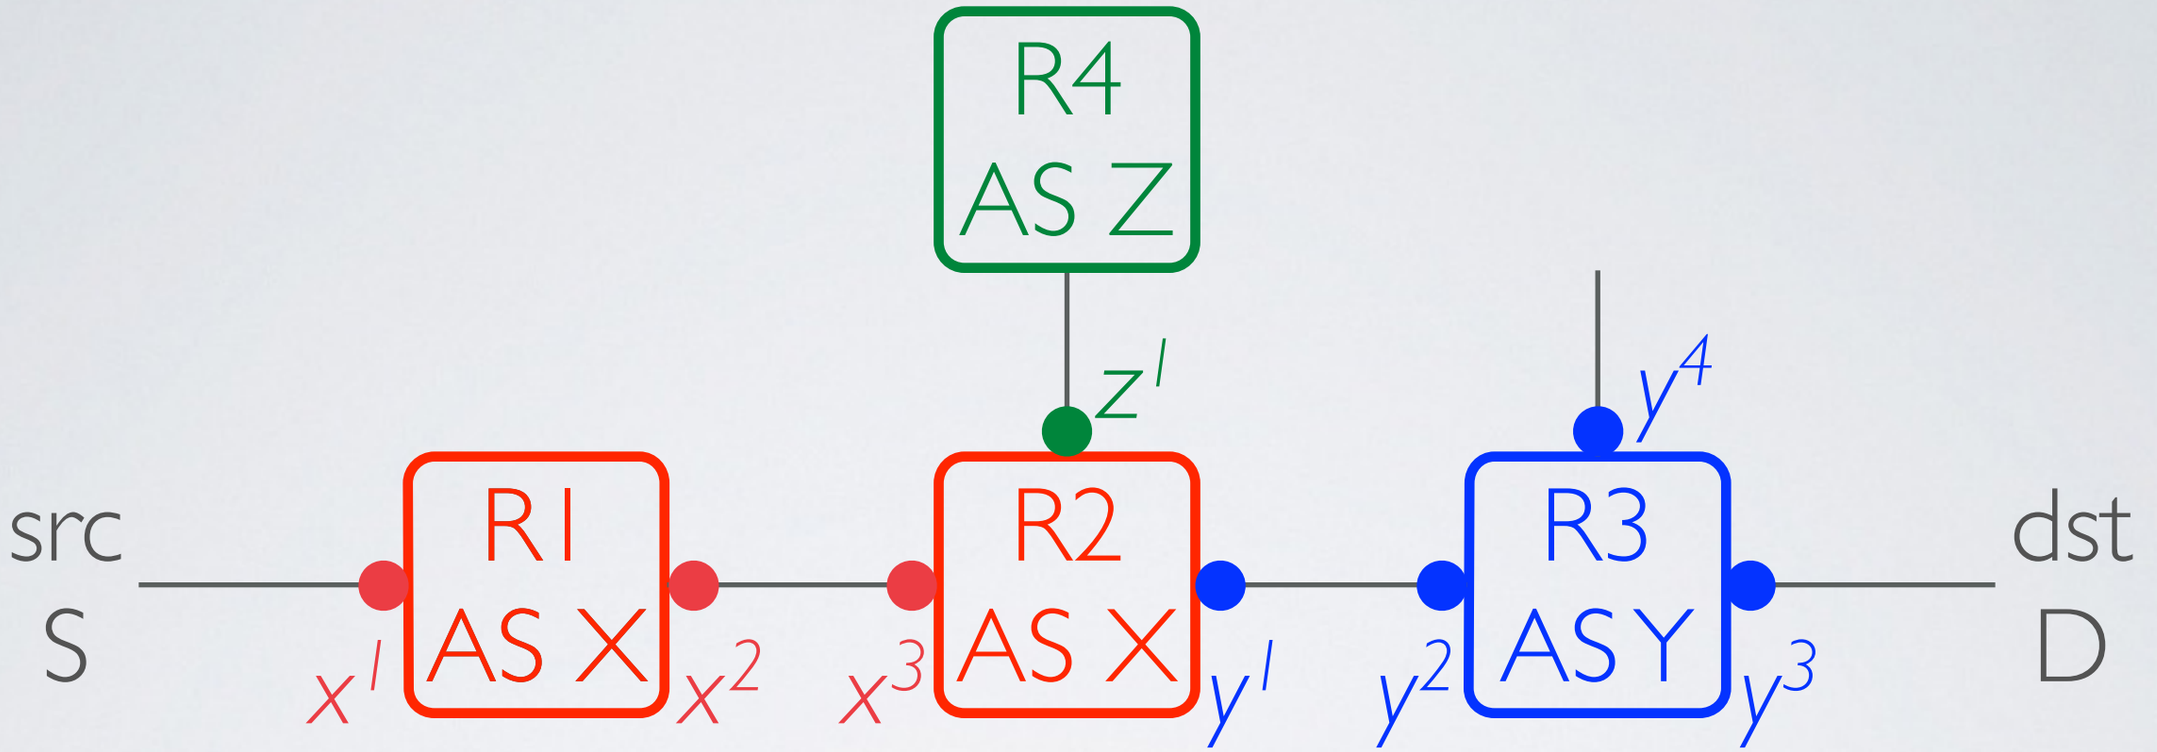

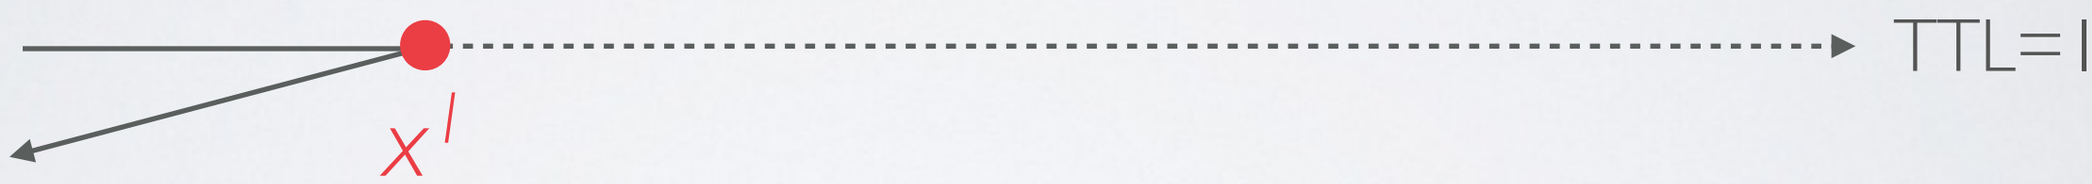

ASSUMED TRACEROUTE BEHAVIOR

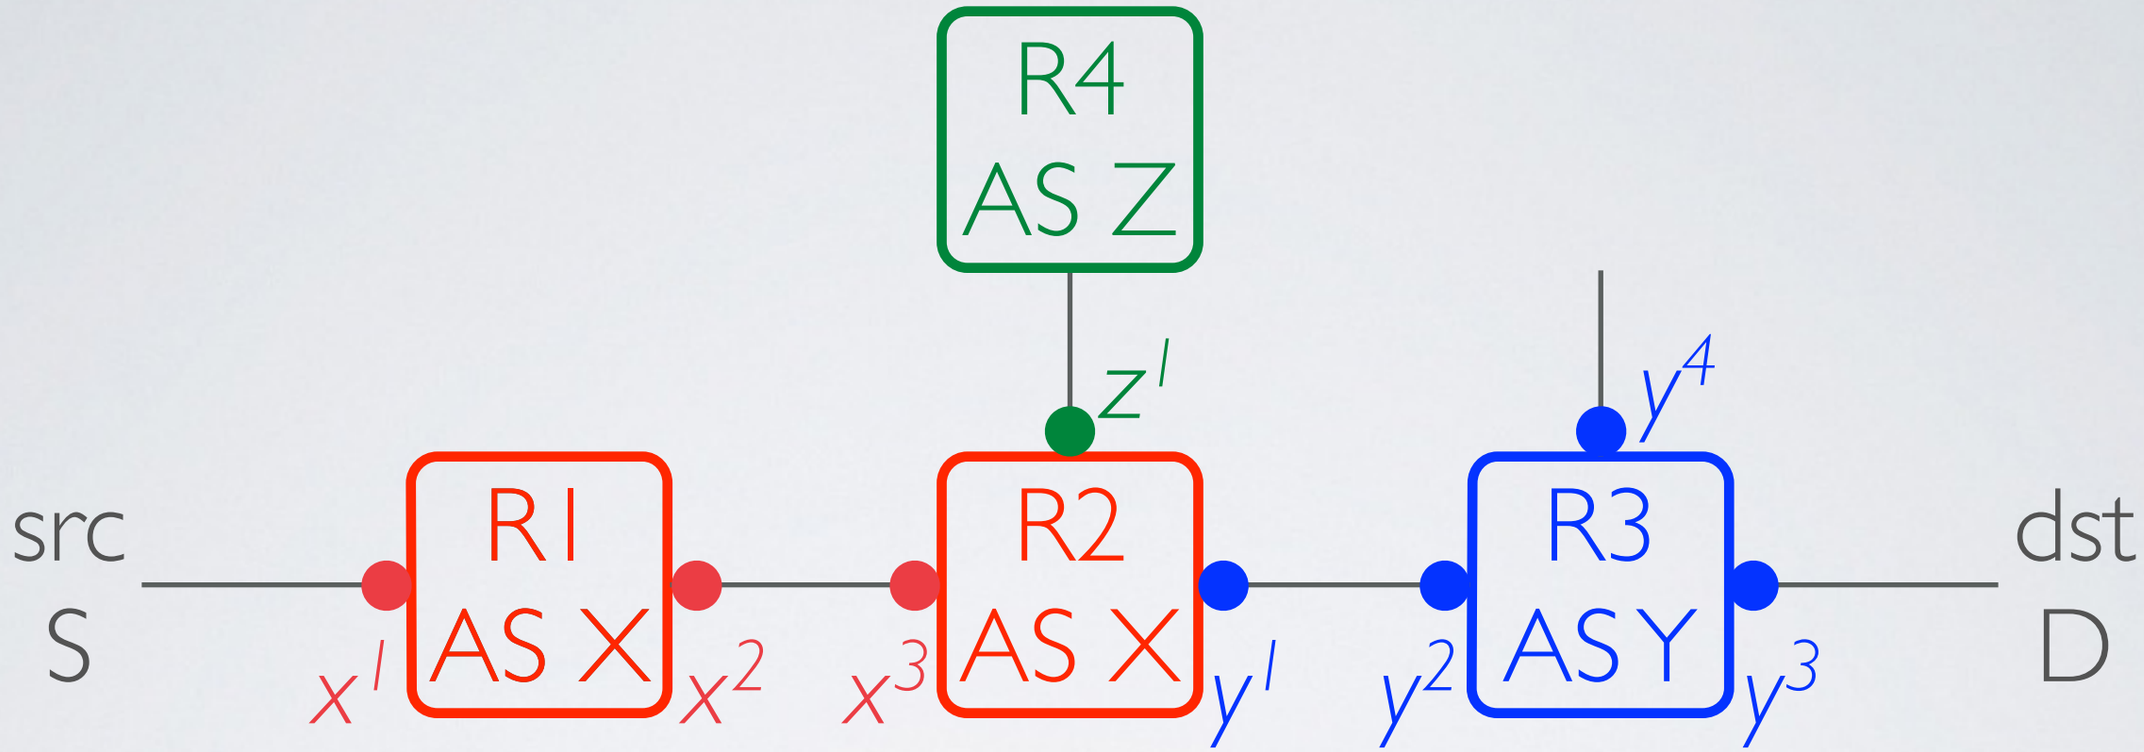

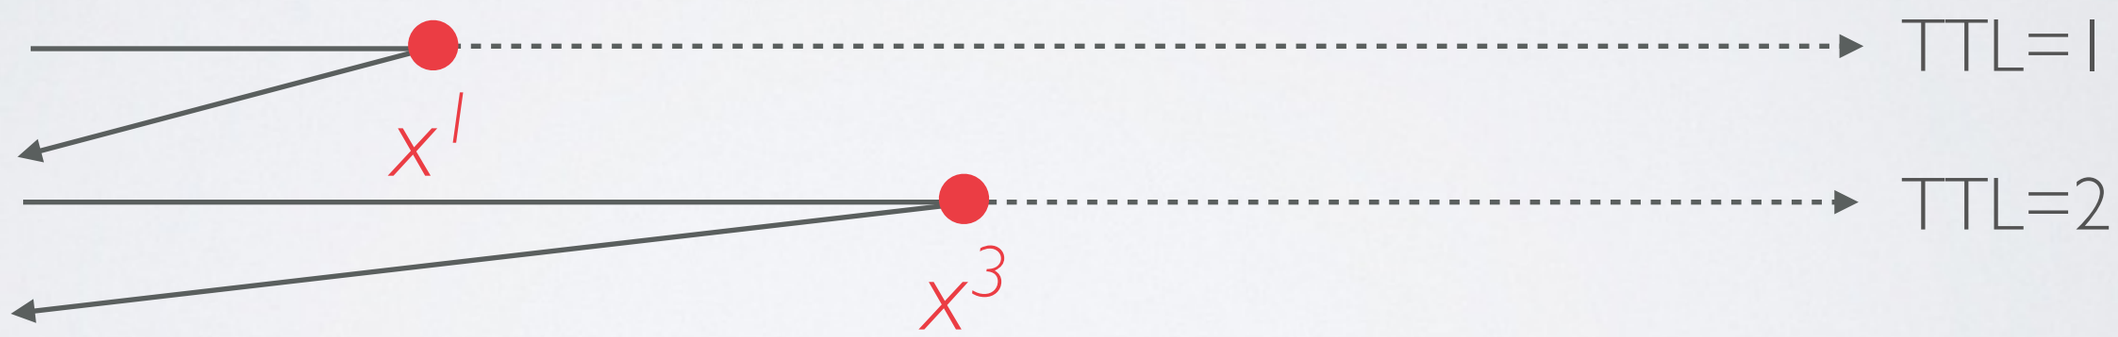

ASSUMED TRACEROUTE BEHAVIOR

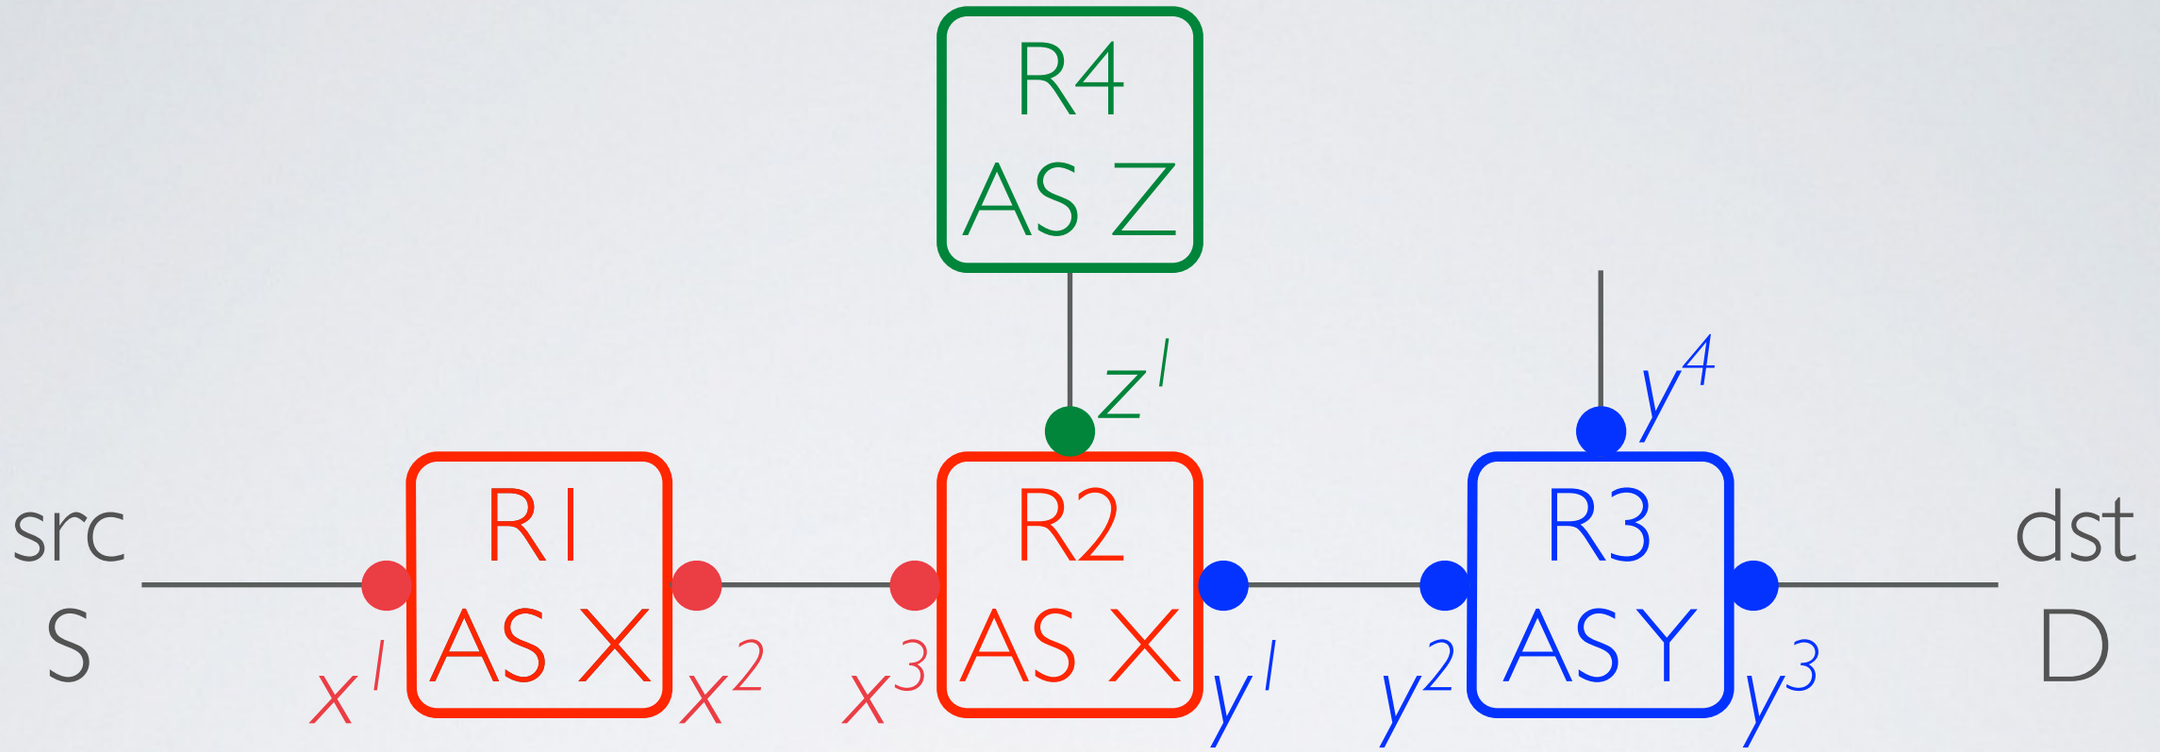

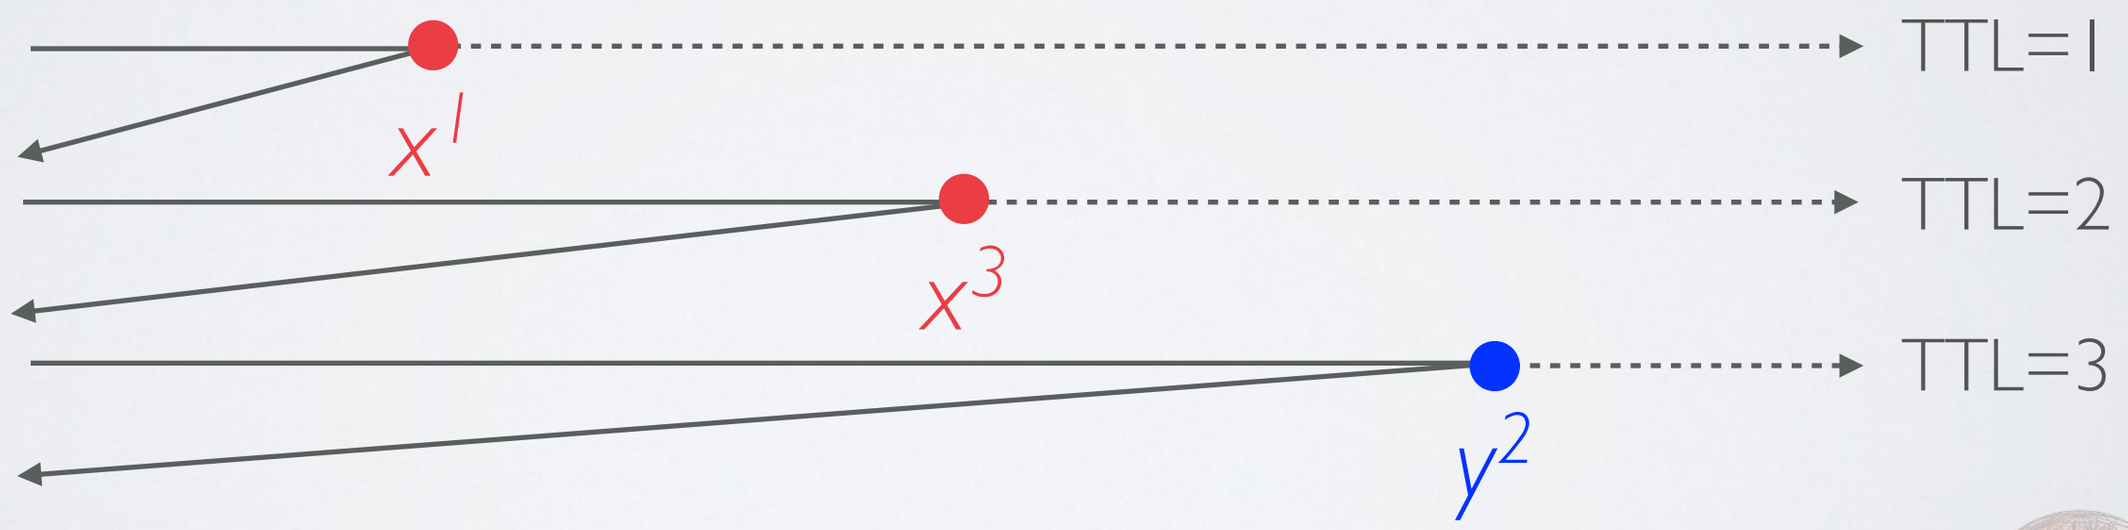

ASSUMED TRACEROUTE BEHAVIOR

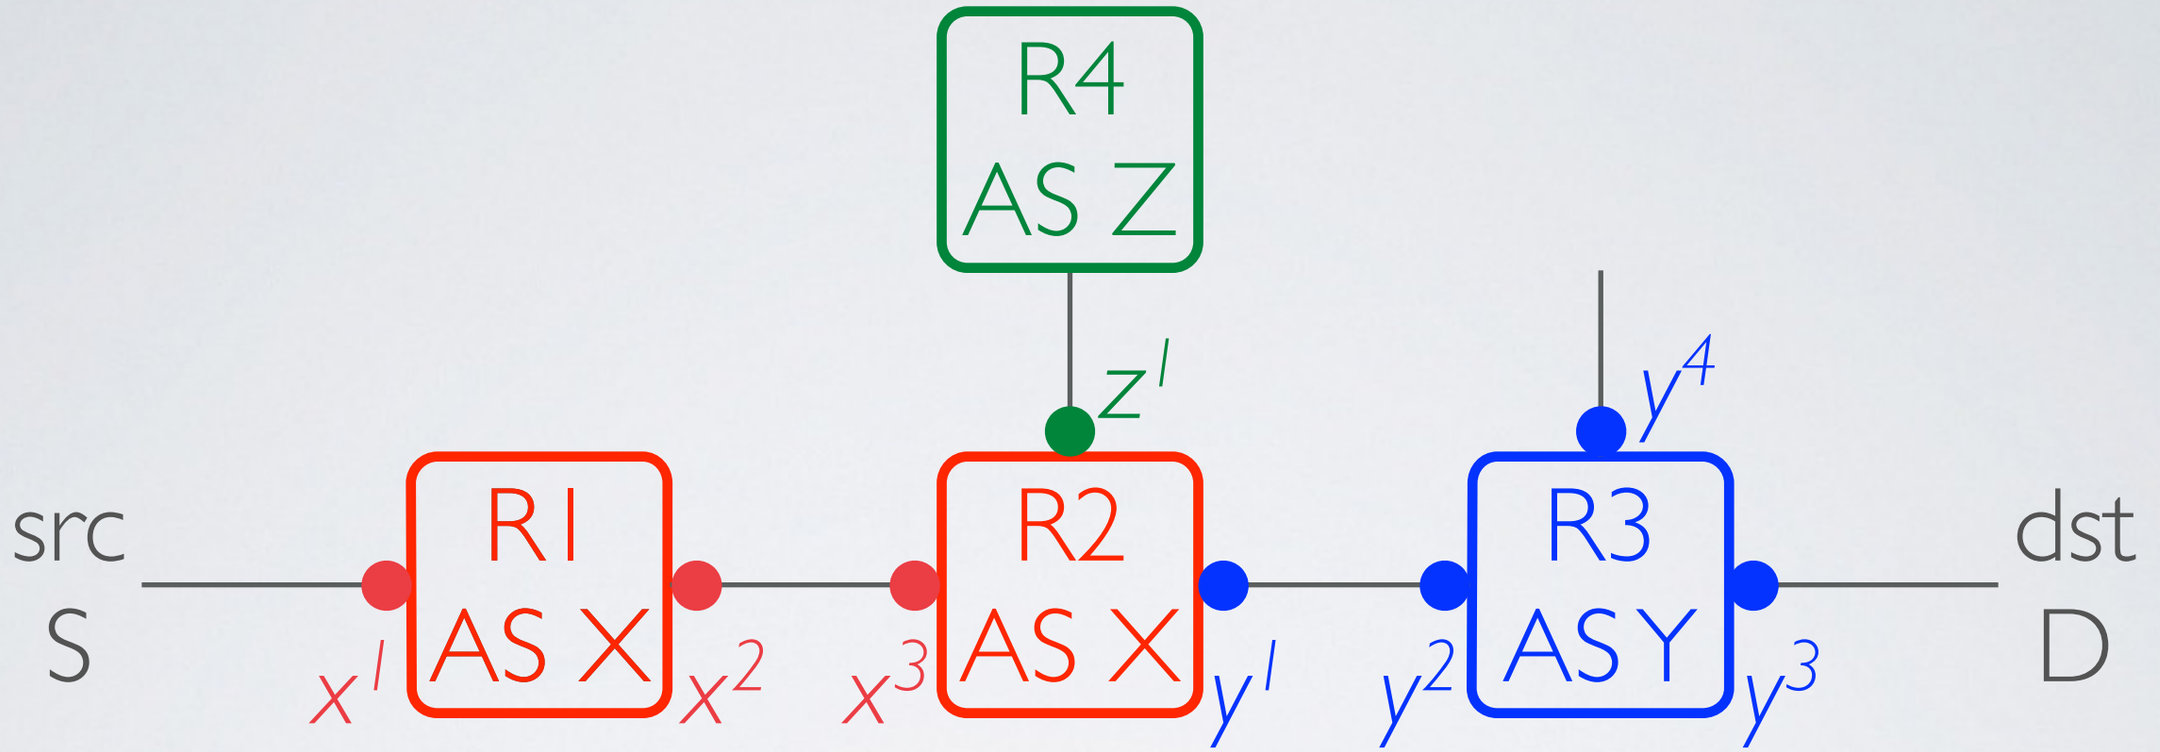

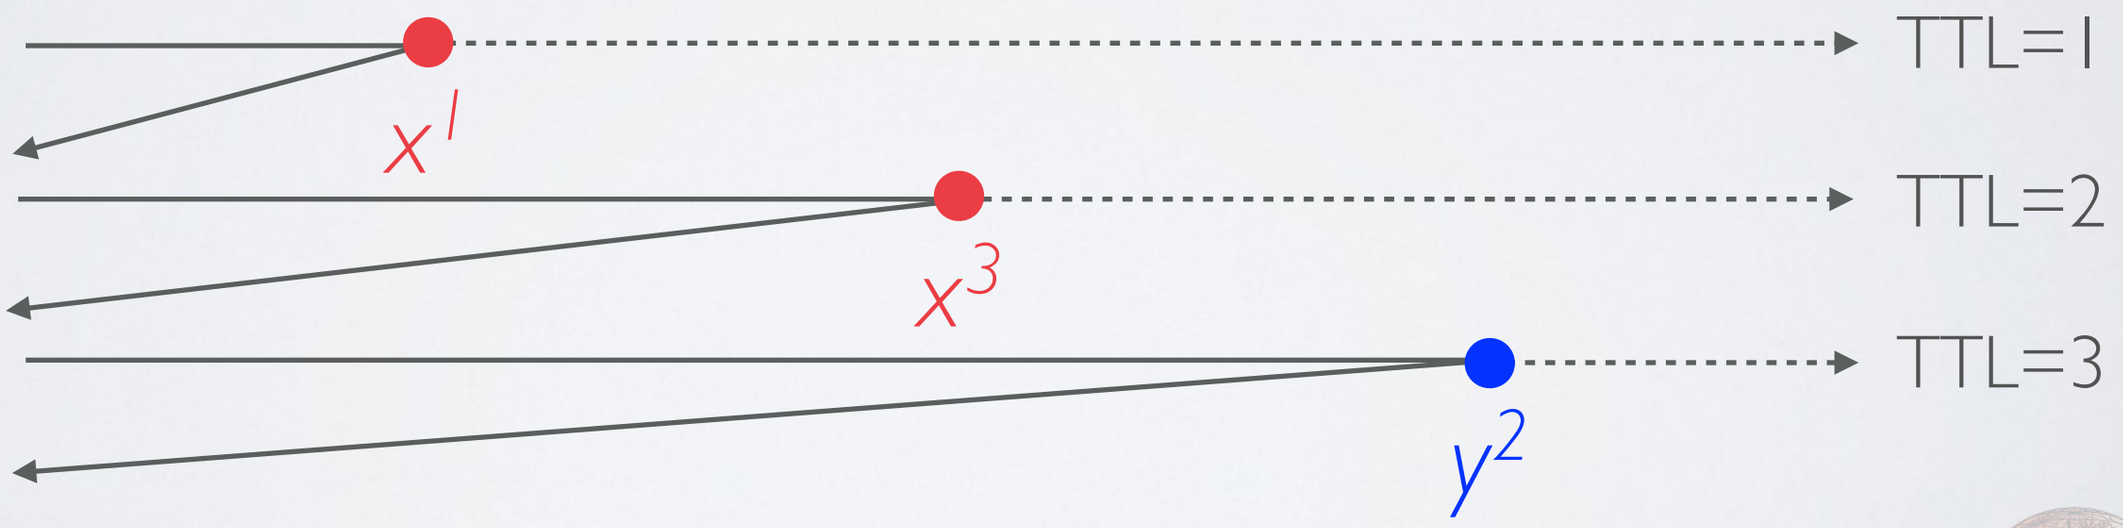

AS path inference: X Y

HOW SHOULD A ROUTER BEHAVE?

(RFC 1812)

4.3.2.4 ICMP Message Source Address

> Except where this document specifies otherwise, the IP source address in an ICMP message originated by the router MUST be one of the IP addresses associated with the physical interface over which the ICMP message is transmitted. If the interface has no IP addresses associated with it, the router's router-id (see Section [5.2.5]) is used instead.

*That is, the address of the **out-bound** interface from which the ICMP message is sent*

OFF-PATH AND THIRD-PARTY ADDRESSES

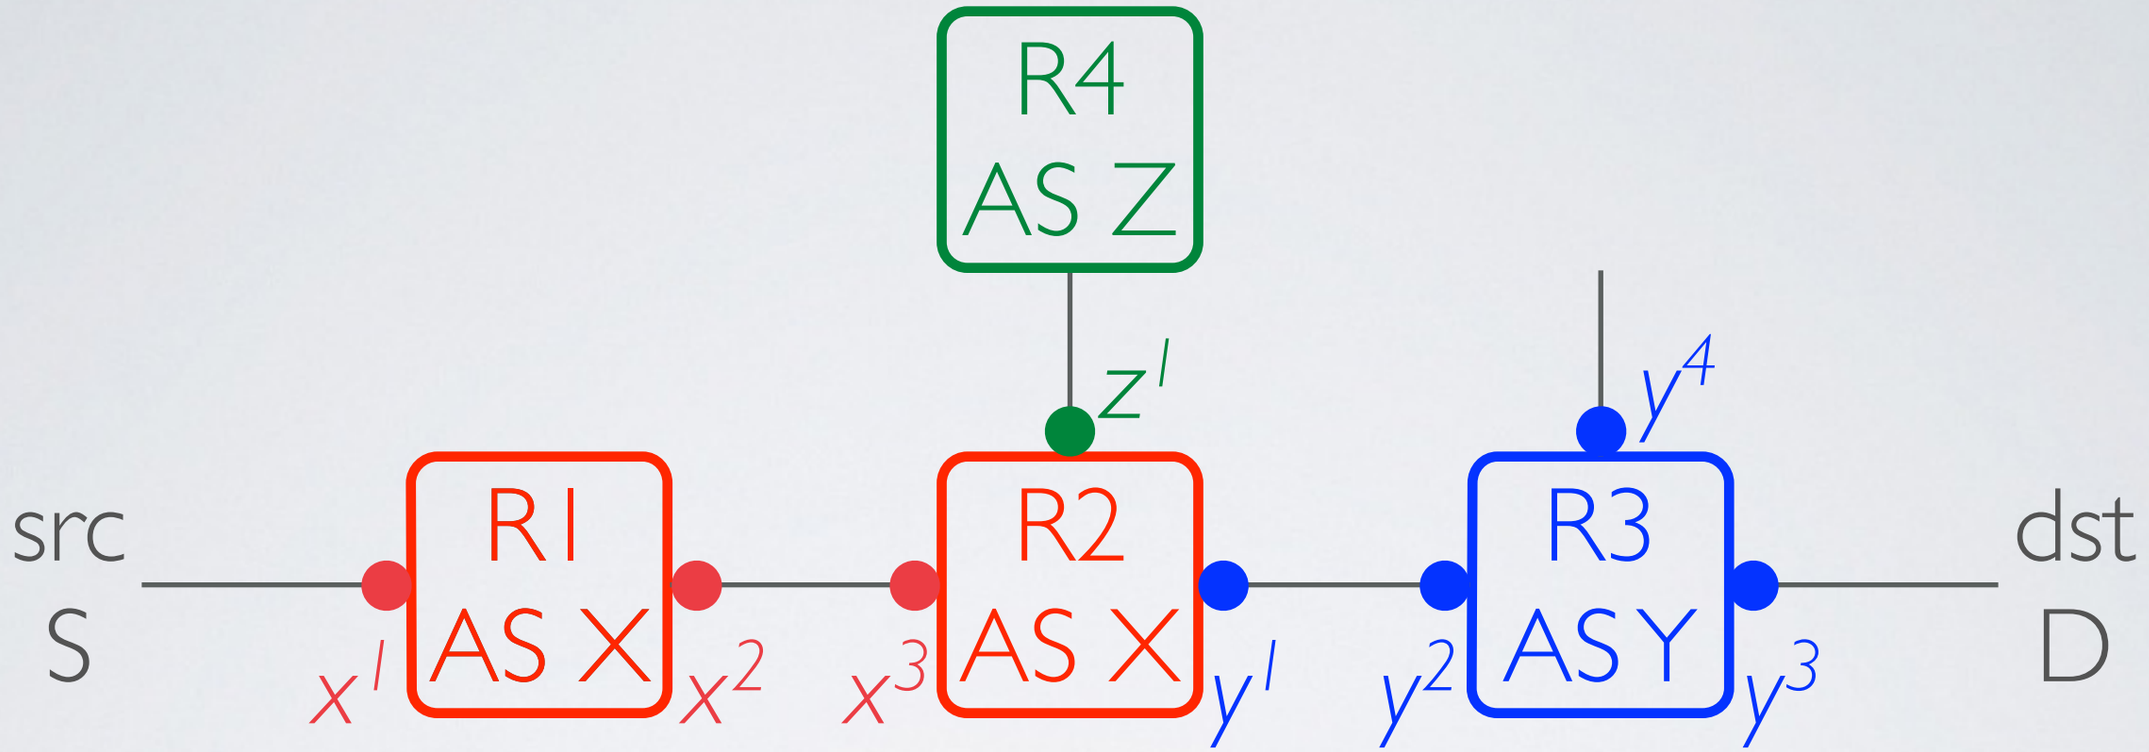

OFF-PATH AND THIRD-PARTY ADDRESSES

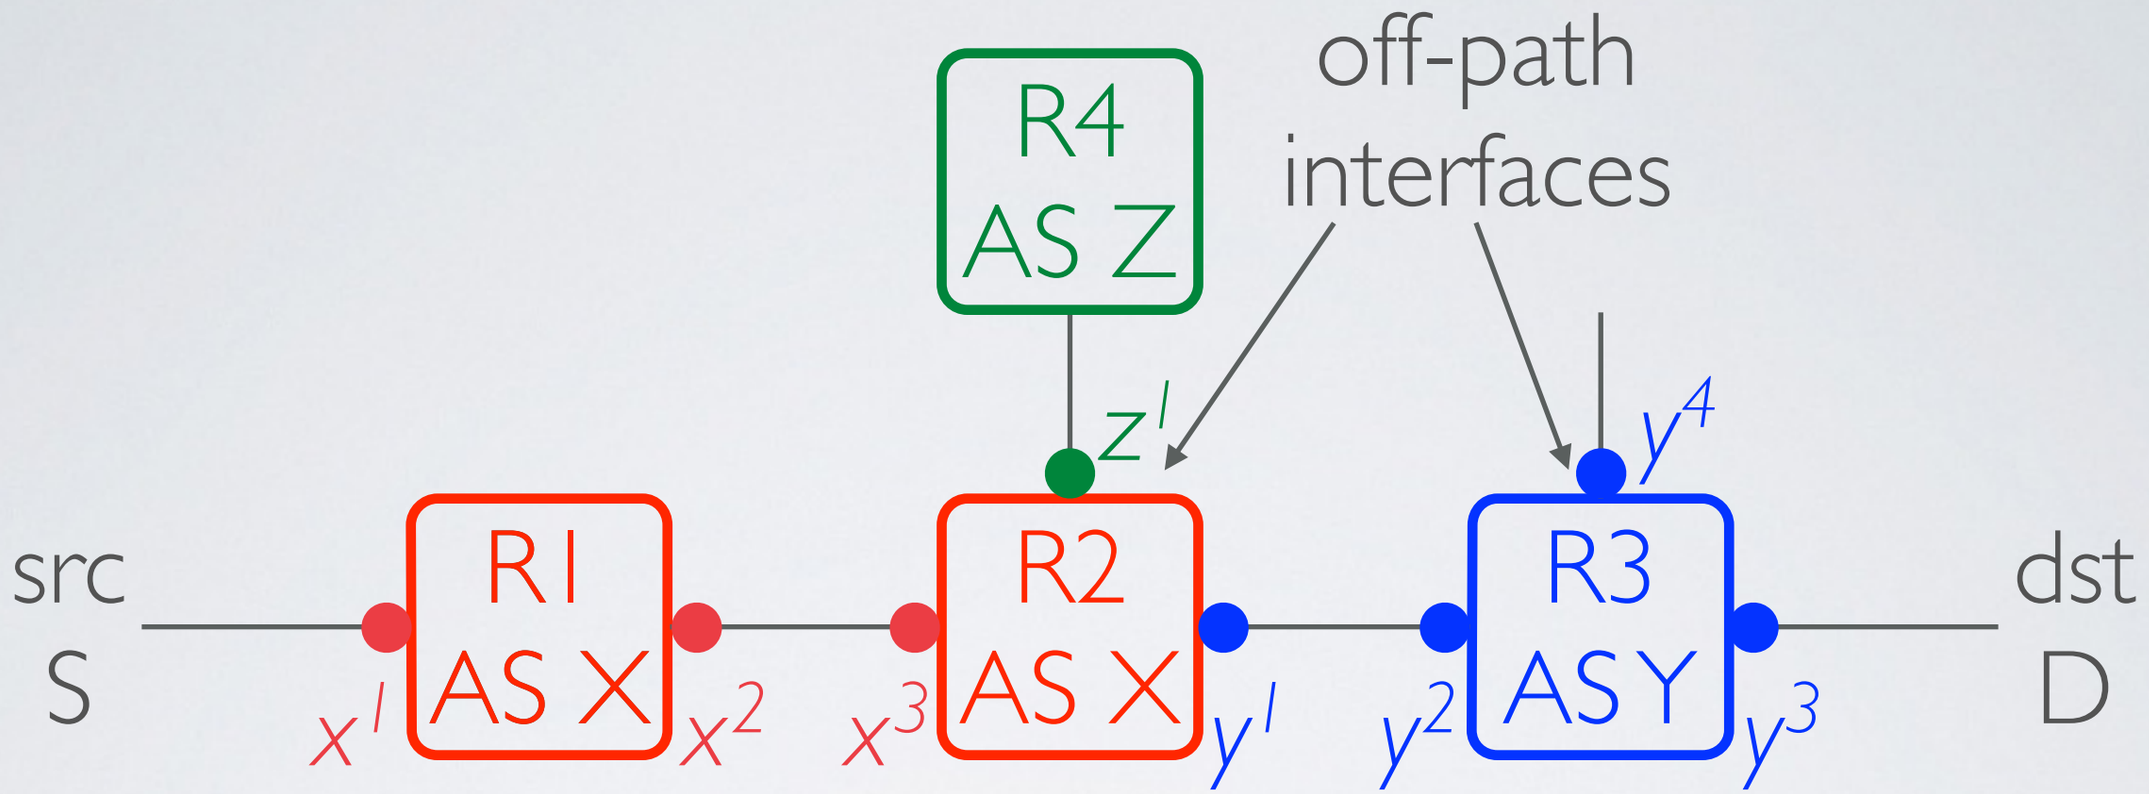

OFF-PATH AND THIRD-PARTY ADDRESSES

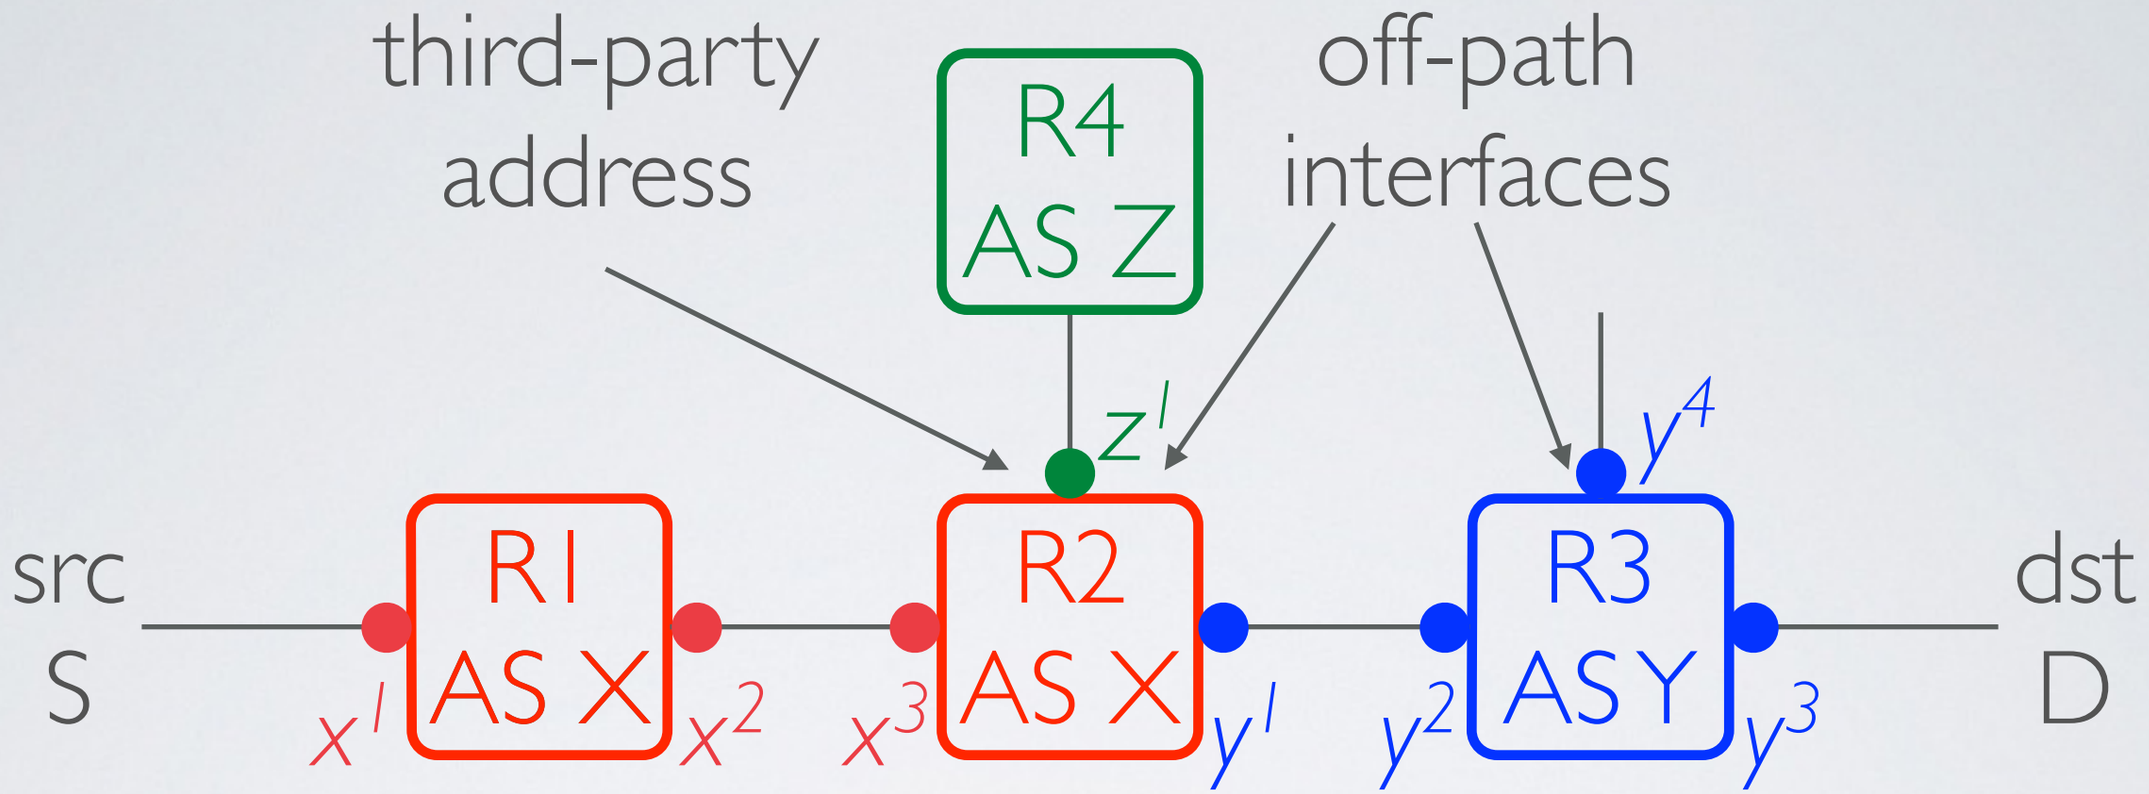

OFF-PATH AND THIRD-PARTY ADDRESSES

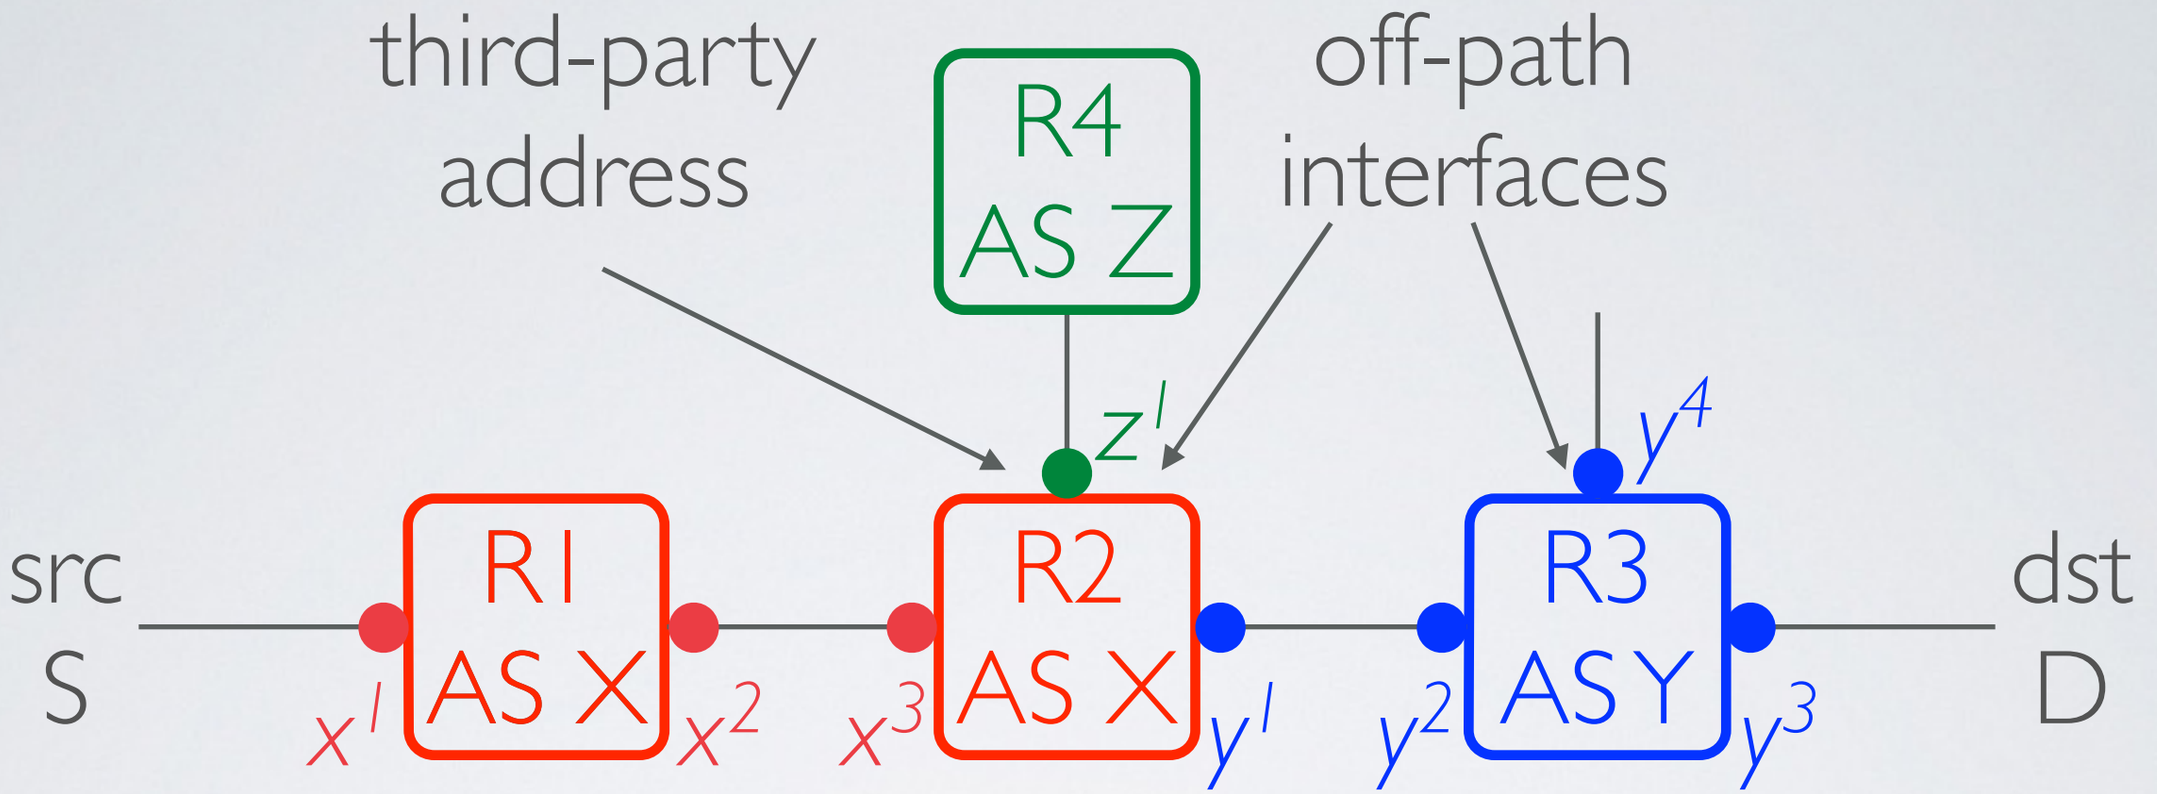

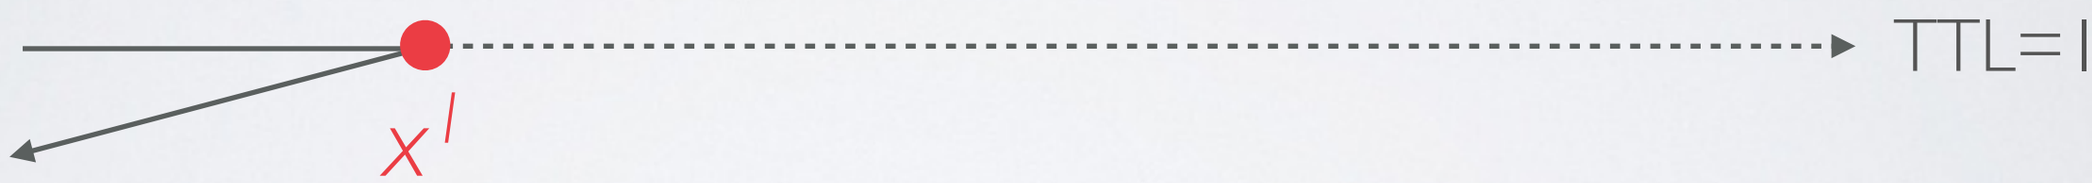

OFF-PATH AND THIRD-PARTY ADDRESSES

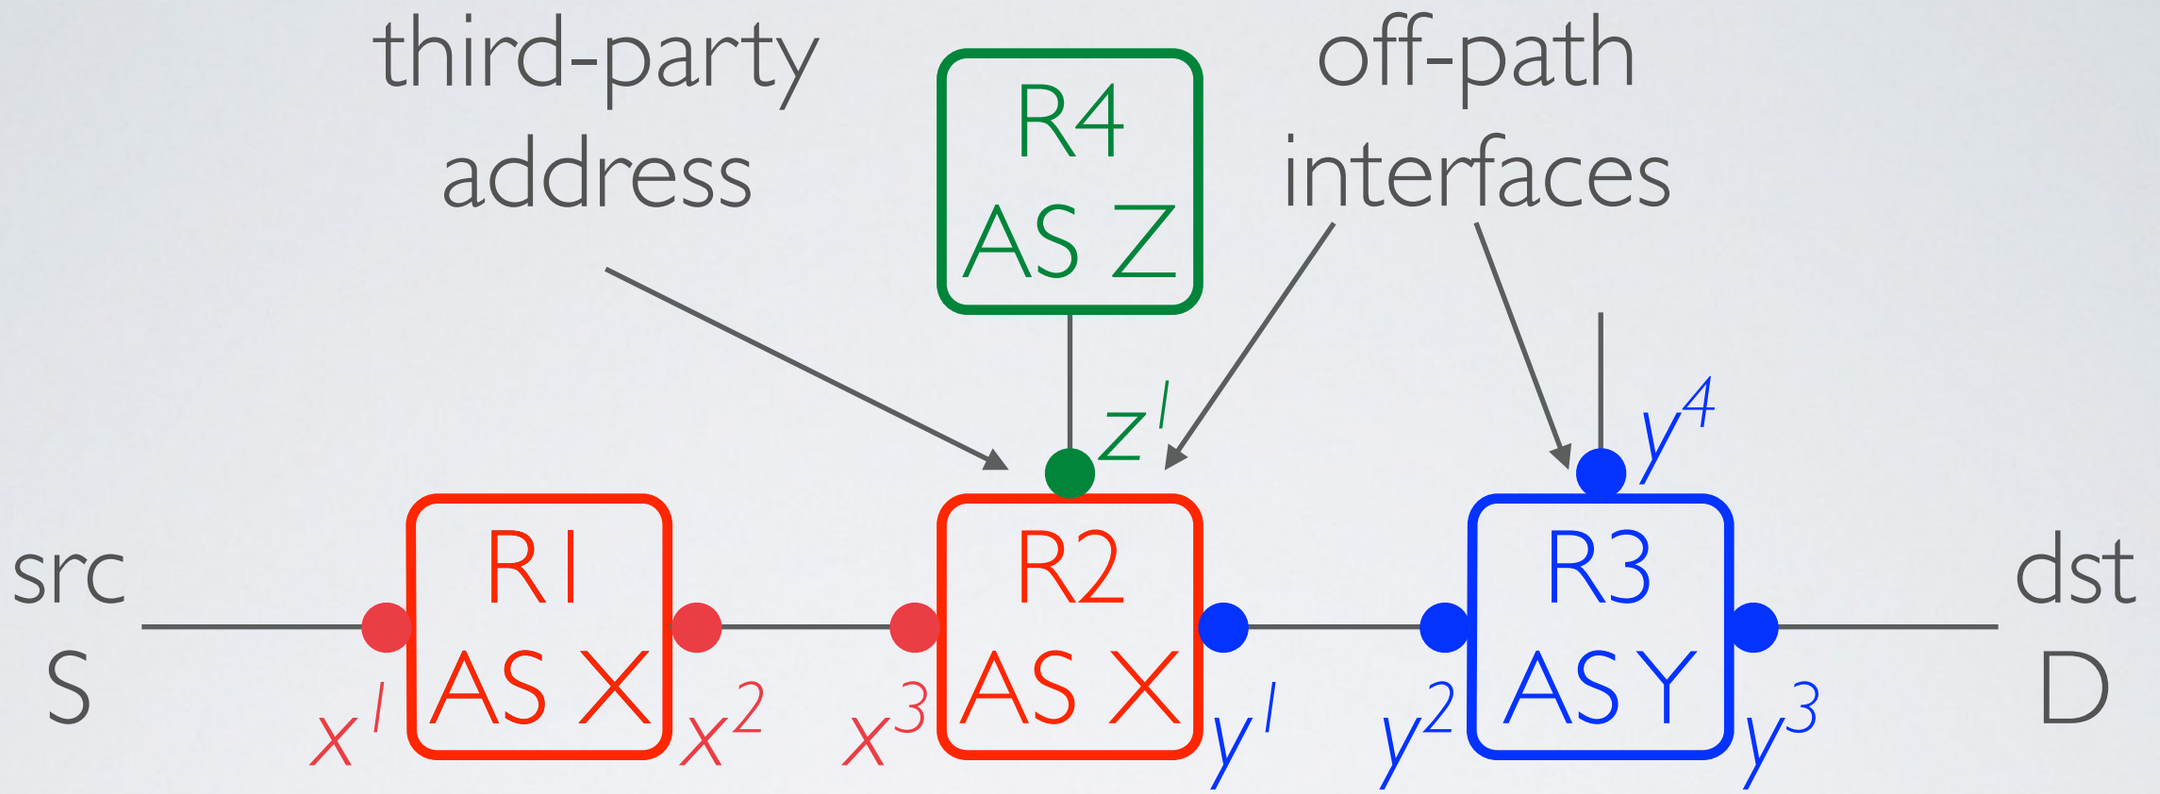

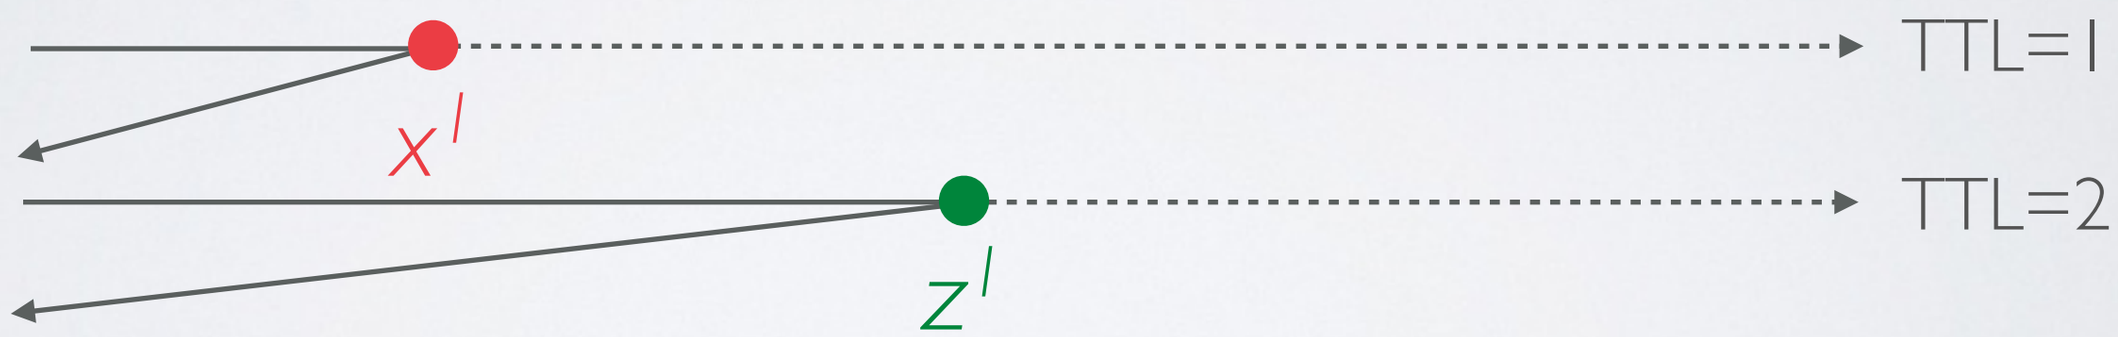

OFF-PATH AND THIRD-PARTY ADDRESSES

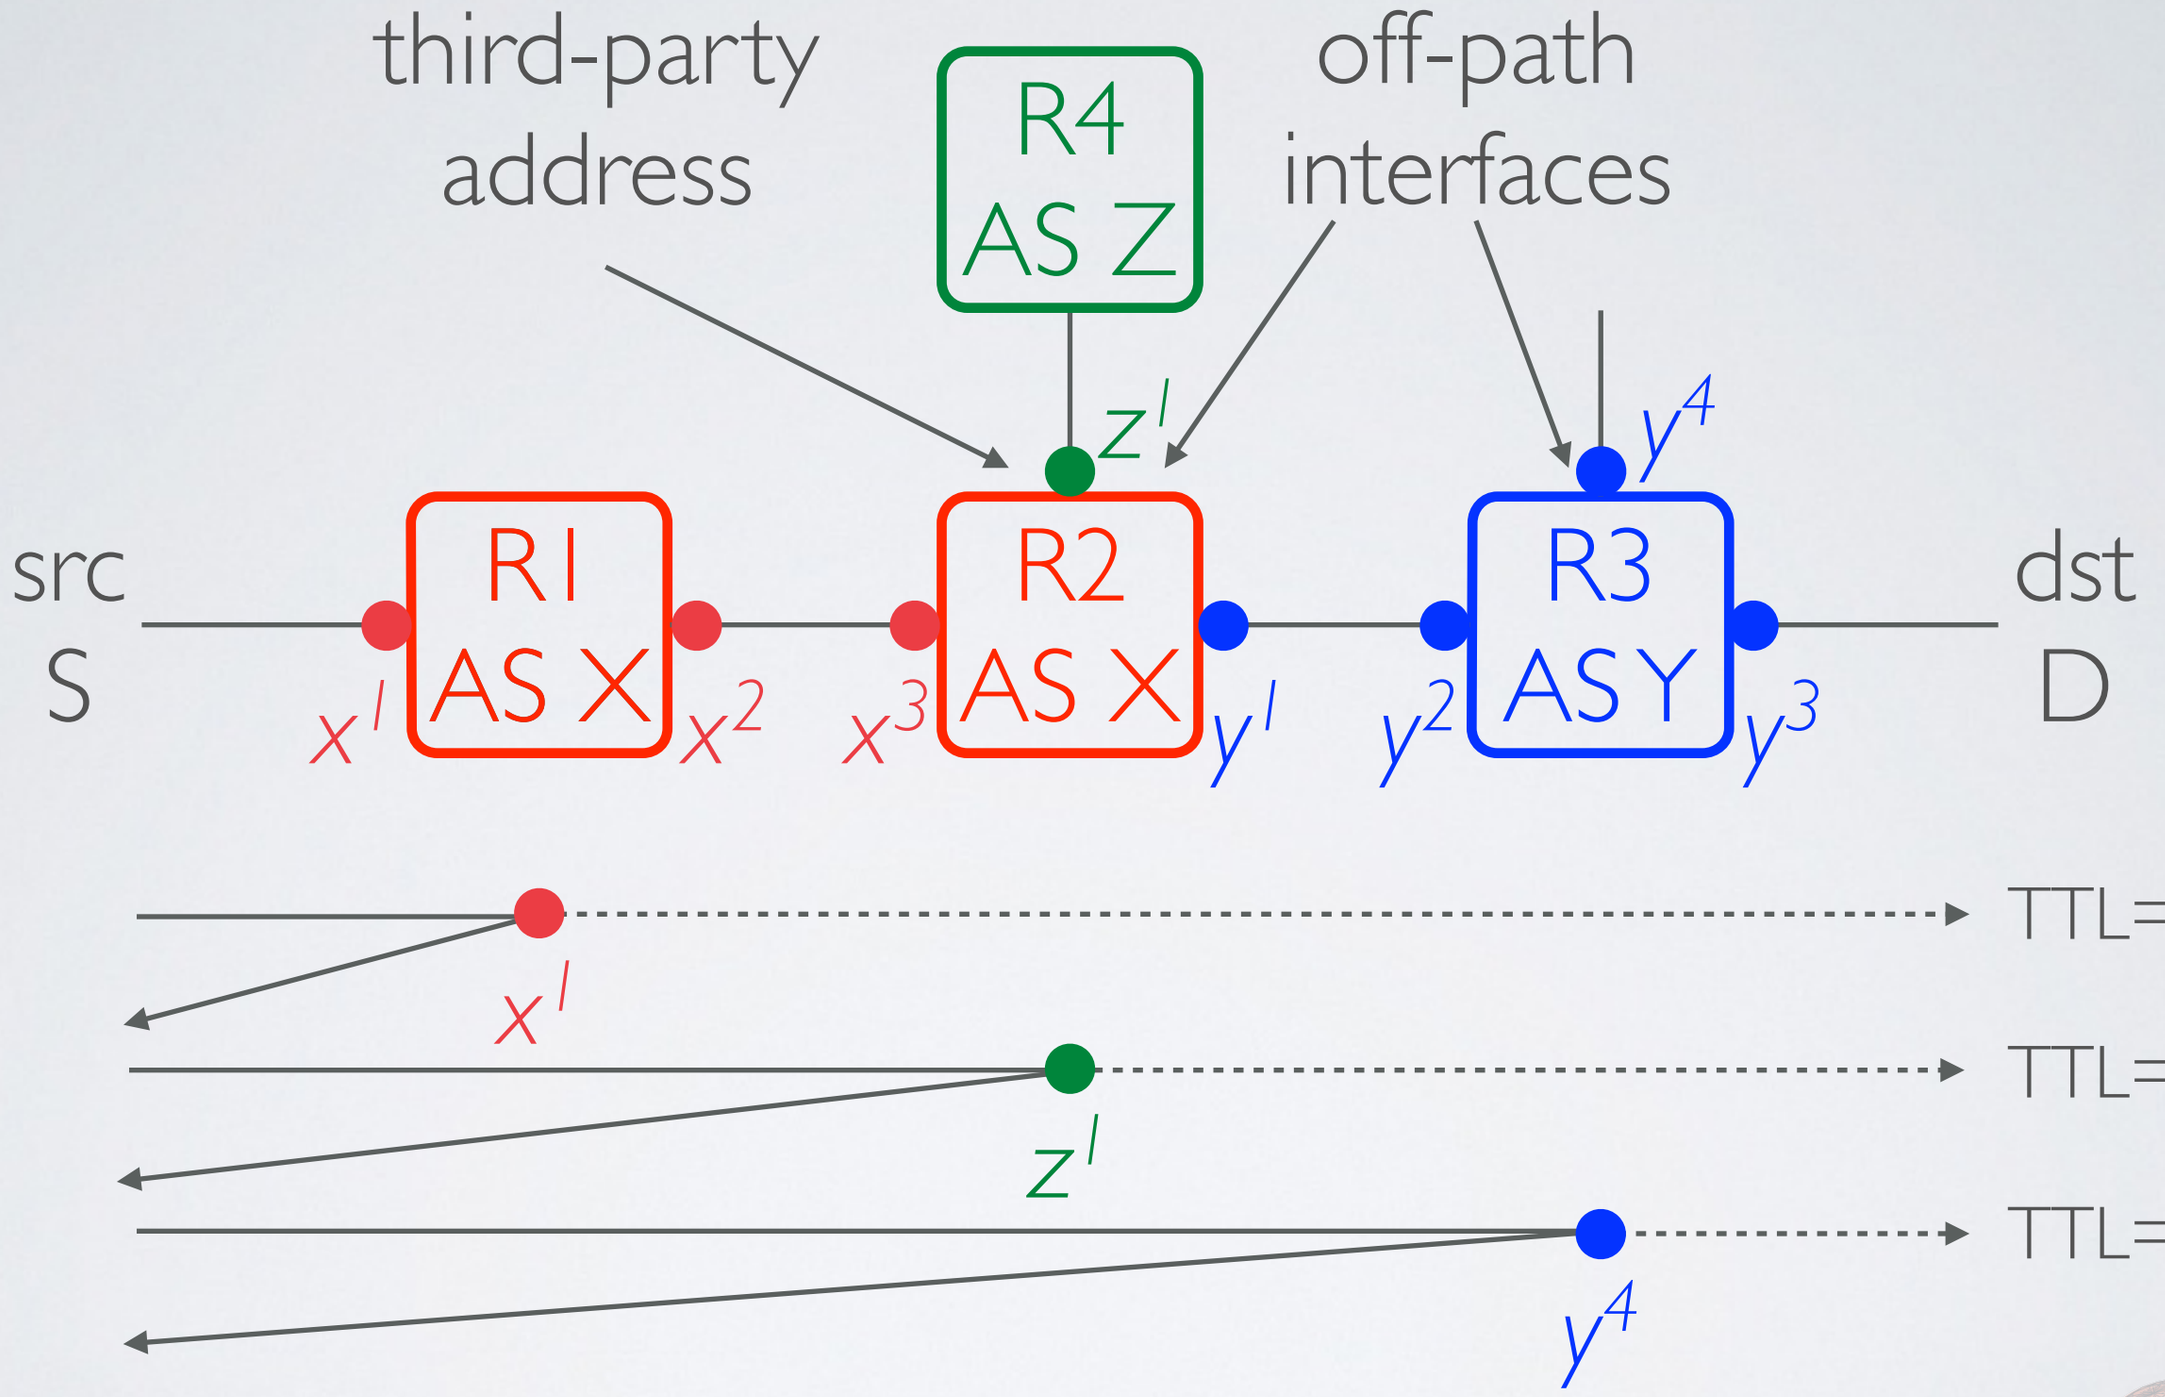

OFF-PATH AND THIRD-PARTY ADDRESSES

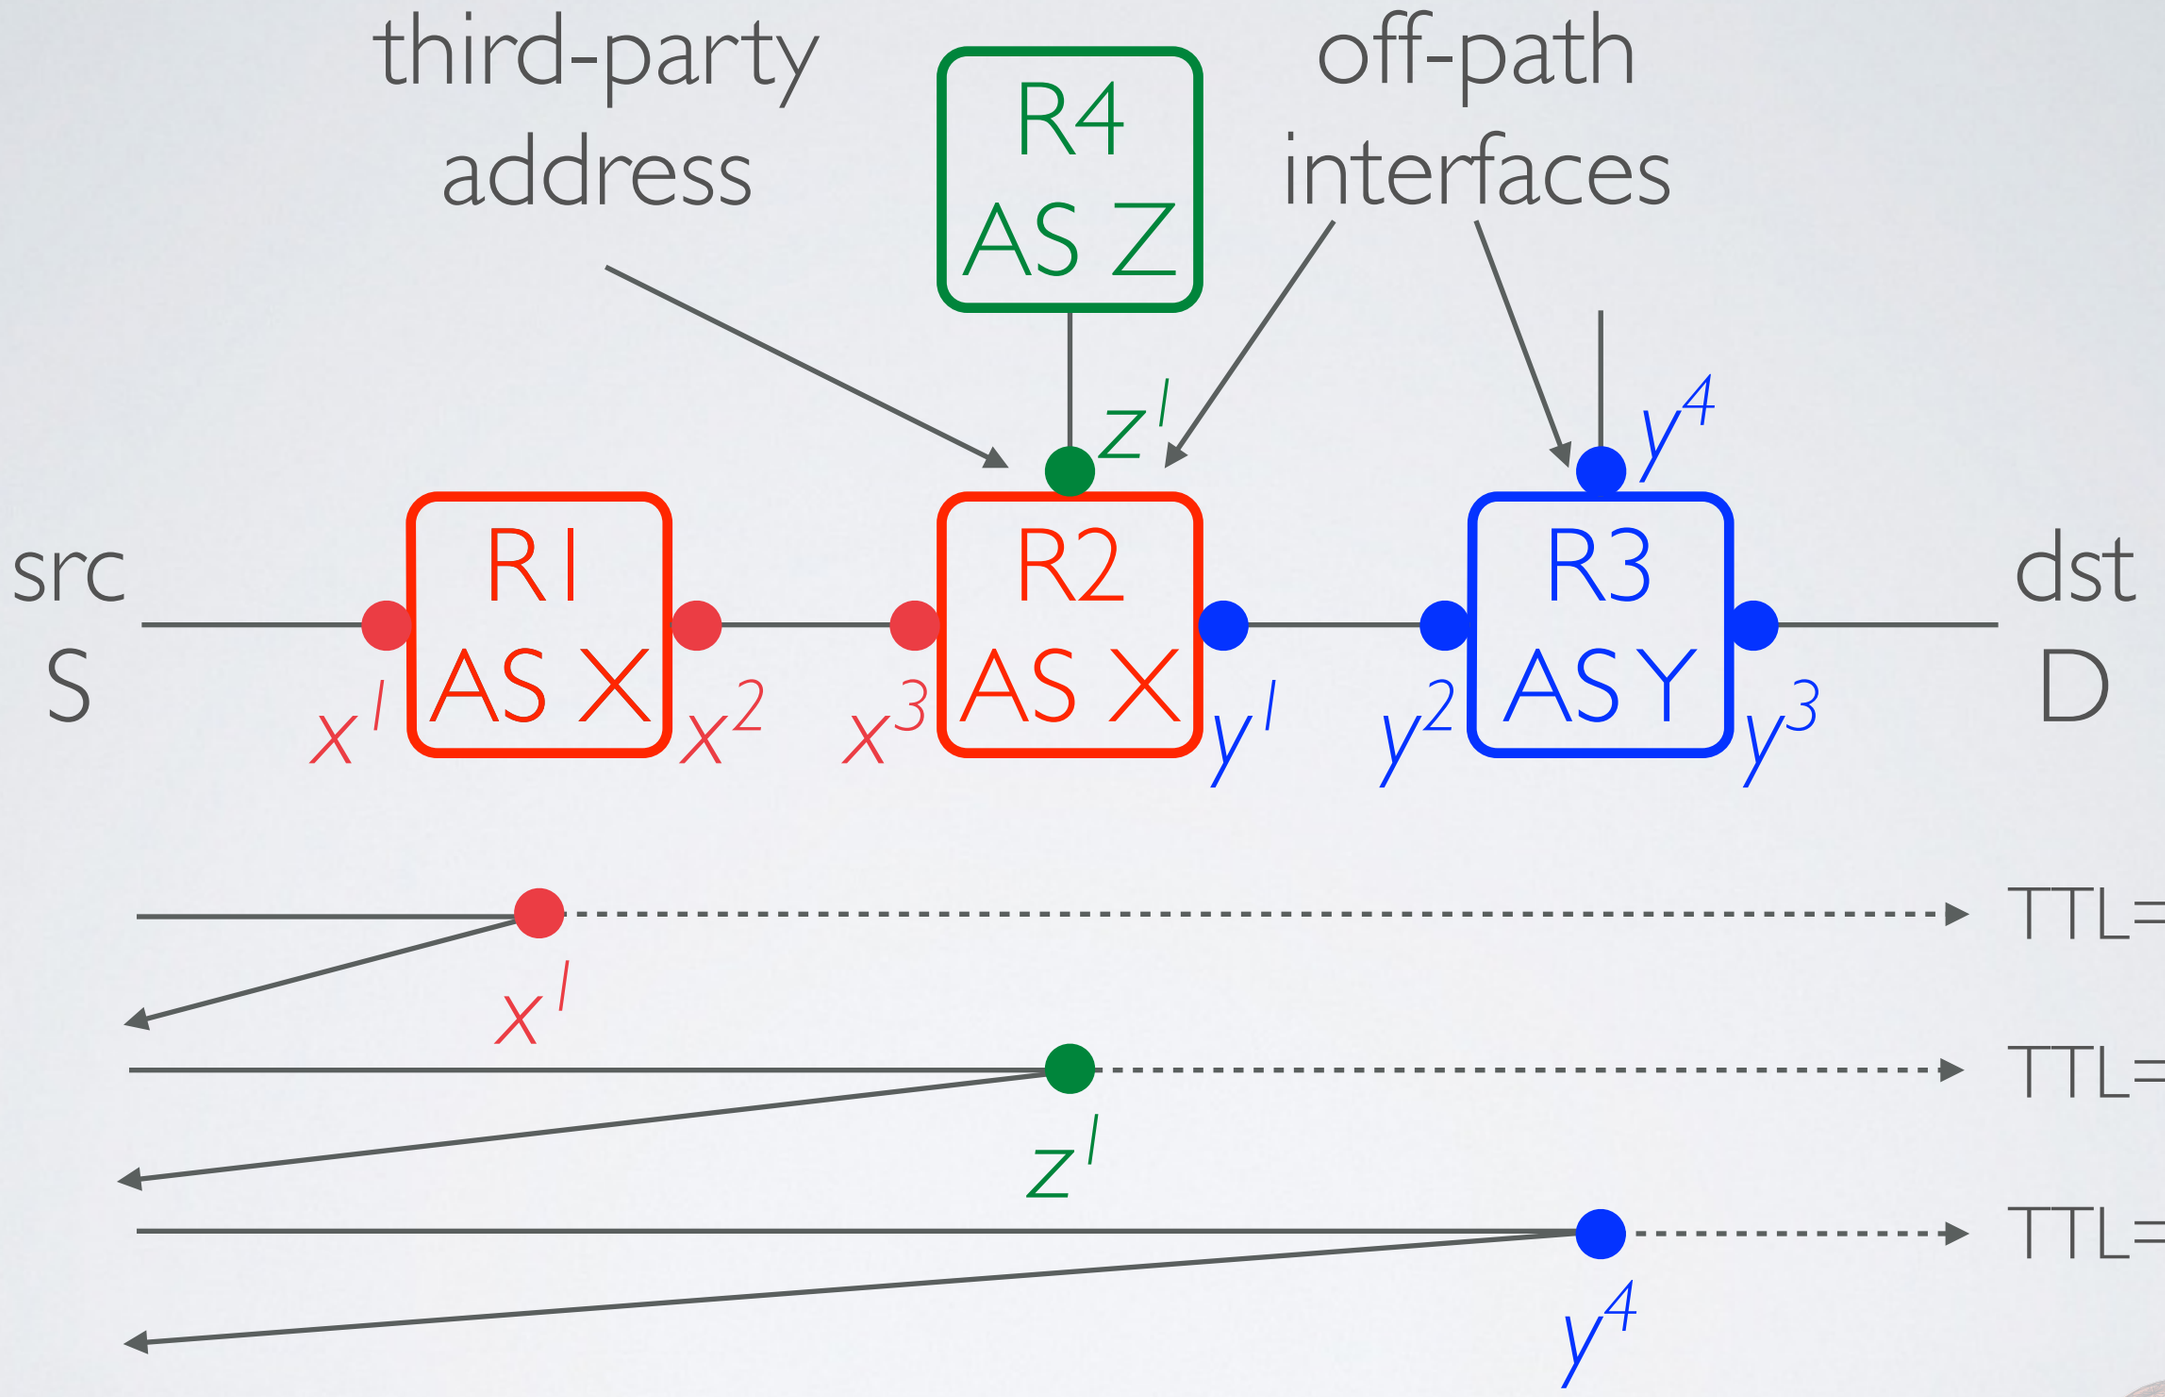

AS path inference: X Z Y

FINDINGS OF MARCHETTA ET AL.

(PAM2013)

- Most classifiable addresses in traceroute paths are off-path
- Consecutive off-path addresses are common
  - More than half of off-path sequences were at least 3 hops
- Presence of off-path addresses in traceroute much more widespread than previously believed

WHAT MOTIVATED OUR WORK?

- If technique and results from Marchetta *et al.* PAM2013 are correct:
  - traceroute is unfit for purpose, operationally and in research.
  - their technique would help us make more solid inferences.
- But: no validation reported.
- Our goal was to assess the correctness of their technique and findings.

IP PRE-SPECIFIED TIMESTAMPS

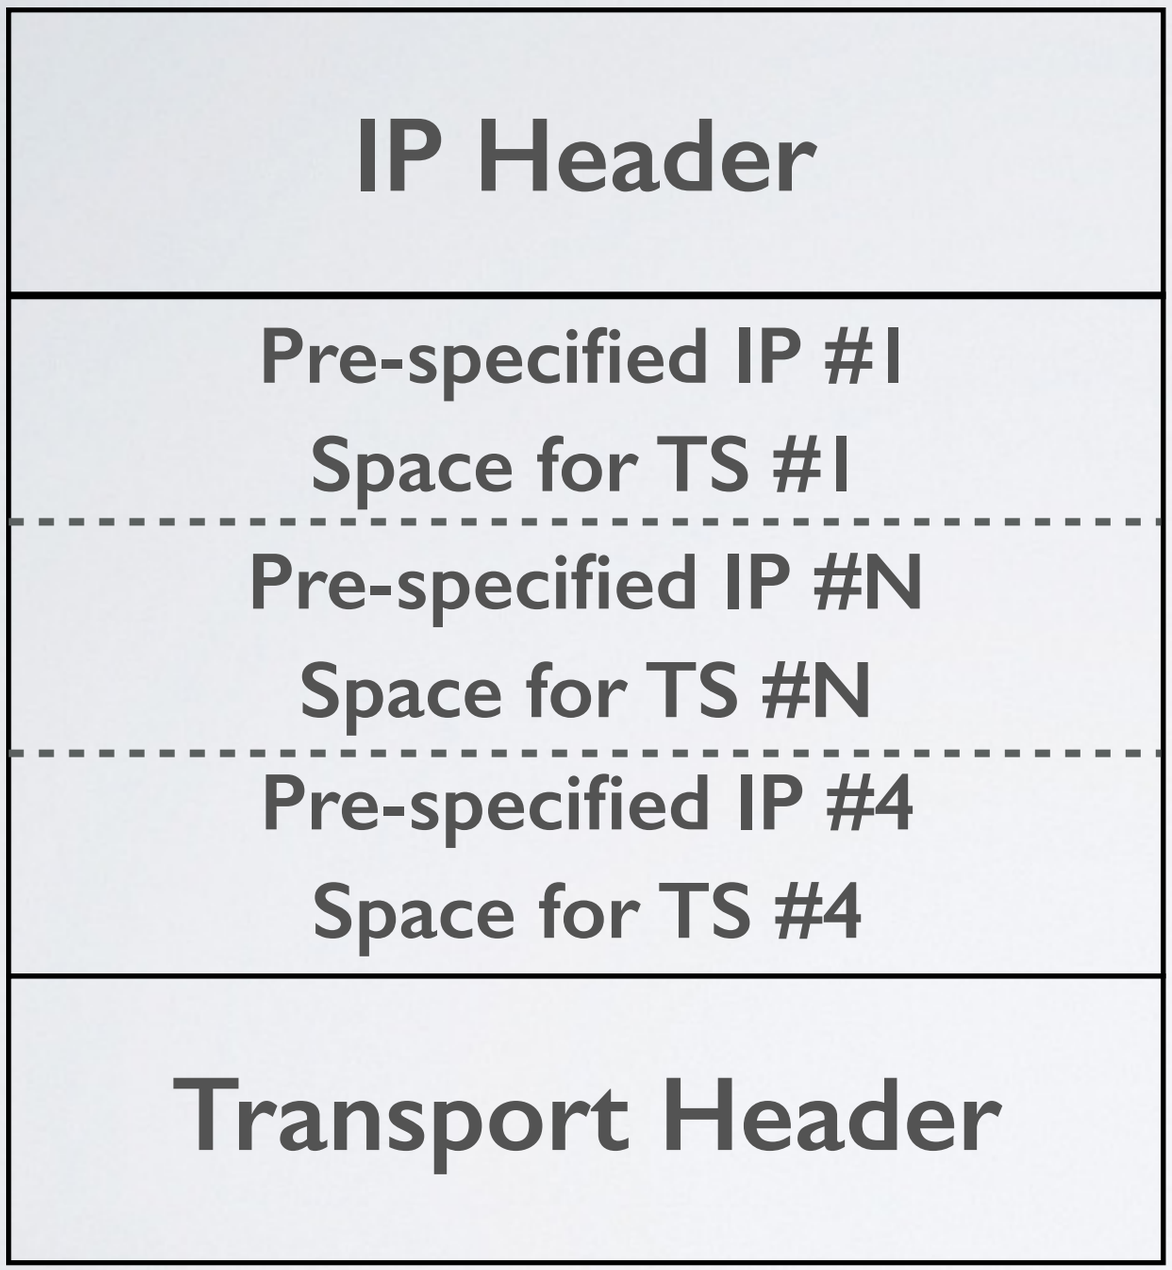

Routers should embed a timestamp if the specified interface is visited

IP Options

Sherry *et al.* notation: a packet sent to Z that asks the routers with addresses A, B, C, D to embed timestamps is written as Z|ABCD

CRITICAL ASSUMPTION

(Marchetta et al. PAM 2013)

If a router **does not embed** a timestamp for a specified IP address when forwarding a packet to a destination *X*, but **does embed** timestamps when packets are sent to the router, then an address observed in traceroute towards *X* was **not in the forwarding path**.

TP-TRACEROUTE ALGORITHM

(Marchetta et al. PAM 2013)

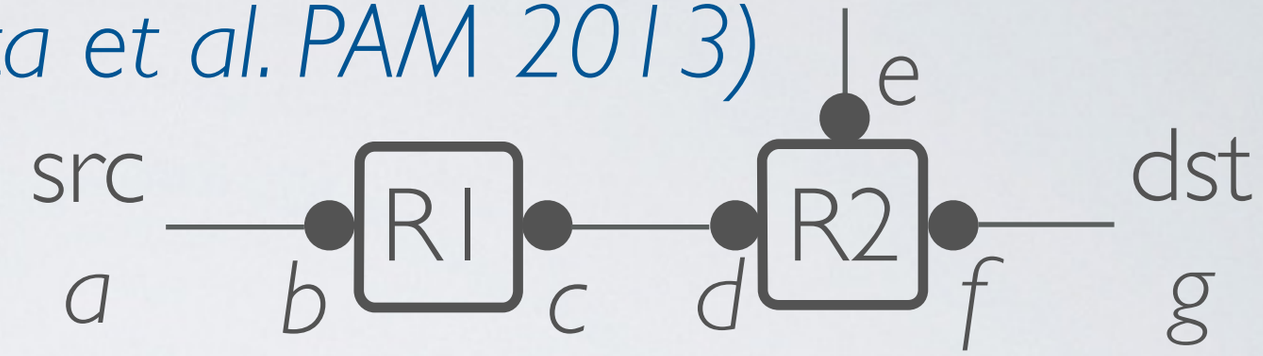

caida

TP-TRACEROUTE ALGORITHM

(Marchetta et al. PAM 2013)

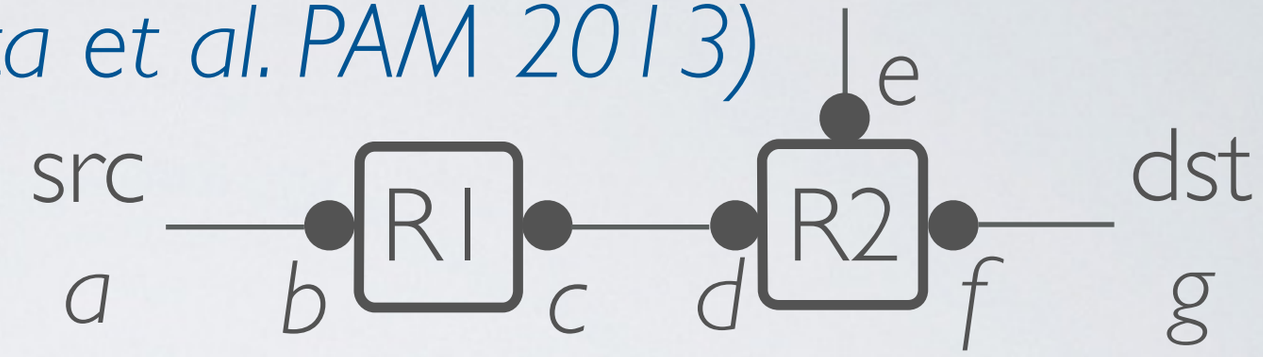

1. UDP Timestamp g|gggg

Port unreach w/ IP timestamp option quoted

If no response, or no timestamp quote, then stop.

TP-TRACEROUTE ALGORITHM

(Marchetta et al. PAM 2013)

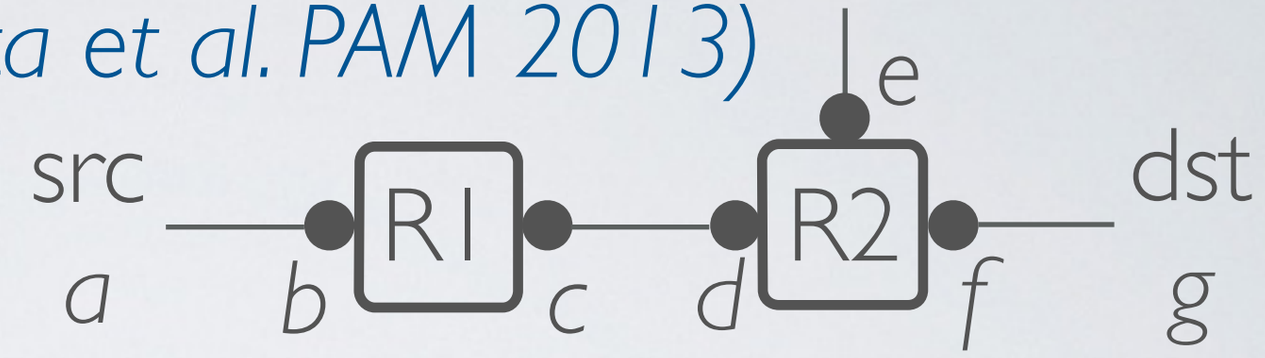

1. UDP Timestamp g|gggg

Port unreach w/ IP timestamp option quoted

2. UDP Traceroute g

b

e

g

Infer the forward IP path towards g.

TP-TRACEROUTE ALGORITHM

(Marchetta et al. PAM 2013)

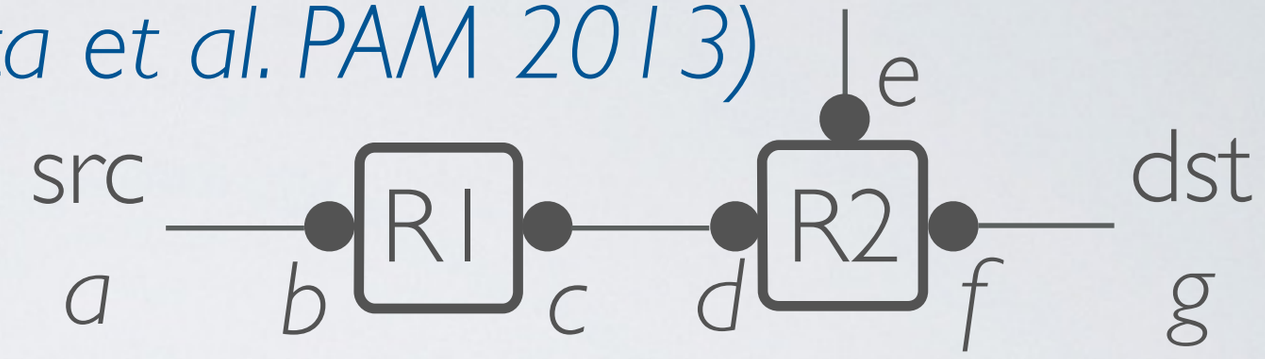

1. UDP Timestamp g|gggg

2. UDP Traceroute g

3. ICMP Timestamp b|bbbb

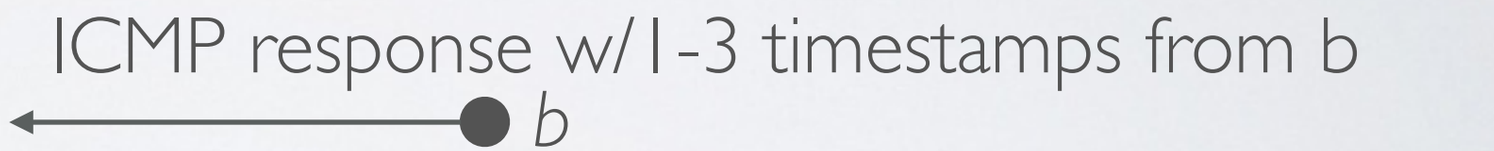

4. ICMP Timestamp e|eeee

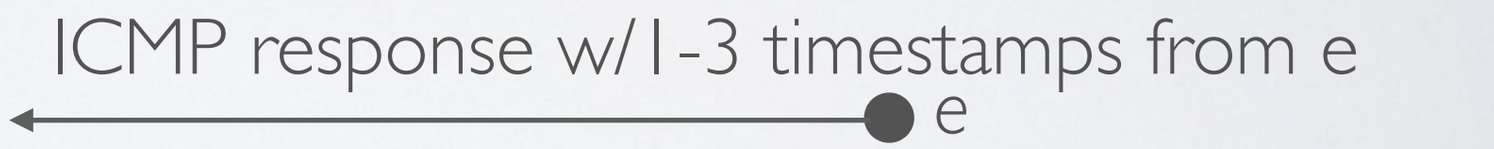

If no response,
or 0 timestamps, (does not support option)
or 4 timestamps, (will embed timestamp regardless of visiting interface)
then tptraceroute cannot infer if the interface was visited or not.

TP-TRACEROUTE ALGORITHM

(Marchetta et al. PAM 2013)

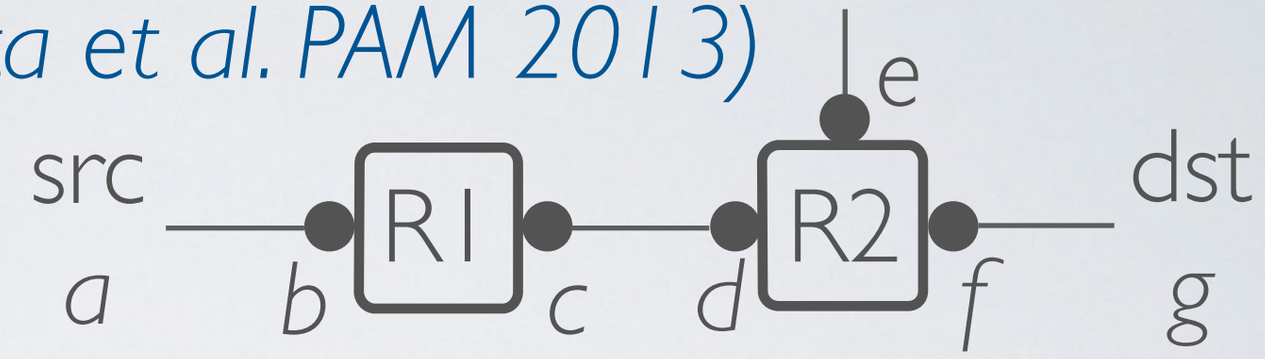

1. UDP Timestamp g|gggg

Port unreach w/ IP timestamp option quoted

2. UDP Traceroute g

b

e

g

3. ICMP Timestamp b|bbbb

ICMP response w/ 1-3 timestamps from b

b

4. ICMP Timestamp e|eeee

ICMP response w/ 1-3 timestamps from e

e

5. UDP Timestamp g|bbbb

Port unreach w/ 1-3 timestamps from b

g

Because 1-3 timestamps were inserted by b, infer it is on-path

TP-TRACEROUTE ALGORITHM

(Marchetta et al. PAM 2013)

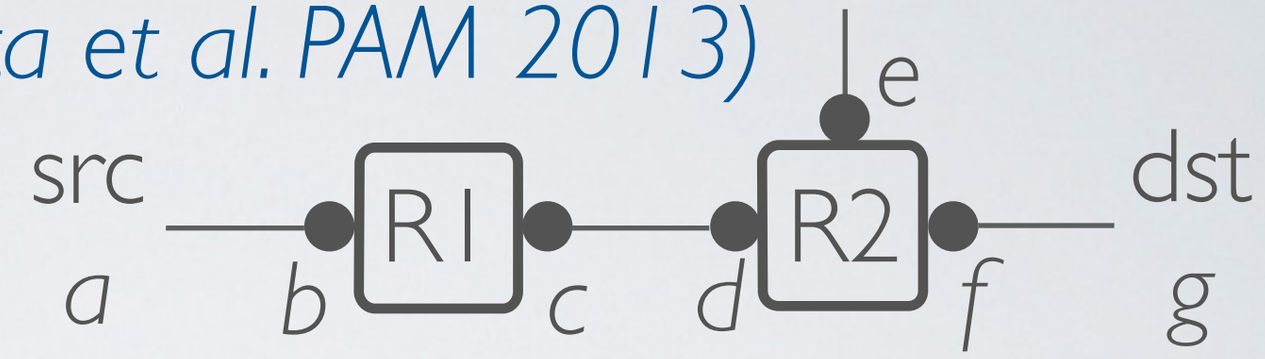

1. UDP Timestamp g|gggg — Port unreach w/ IP timestamp option quoted
2. UDP Traceroute g — *b*, e, g
3. ICMP Timestamp b|bbbb — ICMP response w/1-3 timestamps from b
4. ICMP Timestamp e|eeee — ICMP response w/1-3 timestamps from e
5. UDP Timestamp g|bbbb — Port unreach w/1-3 timestamps from b
6. UDP Timestamp g|eeee — Port unreach w/0 timestamps from e

Because 0 timestamps were inserted by e, infer it is off-path

CROSS VALIDATION OF TP-TRACEROUTE

- Limit ourselves to interfaces we infer are the in-bound interface on a router:
  - **For those interfaces, what inference does tp-traceroute make?**

CROSS VALIDATION OF TP-TRACEROUTE

- Limit ourselves to interfaces we infer are the in-bound interface on a router:
  - **For those interfaces, what inference does tp-traceroute make?**

CROSS VALIDATION OF TP-TRACEROUTE

- Limit ourselves to interfaces we infer are the in-bound interface on a router:
  - **For those interfaces, what inference does tp-traceroute make?**

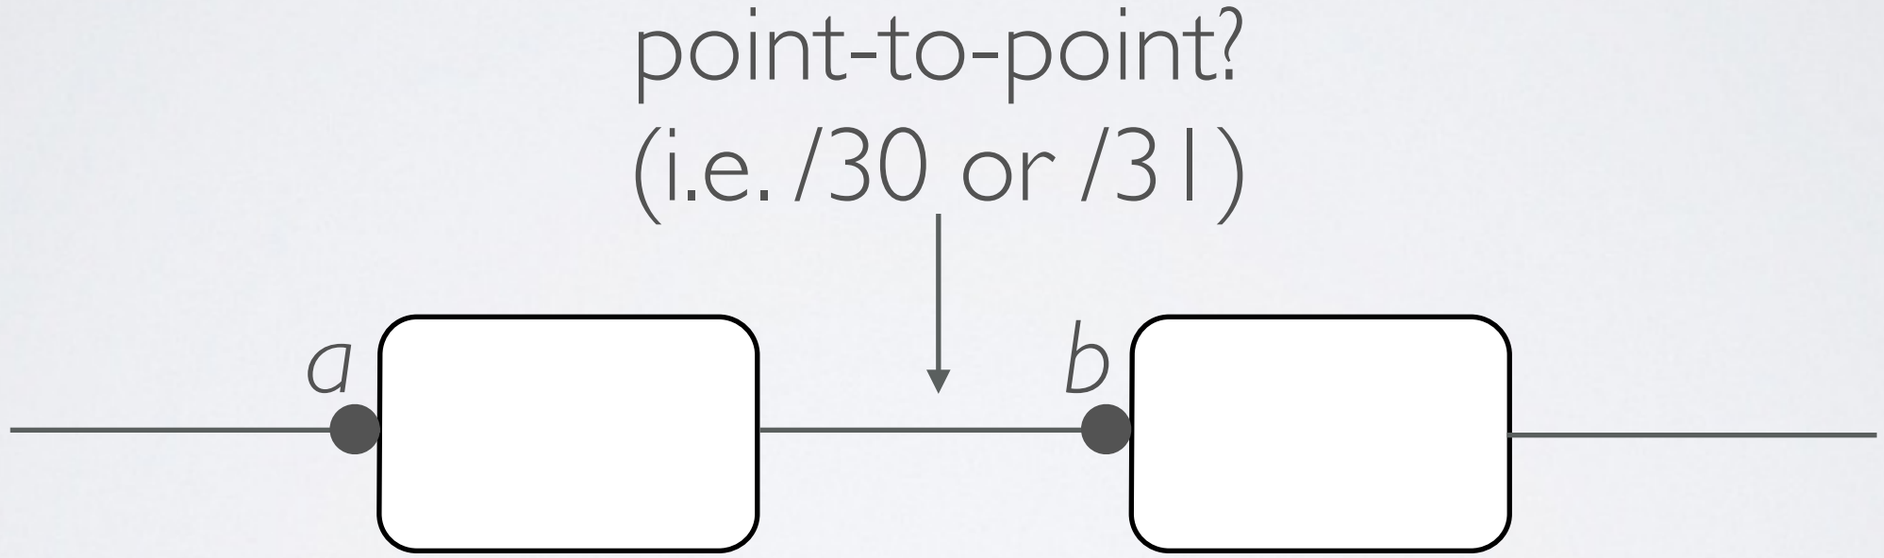

CROSS VALIDATION OF TP-TRACEROUTE

- Limit ourselves to interfaces we infer are the in-bound interface on a router:
  - **For those interfaces, what inference does tp-traceroute make?**

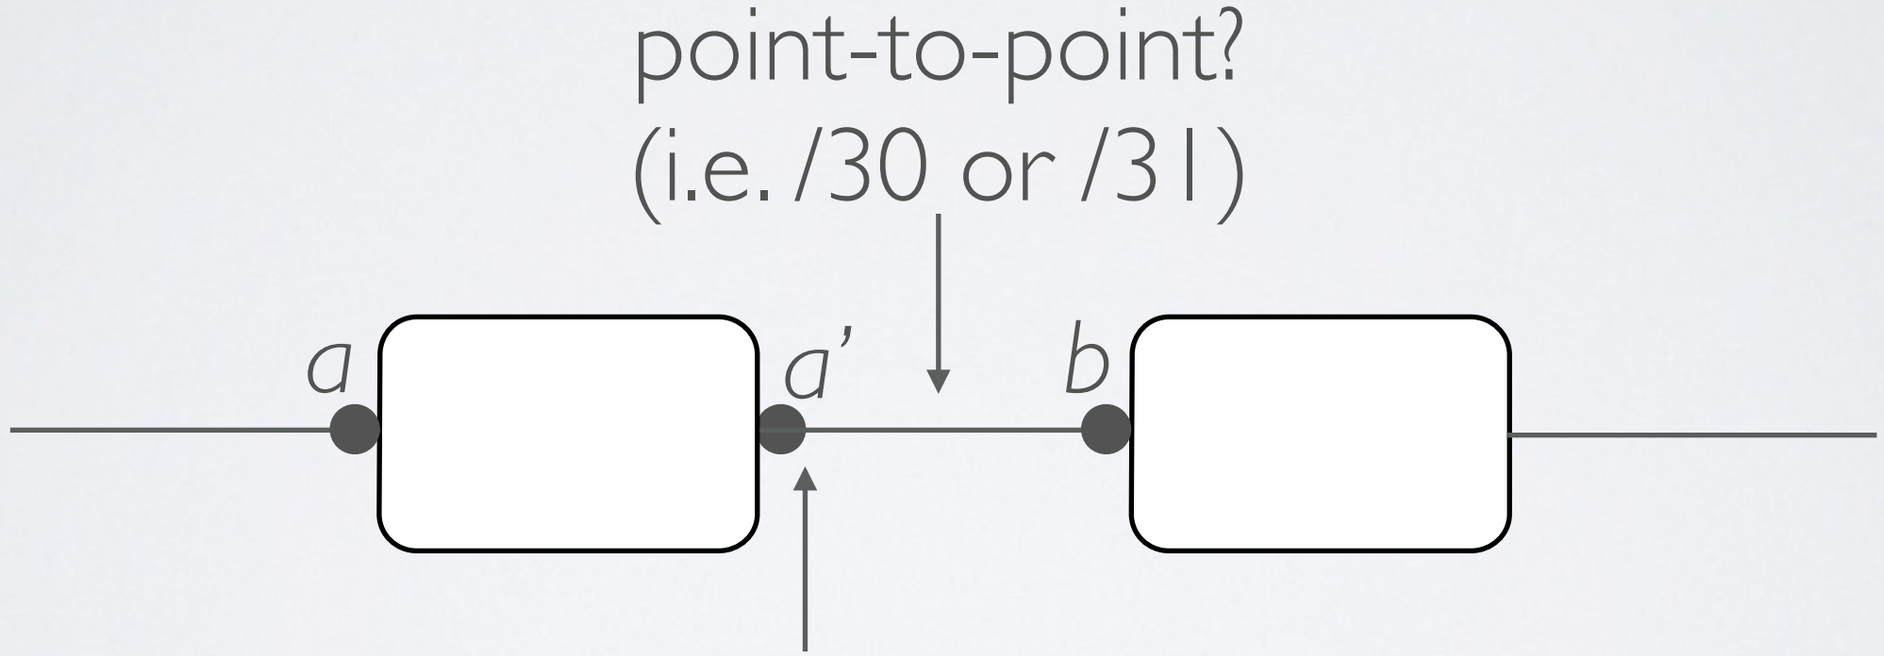

caida

ARE *a*+*a*' ON THE SAME ROUTER?

- Also known as **alias resolution**

  - **extensive validation history:**
    Rocketfuel SIGCOMM2002, **Radargun** IMC2008, **MIDAR** ToN2013

- Two techniques used in this work:

  - repeated **Ally**-style tests

    - using ICMP-echo, TCP-ack, and UDP probes

    - monotonic IPID sequence from non-overlapping probes/replies to *a* and *a*', repeated every 10 minutes for an hour to allow divergence

  - one-off **Mercator** test (if necessary)

    - responses to probes to *a* and *a*' come from common source address

OUR METHOD

- Eight CAIDA Archipelago (Ark) vantage points (VPs)
- Each obtained 10K traceroutes to **responding** destinations chosen at random from ISI Census data
  - for each hop, classify as on- or off-path using **tp-traceroute** (Marchetta *et al.*) technique
  - for each link, infer if point-to-point (our cross-validation) using **prefixscan** in scamper

OUR DATA

- 197,335 IP-level links
- 81,315 inferred point-to-point IP links (**interface B on-path**)
  - In our data, between 77.1% and 90.0% of interfaces in traceroute on point-to-point links were classified by tp-traceroute **as off-path**

Pre-specified IP Timestamps are an unreliable primitive to determine if an address is on- or off-path

WHAT FRACTION OF INTERFACES IN TRACEROUTE ARE IN-BOUND?

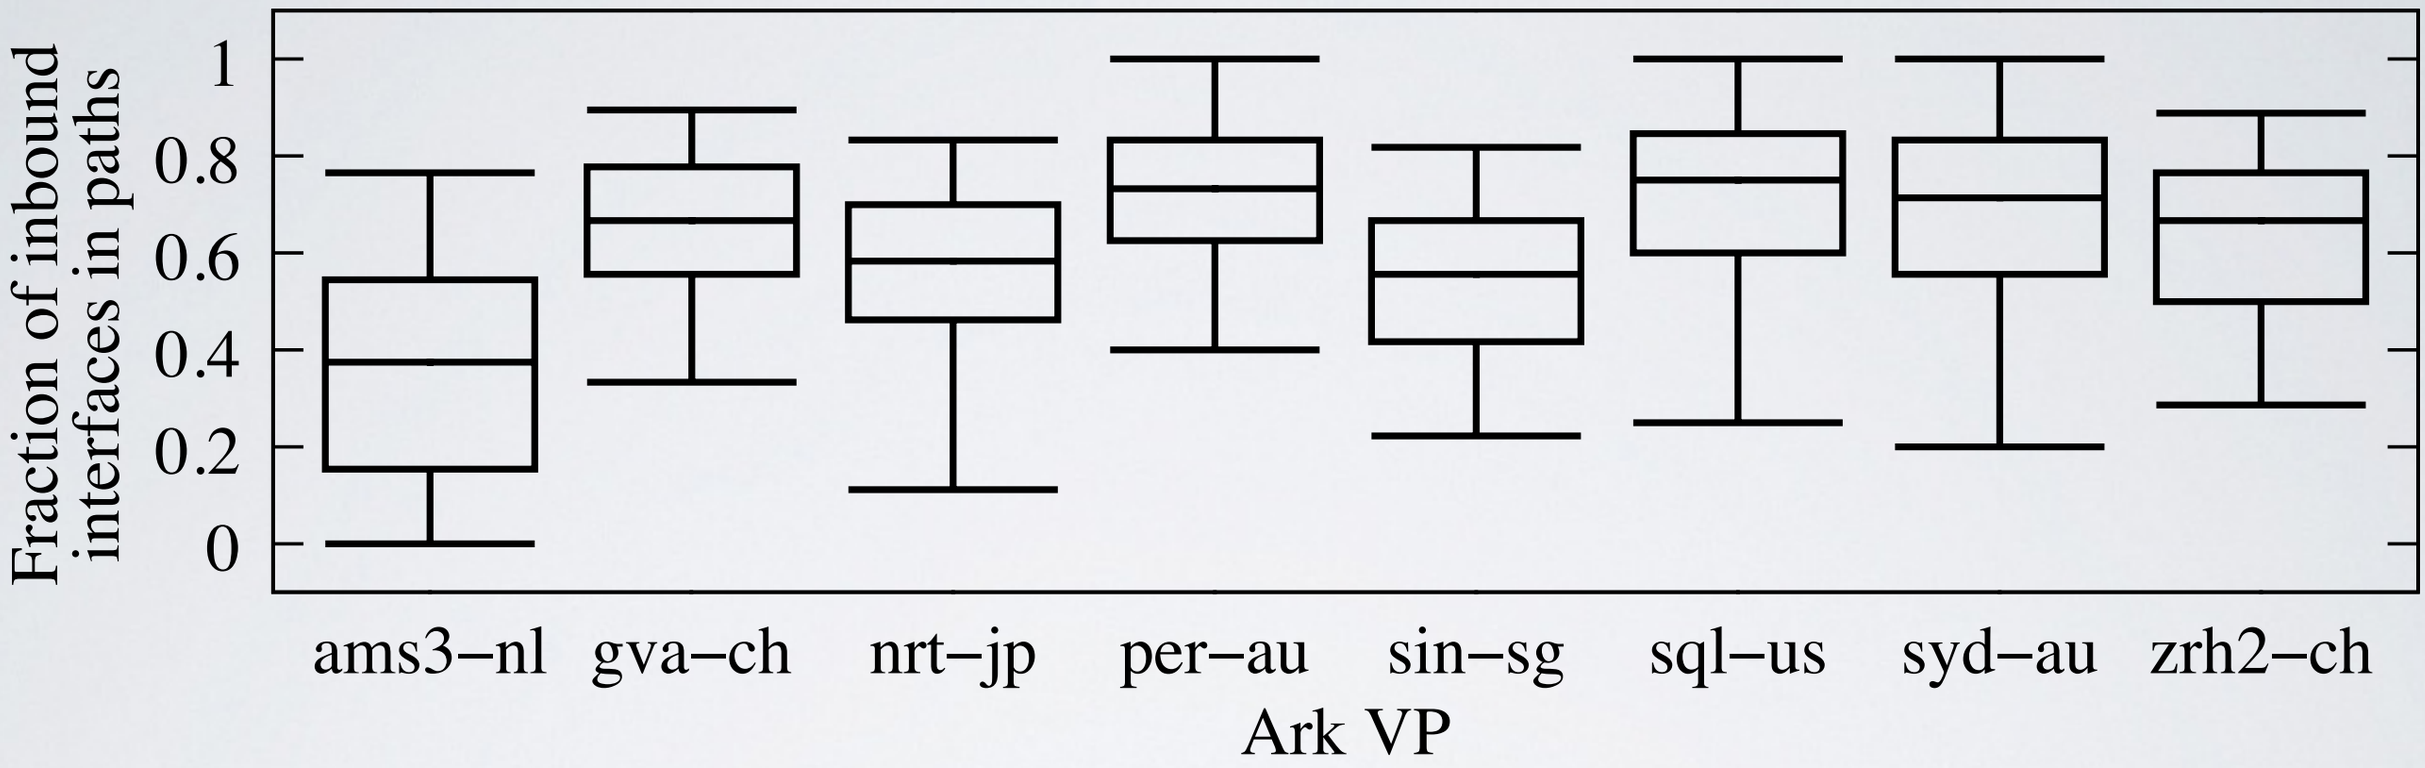

For 7 of 8 VPs, more than half of the interfaces observed in a traceroute were in-bound.

Lower-bound: these are just the routers we could resolve for aliases.

LIMITATION

(multiple point-to-point links can exist between routers, and the address observed in traceroute might be off-path)

LIMITATION

(multiple point-to-point links can exist between routers, and the address observed in traceroute might be off-path)

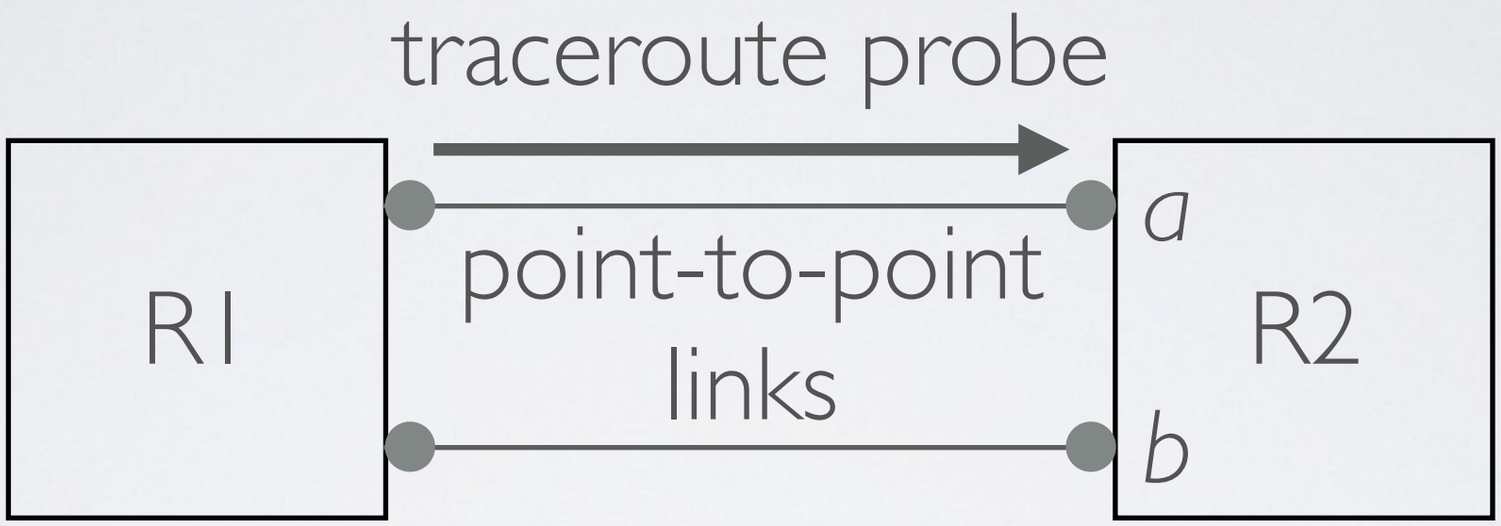

LIMITATION

(multiple point-to-point links can exist between routers, and the address observed in traceroute might be off-path)

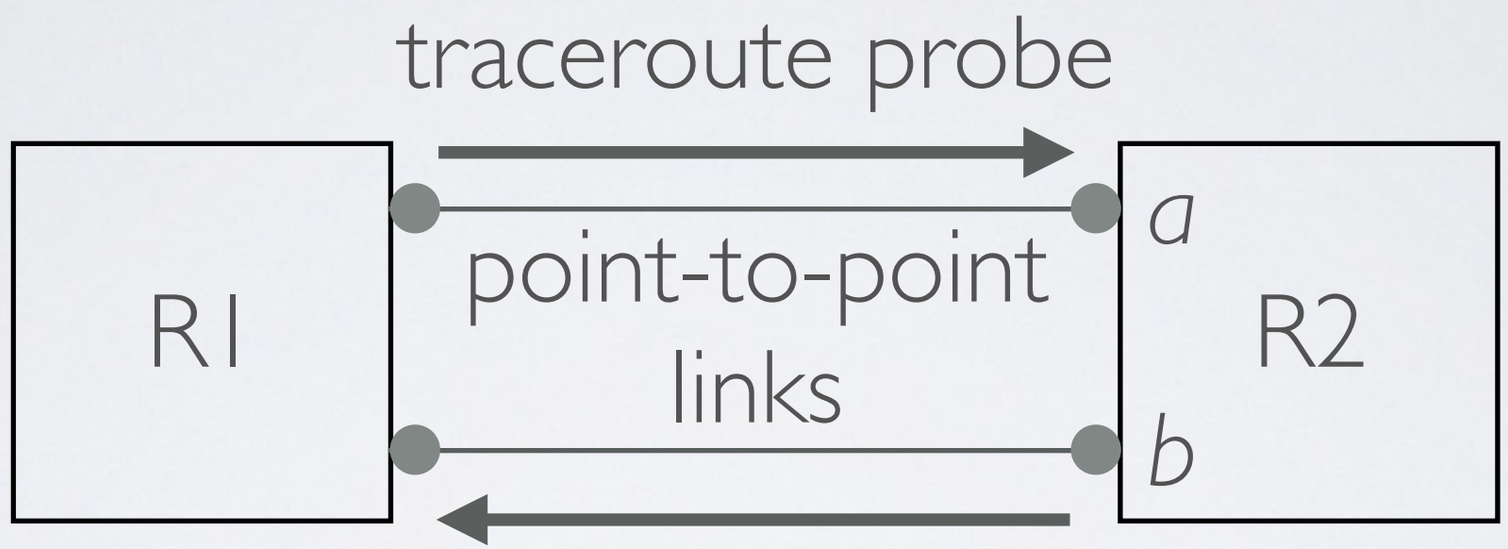

SUMMARY

- Traceroute has an important role in overcoming the visibility issue of AS topology data because we have no other way of uncovering some peerings.

- Traceroute-derived AS paths are messy to work with due to artifacts in IP2AS mappings.

- Presence of off-path addresses not as wide-spread as suggested by Marchetta *et al.* in PAM2013.

- Deriving a technique that accurately infers AS links from traceroute paths remains an important and currently unsolved problem

CONFIDENCE IN ALIAS INFERENCES

- Alias resolution has an extensive validation history
  - **Rocketfuel** SIGCOMM2002, **Radargun** IMC2008, **MIDAR** ToN2013
- IP-ID techniques are generally reliable when they declare two addresses as aliases

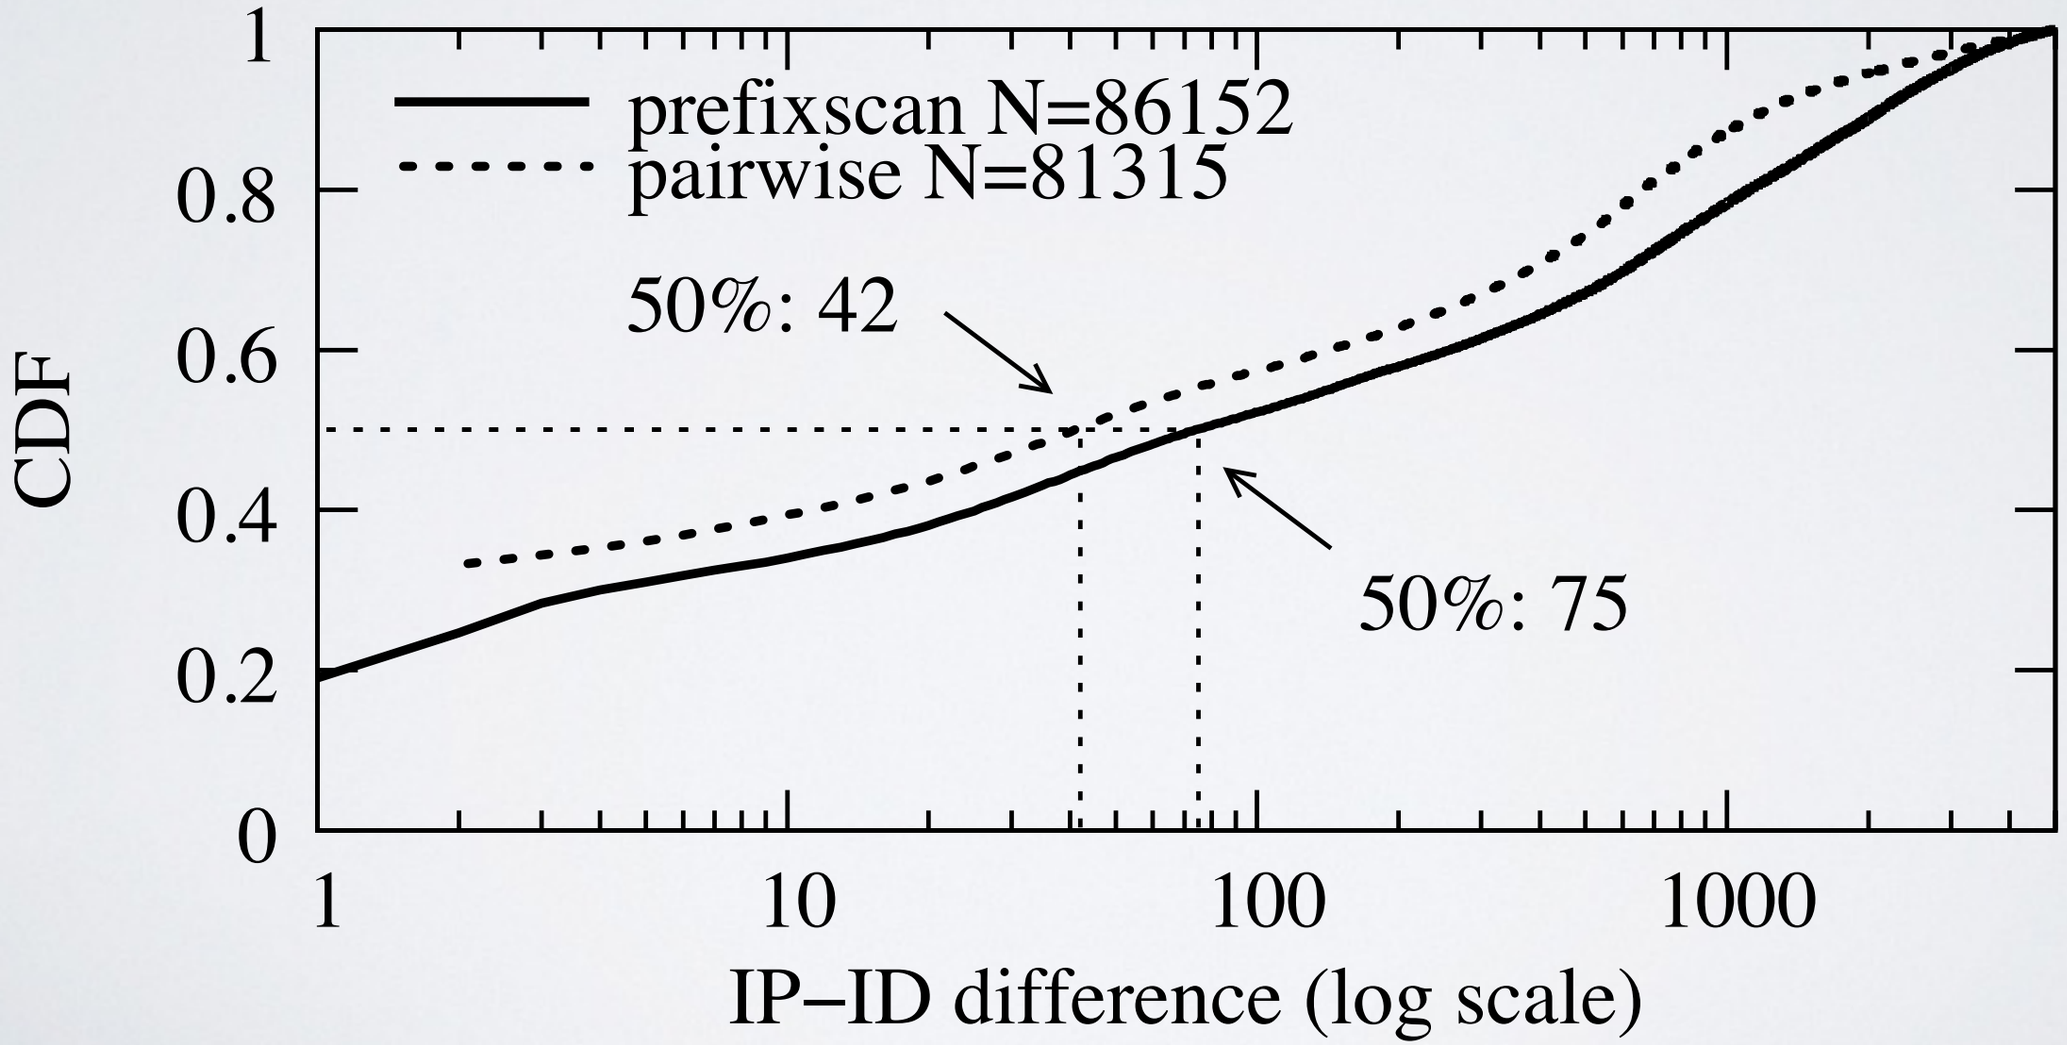

SAMPLING OF INTERFACES

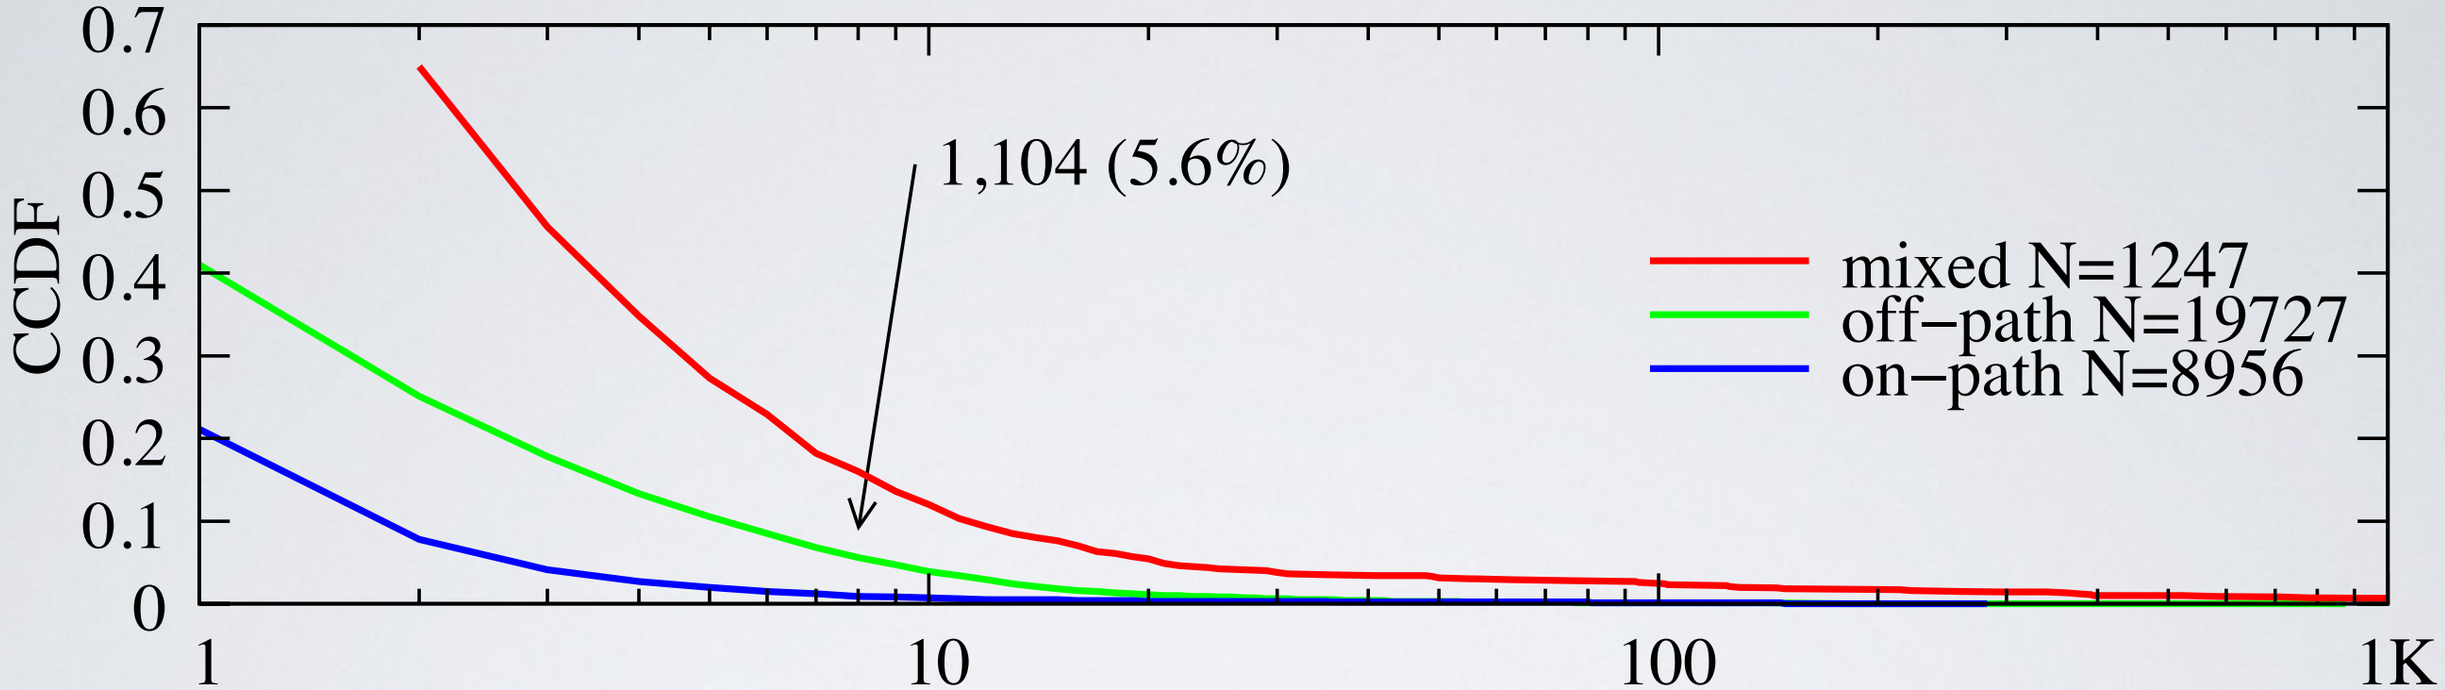

5.6% of interfaces that we inferred as on-path were traversed in at least 8 source-destination peers with UDP G|BBBB packets.

We are at least 99% confident the previous hop did not load balance UDP G|BBBB packets on a path avoiding B.

caida

MIXED CLASSIFICATIONS

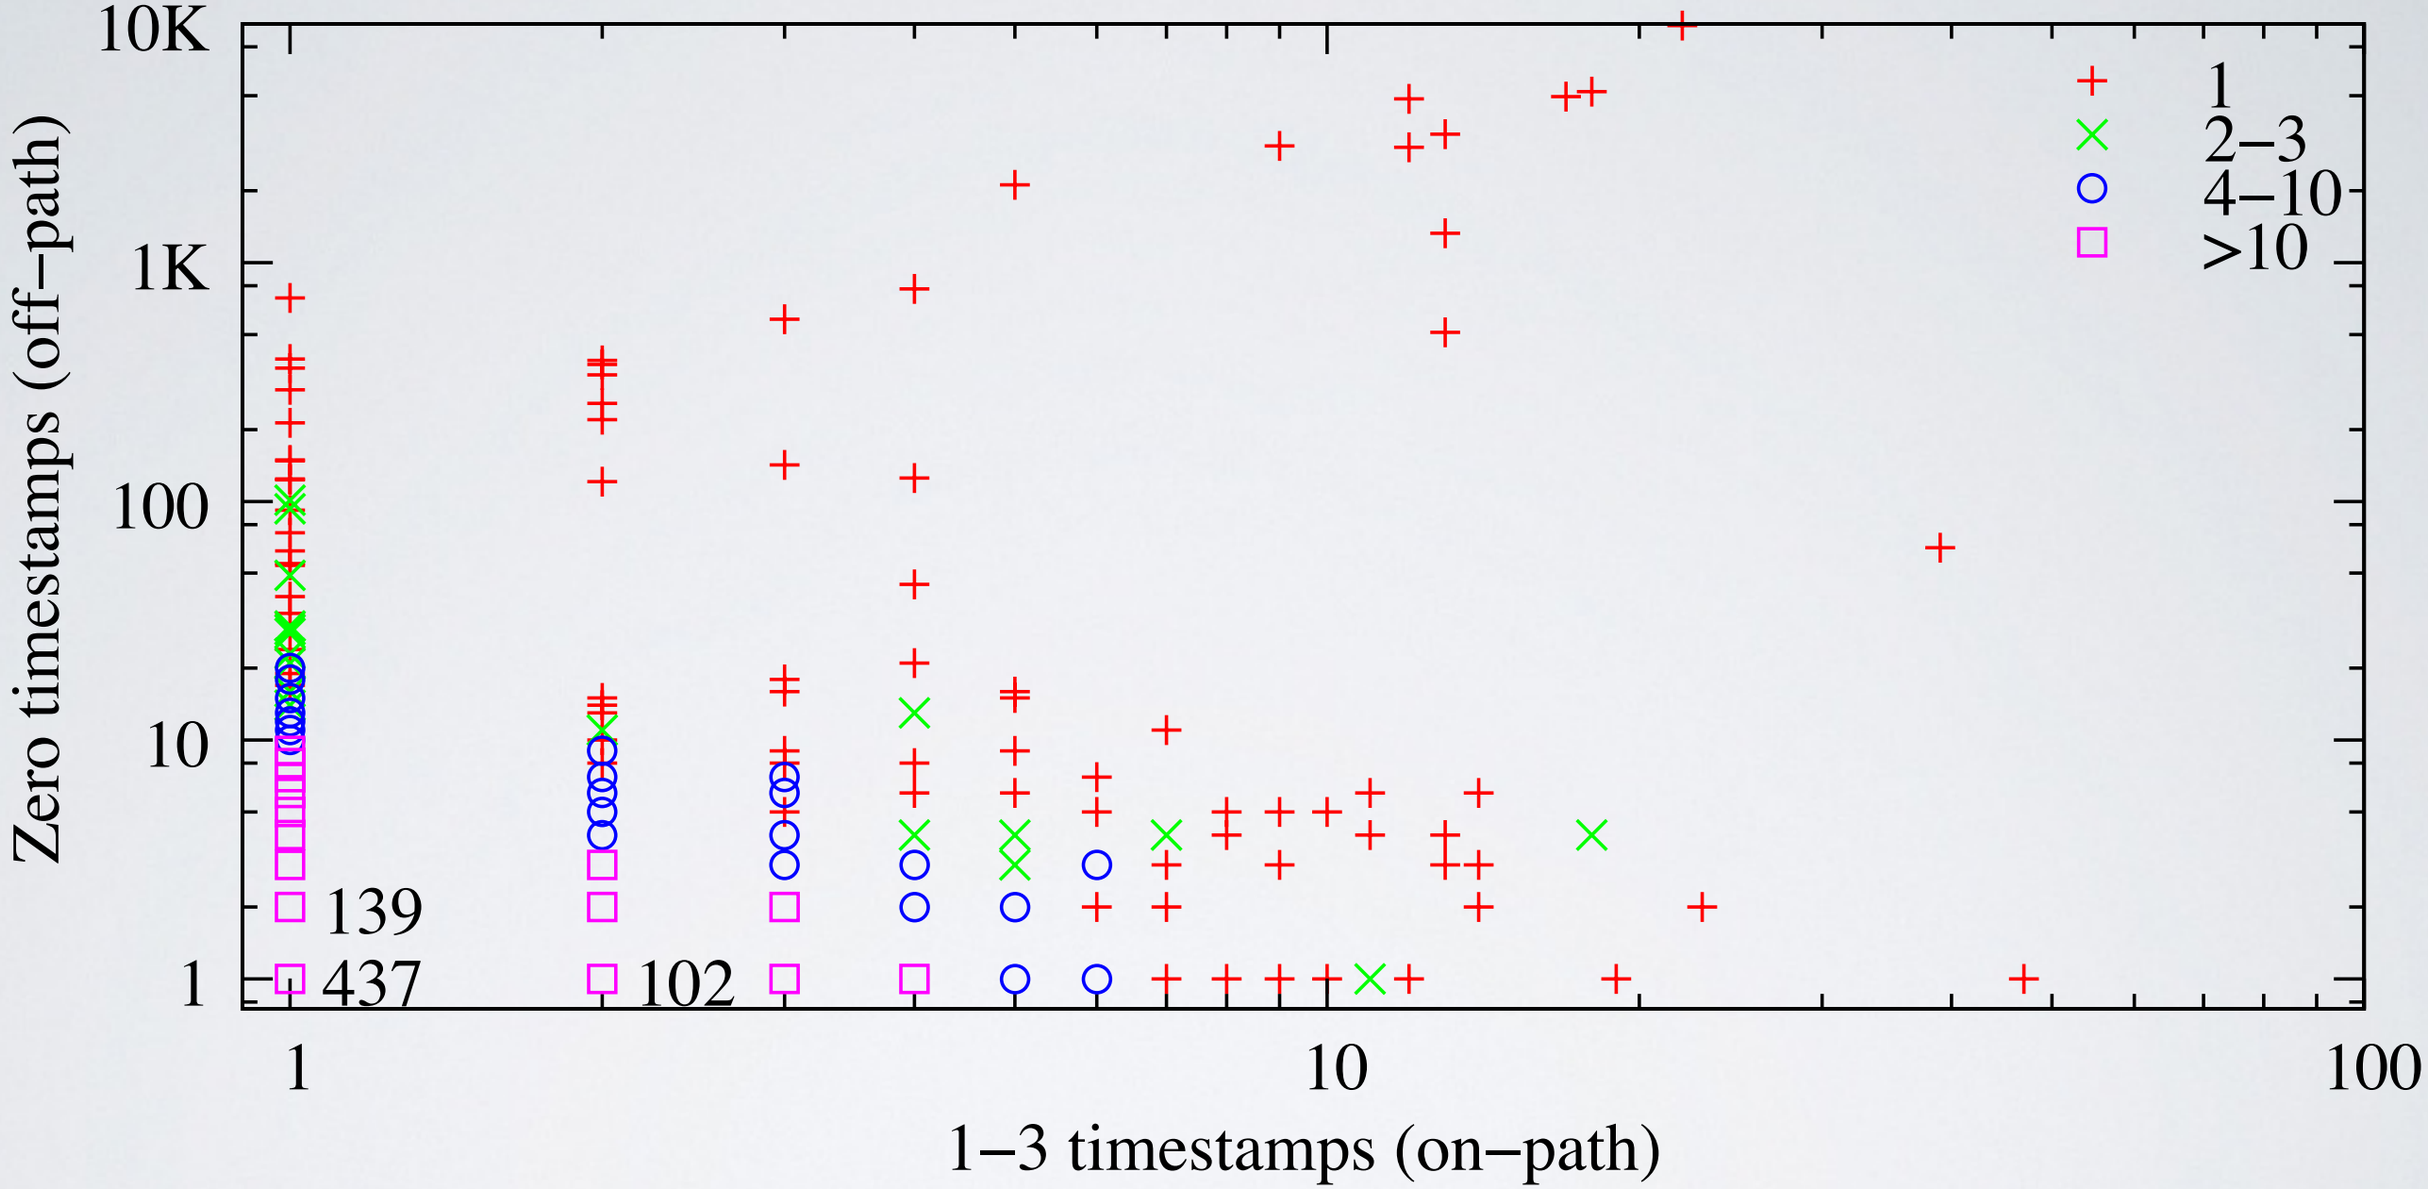

Most interfaces with mixed behavior appeared as on-path for just one source-destination pair