Touring the Internet in a TCP Sidecar

Rob Sherwood   Neil Spring

University of Maryland

IMC 2006

Topology Discovery: Along for the Ride

Goal: Internet's complete router-level topology

- Challenges:
  - Accuracy: Noisy data creates false links, nodes
  - Completeness: Sections difficult to probe
  - Validation: No complete map exists
- Tools/Contributions:
  - Passenger: augment probes with IP RR option
  - Sidecar: attach probes to TCP streams
- Previous work:
  - Limited coverage
  - Predominantly traceroute-based

Understanding Traceroute

Tool that uses TTL-limited probes to discover routers along the path from source to destination.

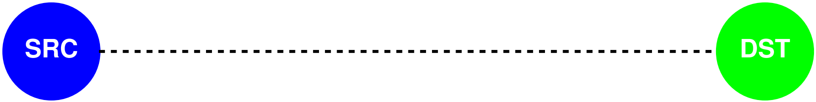

Understanding Traceroute

Tool that uses TTL-limited probes to discover routers along the path from source to destination.

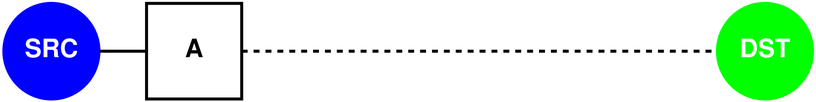

Understanding Traceroute

Tool that uses TTL-limited probes to discover routers along the path from source to destination.

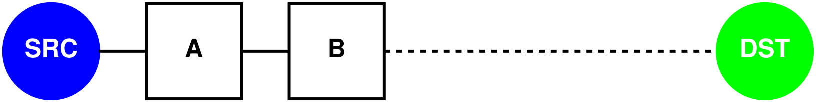

Understanding Traceroute

Tool that uses TTL-limited probes to discover routers along the path from source to destination.

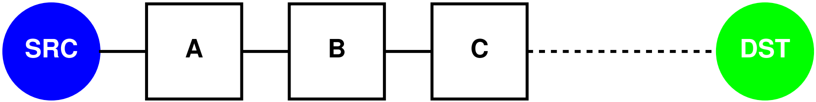

Understanding Traceroute

Tool that uses TTL-limited probes to discover routers along the path from source to destination.

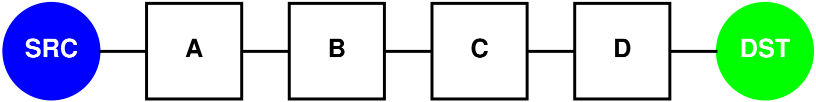

Understanding Traceroute

Tool that uses TTL-limited probes to discover ~~routers~~ interface addresses along the path from source to destination.

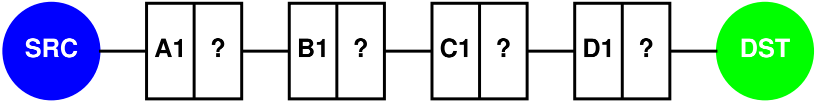

Understanding Traceroute

Tool that uses TTL-limited probes to discover ~~routers~~ interface addresses along ~~the path~~ some set of paths from source to destination.

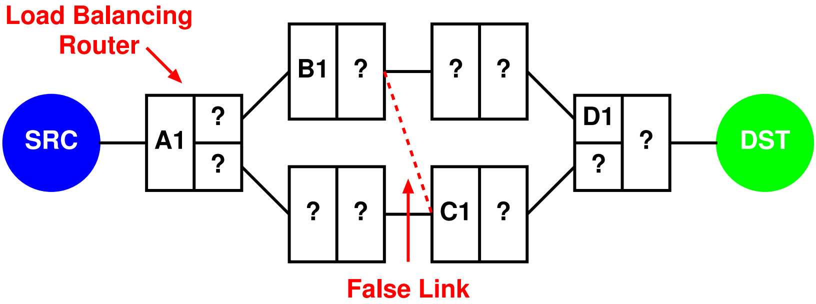

Understanding Traceroute

Tool that uses TTL-limited probes to discover ~~routers~~ interface addresses except hidden routers along ~~the path~~ some set of paths from source to destination.

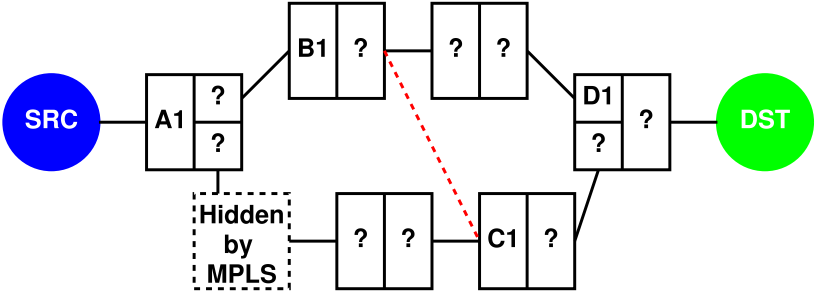

Understanding Traceroute

Tool that uses TTL-limited probes to discover ~~routers~~ interface addresses except hidden routers along ~~the path~~ some set of paths from source to destination unless behind firewall or NAT.

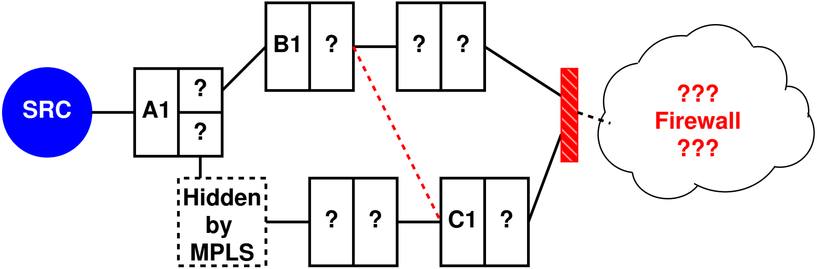

Understanding Traceroute

Tool that uses TTL-limited probes to discover ~~routers~~ interface addresses except hidden routers along ~~the path~~ some set of paths from source to destination unless behind firewall or NAT.

- Will never see layer 2 devices or backup links.

Understanding Traceroute

Tool that uses TTL-limited probes to discover ~~routers~~ interface addresses except hidden routers along ~~the path~~ some set of paths from source to destination unless behind firewall or NAT.

- Will never see layer 2 devices or backup links.
- Will see abuse reports confusing traceroute probes for attack traffic.

Summary of Traceroute's Limitations

- Unresolved aliases: false nodes
- Undetected multiple paths: false links
- Fails to discover hidden routers

- Firewalls/NATs block probes
- Abuse reports limit scope of experiments

- Layer 2 devices and unused/backup links never discovered

Mitigating Traceroute's Limitations

Passenger: IP Record Route Option

- Unresolved aliases: false nodes
- Undetected multiple paths: false links
- Fails to discover hidden routers

- Firewalls/NATs block probes
- Abuse reports limit scope of experiments

- Layer 2 devices and unused/backup links never discovered

Mitigating Traceroute's Limitations

Passenger: IP Record Route Option

- Unresolved aliases: false nodes
- Undetected multiple paths: false links
- Fails to discover hidden routers

Sidecar: Transparently embedded probes

- Firewalls/NATs block probes
- Abuse reports limit scope of experiments

- Layer 2 devices and unused/backup links never discovered

Mitigating Traceroute's Limitations

Passenger: IP Record Route Option

- Unresolved aliases: false nodes
- Undetected multiple paths: false links
- Fails to discover hidden routers

Sidecar: Transparently embedded probes

- Firewalls/NATs block probes
- Abuse reports limit scope of experiments

Not Addressed

- Layer 2 devices and unused/backup links never discovered

Record Route IP Option

RFC791: Record path into packet's IP header

- At most 9 IP addresses recorded

Conventional Wisdom:

- Nine hops is too few to be useful
  - Average path length > 9
- Firewalls drop packets with IP options
  - Additional cause for abuse reports
- IP options increase router processing time

Passenger: TR+RR

Augment traceroute (TR) probes with IP Record Route option (RR):

- RR records the outgoing address
- Prevents false links
- New alias resolution technique

Making RR work:

- 9 hops + PlanetLab = 87-98% of addresses
- Avoid firewalls: Don't probe close to end-hosts
- Destination support not required
  - IP options included in ICMP response

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR : RR array

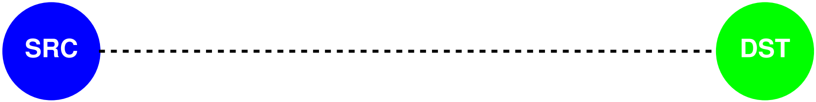

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	$\emptyset$

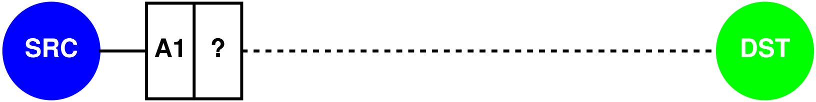

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2

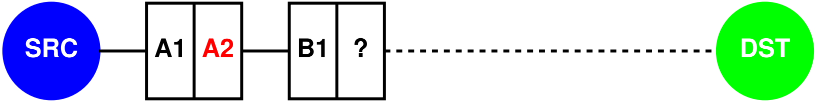

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2
3	C1	:	A3,E2

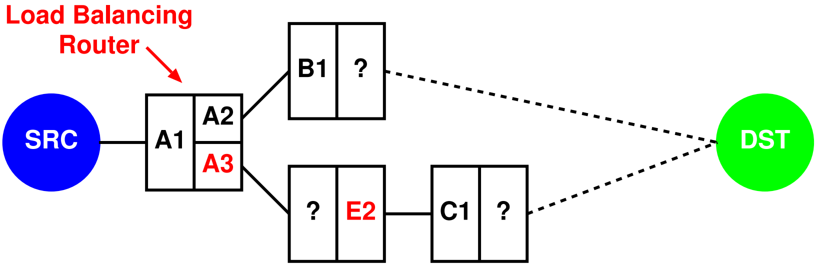

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2
3	C1	:	A3,E2
4	D1	:	A2,B2,F2

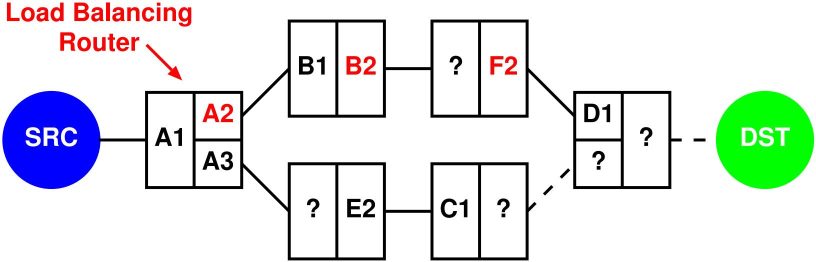

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	$\emptyset$
2	B1	:	A2
3	C1	:	A3,E2
4	D1	:	A2,B2,F2
5	DST	:	A2,B2,F2,D3

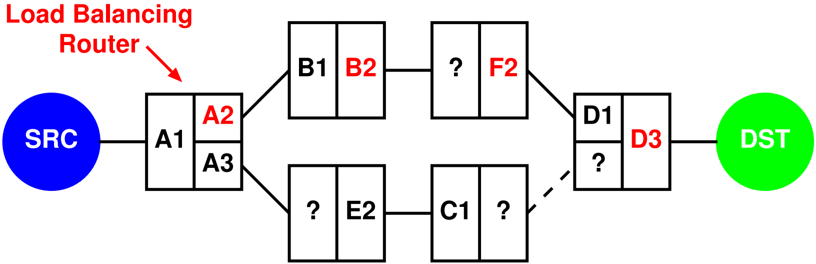

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2
3	C1	:	A3,E2
4	D1	:	A2,B2,F2
5	DST	:	A2,B2,F2,D3

- Alias resolution: ith TR IP is an alias for ith RR IP along the same path

Discovery with TR+RR

- TR records incoming IP
- RR records outgoing IP

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2
3	C1	:	A3,E2
4	D1	:	A2,B2,F2
5	DST	:	A2,B2,F2,D3

- Alias resolution: ith TR IP is an alias for ith RR IP along the same path
- . . . but varying RR implementation make this not true in general

Matching TR and RR IPs

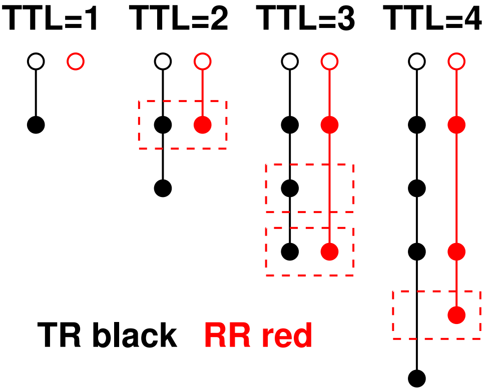

- RR classification is a new problem
- Imperfect classification heuristics in paper
- Formal system subject of continued work

Sidecar: Transparent Probing

- Inject measurement packets into non-measurement TCP streams
- Probes are replayed packets with lower TTL and RR option
  - Transparent w.r.t. end-points
  - No abuse reports triggered by Sidecar probes
- Sidecar probes traverse NATs/Firewalls
- Technique generalizes to non-topology measurements [WORLDS'06]

Preliminary Experiments and Results

Inject Passenger+Sidecar probes into:

- All CoDeeN traffic: May 17-24th
  - 13M (src,dst) pairs
  - 22K IP addresses/891 ASes discovered
  - 65.8% of links corroborated with RR
- Web Crawl
  - PlanetLab × 160K web servers
  - 51M (src,dst) pairs
  - 375K IP addresses/8,739 ASes discovered
  - 69.1% of links corroborated with RR

Conclusions and Future Work

- RR option prematurely dismissed
  - Prevents false link assertions
  - Discovers hidden routers
  - New alias resolution technique
- New problem of RR implementation classification
- Sidecar: new technique for unobtrusive measurements
  - Less abuse reports → more destinations
- Future work:
  - Better router classification
  - More traffic sources

Implementation Diversity

- Matching RR to TR addresses is difficult
- Most routers don't RR when expiring probes

TTL	TR	:	RR array
1	A1	:	∅

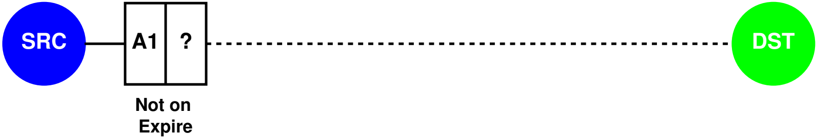

Implementation Diversity

- Matching RR to TR addresses is difficult
- Most routers don't RR when expiring probes but some do

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2,B2

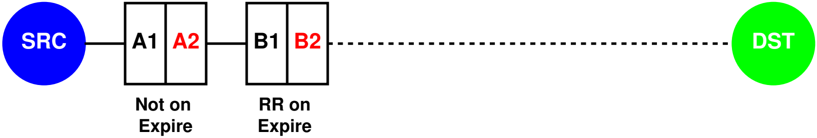

Implementation Diversity

- Some implement RR but not TR
- Some don't implement RR

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2,B2
3	C1	:	A2,B2,F2

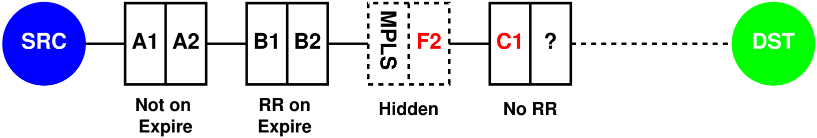

Implementation Diversity

- Some routers drop packets with IP options

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2,B2
3	C1	:	A2,B2,F2
4	?	:	dropped

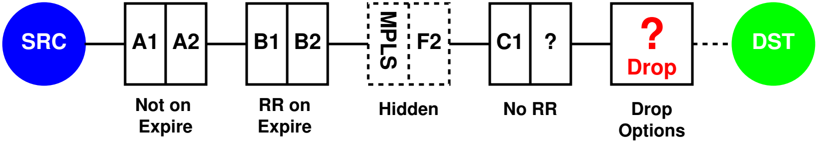

Implementation Diversity

- Alias resolution requires classifying RR behavior

TTL	TR	:	RR array
1	A1	:	∅
2	B1	:	A2,B2
3	C1	:	A2,B2,F2
4	?	:	dropped

Not on Expire	**RR on Expire**	**Hidden**	**No RR**	**Drop Options**

- RR classification is a new problem
- Imperfect classification heuristics in paper
- Formal system subject of future work