Paris traceroute: Measuring more accurate and complete paths

Brice Augustin
Fabien Viger, Xavier Cuvellier, Matthieu Latapy, Clémence Magnien, Timur Friedman and Renata Teixeira

Laboratoire LIP6 – CNRS
Université Pierre et Marie Curie – Paris 6

What's wrong with traceroute?

```
-bash-3.00$ traceroute -n www.google.com
 1  70.87204.1  8.323 ms  0.797 ms  1.066 ms
 2  70.84.160.130  0.471 ms  0.262 ms *
 3  70.85.127.109  0.299 ms  0.258 ms  0.256 ms
 4  70.87.253.17  0.302 ms  0.206 ms *
 5  208.172.139.129  0.569 ms  0.556 ms  0.480 ms
 6  204.70.193.193  28.347 ms 204.70.192.49  0.694 ms *
 7  208.172.97.170  28.380 ms 204.70.193.185  28.378 ms 208.172.97.170  28.374 ms
 8  208.172.99.94  28.356 ms 208.172.108.6  28.483 ms 208.172.99.94  28.444 ms
 9  72.14.238.57  30.792 ms  30.674 ms 208.172.108.6  28.437 ms
10  72.14.238.151  31.371 ms 72.14.238.57  30.653 ms  30.718 ms
11  66.249.95.194  40.722 ms 72.14.238.151  31.237 ms 66.249.95.194  40.870 ms
12  216.239.51.104  31.390 ms 72.14.238.190  40.858 ms 216.239.51.104  31.357 ms
```

What's wrong with traceroute?

```
-bash-3.00$ traceroute -n www.google.com
 1  70.87204.1  8.323 ms  0.797 ms  1.066 ms
 2  70.84.160.130  0.471 ms  0.262 ms *
 3  70.85.127.109  0.299 ms  0.258 ms  0.256 ms
 4  70.87.253.17   02 m    206 ms *
 5  208.172.139.12  0.56   0.556 ms  0.480 m
    .70    193  28.347 ms 20   192.49   0.60  ms *
    .17    170  28.380 ms 20   .193     28.378 ms   172.97.170  28.  7 ms
 8  208.172.99.94    356   208.172.108.  8.483 ms 20   2.99.94  28.444 ms
 9  72.14.238.57    92 m  0.674 ms 208.172.108.6  28.437 ms
10  72.14.238.151  31.371 ms 72.14.238.57  30.653 ms  30.718 ms
11  66.249.95.194  40.722 ms 72.14.238.151  31.237 ms 66.249.95.194  40.870 ms
12  216.239.51.104  31.390 ms 72.14.238.190  40.858 ms 216.239.51.104  31.357 ms
```

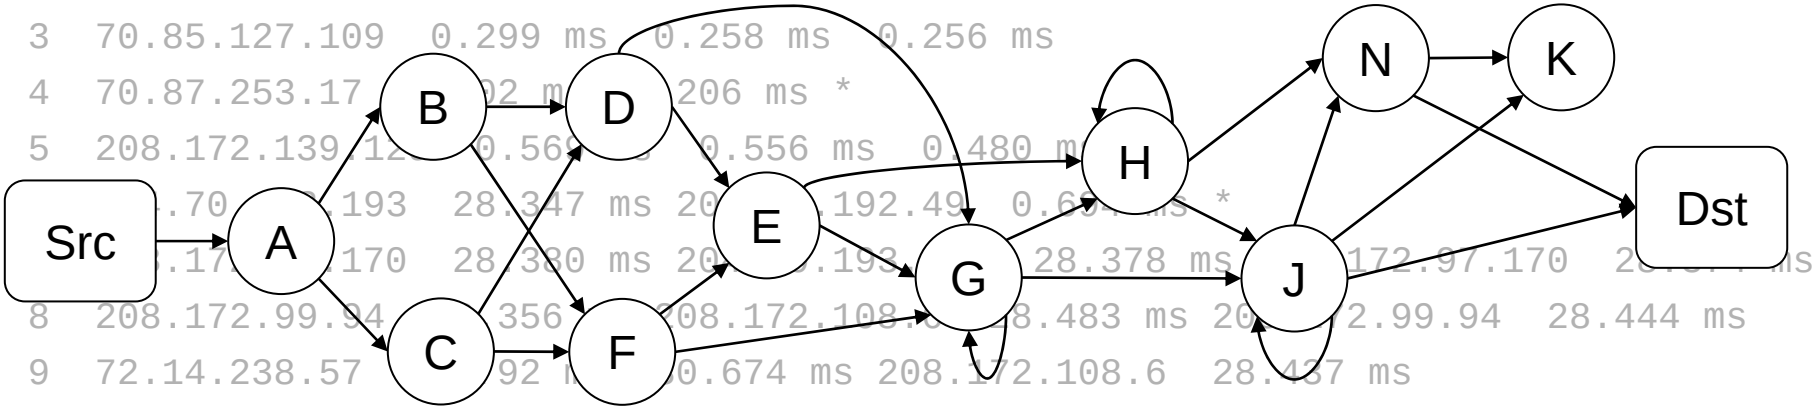

Findings

- Identified traceroute deficiencies on load balanced paths
  - Measured paths are **inaccurate** and **incomplete**
  - May diagnose an incorrect path
- Many routers have load balancing capabilities
  - Per-packet, per-flow, per-destination
- Built a new traceroute: **Paris traceroute**

Classic vs Paris traceroute

Classic traceroute:

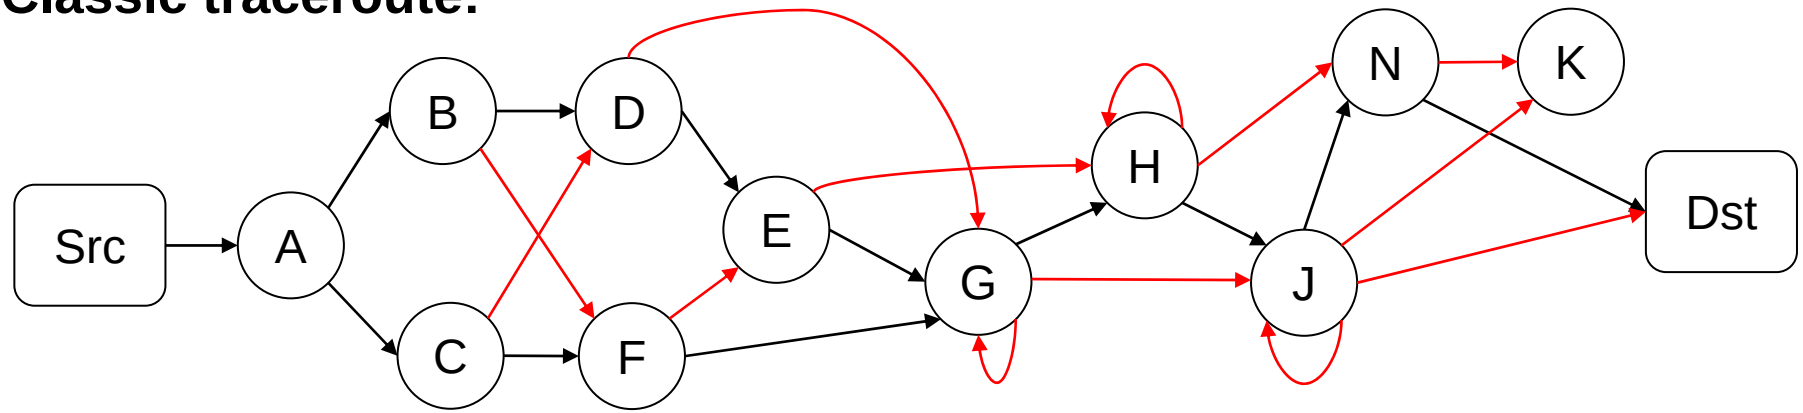

Paris traceroute:

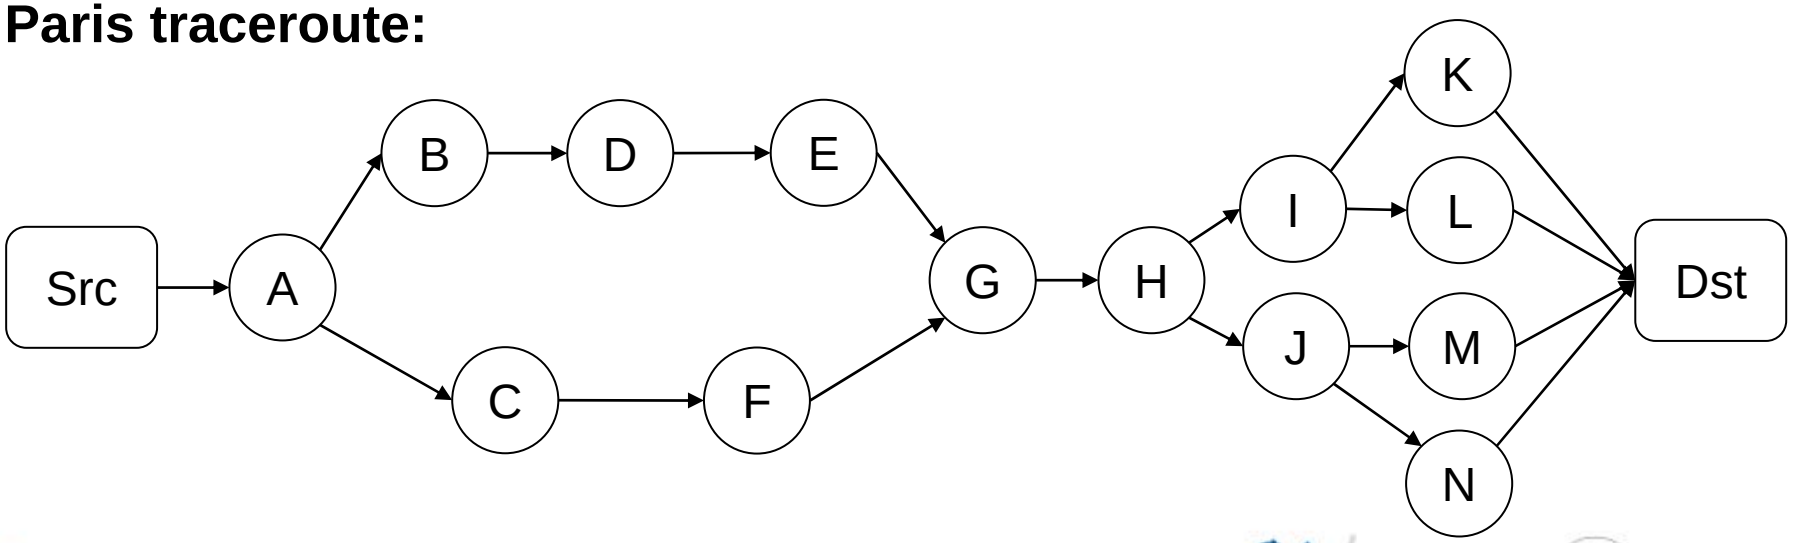

Traceroute under load balancing

Actual network:

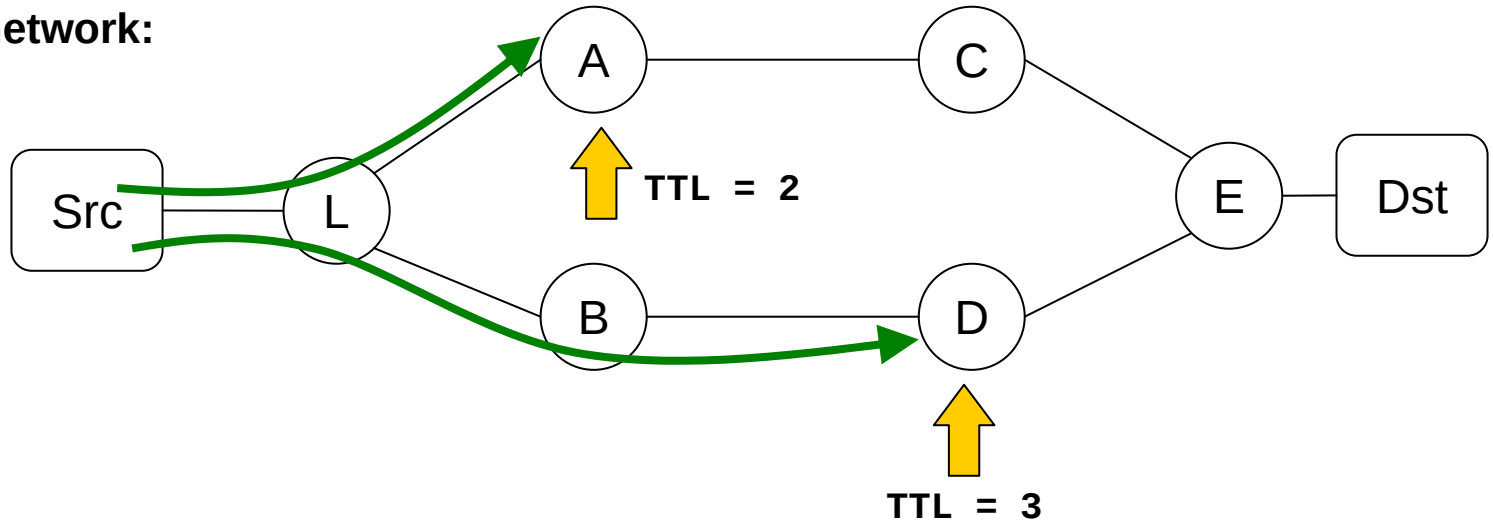

Inferred path:

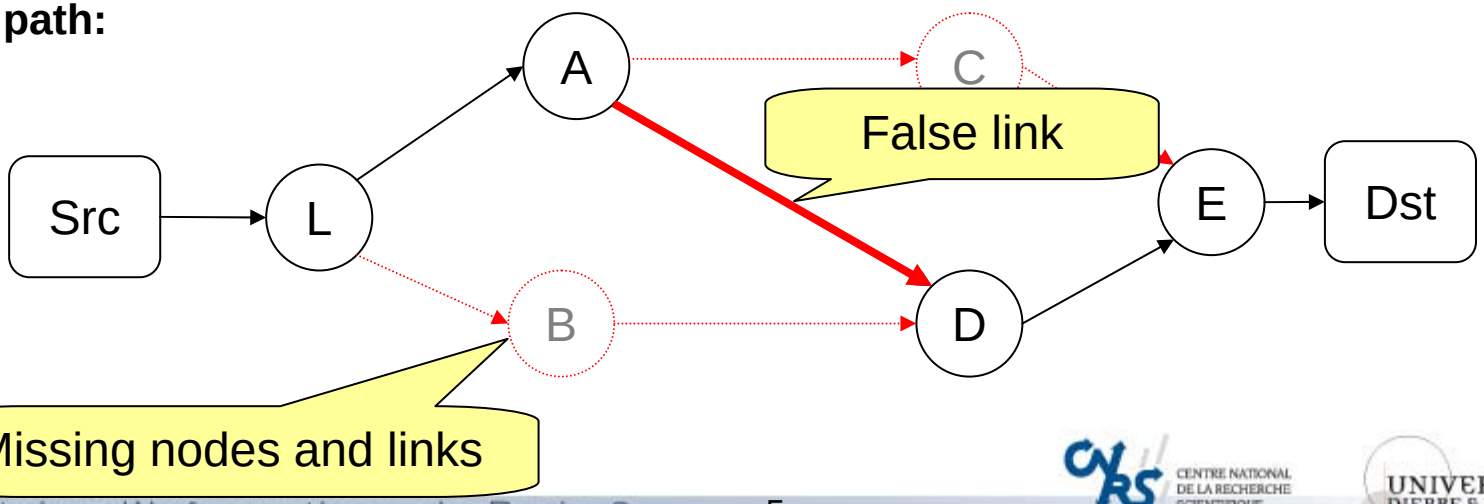

Hard to diagnose aberrant paths

Actual network:

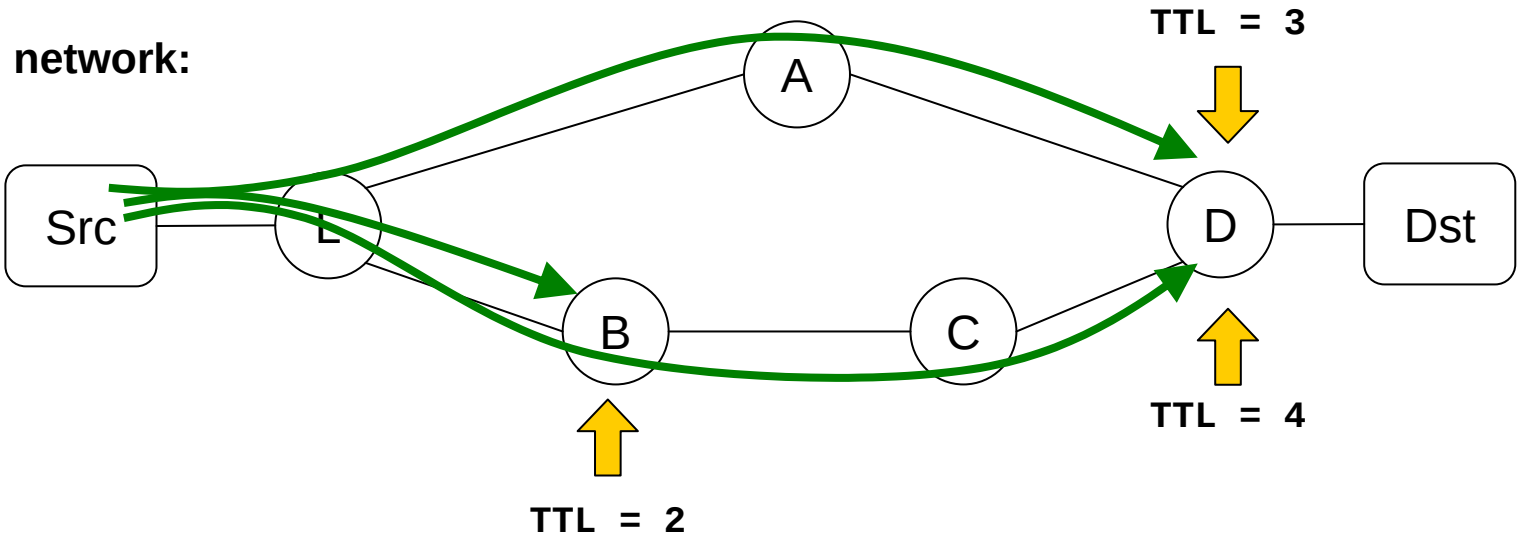

Inferred path:

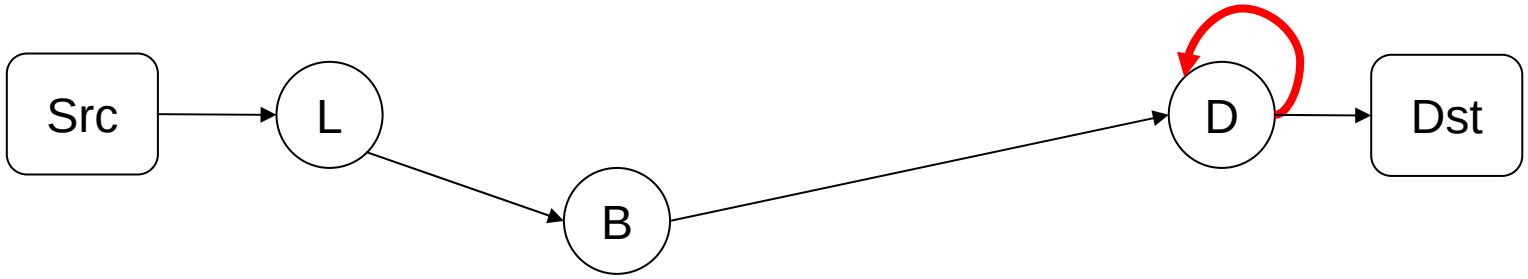

Hard to diagnose unstable paths

Actual network:

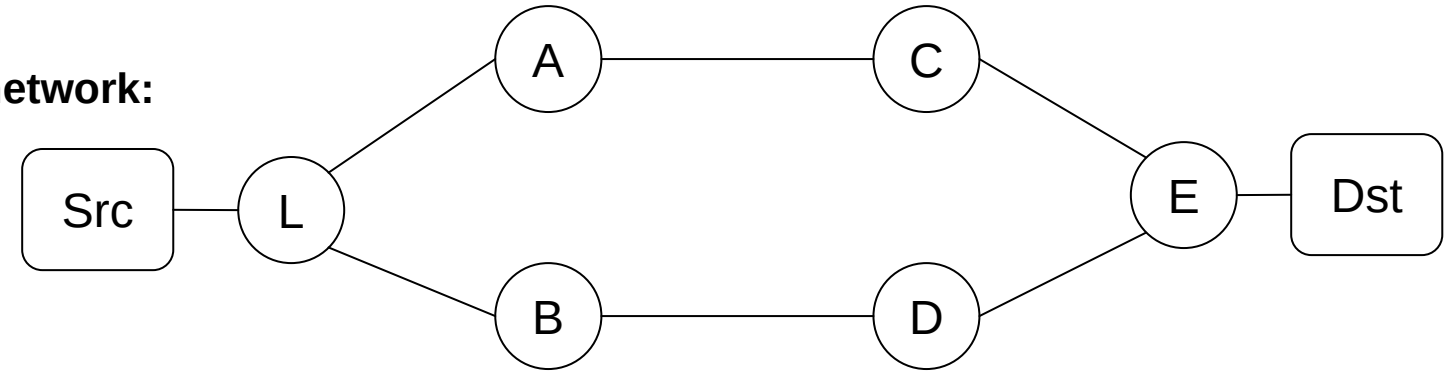

Inferred path:

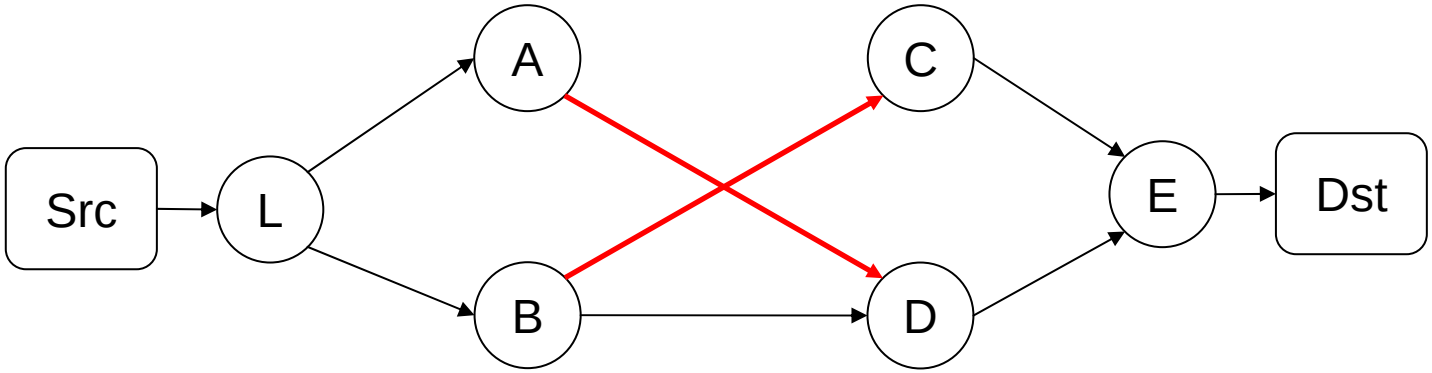

Problems happen even under per-flow load balancing

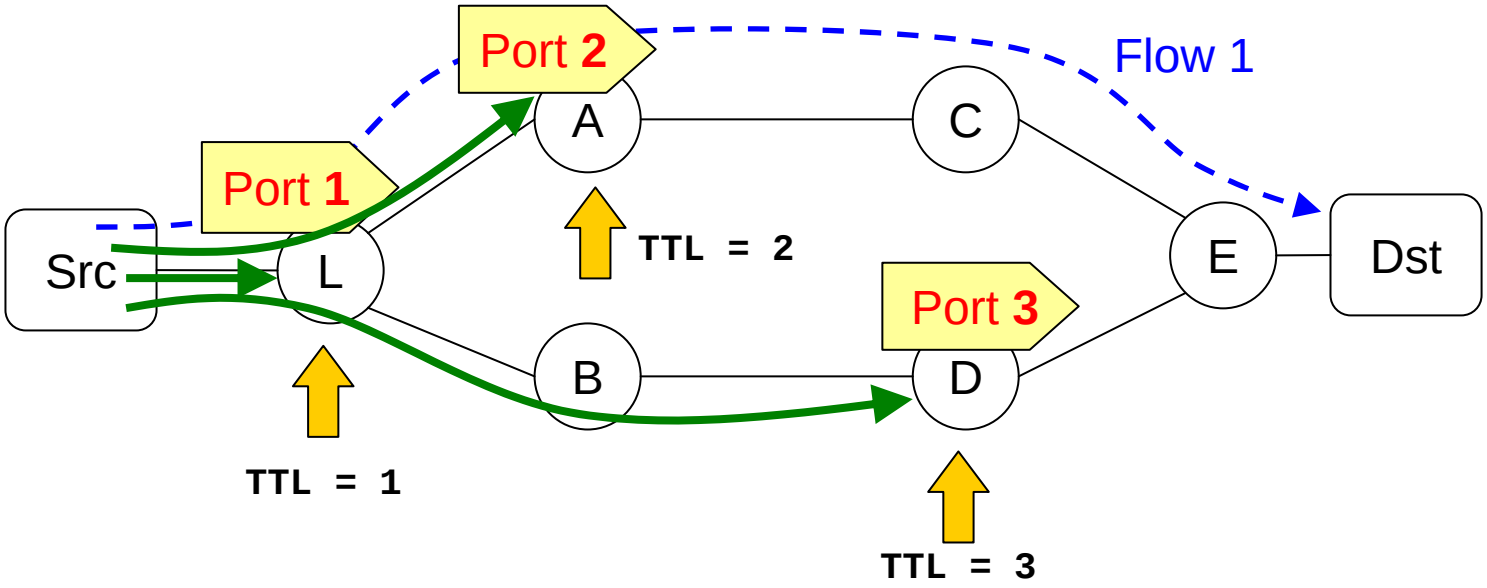

- Traceroute uses the destination port as identifier
- Per-flow load balancers use the destination port as part of the flow identifier

Paris traceroute: Tracing a single path

- Solves the problem with per-flow load balancing
  - Packets have the same flow identifier
- Works with UDP, TCP and ICMP

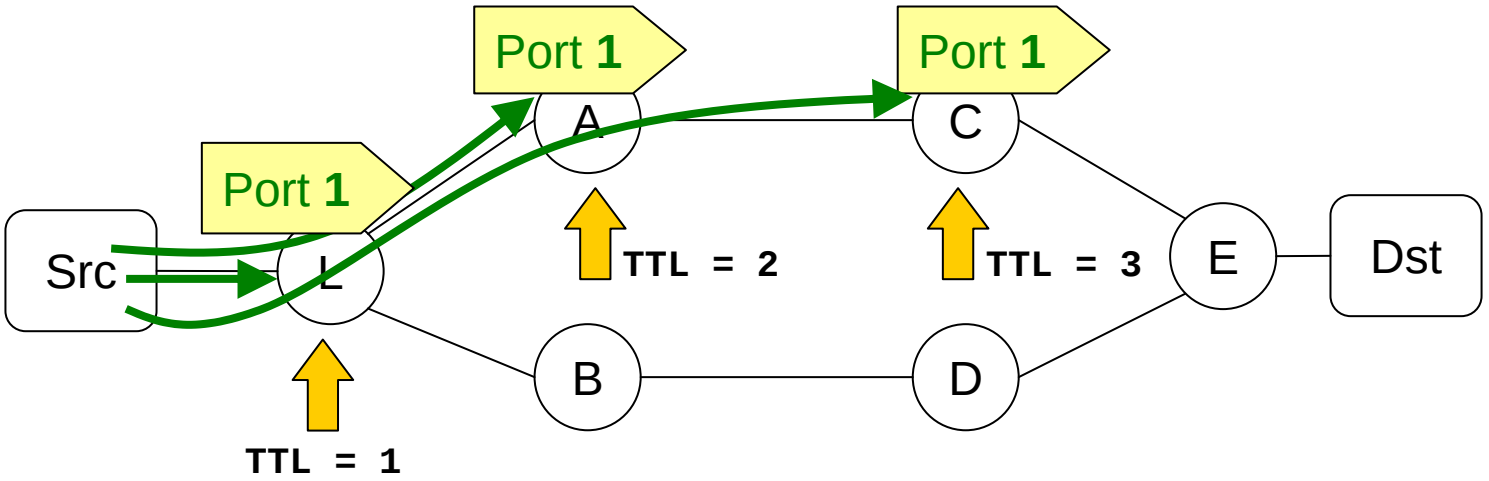

Identifying the probes

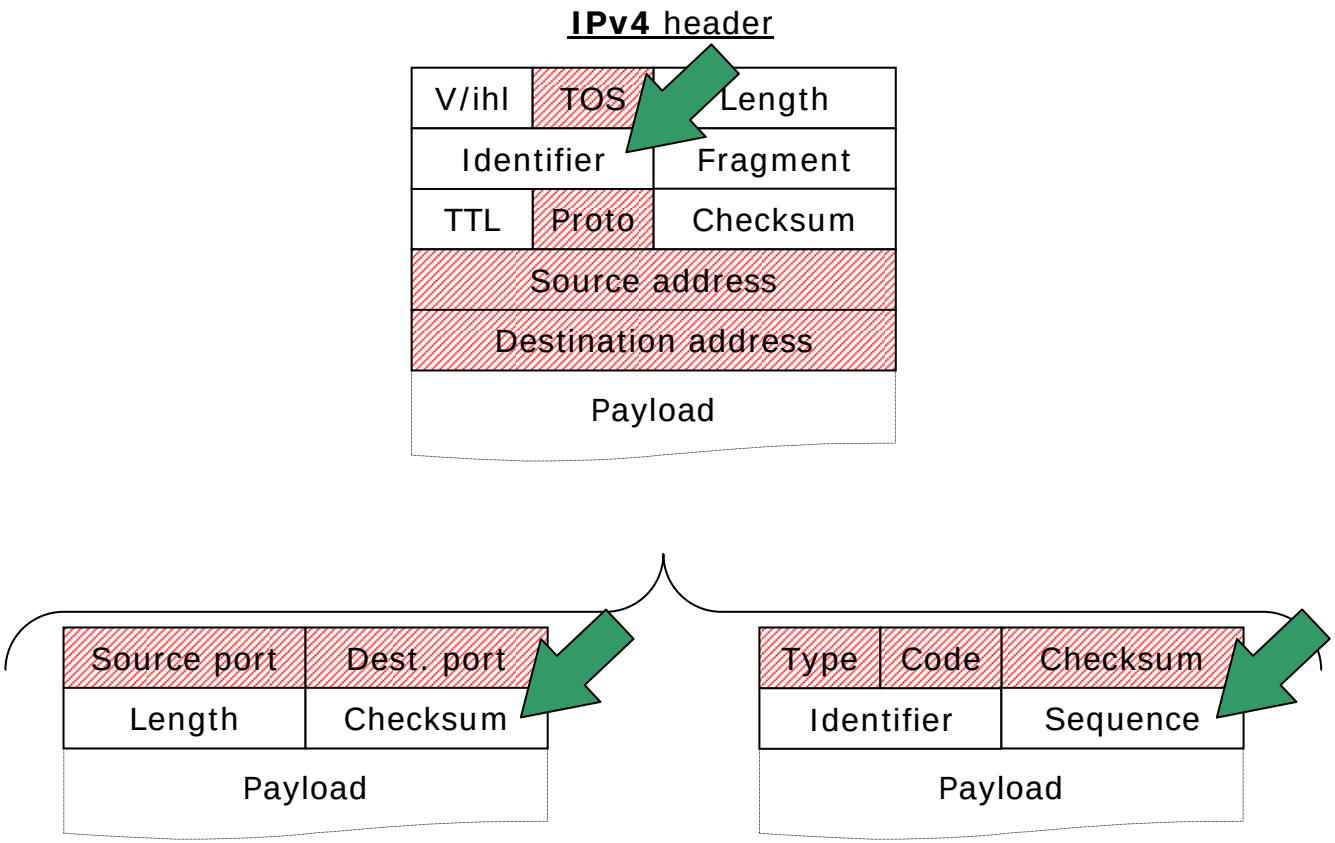

Paris traceroute: Tracing all the paths

- Change the probing strategy
- At each hop:
  - Send packets with a different flow identifier
  - Send enough probes to enumerate all interfaces with a high degree of confidence
  - Classify load balancers: per-flow or per-packet

Interfaces enumeration

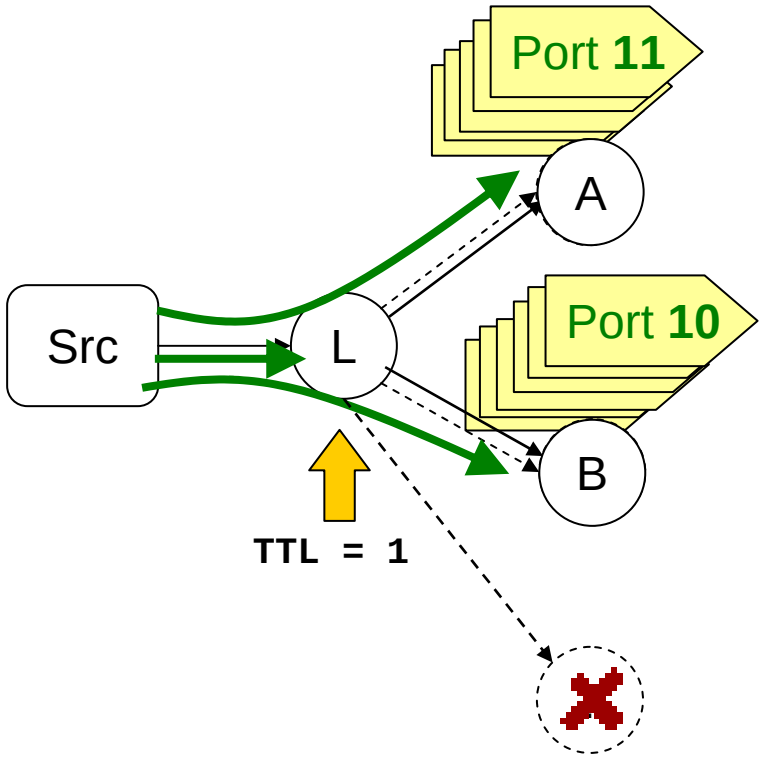

Steps:

- Interfaces after L?
- Suppose 2 interfaces
- Send 6 packets through L
- Responses from 2 interfaces
- Suppose actually 3 interfaces
- Send 5 more packets through L
- No third interface
- Stop probing

Probing overhead per hop

- Classic traceroute: 3 packets by default
- Paris traceroute: at least 6 packets to rule out load balancing

# interfaces	1	2	3	4	5	6	…	15	16
# packets	6	11	16	21	27	33	…	90	96

- Up to 96 probes (up to 16 responding interfaces in our traces)

Paris traceroute output

```
-bash-3.00$ paris-traceroute –n –a exh www.amazon.com
 1  132.227.74.1
 2  132.227.106.254
[...]
 7  194.68.129.201
 8  64.125.23.13:0,1,3,6,9,10  64.125.23.9:2,4,5,7,8
 9  64.125.27.225:0,1,3,6,9,10  64.125.27.165:2,4,5,7,8
10  64.125.27.57:0,1,3,6,9,
11  64.125.28.126:0,1,3,6,9,10  64.125.29.109:2,4,5,7,8
12  64.125.29.230
[...]
```

Single path, no load balancing

Per-flow load balancer

2 load-balanced paths

List of flow identifiers

Load balancing is common

- Measurements from 15 sources to 70,000 destinations
- Paths affected by load balancing:
  - 39% by per-flow
  - 2% by per-packet
  - 70% by per-destination
- Many Tier-1s use load balancing

Load balancing causes anomalies

From our LIP6 vantage point:

- Diamonds appear in 30% of the destinations
  - Paris traceroute removes 10,662 from 19,159 (**56%**)
- Loops appear in 4.5% of the measured routes
  - Paris traceroute removes 5,047 from 5,795 (**87%**)
- Cycles appear in 0.25% of the measured routes
  - Paris traceroute removes 3,886 from 5,674 (**68%**)
- Other causes
  - Routing changes
  - NAT boxes
  - Buggy routers
  - Per-packet load balancing

Load-balanced paths

- Generally short, narrow and symmetric
- Some are extremely long
  - More than 10 hops
- Some others are very wide
  - Up to 16 responding interfaces
- Parallel paths with different hop counts

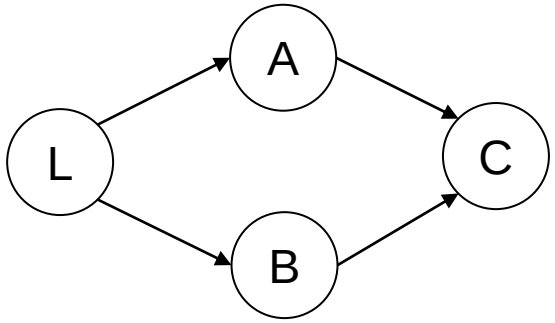

Conclusion

- Vast deployment of load balancing
- Classic traceroute discovers inaccurate and incomplete paths
- Paris traceroute reports more accurate and complete paths

More information

www.paris-traceroute.net

Perspectives

- Measure “native path diversity” in the internet (submitted to IMC2007)
- Handle some probing subtleties
- Simple extensions to detect:
  - Per-destination load balancing
  - Uneven load balancing
- Return path diversity

Backup slides/making of

Hop by hop progression

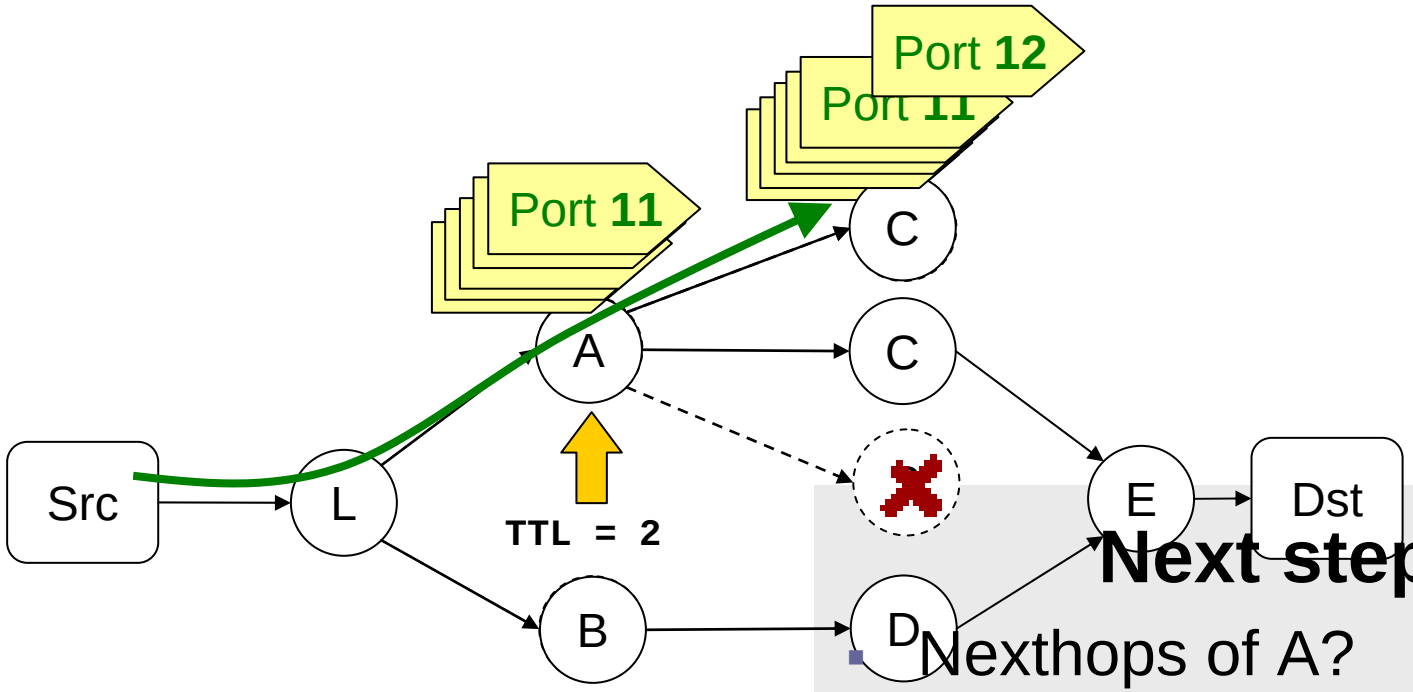

Next steps:

- Nexthops of A?
- Find enough packets traversing A
- Enumerate interfaces

Load balancer classification

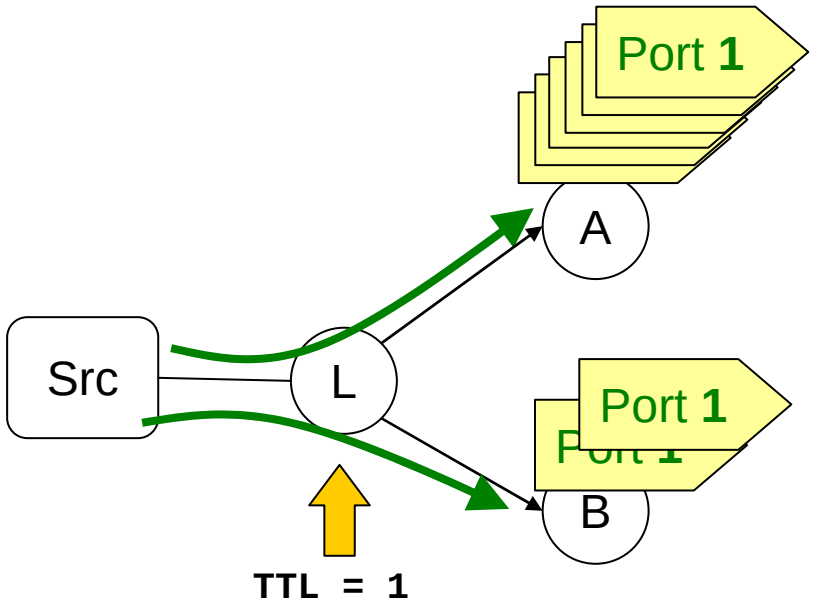

Steps:

Per-packet
or
per-flow

?

- **Suppose per-packet**
- **Send 6 identical packets**
- **Responses from 2 interfaces**

What's wrong with traceroute?

```
-bash-3.00$ traceroute -n www.google.com
 1  70.87.204.1  8.323 ms  0.797 ms  1.066 ms
 2  70.84.160.130  0.471 ms  0.262 ms *
 3  70.85.127.109  0.299 ms  0.258 ms  0.256 ms
 4  70.87.253.17  0.302 ms  0.206 ms *
 5  208.172.139.129  0.569 ms  0.556 ms  0.480 ms
 6  204.70.193.193  28.347 ms 204.70.192.49  0.694 ms *
 7  208.172.97.170  28.380 ms 204.70.193.185  28.378 ms 208.172.97.170  28.374 ms
 8  208.172.99.94  28.356 ms 208.172.108.6  28.483 ms 208.172.99.94  28.444 ms
 9  72.14.238.57  30.792 ms  30.674 ms 208.172.108.6  28.437 ms
10  72.14.238.151  31.371 ms 72.14.238.57  30.653 ms  30.718 ms
11  66.249.95.194  40.722 ms 72.14.238.151  31.237 ms 66.249.95.1  40.870
12  216.239.51.1  31.390  72.14.238.190  40.858 ms 216.239.5  31.
```

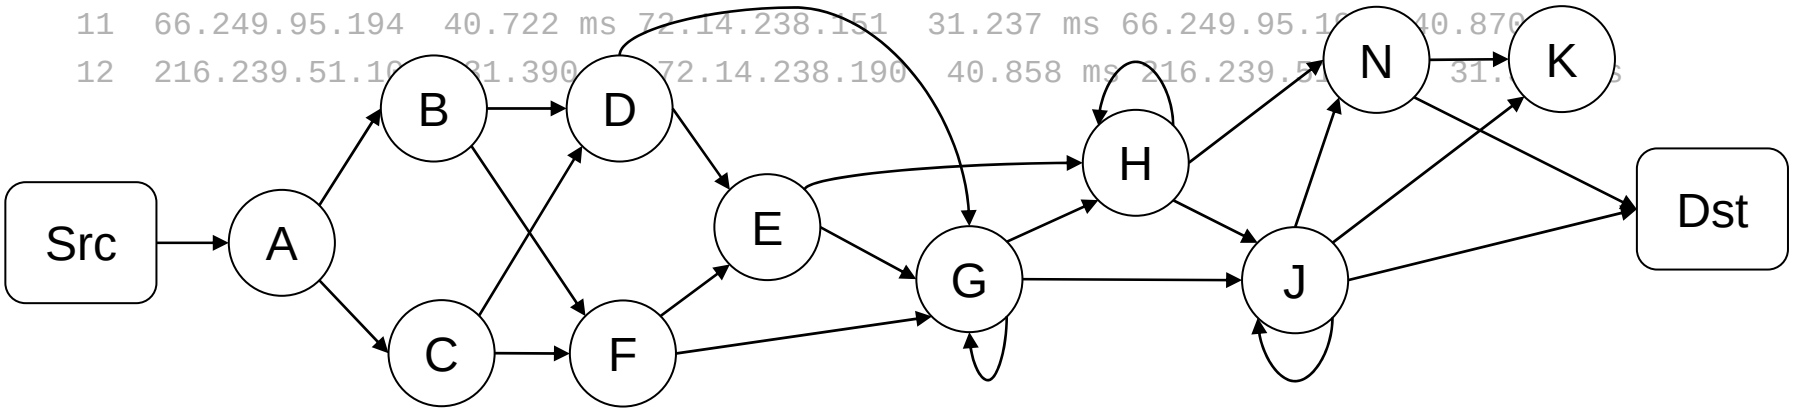

Introduction

- Traceroute measures a path between two hosts in an IP network
- It is widely used by:
  - Network operators
  - Networking researchers
  - Geeks/Computer enthusiasts

Hard to diagnose aberrant paths

Inferred path:

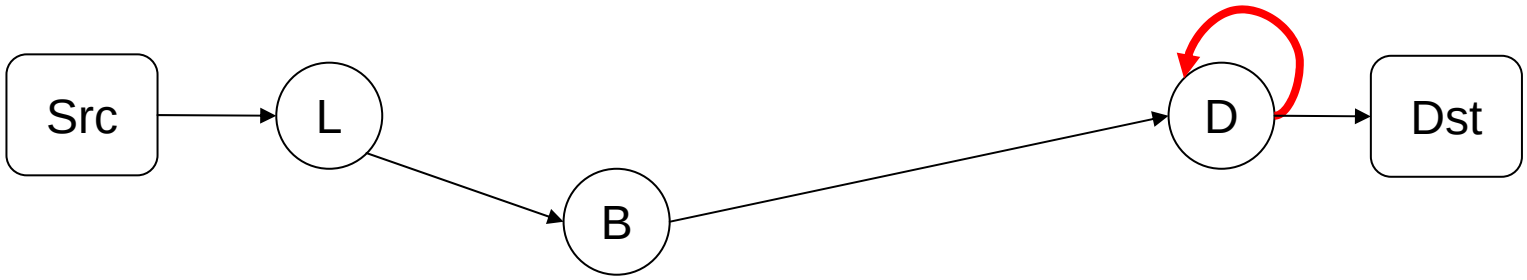

Actual network:

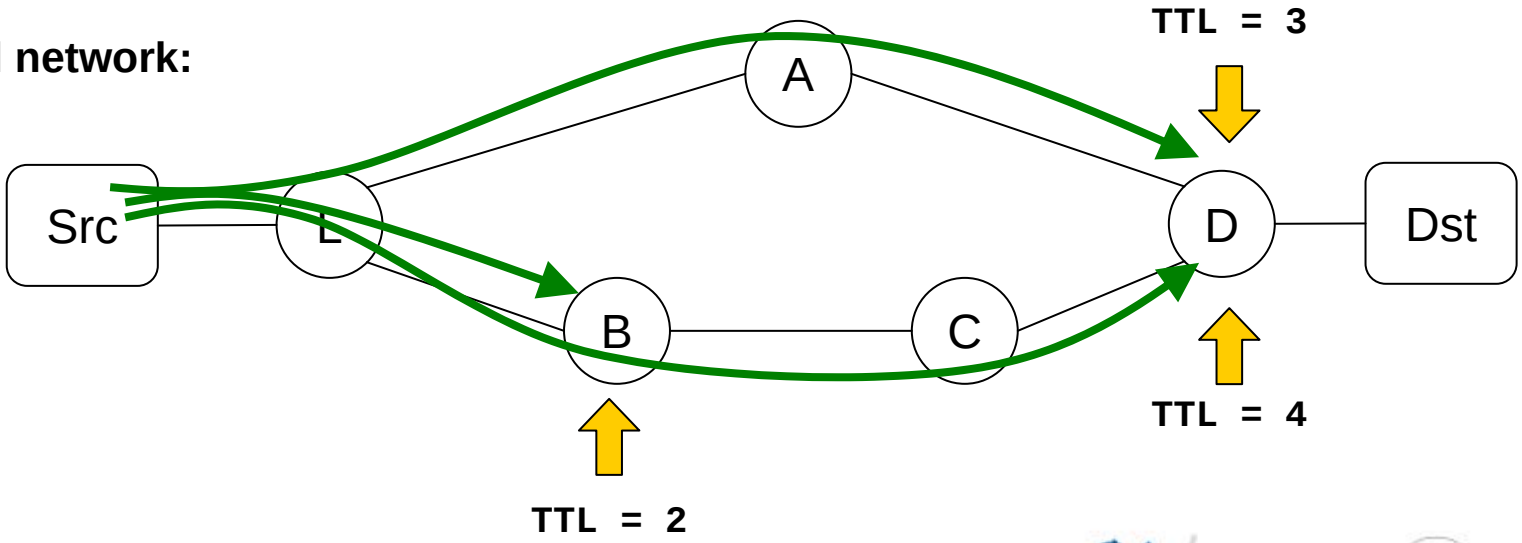

Hard to diagnose unstable paths

Inferred path:

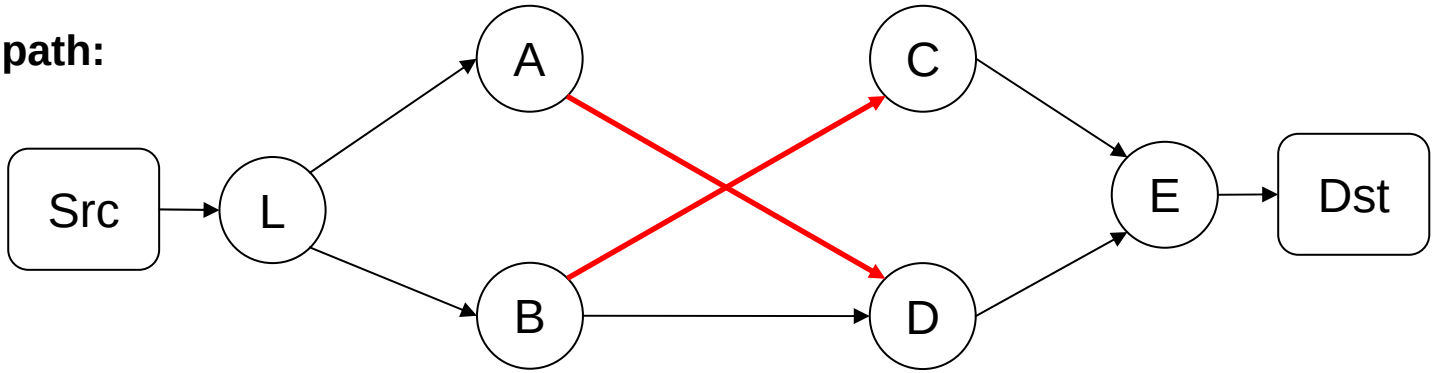

Actual network:

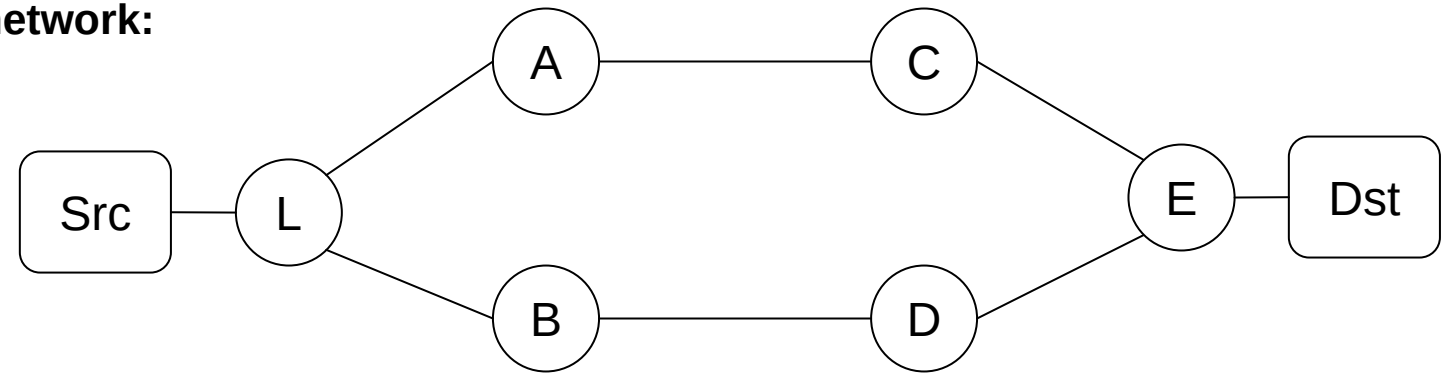

Introduction