# Measured Impact of Tracing Straight

Matthew Luckie, David Murrell

WAND Network Research Group
Department of Computer Science
University of Waikato

# The Problem

- Until recently, load balancing was ignored by macroscopic traceroute collection software
- Negative impacts include
  - inference of false loops — Augustin *et al.* IMC '06
  - lower destination reachability — Luckie *et al.* IMC '08
  - inference of false links

# The Problem

- Classic traceroute: probes toward a destination can take different paths
  - Augustin *et al.* IMC'07: 39% of paths have at least one per-flow load-balancer
- Link inferences from classic traceroute data on questionable ground
  - Assumption is the interfaces represent distinct routers which are connected
  - Traceroute technique did not evolve with the Internet
- Lots of traceroute data collected using classic technique
  - This work is about trying to quantify false link inference rate
  - Can't go back and fix pre-2006 data

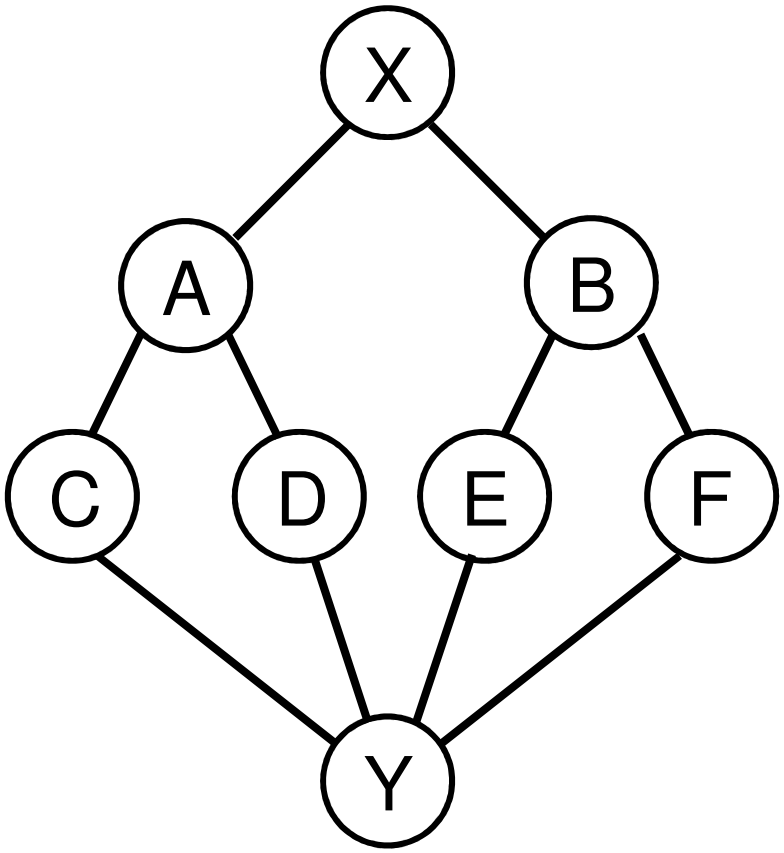


# Interface graph

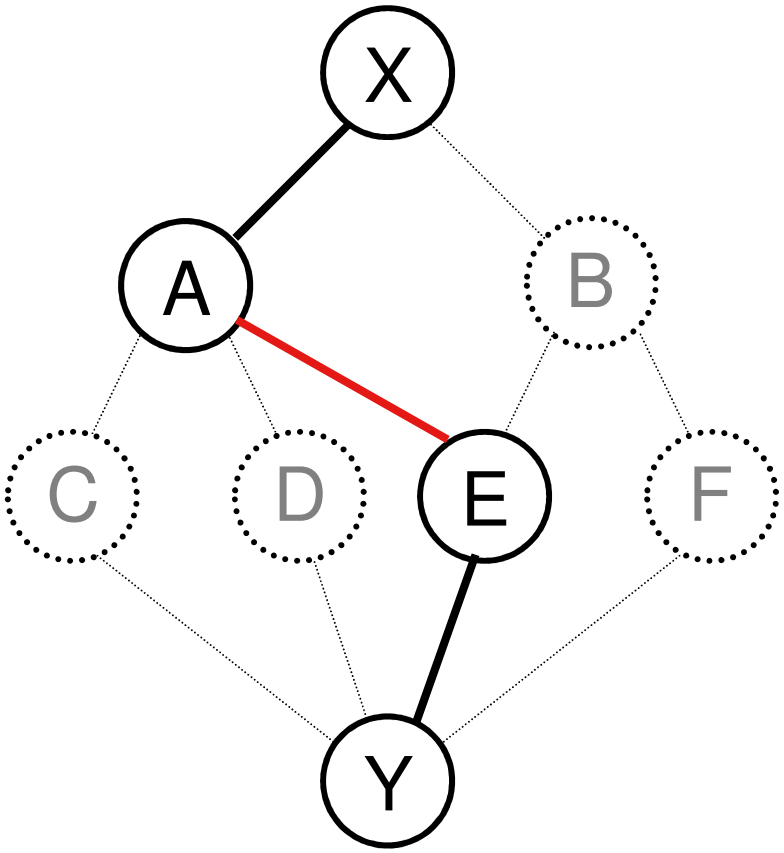


Interface graph Inferred using classic traceroute

A-E is an **artifact** of classic traceroute

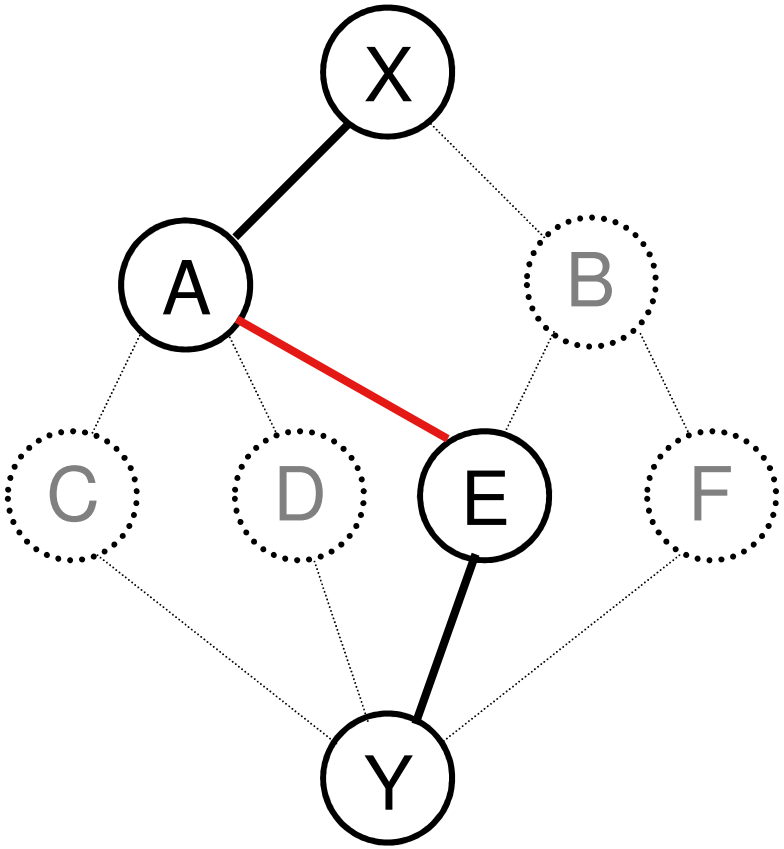


Interface graph revealed using classic traceroute

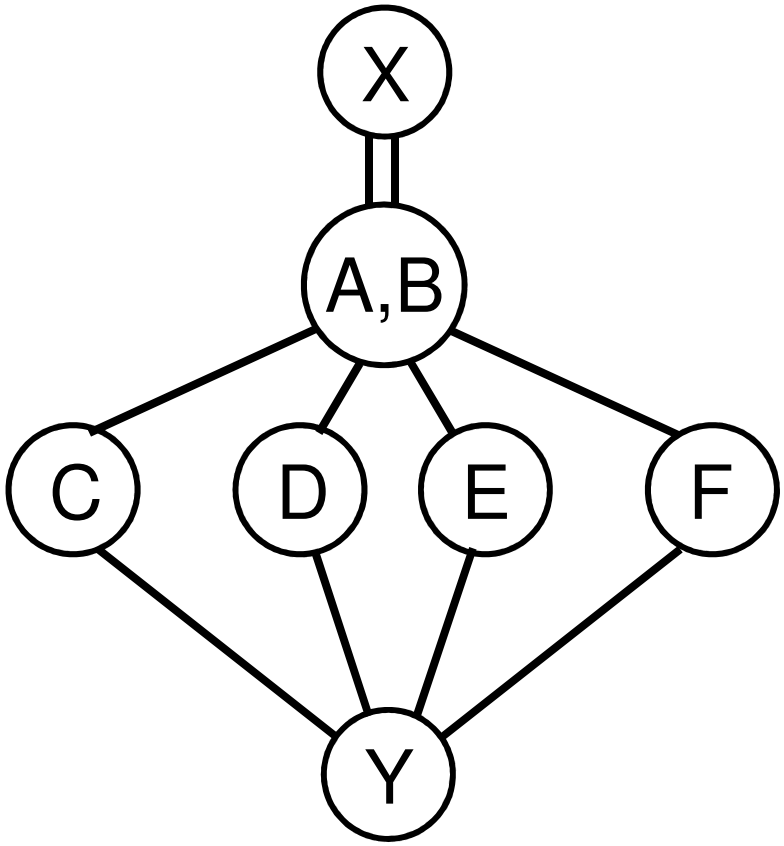


A-E does not introduce a false link if A+B are aliases in the router topology

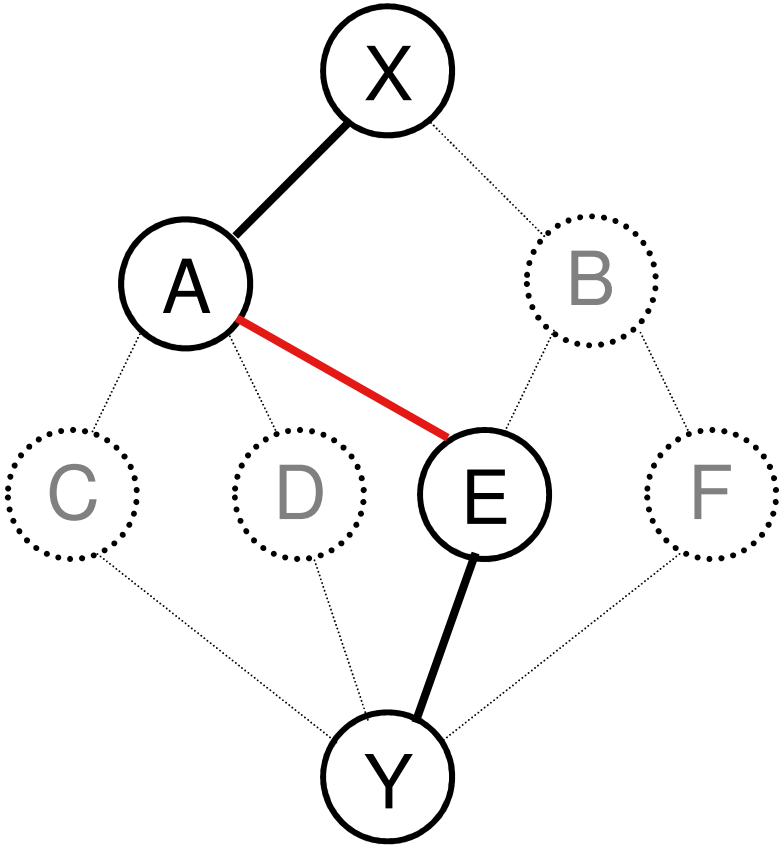


Interface graph revealed using classic traceroute

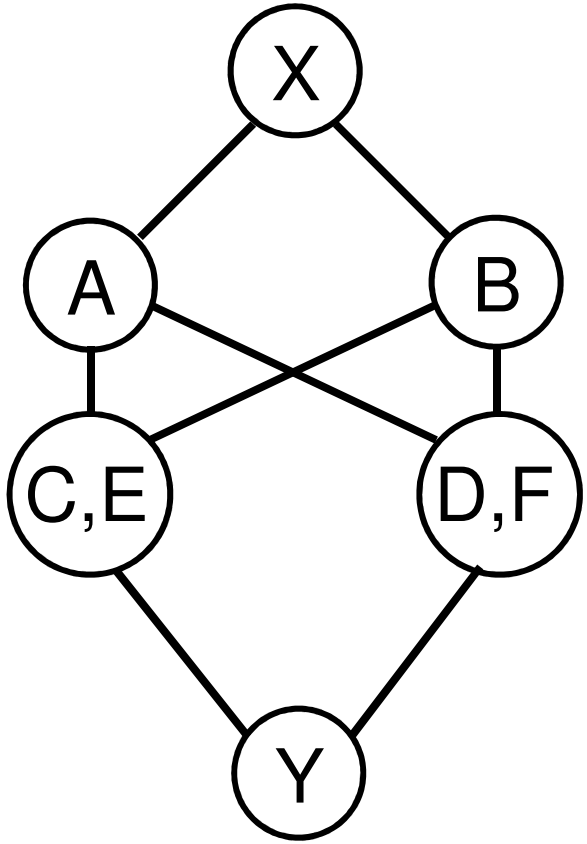


A-E does not introduce a false link if C+E are aliases in the router topology

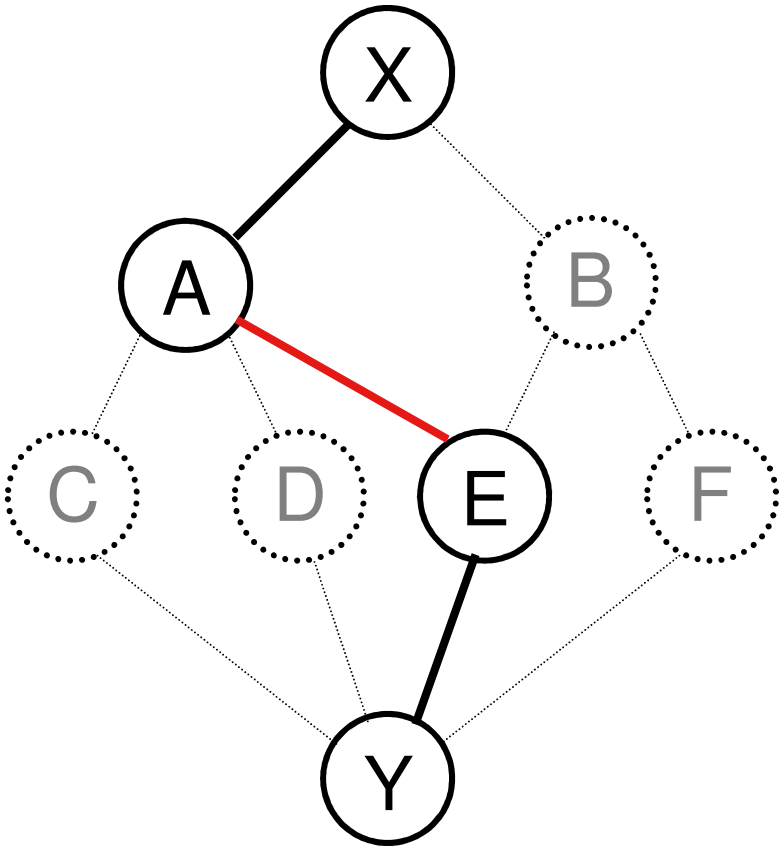


Interface graph revealed using classic traceroute

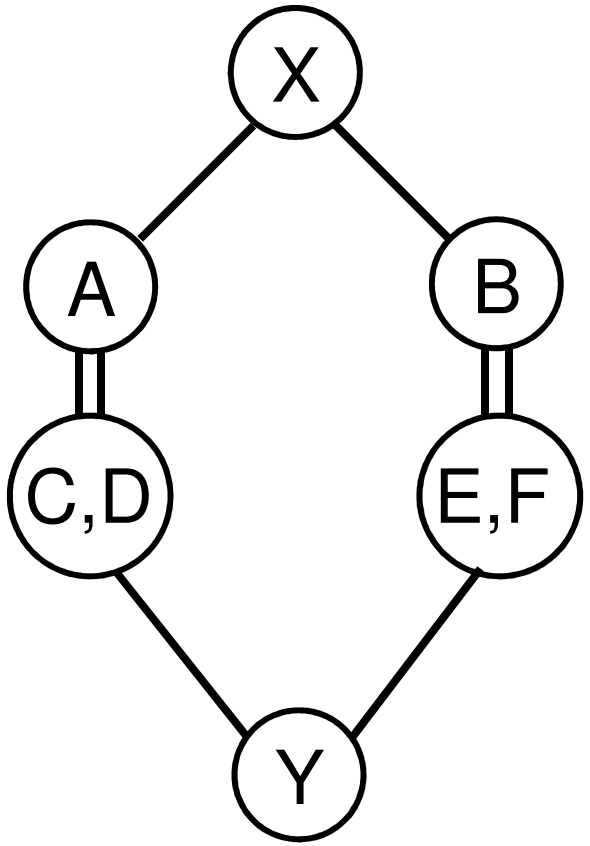


A-E could introduce a false link in this router topology

# Not all load balancing could produce an artifact

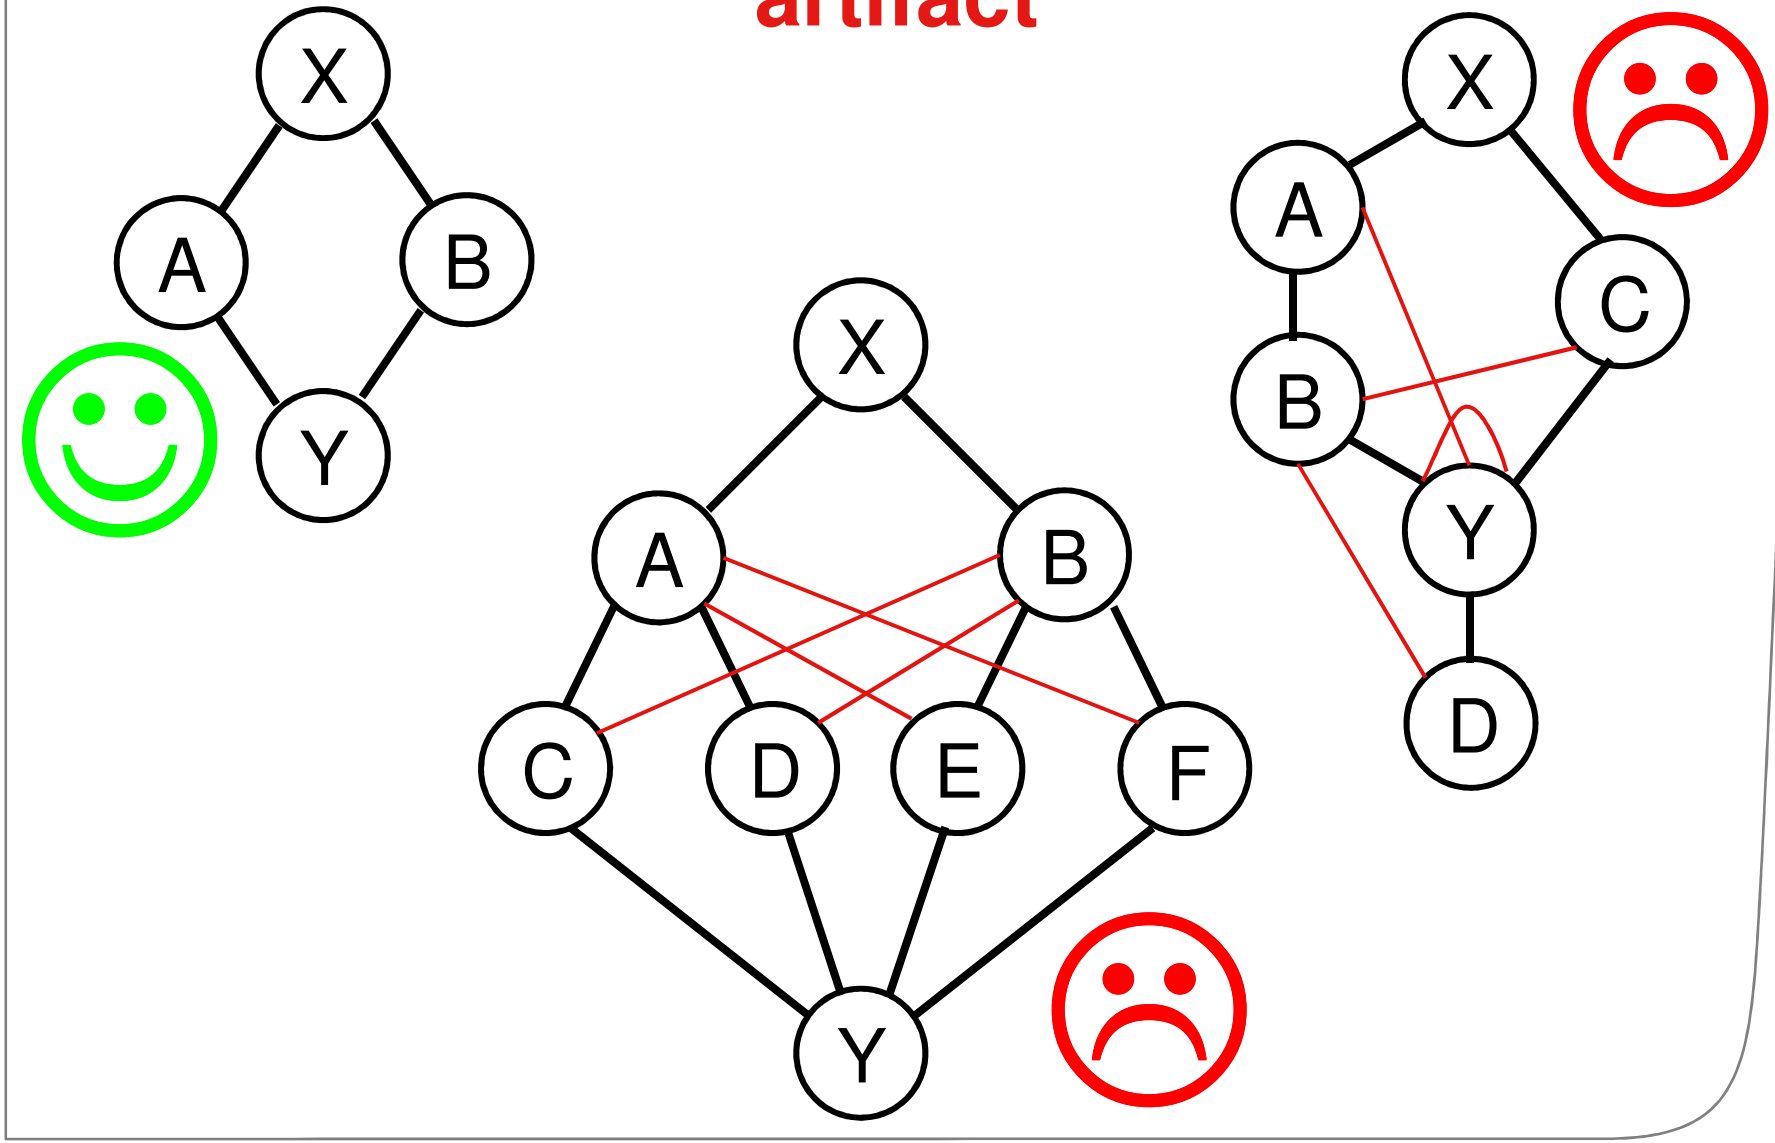

# Methodology

- 12 ark monitors, each probing a list of 19116 addresses
  - derived from BGP data from routeviews on 18th Jan 2010
  - each list contains addresses from distinct prefixes
- For each destination:
  1. Identify artifact links in classic traceroute data
     - use Multipath Detection Algorithm (MDA) to infer all possible links towards a destination to 99% confidence. Augustin *et al.* IMC'07
     - artifact links are those that appear in the output of classic traceroute but not in MDA traceroute.
  2. Determine if the artifact could introduce a false router link

# Methodology

- To guard against false positives as a result of path changes, we use the following procedure
  1. Initial Paris traceroute
  2. Classic traceroute
  3. MDA traceroute to 99% confidence
  4. Final Paris traceroute
- Traces 1 and 4 have to agree, otherwise we discard.
  - Roughly ~1k traces from each VP's ~19k traces were discarded.

# Methodology

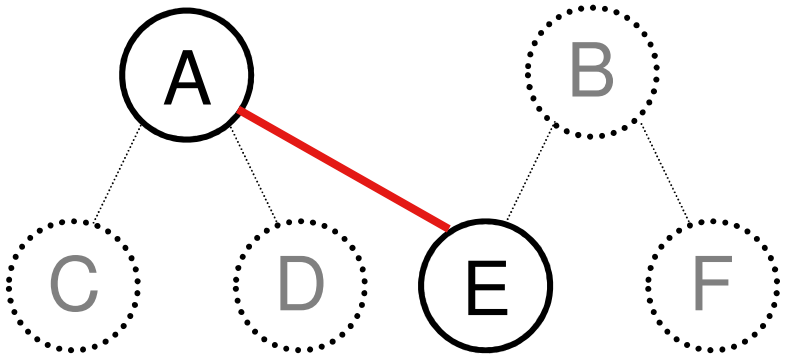


- Determine if an artifact could introduce a false link
  - Use Ally technique where incrementing IPID behaviour is observed for both addresses (Spring *et al.* 2002)
  - If (A,B) or (E,C) or (E,D) are aliases, then the artifact does not introduce a false router adjacency (classification: valid)
  - If (A, B, C, D, E) all exhibit incrementing IPID but no alias is found, we reason the artifact could introduce a false link
  - Otherwise artifact is unclassified.
- Assumption (validated) is that IPID-based alias resolution can rule on whether or not two IP addresses are aliases if incrementing IPID values are observed.

# Summary of data

- Each vantage point (VP) saw roughly the same overall raw link counts:
  - ~52k links MDA-icmp
  - ~55k links MDA-udp
  - ~46k links traceroute

# Result: traces with at least one artifact

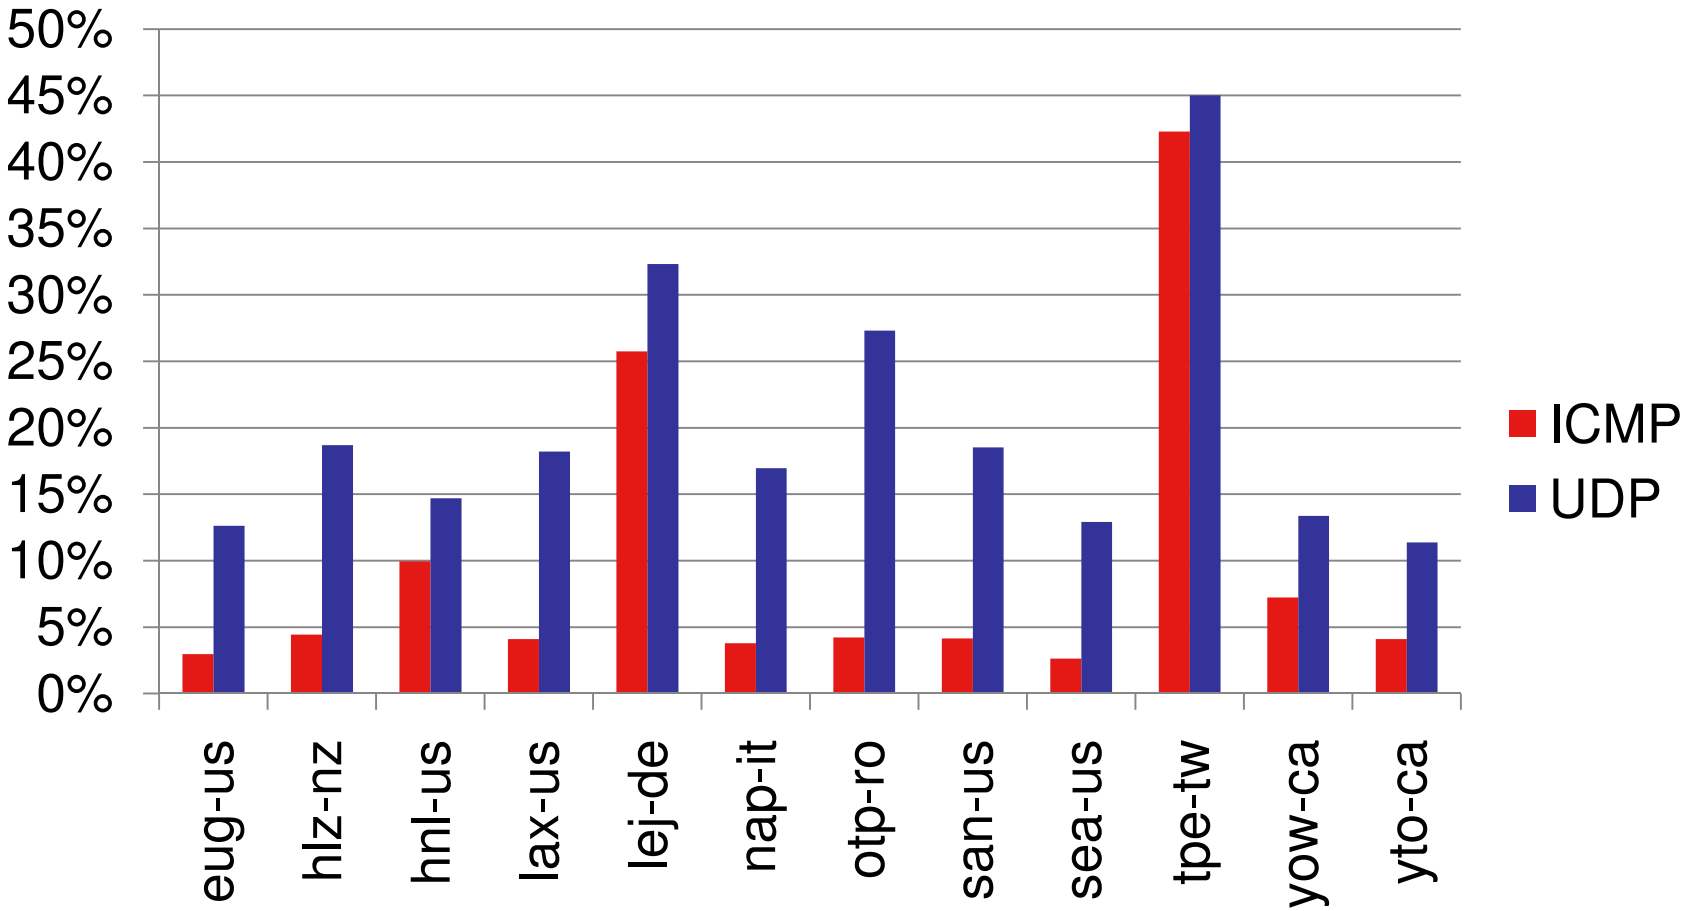


ICMP: 4.9% of traces (on average) have at least one artifact
UDP: 16.5% of traces (on average) have at least one artifact
Note: figures exclude tpe-tw and lej-de

# Result: unique artifact links, ICMP

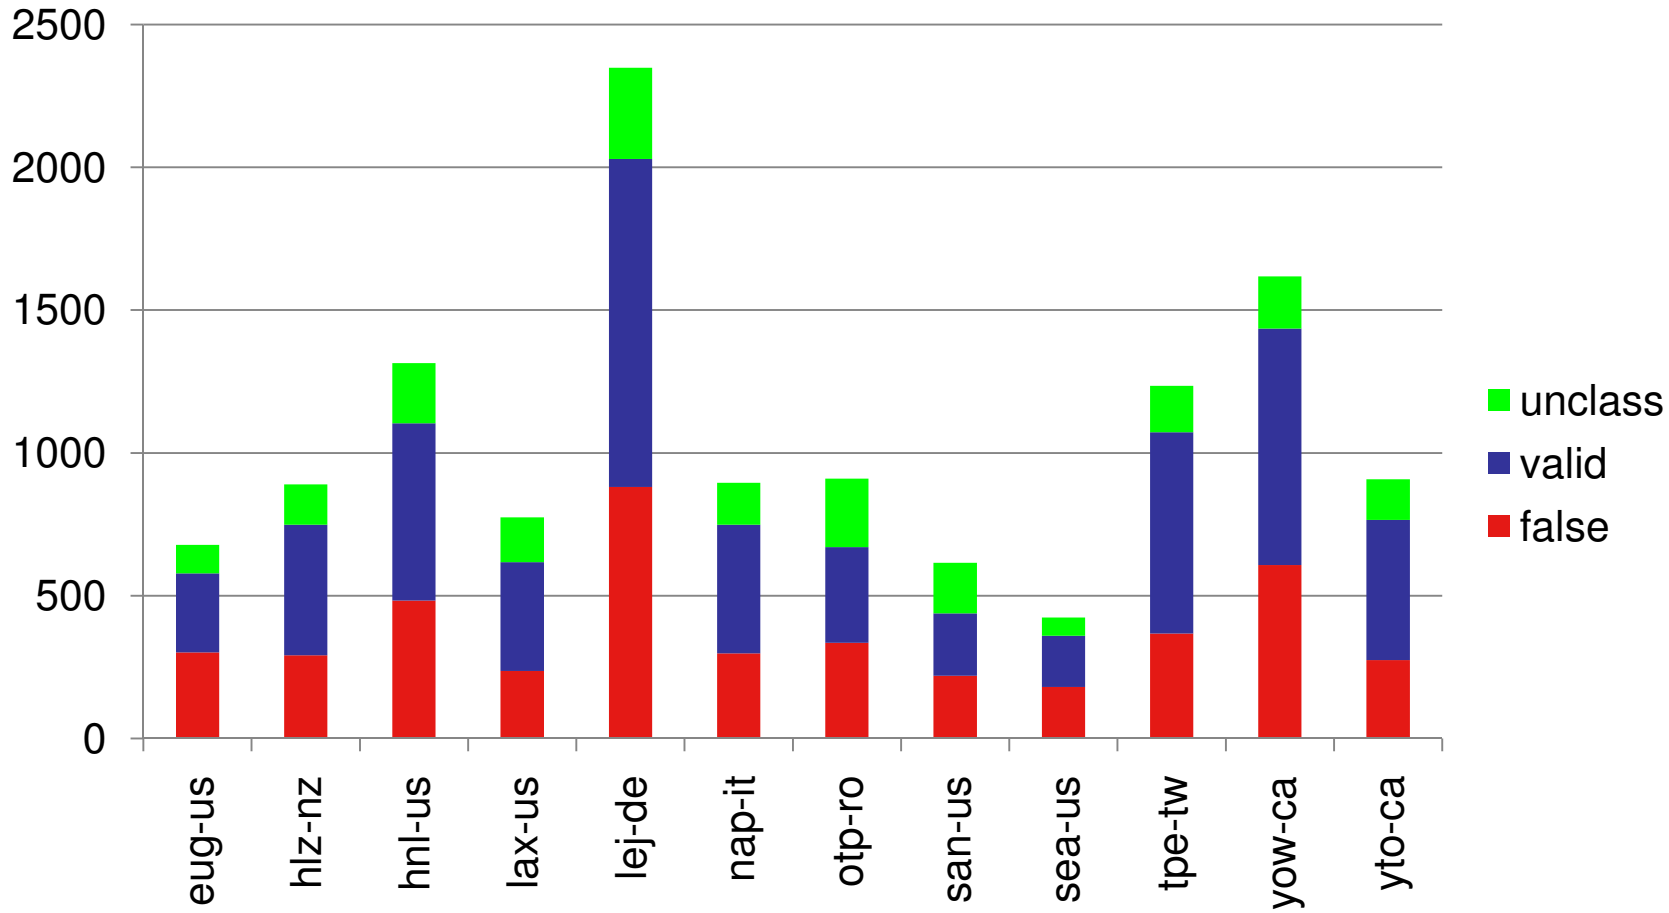


On average, 372 links per VP are classified as false links.
0.8% of each VP's link set.
9618 total artifacts. 3359 (35%) false links.

# Result: unique artifact links, UDP

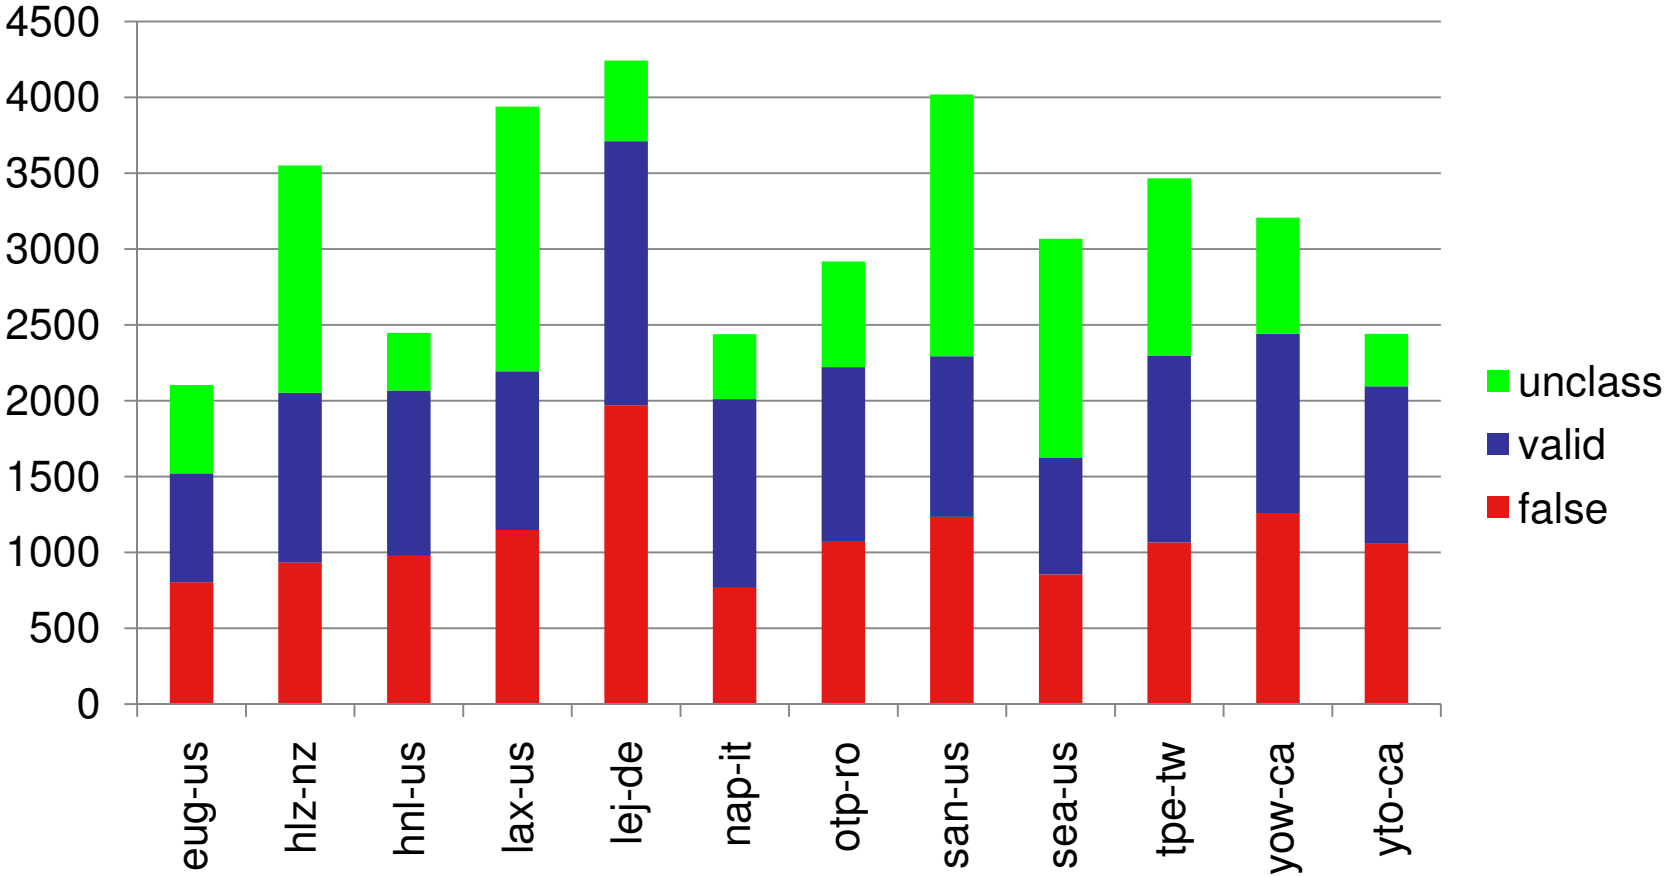


On average, 1094 links per VP are classified as invalid.
2.4% of each VP's link set.
27570 total artifacts. 10653 (39%) false links.

# Artifact link impact: analytical alias resolution

- Analytical alias resolution rule: two addresses in a traceroute path can't be aliases if there are no loops.

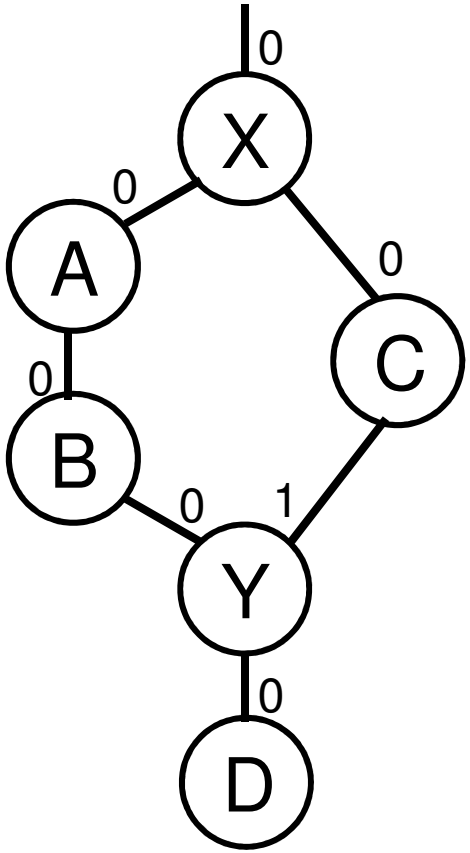


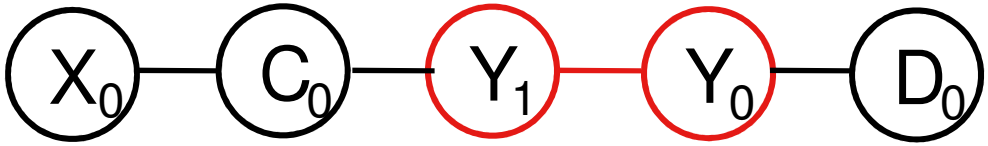


| | | |
|---|---|---|
| ICMP: | 598 / 3359 | (18%) |
| UDP: | 1374 / 10653 | (13%) |

# Validating use of Ally technique to classify artifact links

- Use ping –R and ICMP-Paris traceroute towards destination
  - RR IP address *usually* from egress interface
  - ICMP time-exceeded IP address *usually* from ingress interface
- Infer addresses used in a sequence of /30 subnets
- Identify ICMP/TCP/UDP IP-ID behaviour for each
- Resolve for aliases using Ally for pairs of addresses with incrementing IPID values.

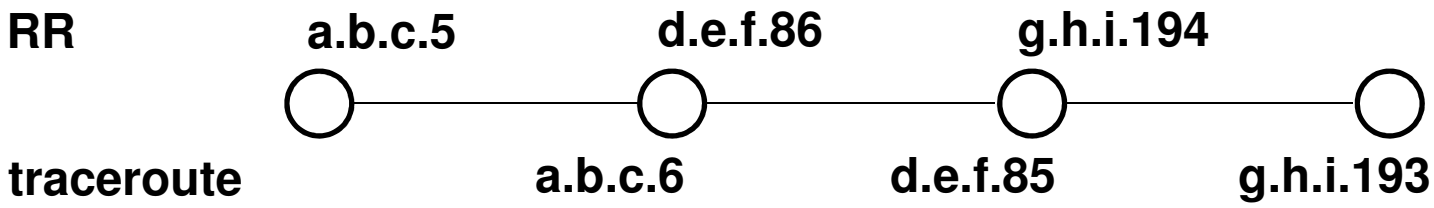

# Validating use of Ally technique: results

- 17 ark monitors, 128237 RR/trace pairs.
- 16285 pairs of likely /30 aliases tested
- Classification obtained for 12200 (75%)
  - Others did not have an incrementing IP-ID for UDP/ICMP/TCP
  - A few targets were classified but then unresponsive to Ally.
- 468 (3.8%) pairs of not-aliases
  - 4.68.110.66 in 156 pairs : structural rejection of /30 inferences, infer that Ally is correct.
  - 64.57.29.98 in 252 pairs : structural rejection of /30 inferences, infer that Ally is correct.
  - Have not investigated other 0.5%
- Result taken: safe to use Ally to rule on aliases where it is usable.

# Summary

- Small but measurable impact on graph produced
  - 0.8% of links in classic ICMP traceroute are invalid per VP
  - 2.4% of links in classic UDP traceroute are invalid per VP
- Classic ICMP-echo approach not as affected by per-flow load balancers as UDP approach.
- Larger problem: heuristics and hacks we use to build IP-layer maps of the Internet sometimes don't hold. VERY HARD TO VALIDATE.

# Future Work, Open Questions

- Future work
  - Extend data collection to use 40 Ark VPs rather than 12
  - Work towards techniques that annotate graphs with the likelihood a traceroute link is valid. Hard, as heuristics fall over in face of incomplete data.
  - Ground truth data
- Open questions
  - How do false links accumulate in traceroute graphs?
  - What is the impact of false links on the graph's properties?
  - What are the limitations of this work's methodology, and can they be addressed?