Deploying MDA Traceroute on RIPE Atlas Probes

Kevin Vermeulen[1], Stephen Strowes[2], Timur Friedman[1]

[1]Sorbonne University, [2]RIPE NCC

Summary

- Multipath Detection Algorithm (MDA) and its limits

- Towards a better MDA:
  - Survey on load balancers
  - Provide heuristics based on data
  - Results

Multipath Detection Algorithm: Definition

- Allows to discover all the paths between a source and a destination, based on paris-traceroute
- Statistical guarantees on the discovered topology
- Potentially sends tens of thousands of packets to discover all the topology
- Makes the worst case hypothesis that every discovered interface could be part of a load balancer

Survey (work in progress)

- 350,000 traceroutes towards destinations from IMPACT IP Hitlist
  - Work divided among 35 PlanetLab nodes as sources
- 100,000 traceroutes computed at the moment (computing still in progress)
- 40% of the traceroutes contained at least one diamond

Survey: diamond lengths

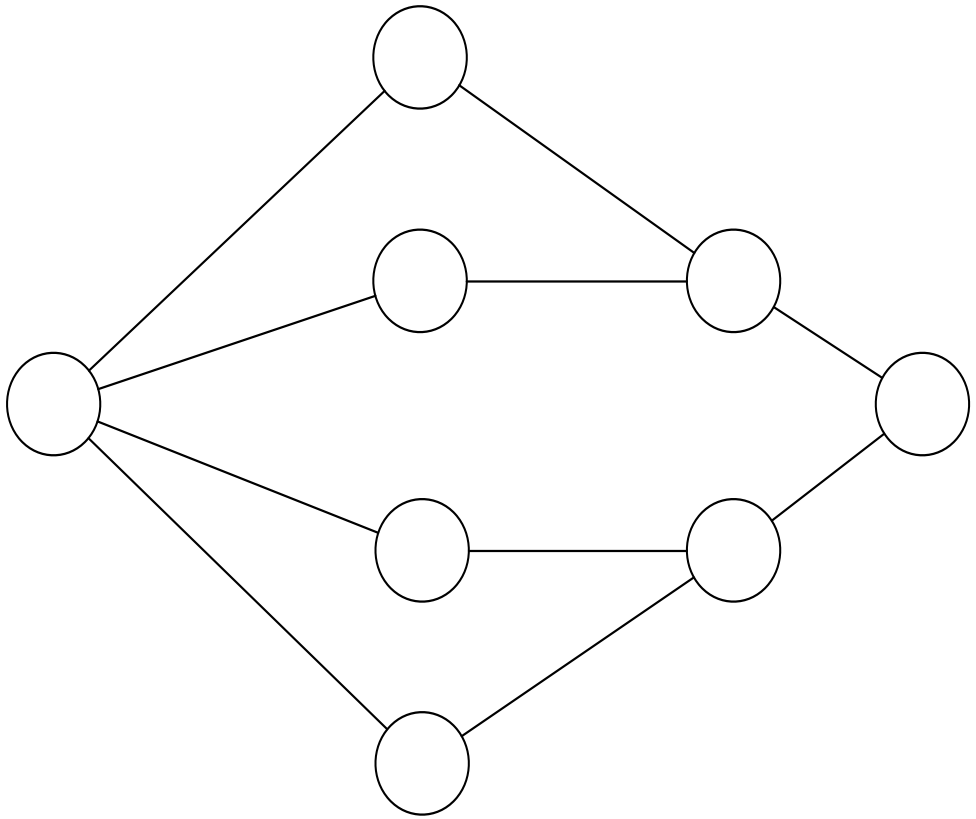

max_length = min_length = 2

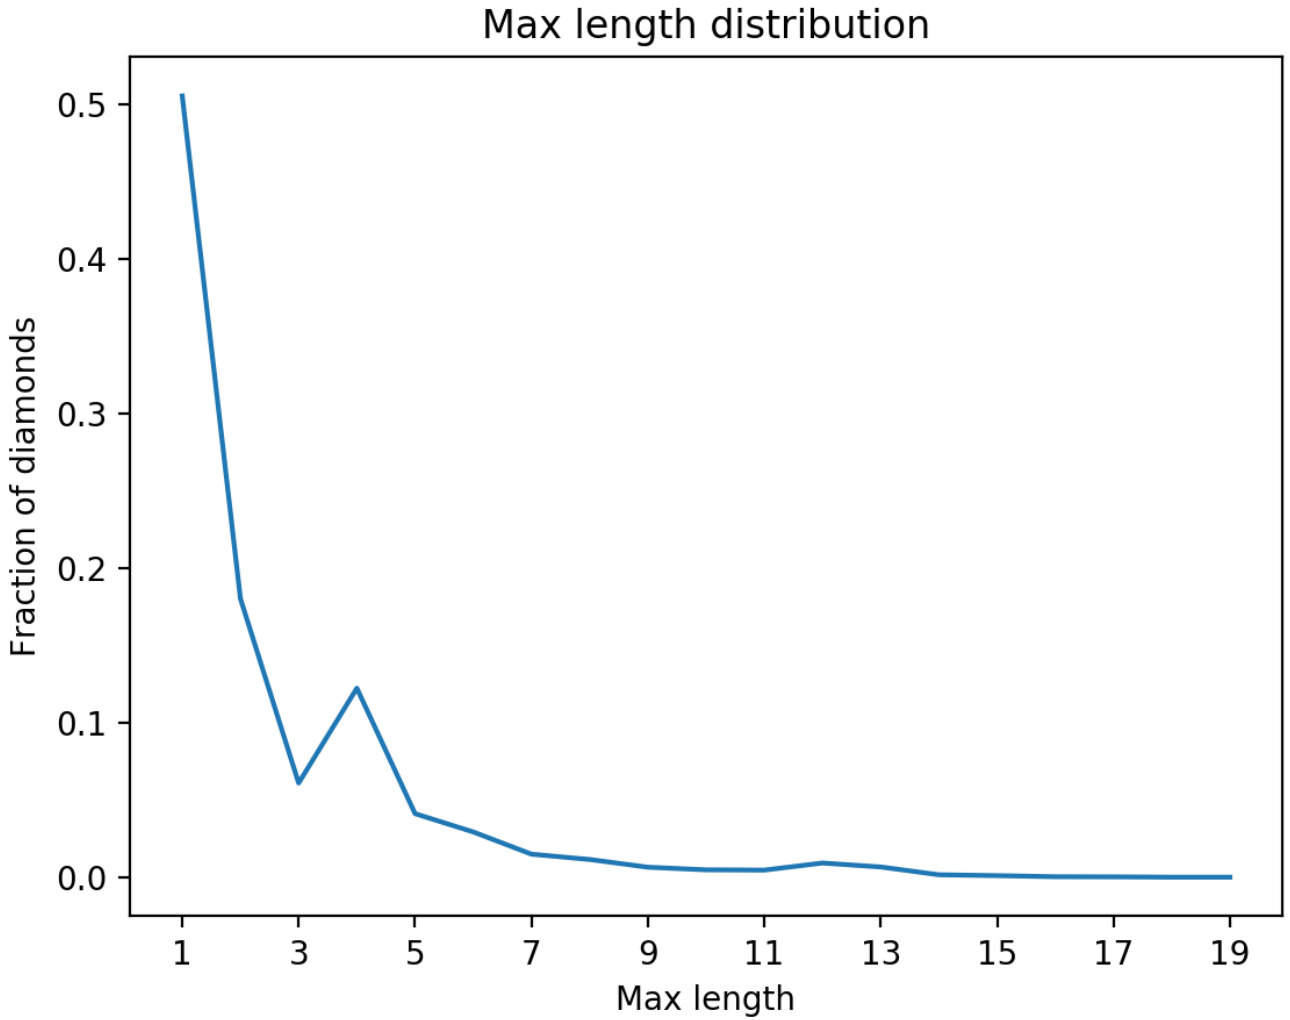

Survey: diamond widths

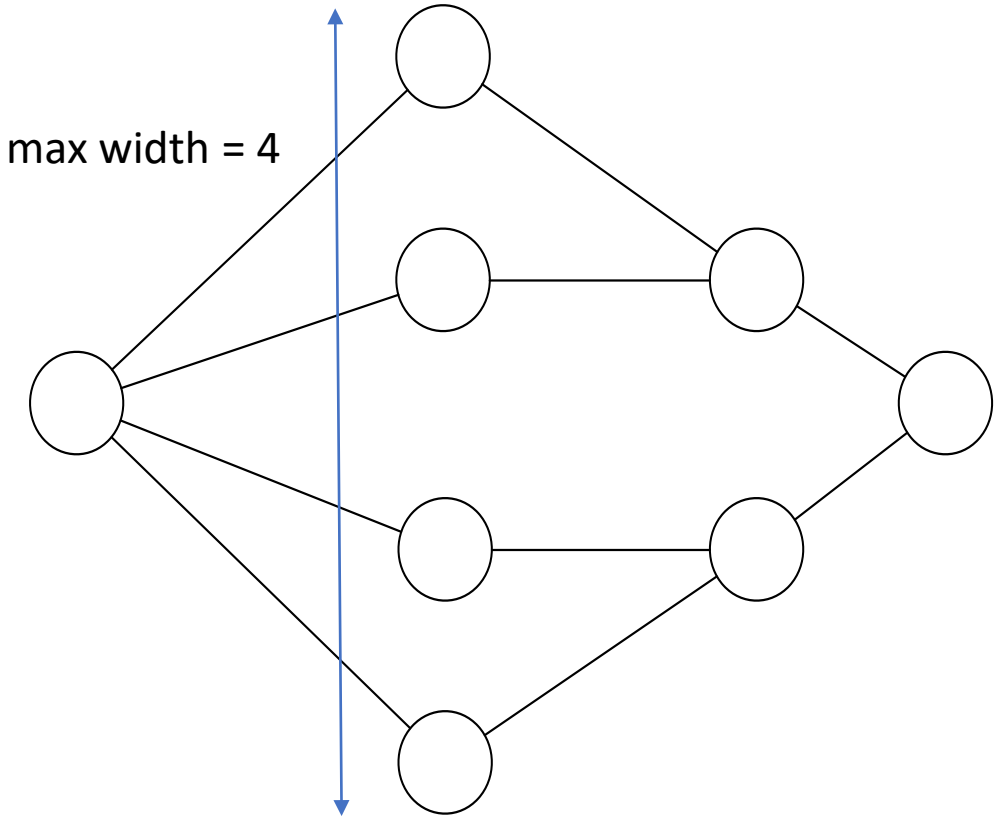

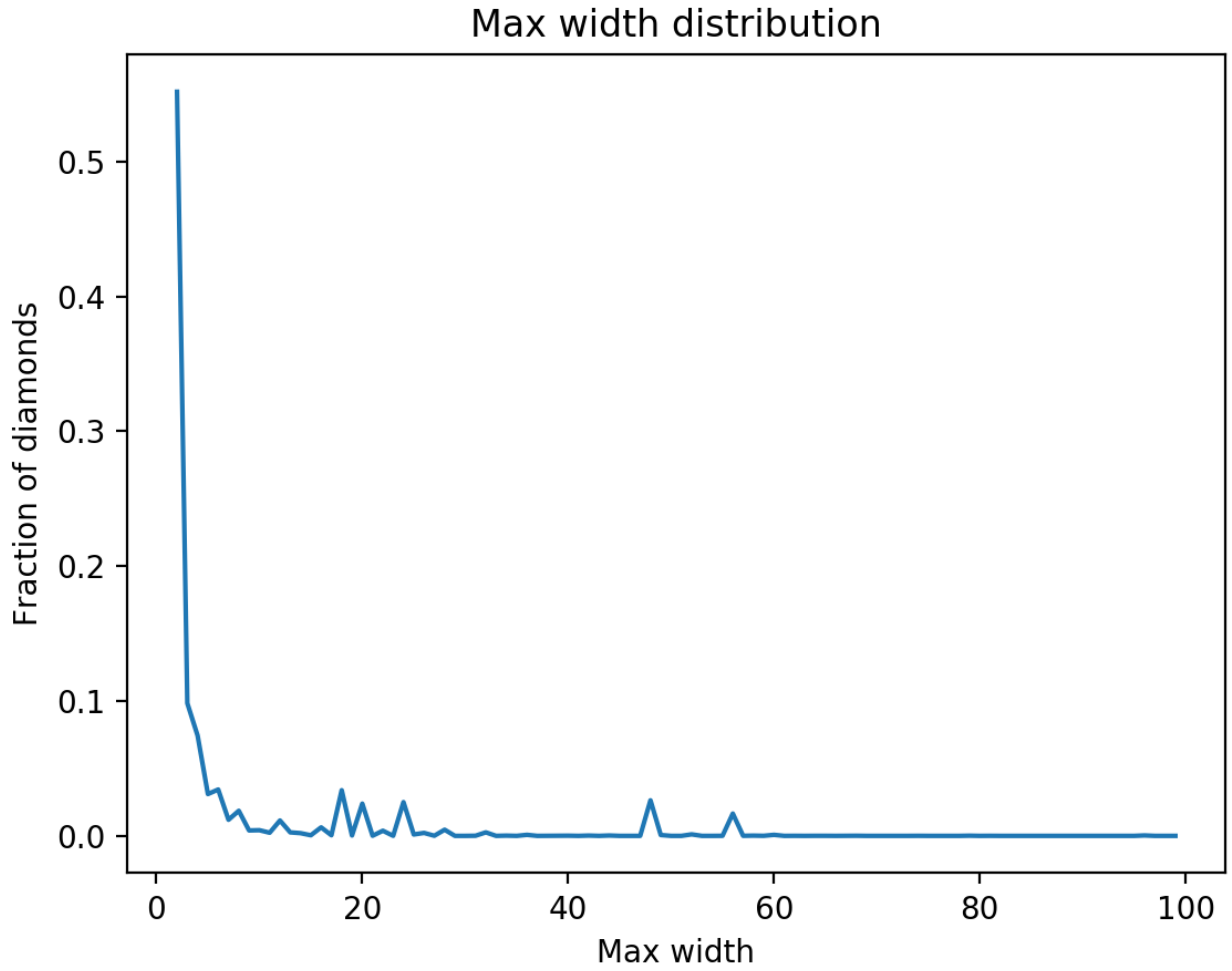

Survey: width asymmetry

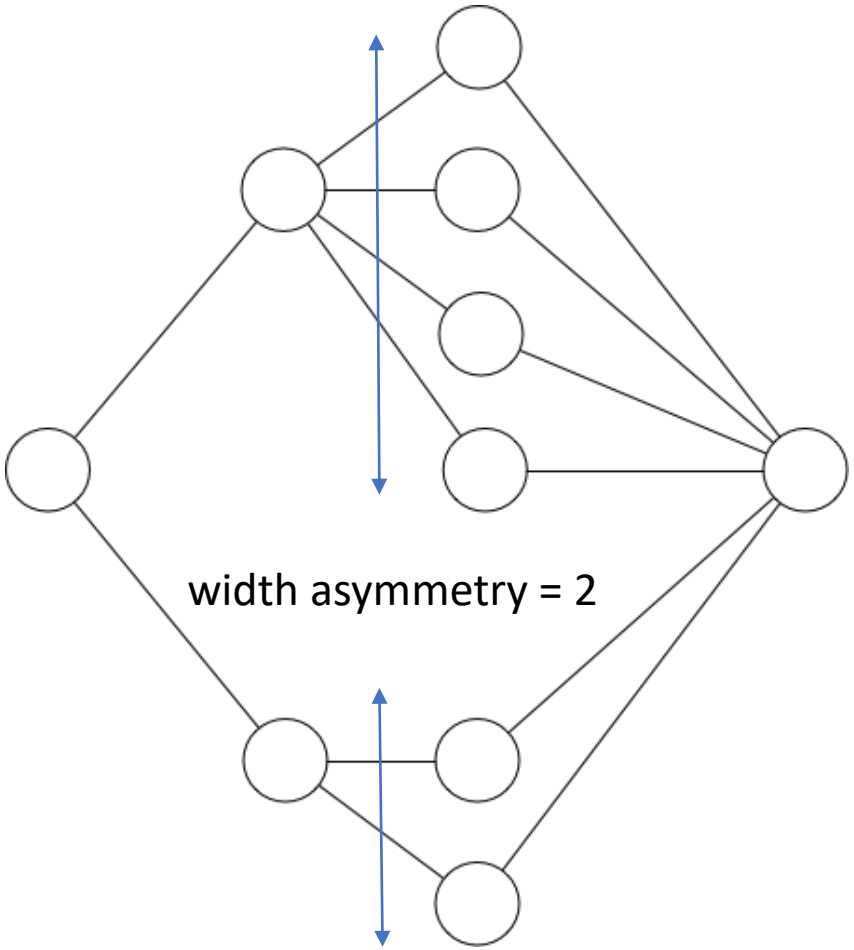

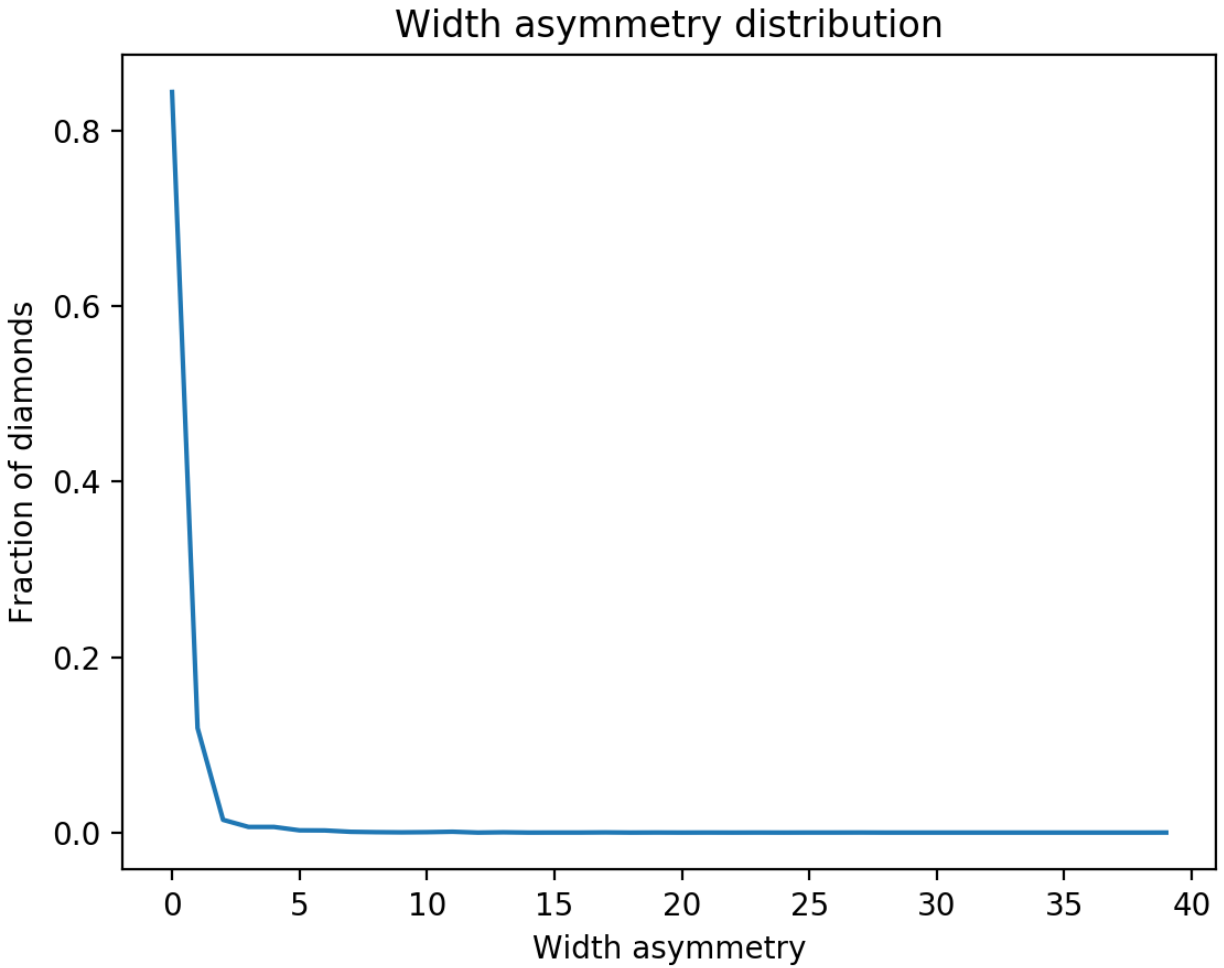

Survey: meshed diamonds

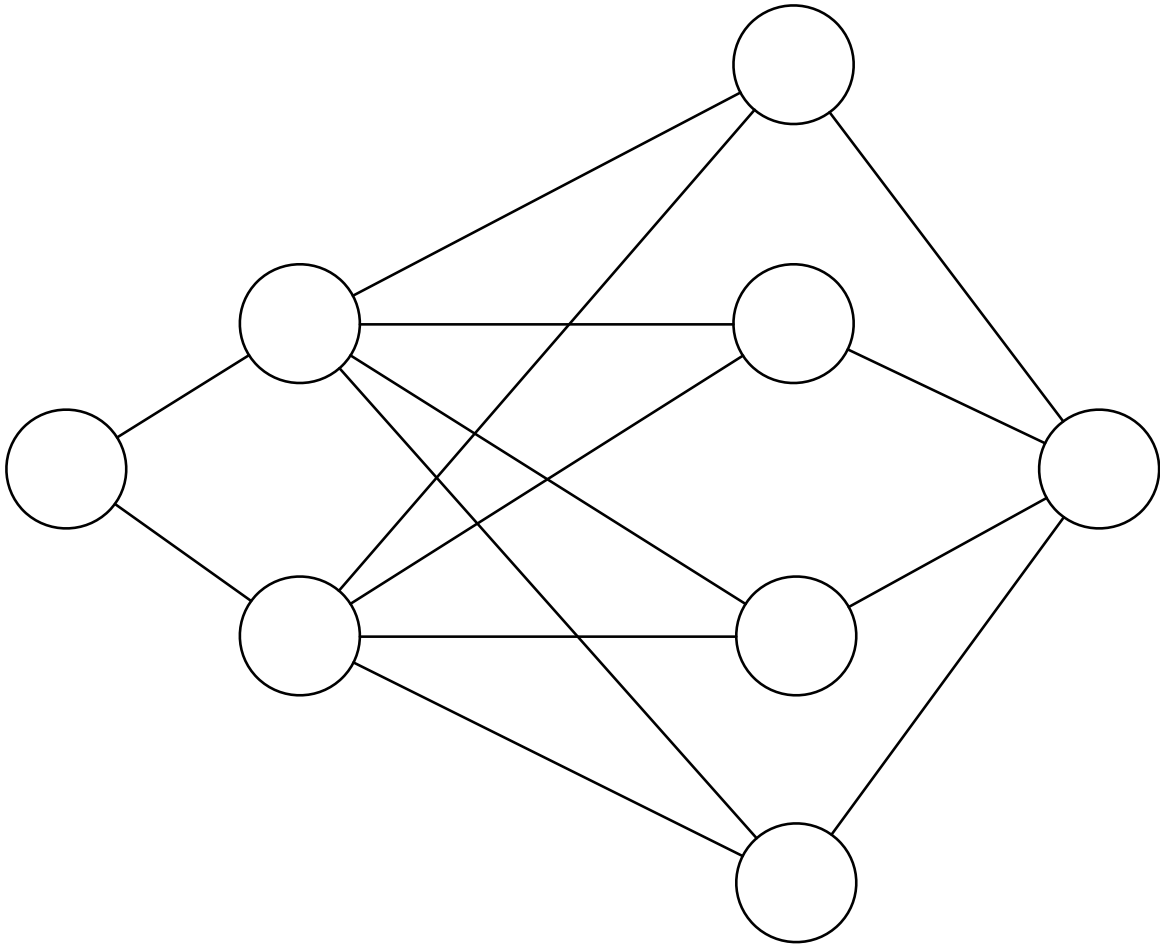

- 15.3 % are meshed diamond
- More meshing metrics are being defined in our ongoing work

The MDA uses 8500 packets to discover this topology!
Source : ple2.planet-lab.eu
Destination : 125.155.82.17

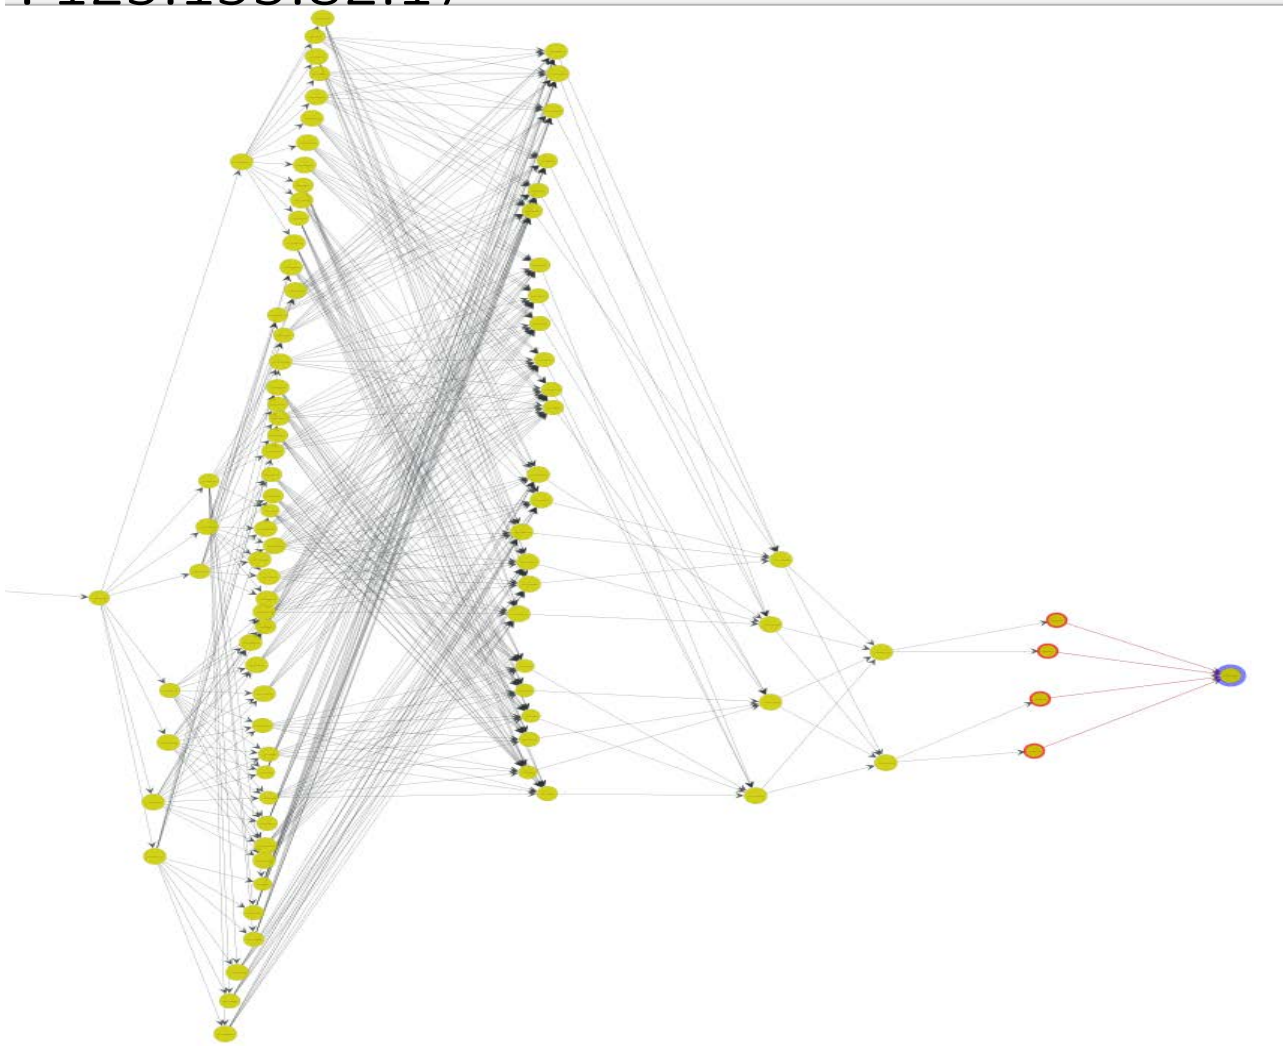

MDA overhead: packets sent

- hop 1: one node discovered
  - Is there a second node?
  - Send n2 = 8 packets to hop 1
  - No more nodes found
- hop 2: six nodes discovered
  - Is there a seventh node?
  - Send n7 = 43 packets
  - No more nodes found
- hop 3: each hop 2 node is a potential branching point!
  - Must find n2 = 8 flows that go to each hop 2 interface
  - Send all 6x8 = 48 packets to hop 3
  - Only 1 node found
- **Can we do better?**

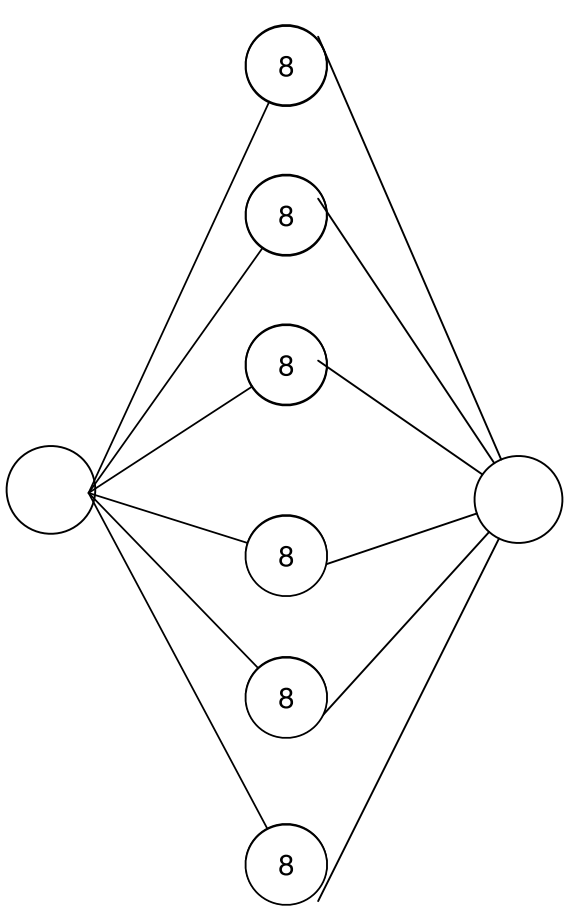

Measured topology

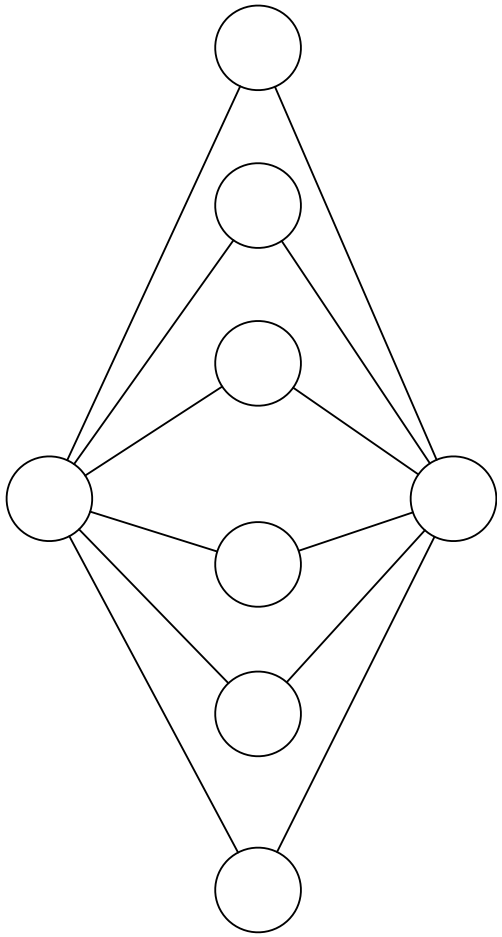

Ground truth

Towards a better MDA

- hop 1 : n2 = 8
- hop 2 : n7 = 43
- hop 3 : n2 = 8
- Assumption: equal probability to reach any of the interfaces at hop 3
- We make this assumption because our survey reveals that most diamonds are symmetric

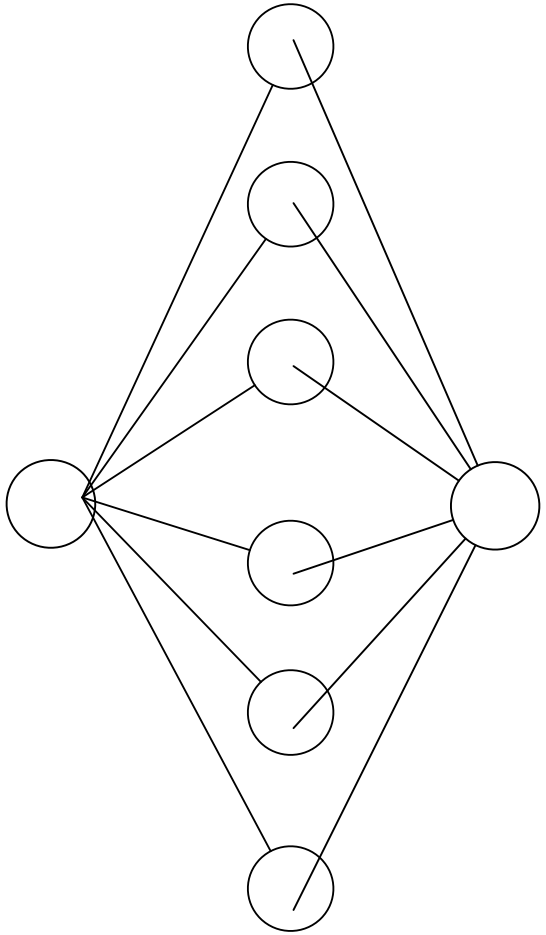

Measured topology

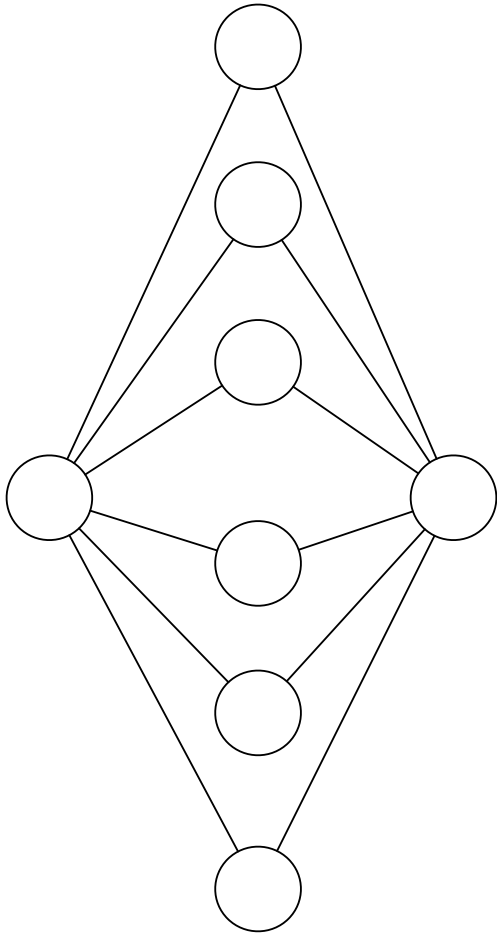

Ground truth

Towards a better MDA

- For each combination of these characterictics: symmetry, asymmetry, meshed, not meshed

-> We provide heuristics to save probes

New vs Classic MDA on a length 1 diamond

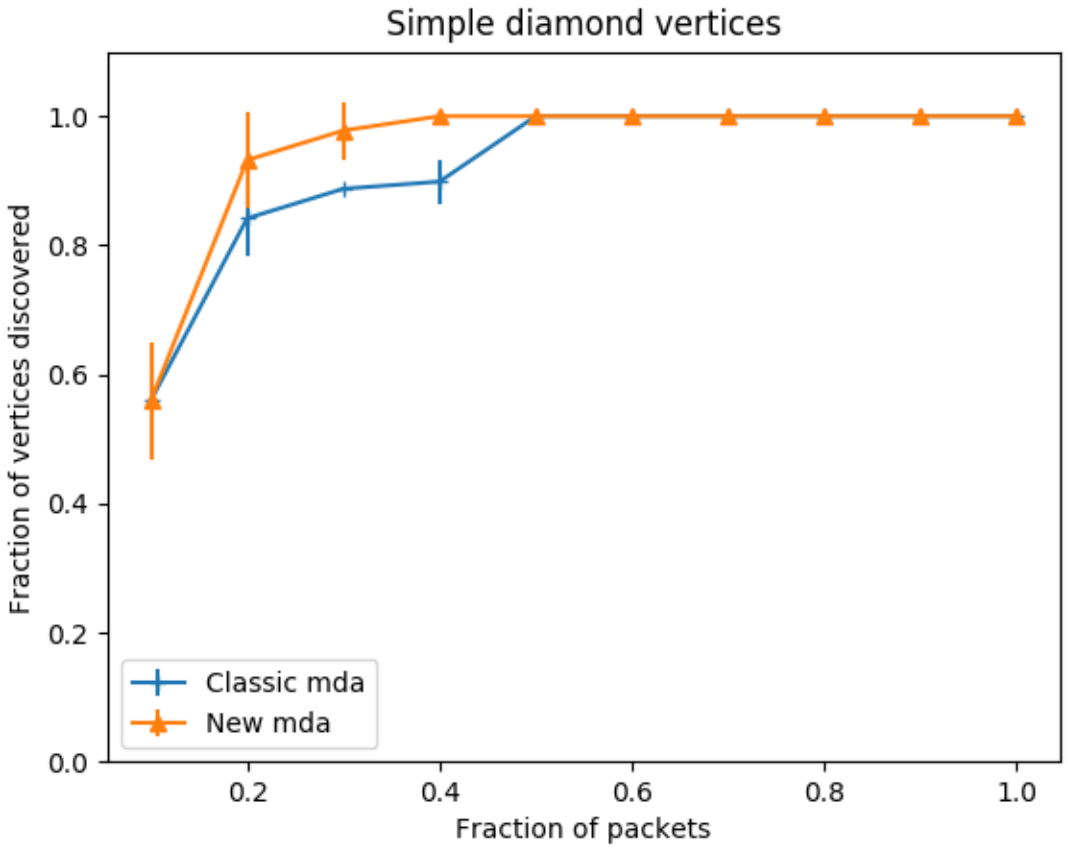

40% of packets are needed to discover all the vertices

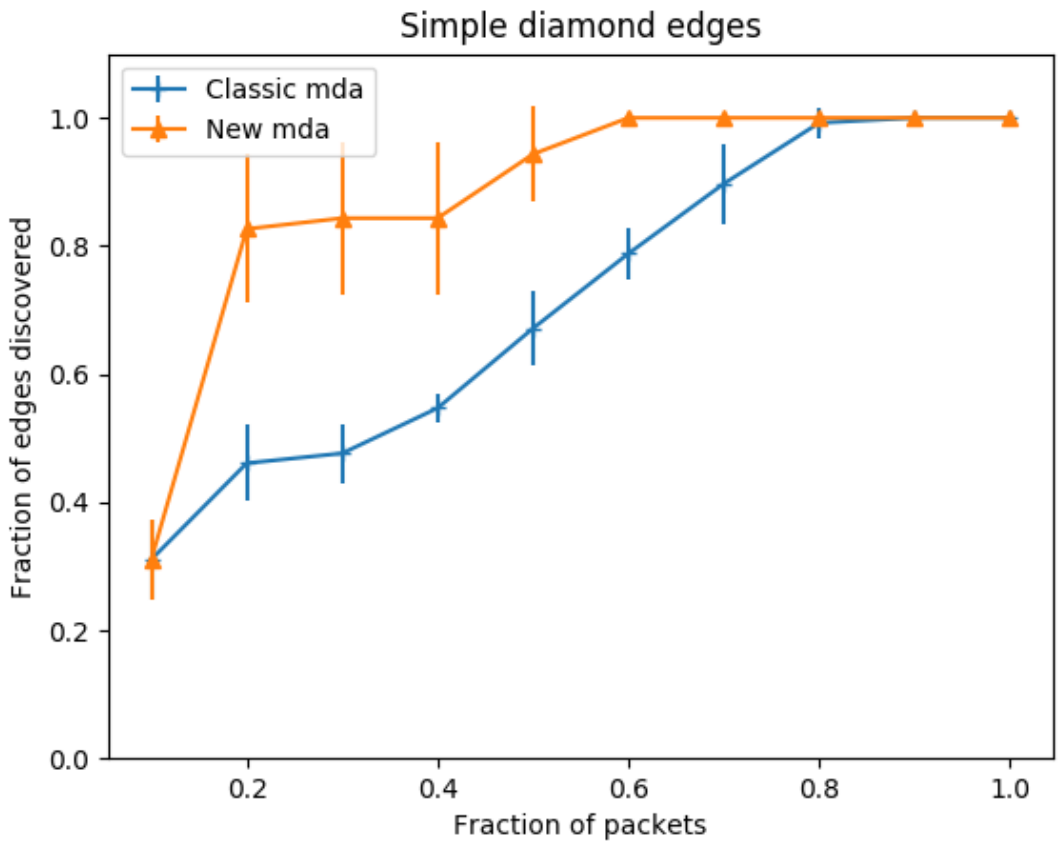

60% of packets are needed to discover all the edges

New vs Classic MDA on a longer symmetric diamond

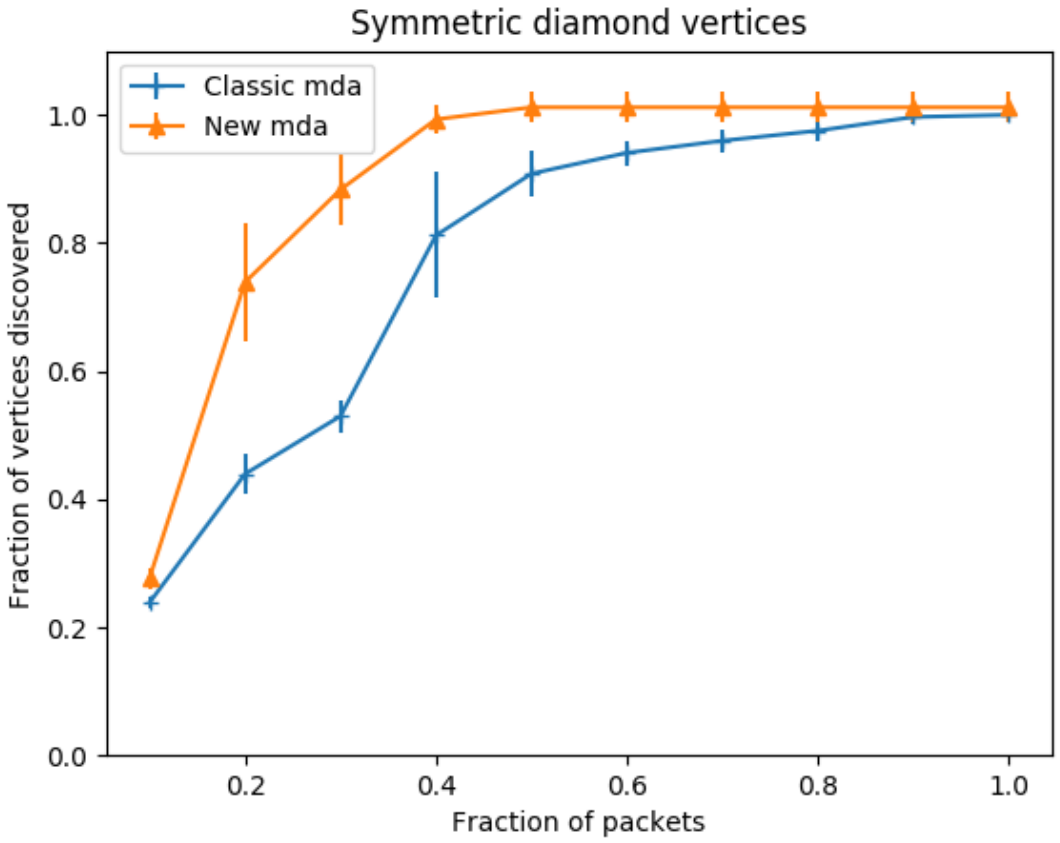

50% of packets are needed to discover all the vertices

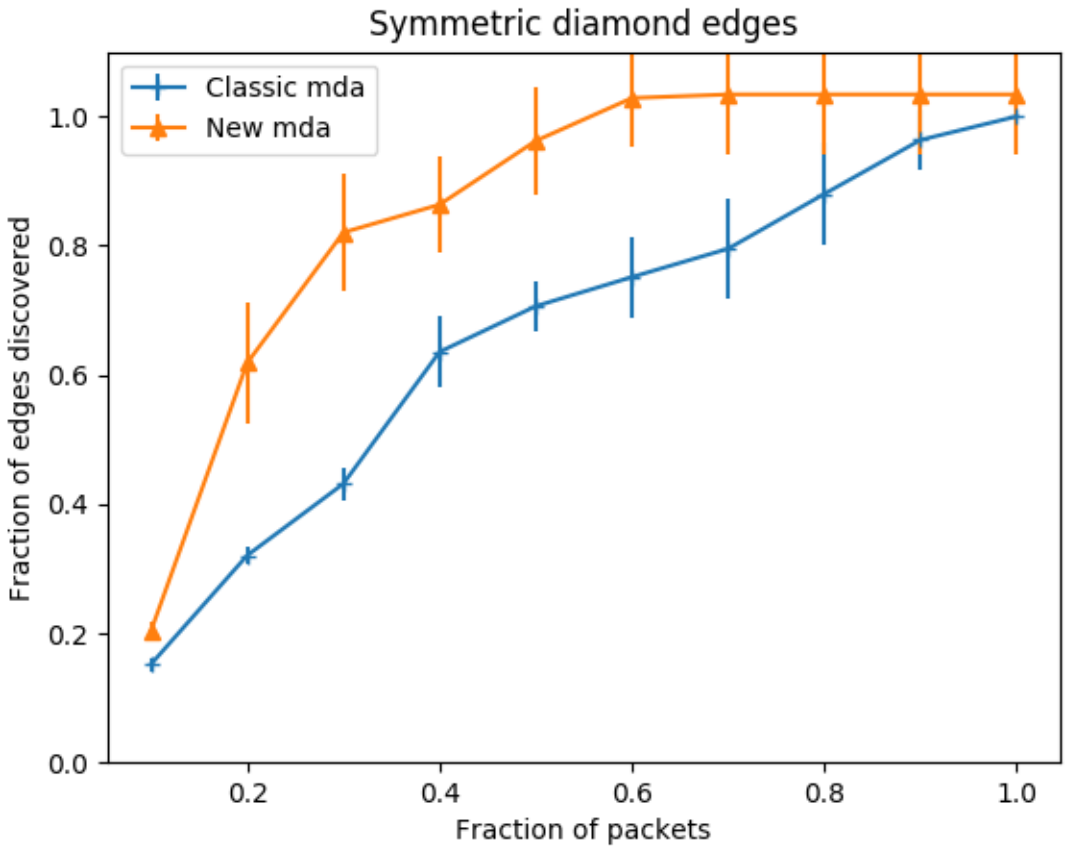

60% of packets are needed to discover all the edges

New vs Classic MDA on a meshed diamond

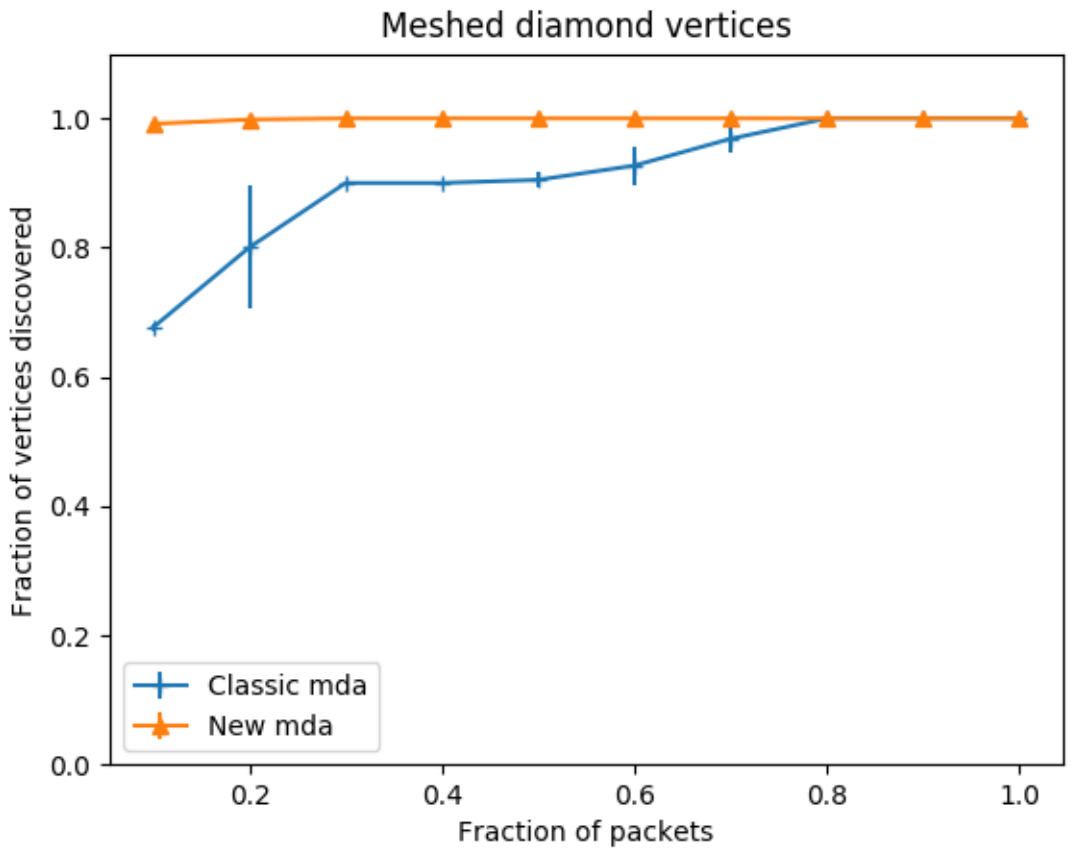

20% of packets are needed to discover all the vertices

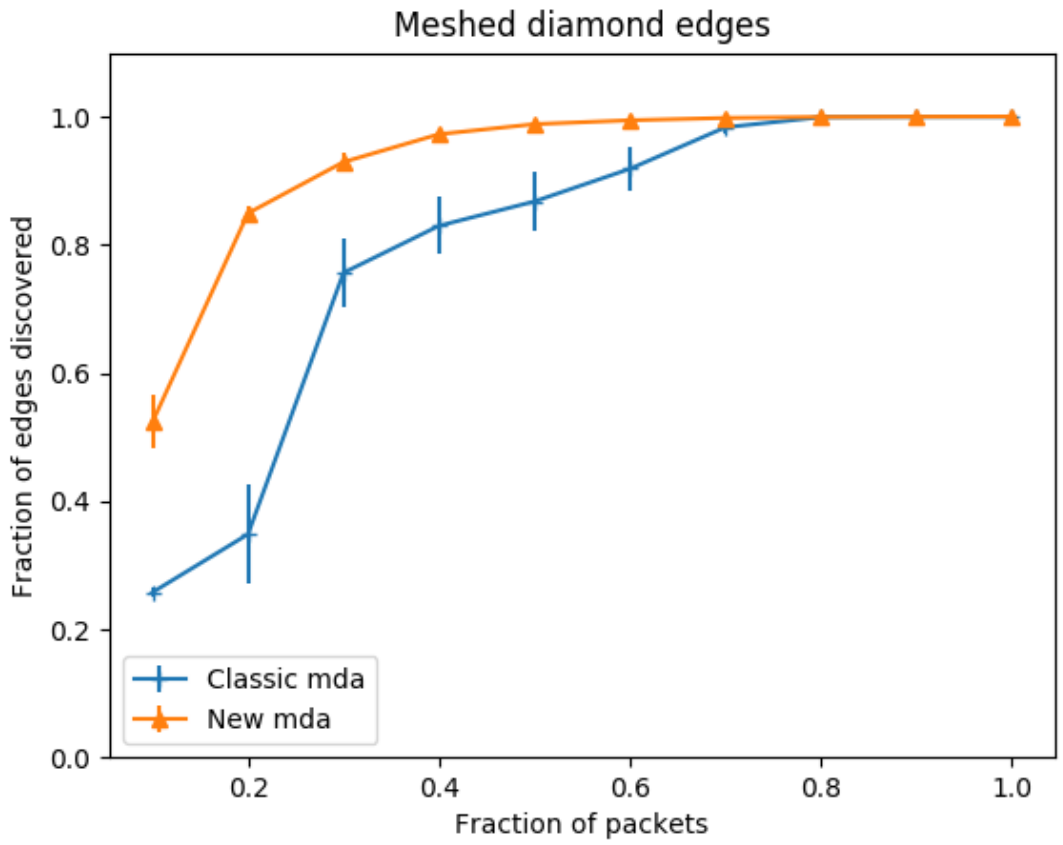

60% of packets are needed to discover all the edges

New vs Classic MDA on an asymmetric diamond

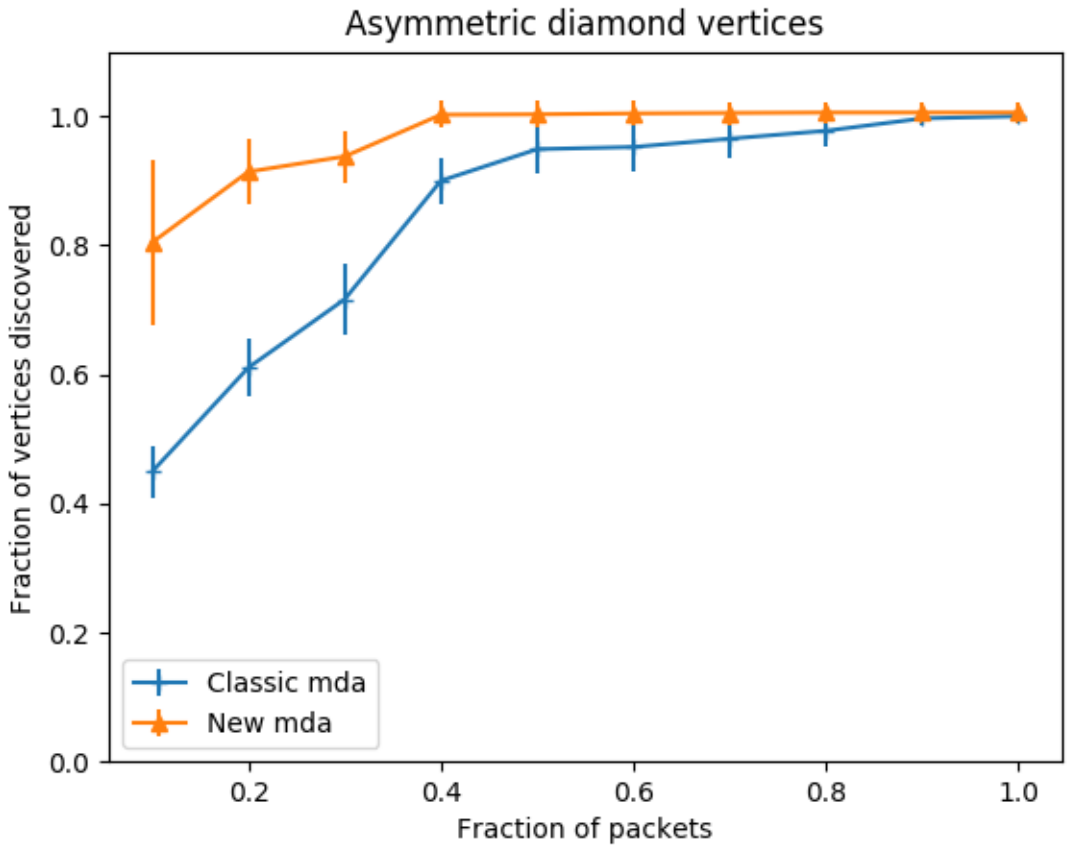

40% of packets are needed to discover all the vertices

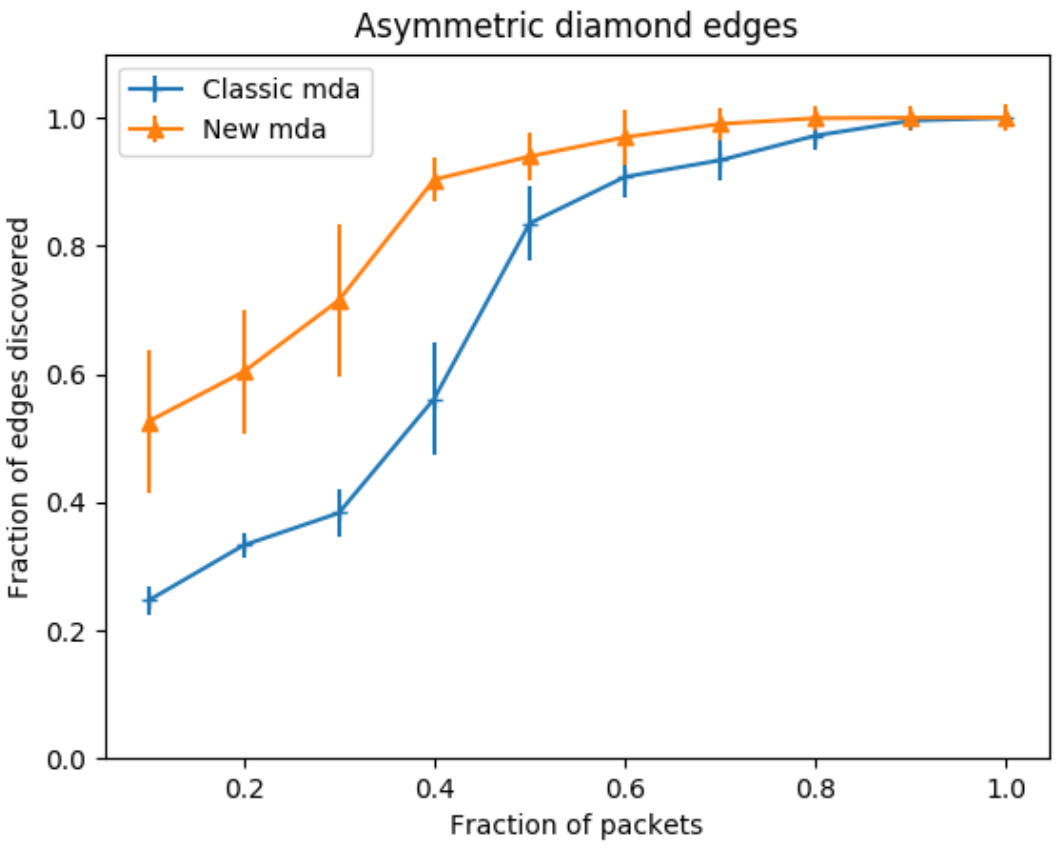

80% of packets are needed to discover all the edges

References

- [1]https://paris-traceroute.net/images/infocom2009.pdf
- [2]https://ant.isi.edu/datasets/all.html
- [3]http://mat.uab.cat/matmat/PDFv2014/v2014n02.pdf

Questions?