Diamond-Miner: Comprehensive Discovery of the Internet's Topology Diamonds

Kévin Vermeulen, Justin P. Rohrer, Robert Beverly,
Olivier Fourmaux, Timur Friedman

USENIX NSDI 2020
26th February 2020

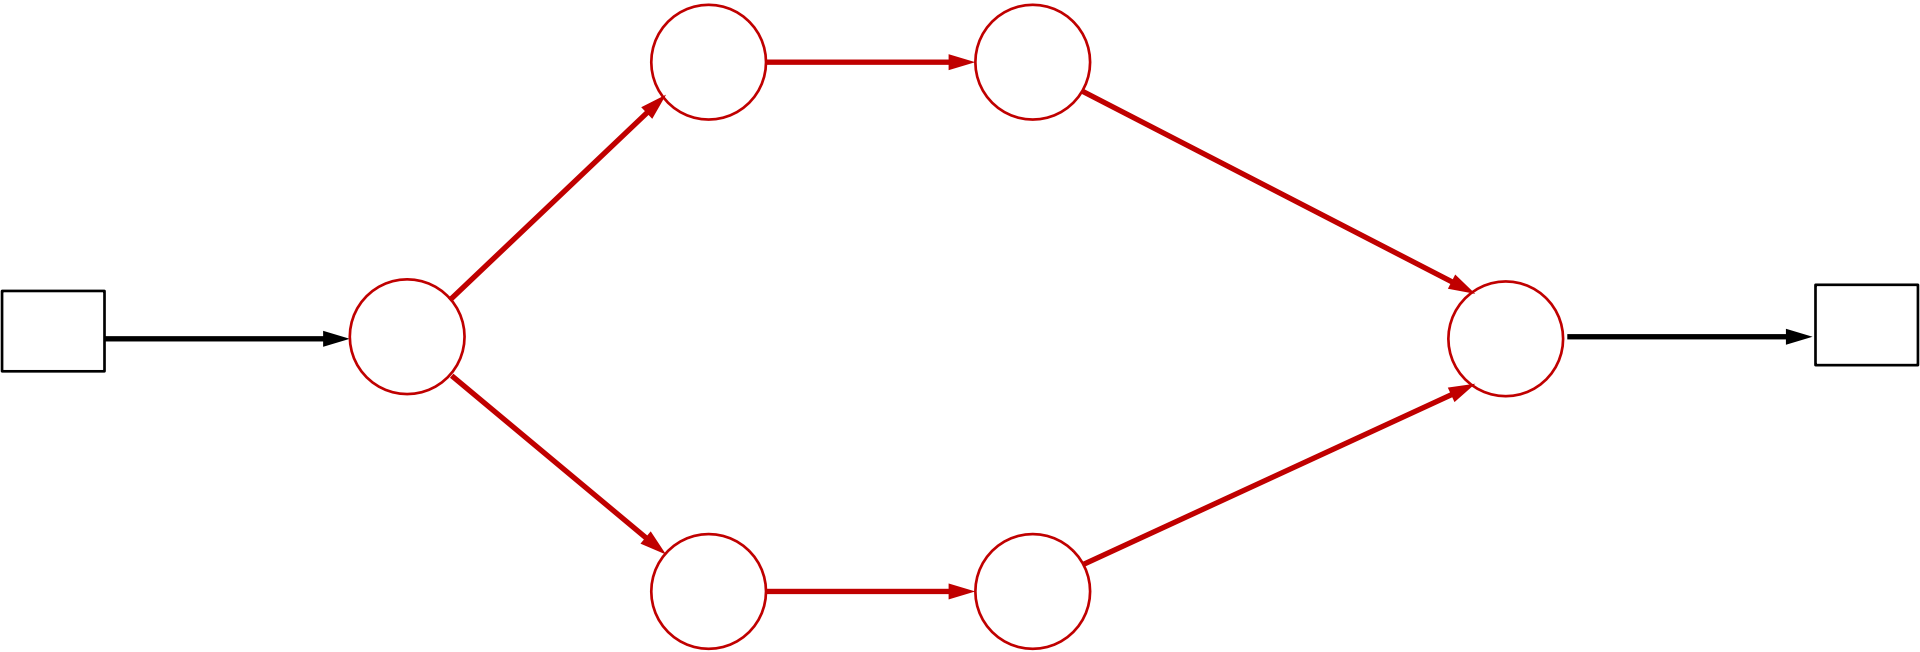

Diamond

Motivation

- IP load balancing is prevalent:
  - Capacity
  - Redundancy
- Existing techniques do not accurrately capture load balanced paths at Internet scale
- Bad idea of true resilience and structure of the Internet
- Diamond-Miner revealed 64% more links than existing Internet maps

Mapping diamonds today

- Single path probing: traceroute, Paris traceroute

- Multipath probing: MDA Paris traceroute

A toy example

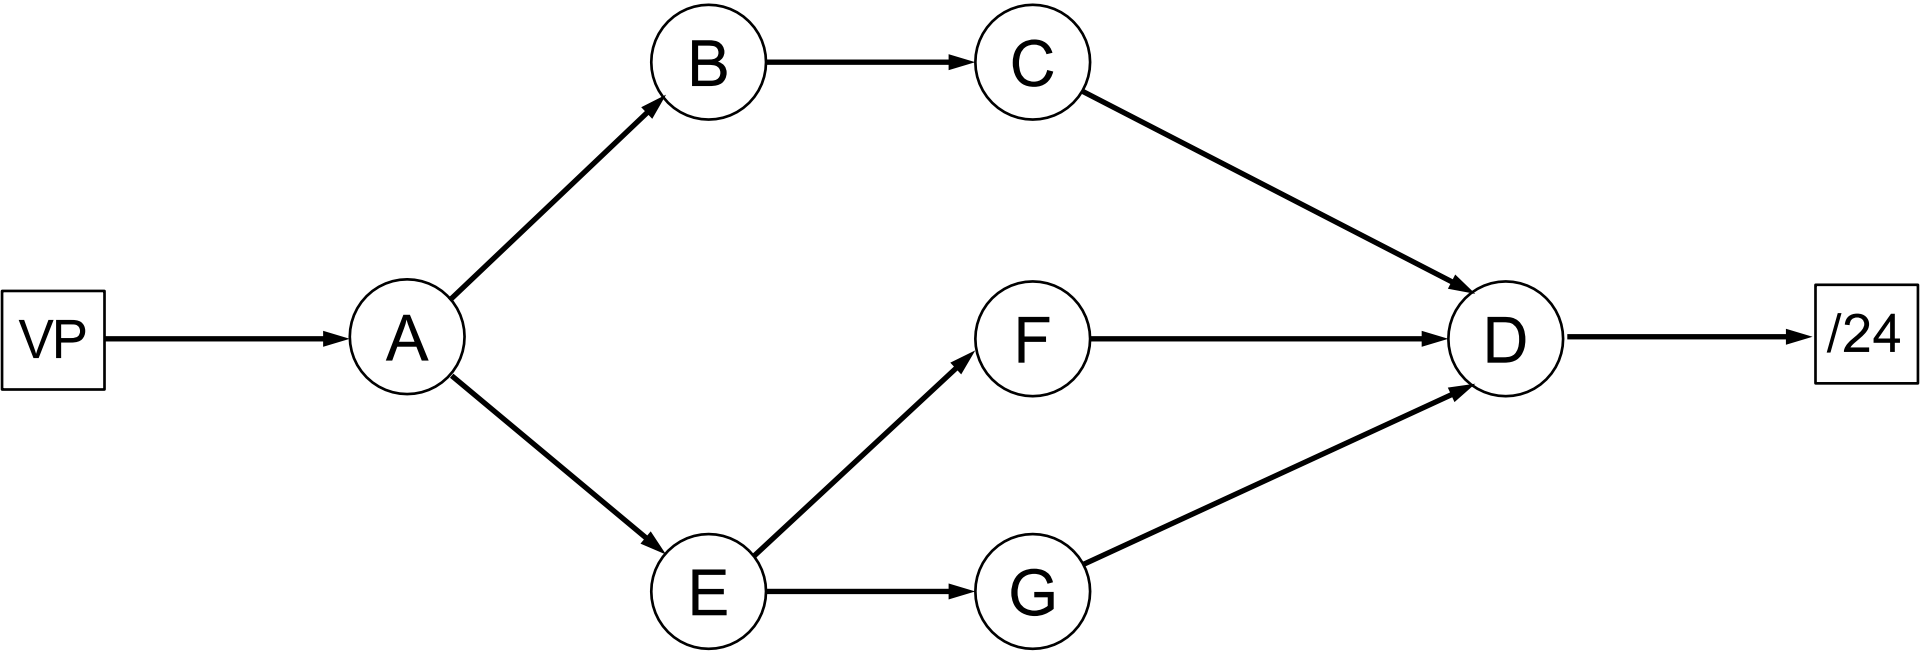

Single path probing: traceroute, an incomplete technique

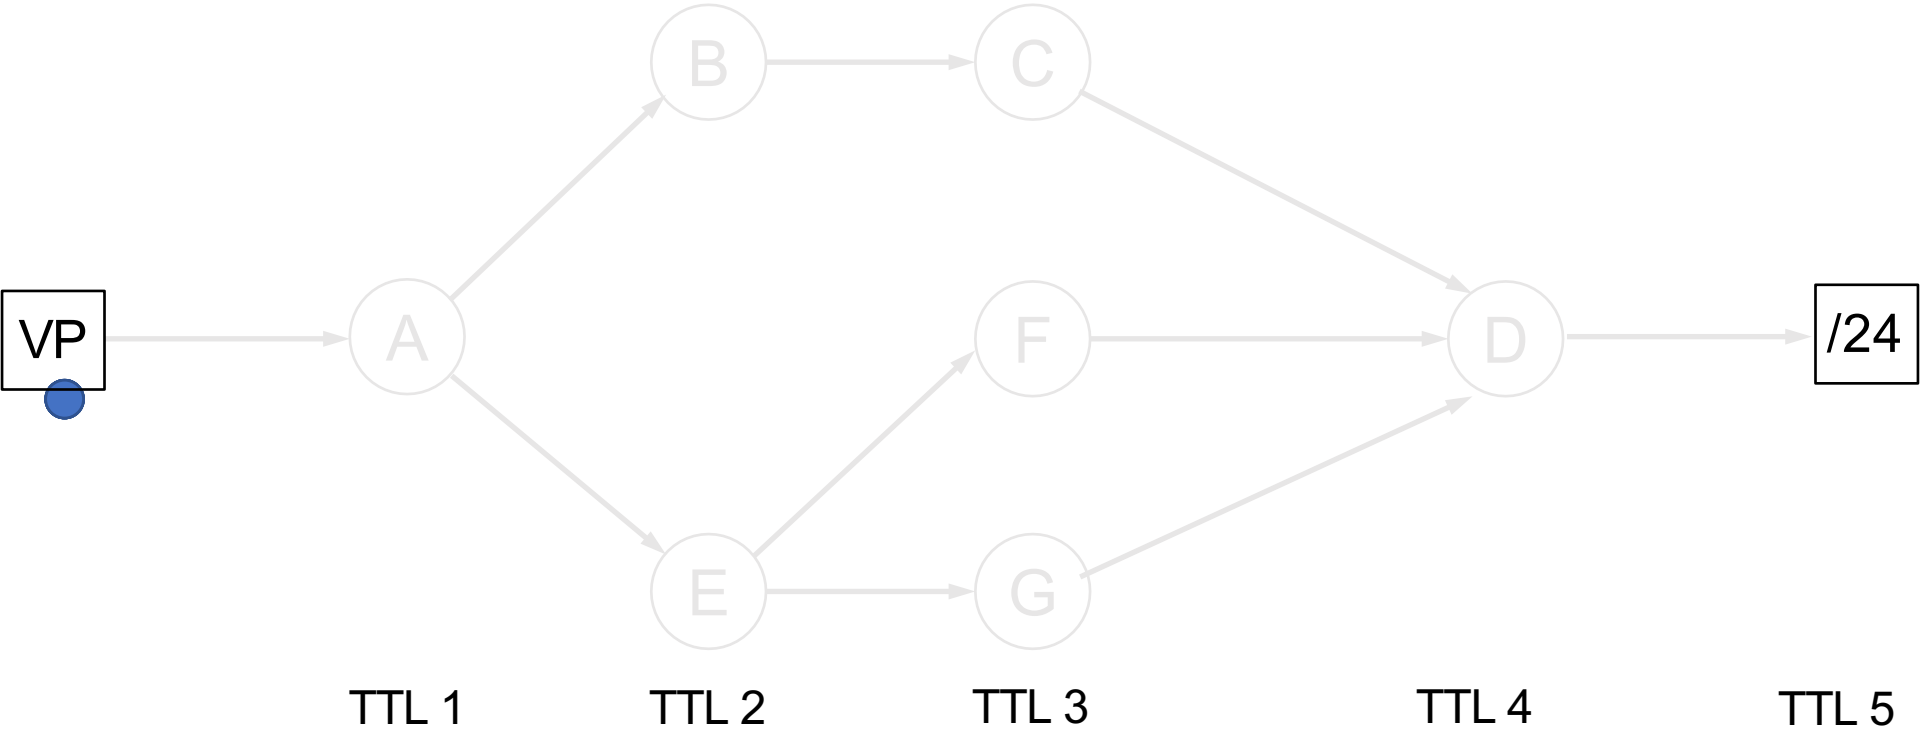

Multipath probing: MDA Paris traceroute, a hop by hop resolving technique

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

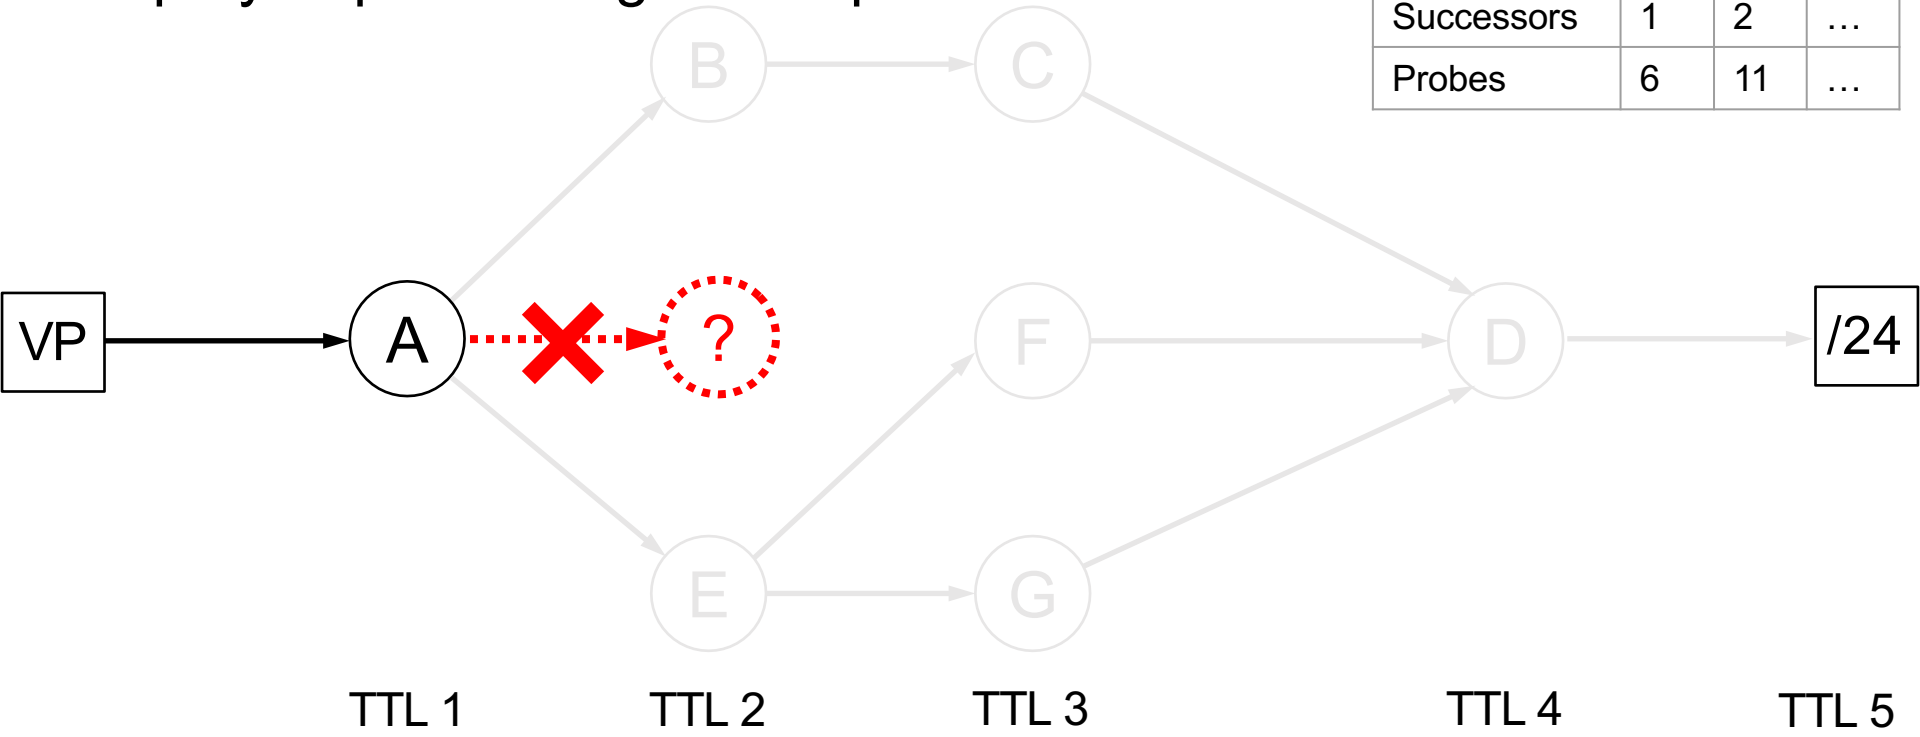

Multipath probing: MDA Paris traceroute, a hop by hop resolving technique

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

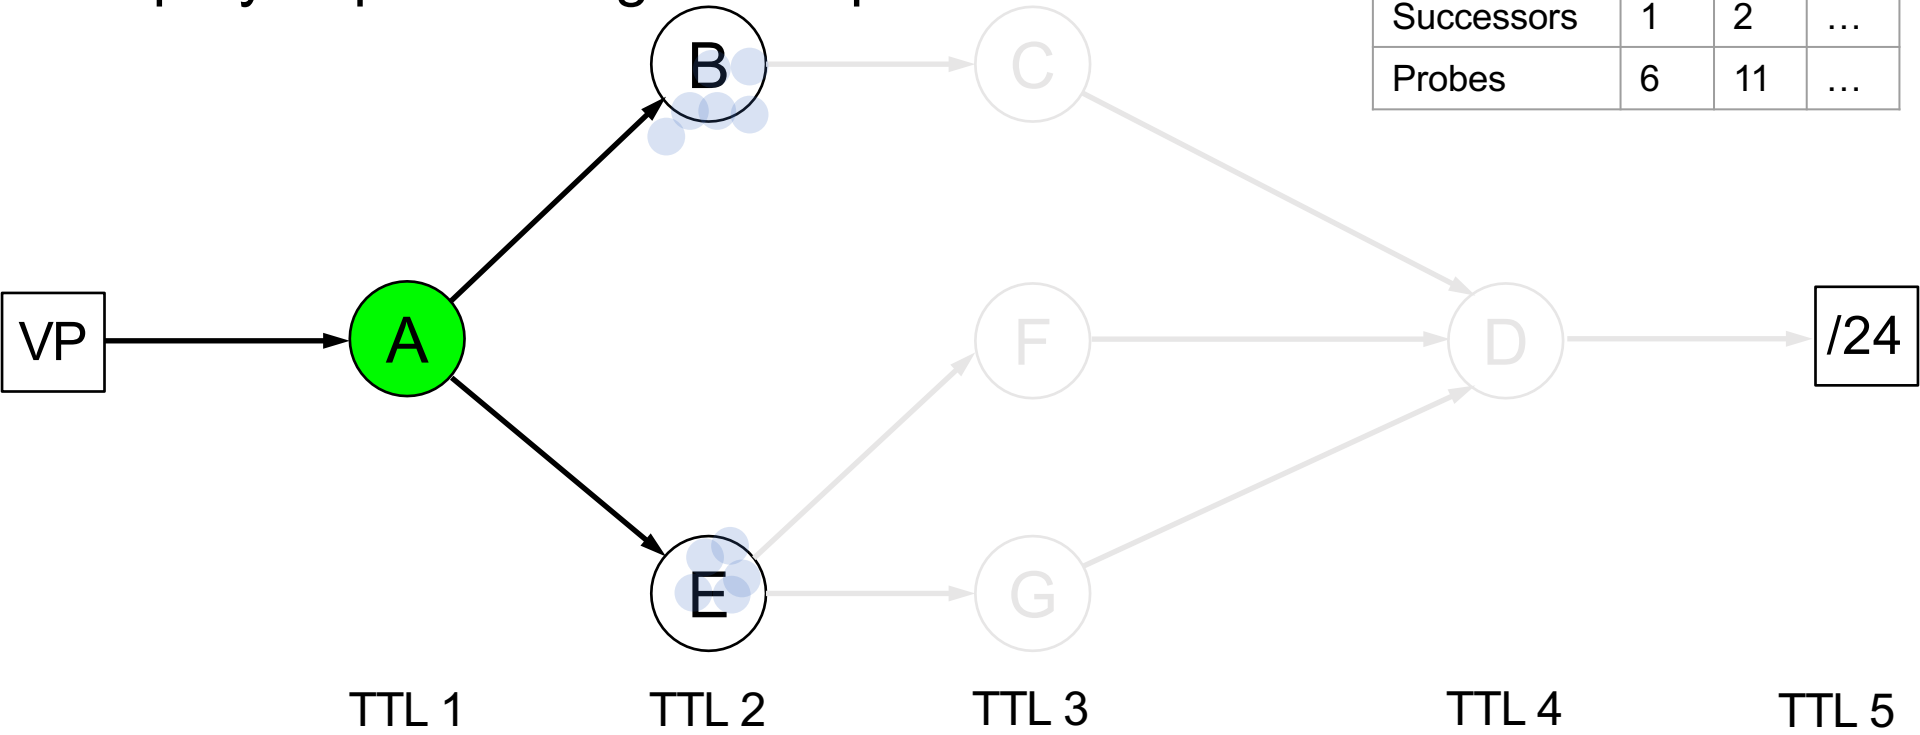

Multipath probing: MDA Paris traceroute, a hop by hop resolving technique

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

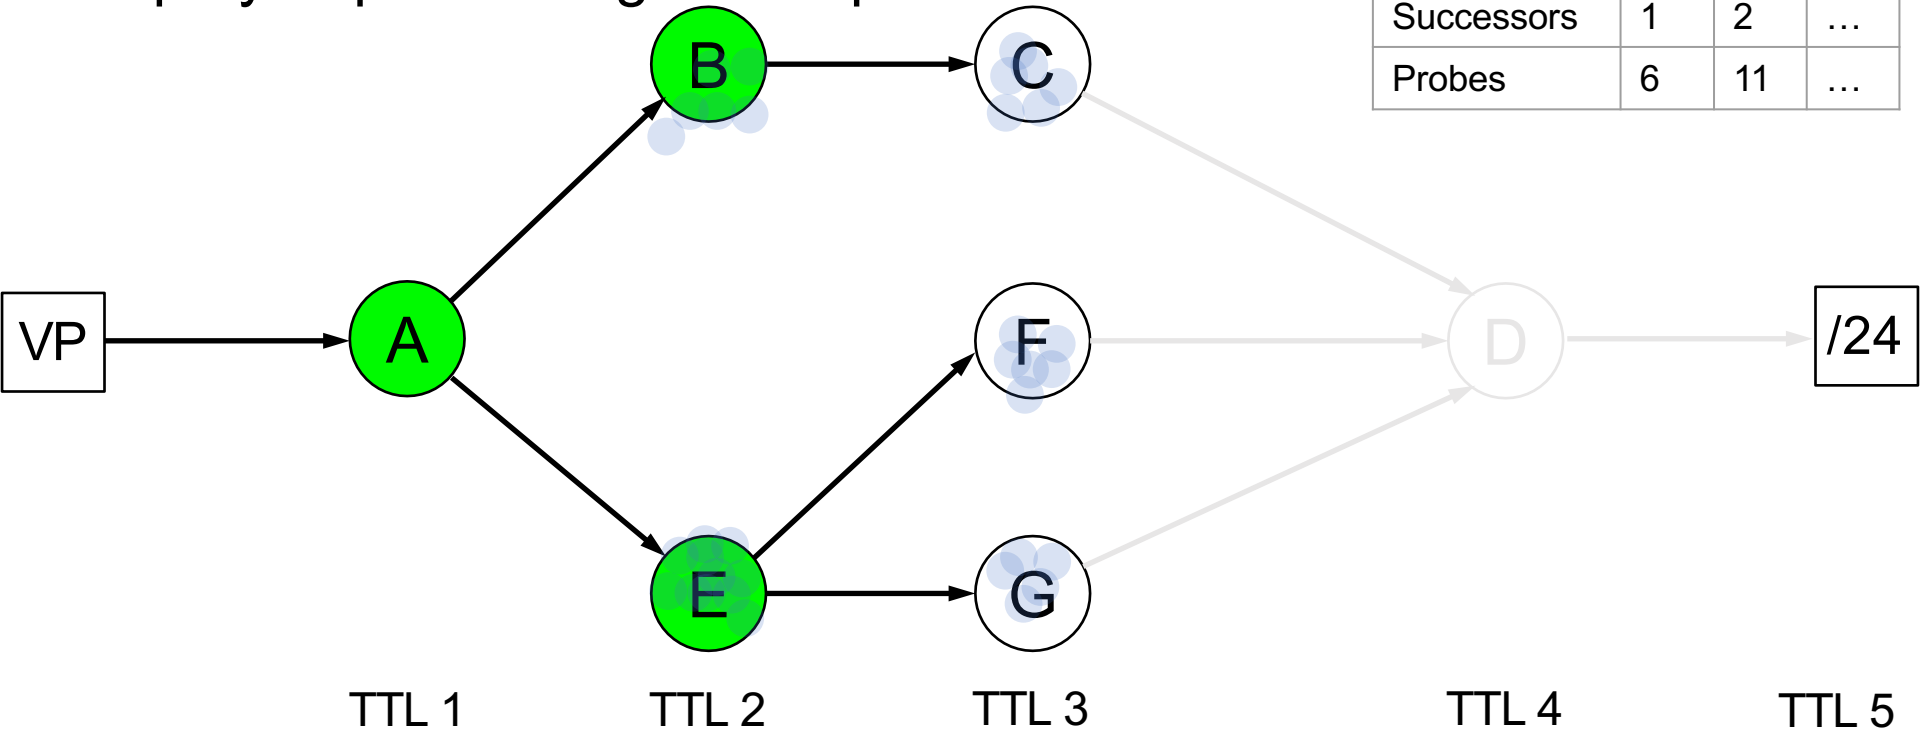

Multipath probing: MDA Paris traceroute, a hop by hop resolving technique

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

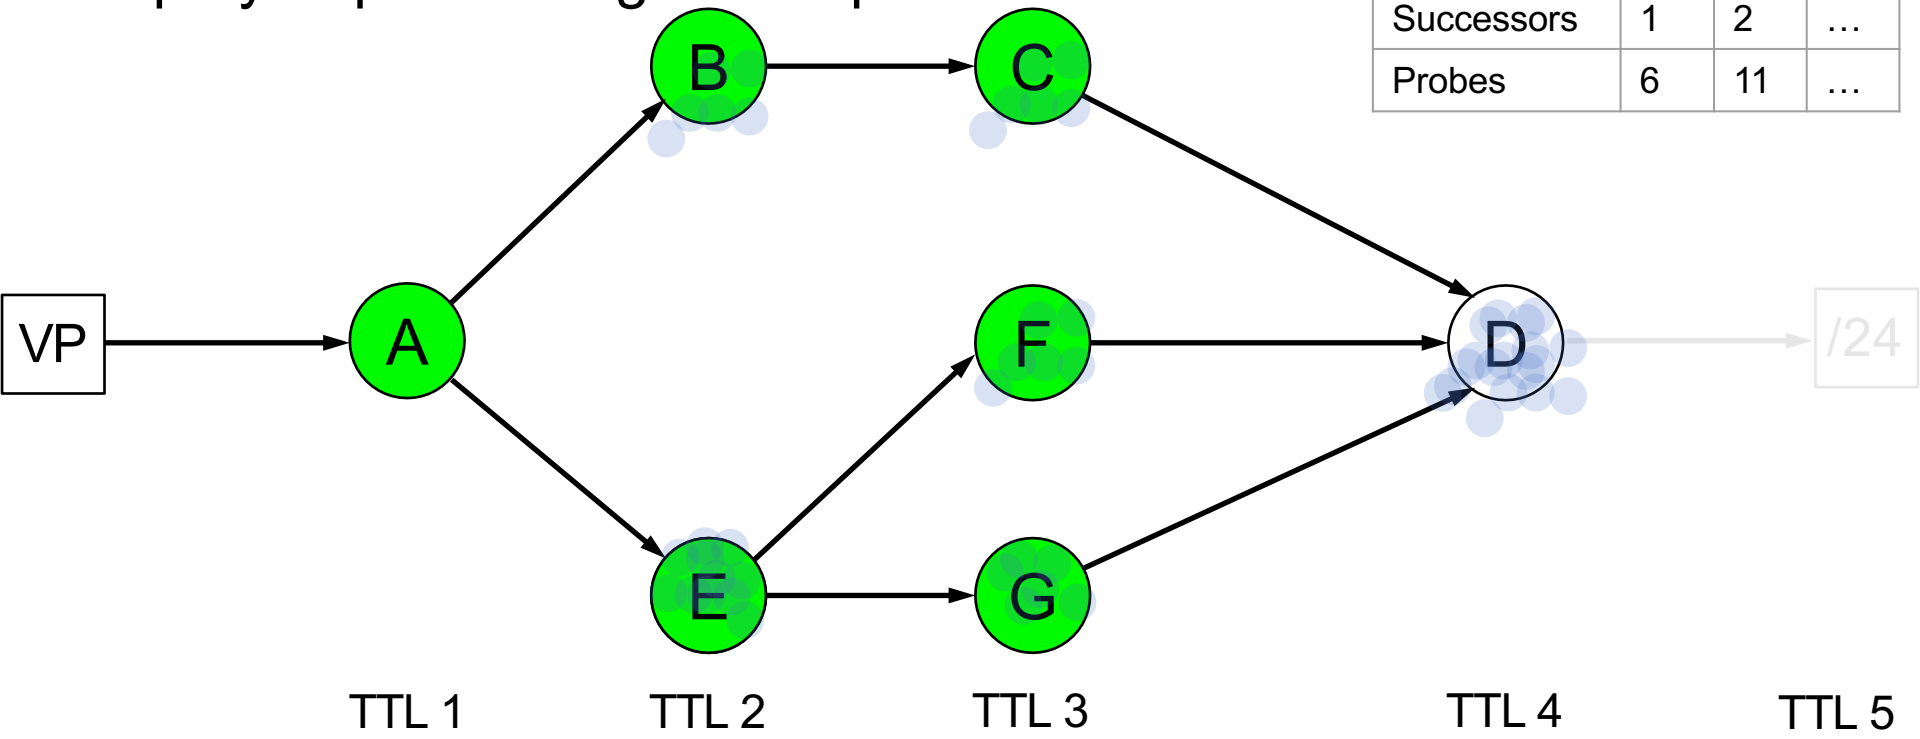

Mapping diamonds today

- Single path probing: traceroute, Paris traceroute

→ **No statistical guarantees**

- Multipath probing: MDA Paris traceroute

→ **No Internet scale**

- How to build a system that provides statistical guarantees at Internet Scale?

Roadmap

- Challenges
- **Diamond-Miner**
- Evaluation
- Conclusion

Contributions

- **Diamond-Miner**: a massively parallelized probing system to map diamonds at **Internet scale** providing **statistical guarantees**

Key Ideas

- No more resolving nodes TTL per TTL, resolves all of the nodes of the topology **concurrently**

- Round based algorithm:
  - Input: topology discovered by the previous round
  - Output: number of probes to send per TTL per destination prefix to achieve statistical guarantees

A toy example

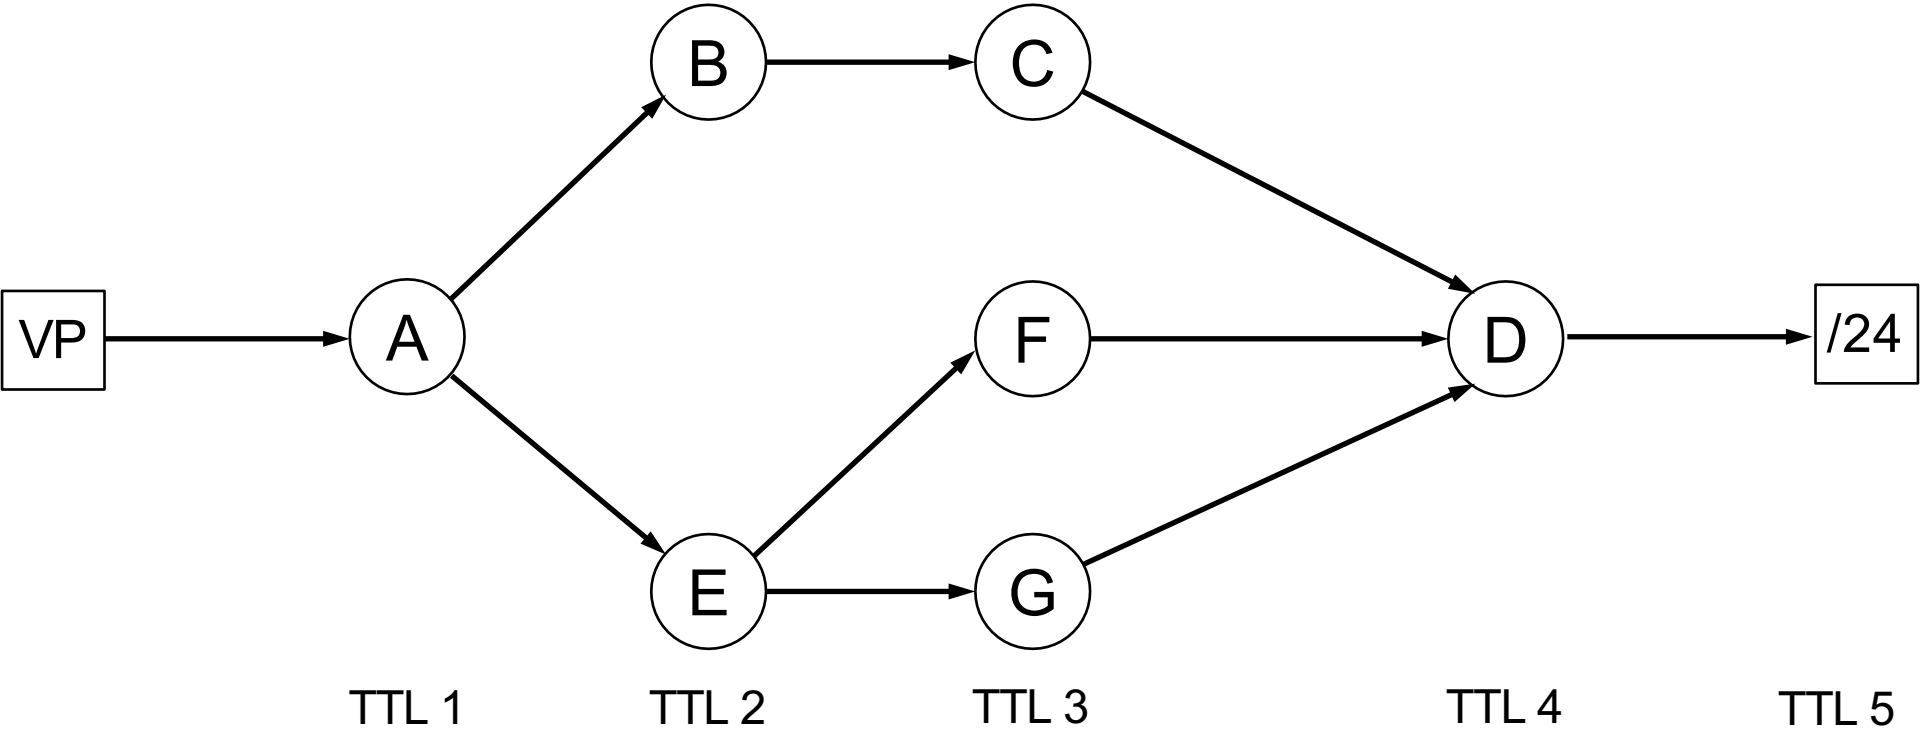

Diamond-Miner

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

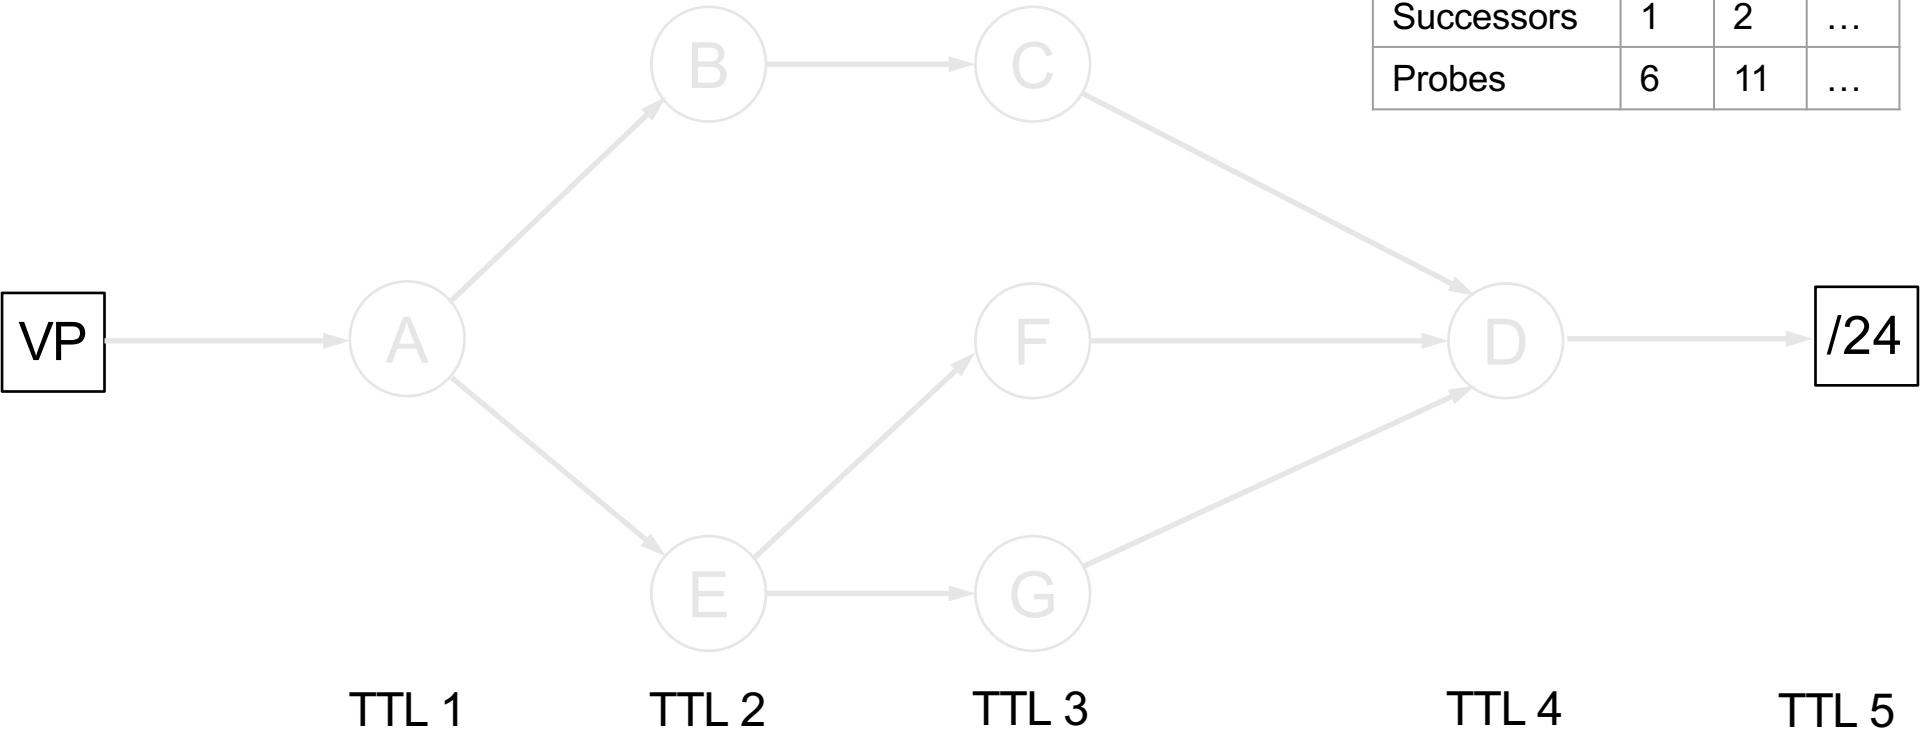

Round 1: send

Diamond-Miner

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

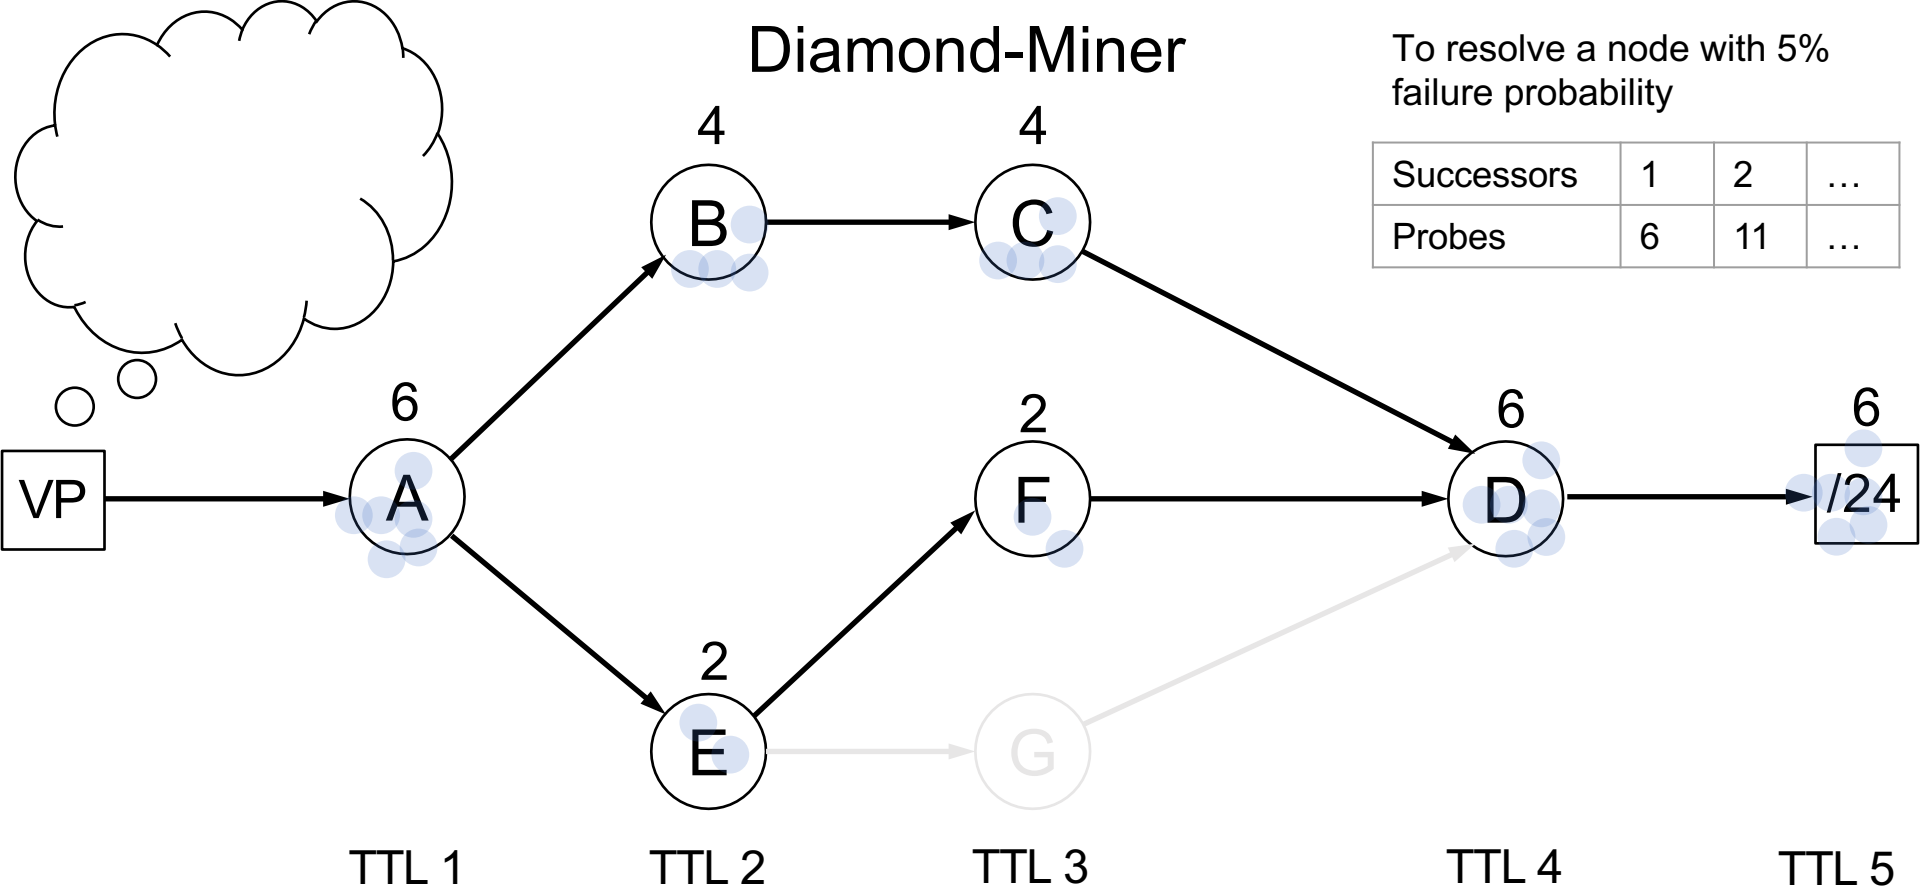

Round 2: compute

Diamond-Miner

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

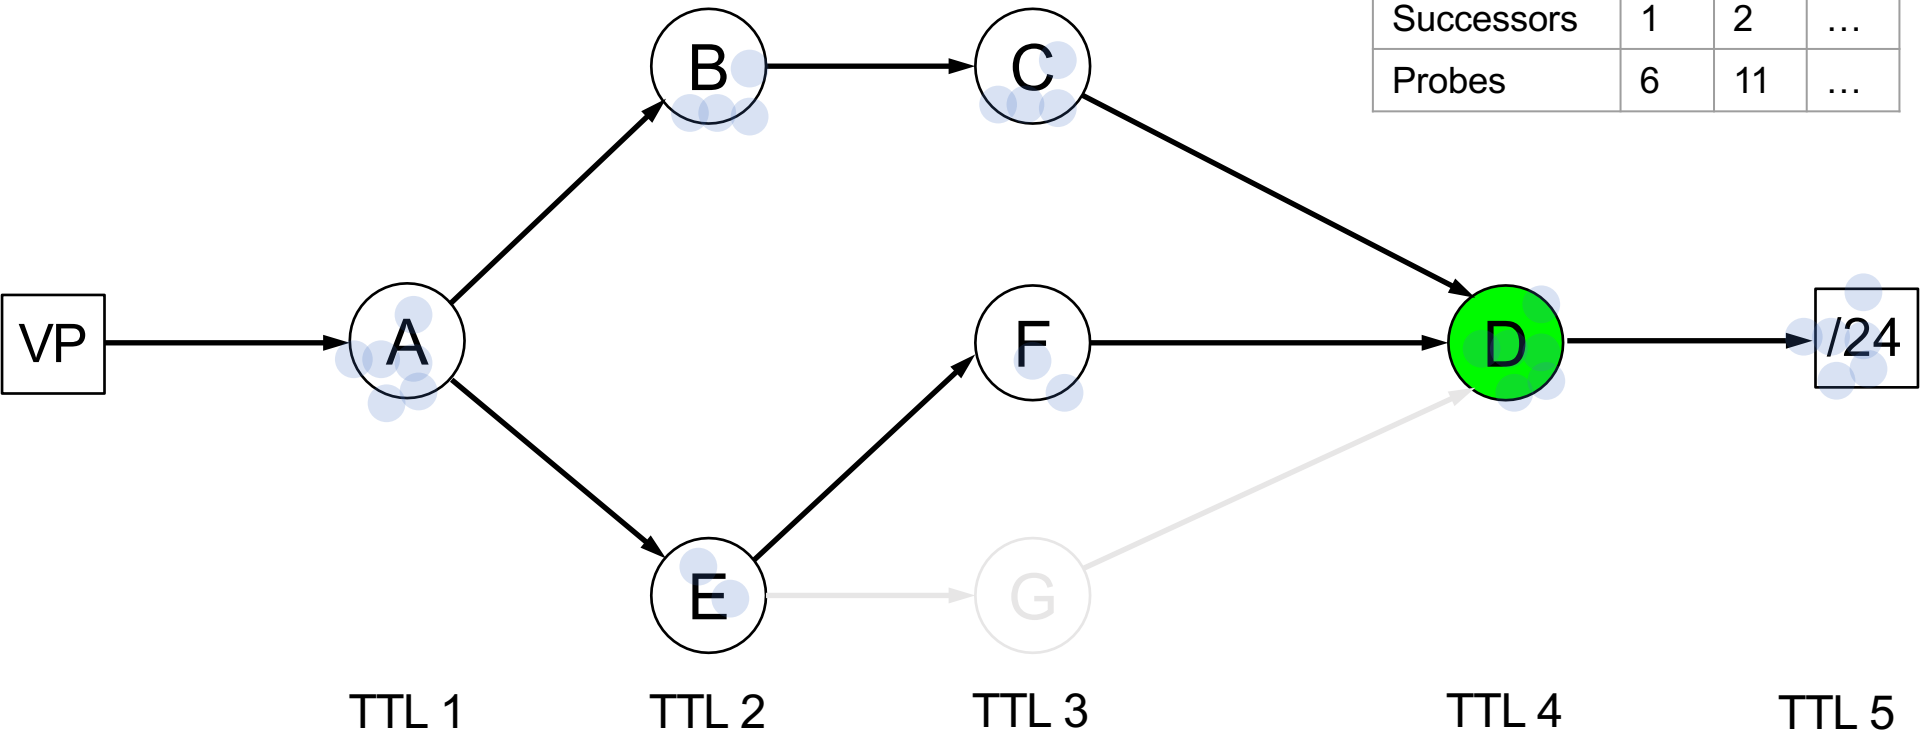

Round 2: send

Diamond-Miner

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

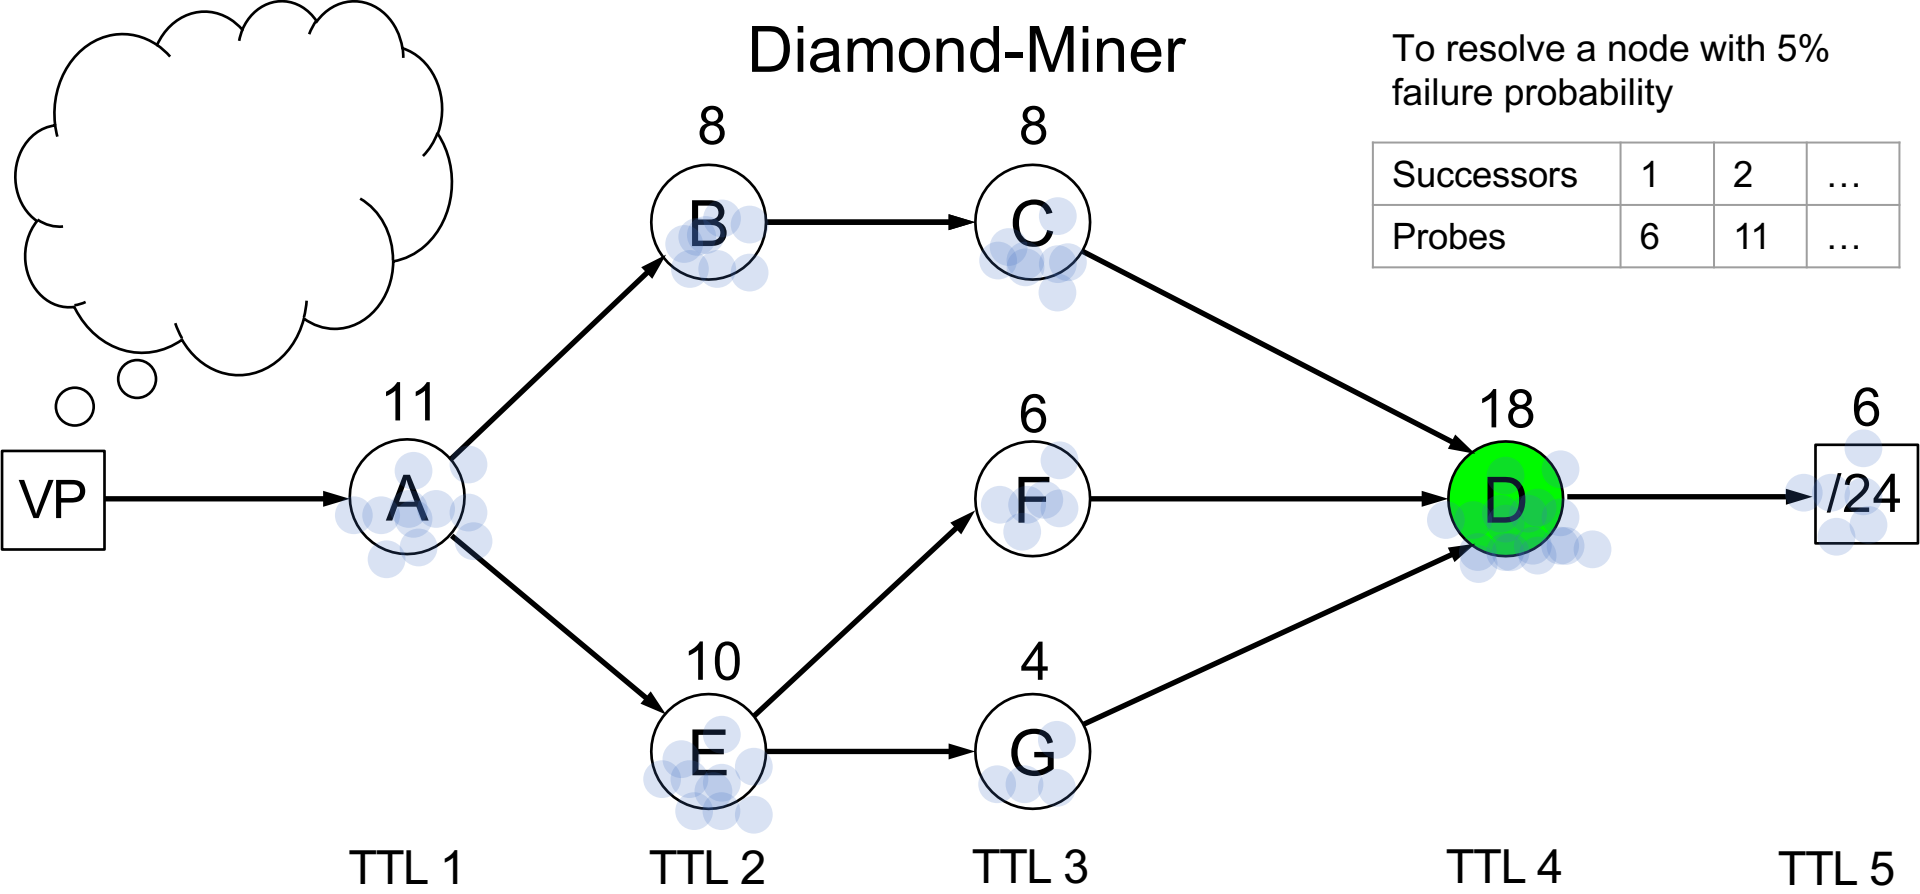

Round 3: compute

Diamond-Miner

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

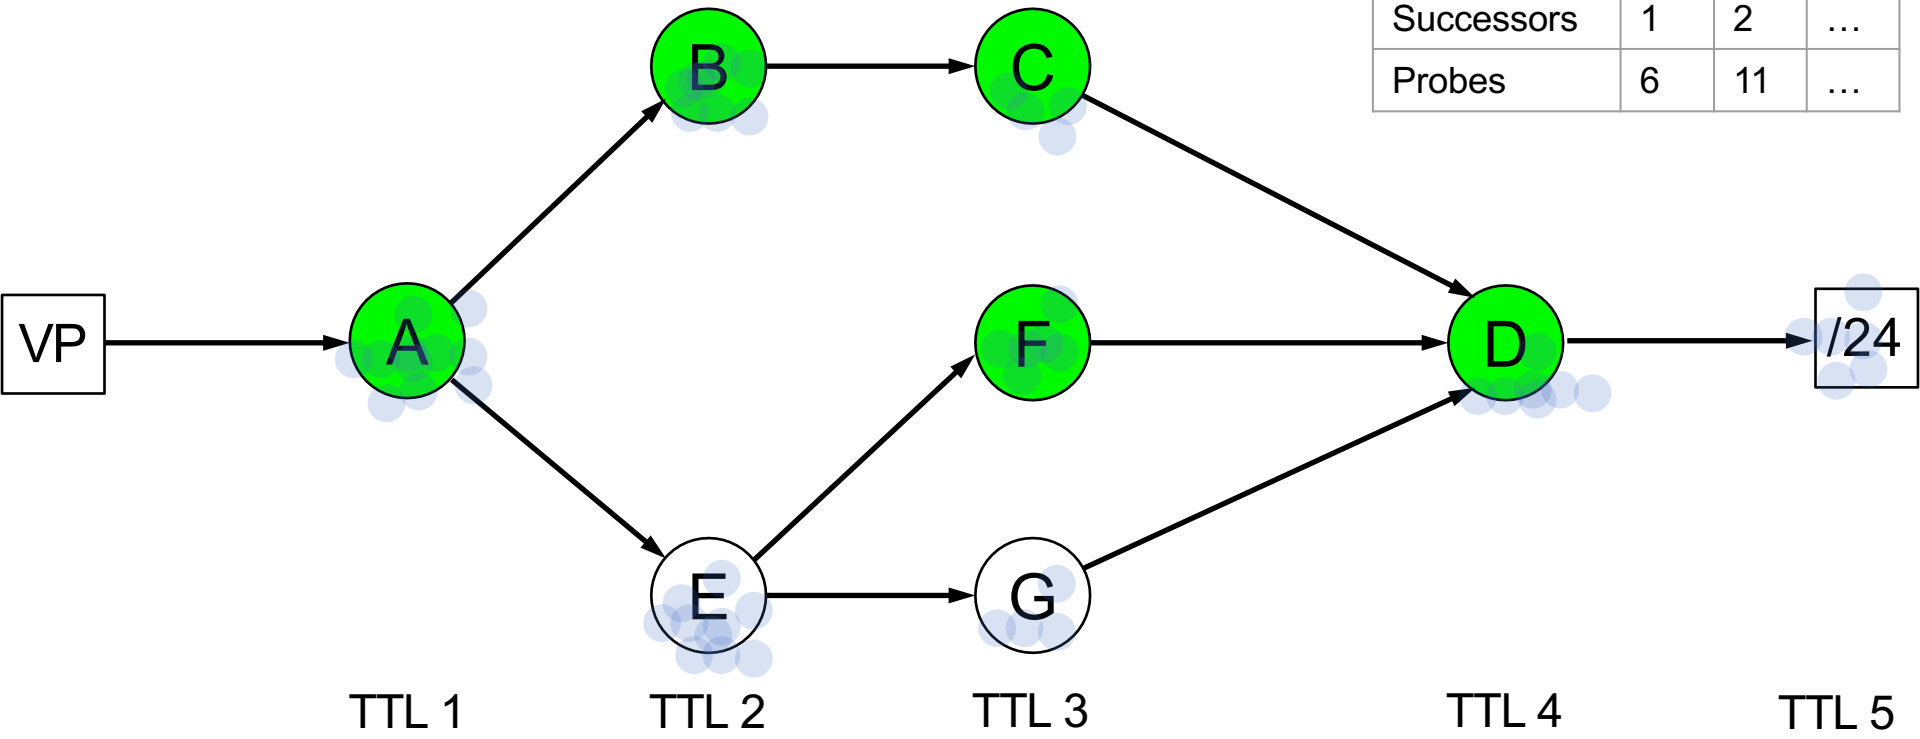

Round 3: send

Diamond-Miner

To resolve a node with 5% failure probability

Successors	1	2	…
Probes	6	11	…

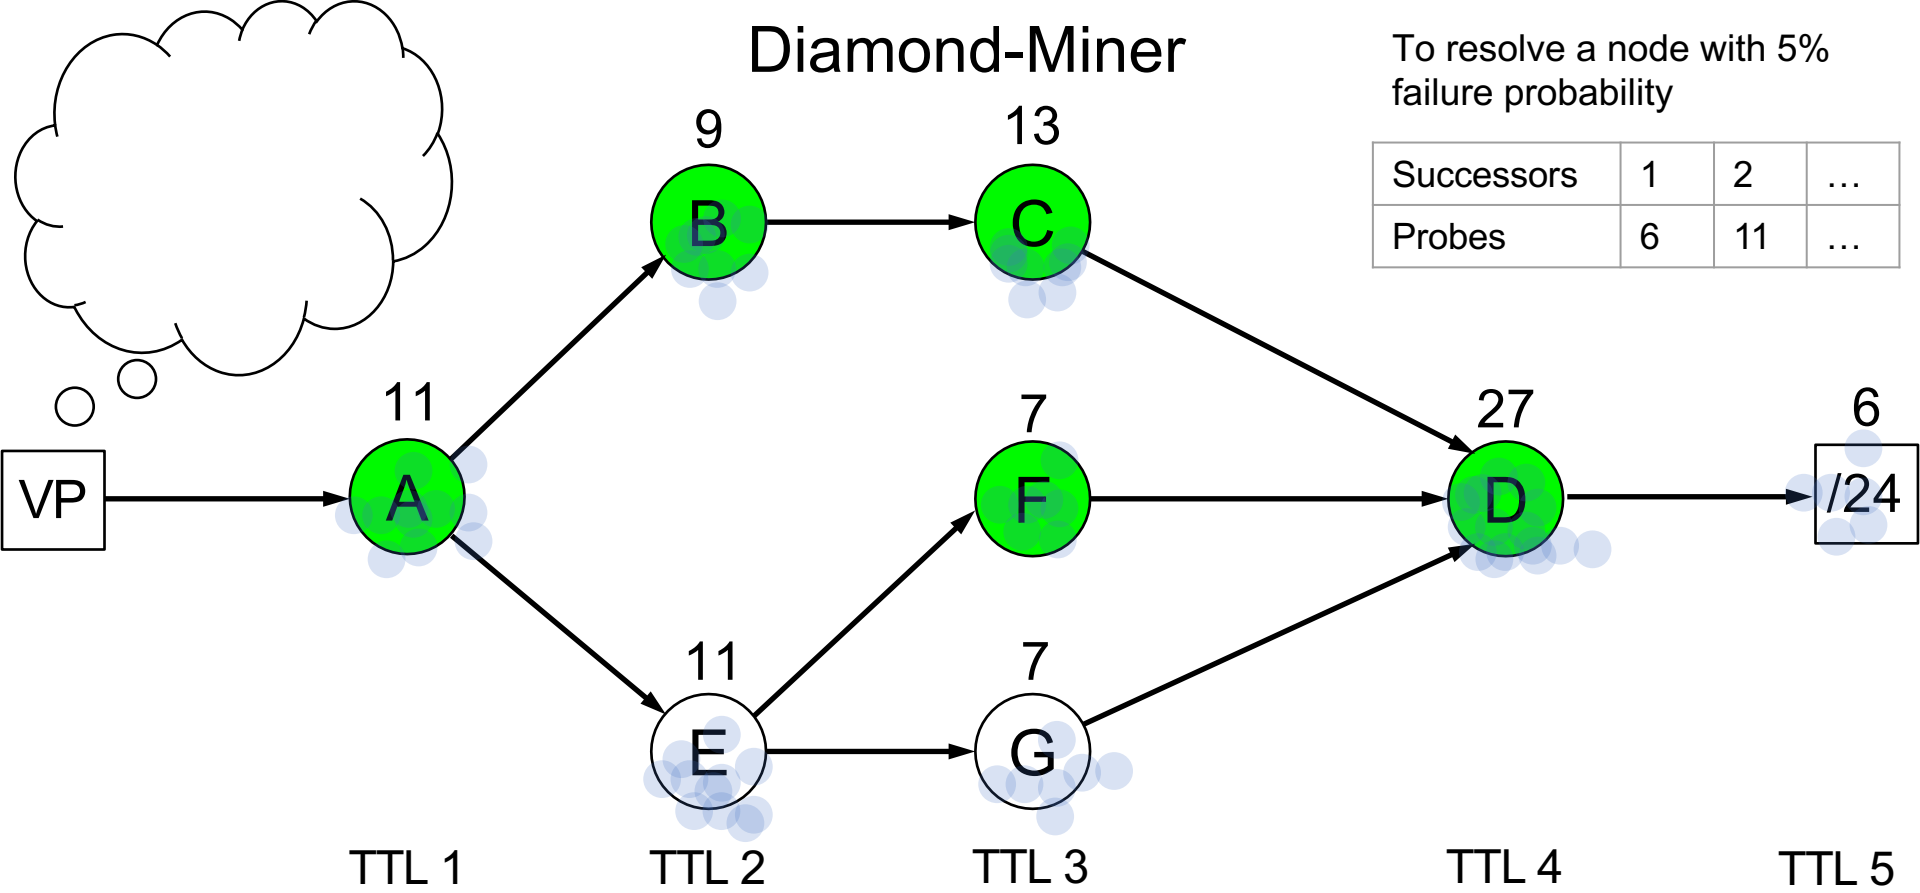

Round 4: compute

Key Ideas

- No more resolving nodes TTL per TTL, **resolves all the nodes concurrently**

- Round based algorithm:
  - Input: topology discovered by the previous round
  - Output: number of probes to send per TTL per destination prefix to achieve statistical guarantees

Scaling Diamond-Miner

- Perform the algorithm on all the /24s in parallel
- No more hop by hop probing constraints

→ **Decrease the time to completion**

Roadmap

- Challenges
- Diamond-Miner
- **Evaluation**
- Conclusion

High probing rate

1 x 100,000 pps

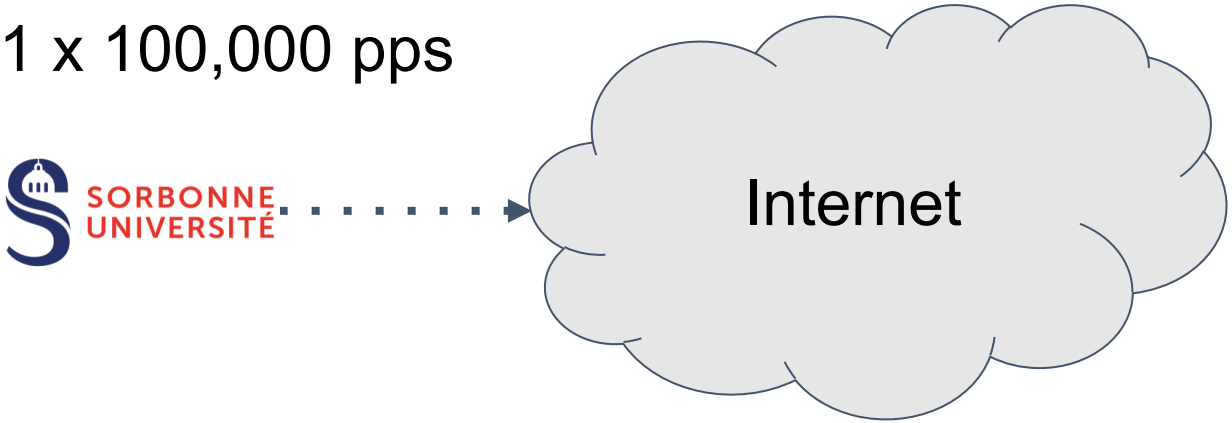

Evaluation (Number of rounds)

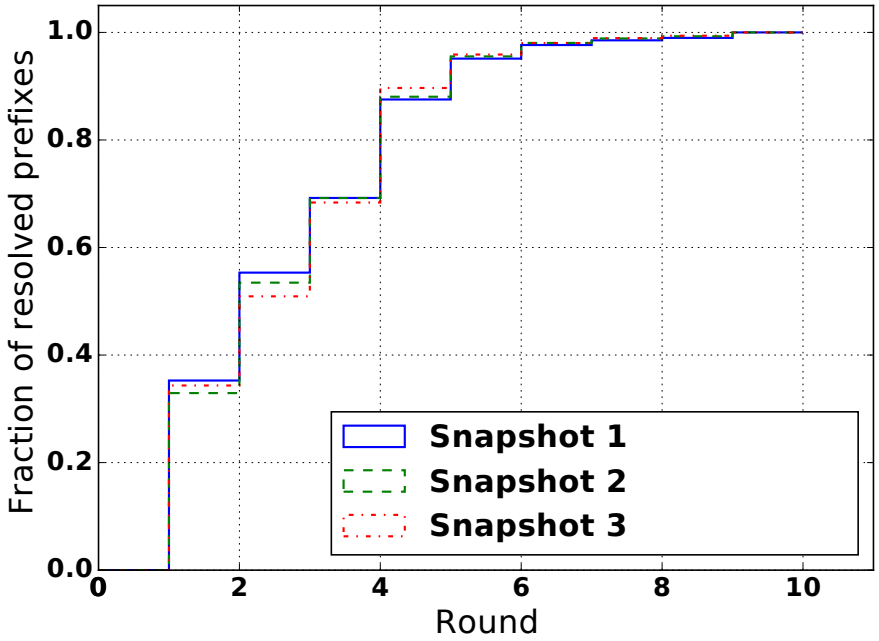

- 10 rounds → > 99% of resolved /24 prefixes

Probes sent and time to completion from a single vantage point for one snapshot

	Probes (billions)	Time to completion
Classic multipath (emulated)	5.9 B	64.3 years
Diamond-Miner	6.6 B	2.5 days

Multiple vantage points

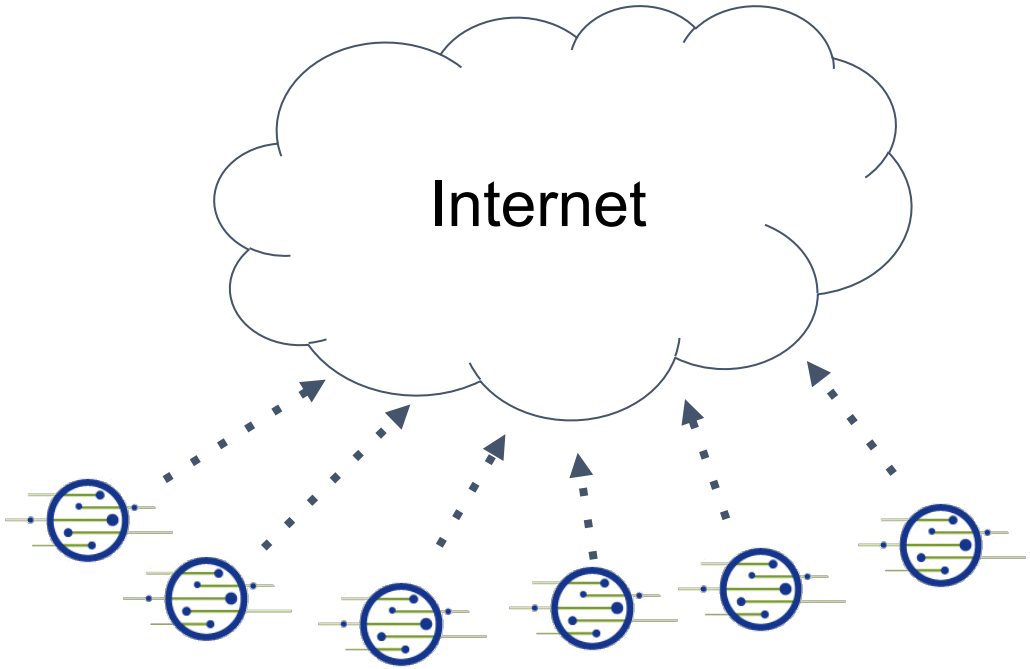

PlanetLab Europe

6 x 10,000 pps

Discovery in one week

	Vantage points	Nodes (millions)	Links (millions)	Probes (billions)
Yarrp	1	0.6	1.3	1.6
D-Miner	1	1.3	4.6	20.1
Yarrp	6	0.8	2.5	1.0
D-Miner	6	1.6	7.1	13.2

Discovery in one week

	Vantage points	Nodes (millions)	Links (millions)	Probes (billions)
D-Miner	6	1.6	7.1	13.2
Ark	~110	1.9	4.3	5.9

Roadmap

- Challenges
- Diamond-Miner
- Evaluation
- Conclusion

Takeaway

- First system capable of tracing diamonds at Internet scale with **statistical guarantees**

- Obtains the **most complete** IP-level topology view from a single server

- All our code is publicly available:
- https://github.com/dioptra-io

Evaluation (Intel Xeon Gold 5122 3.6 GHz, 8 cores)

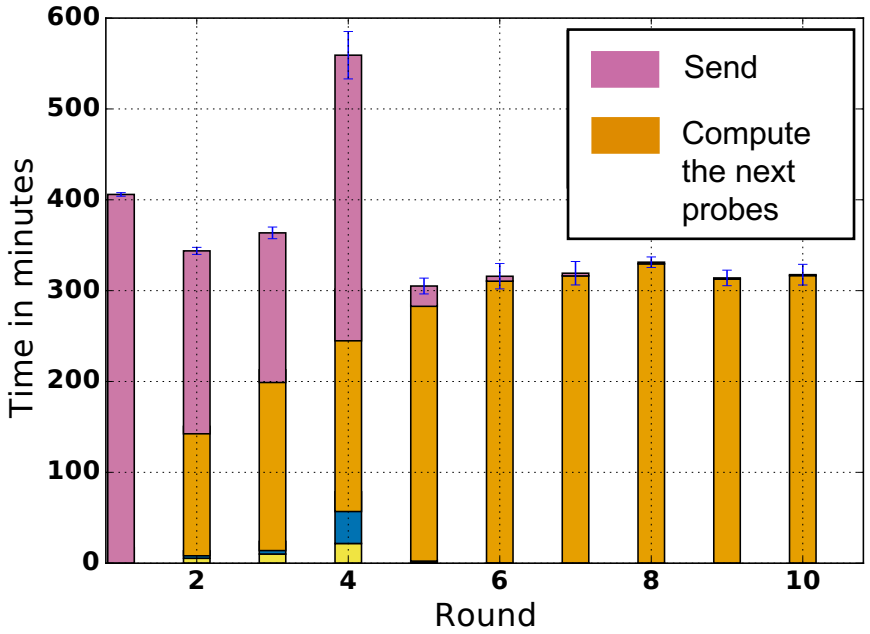

- Most of the time after round 5 is spent in the computation
- **1 snapshot = 1 day**

Motivation

- Resilience
- Security
- Socio-economic
- Basic science!

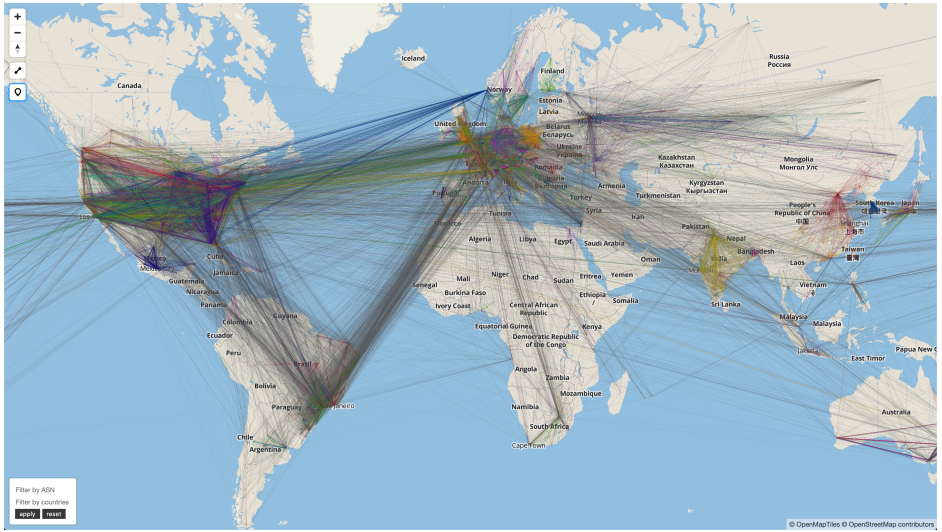

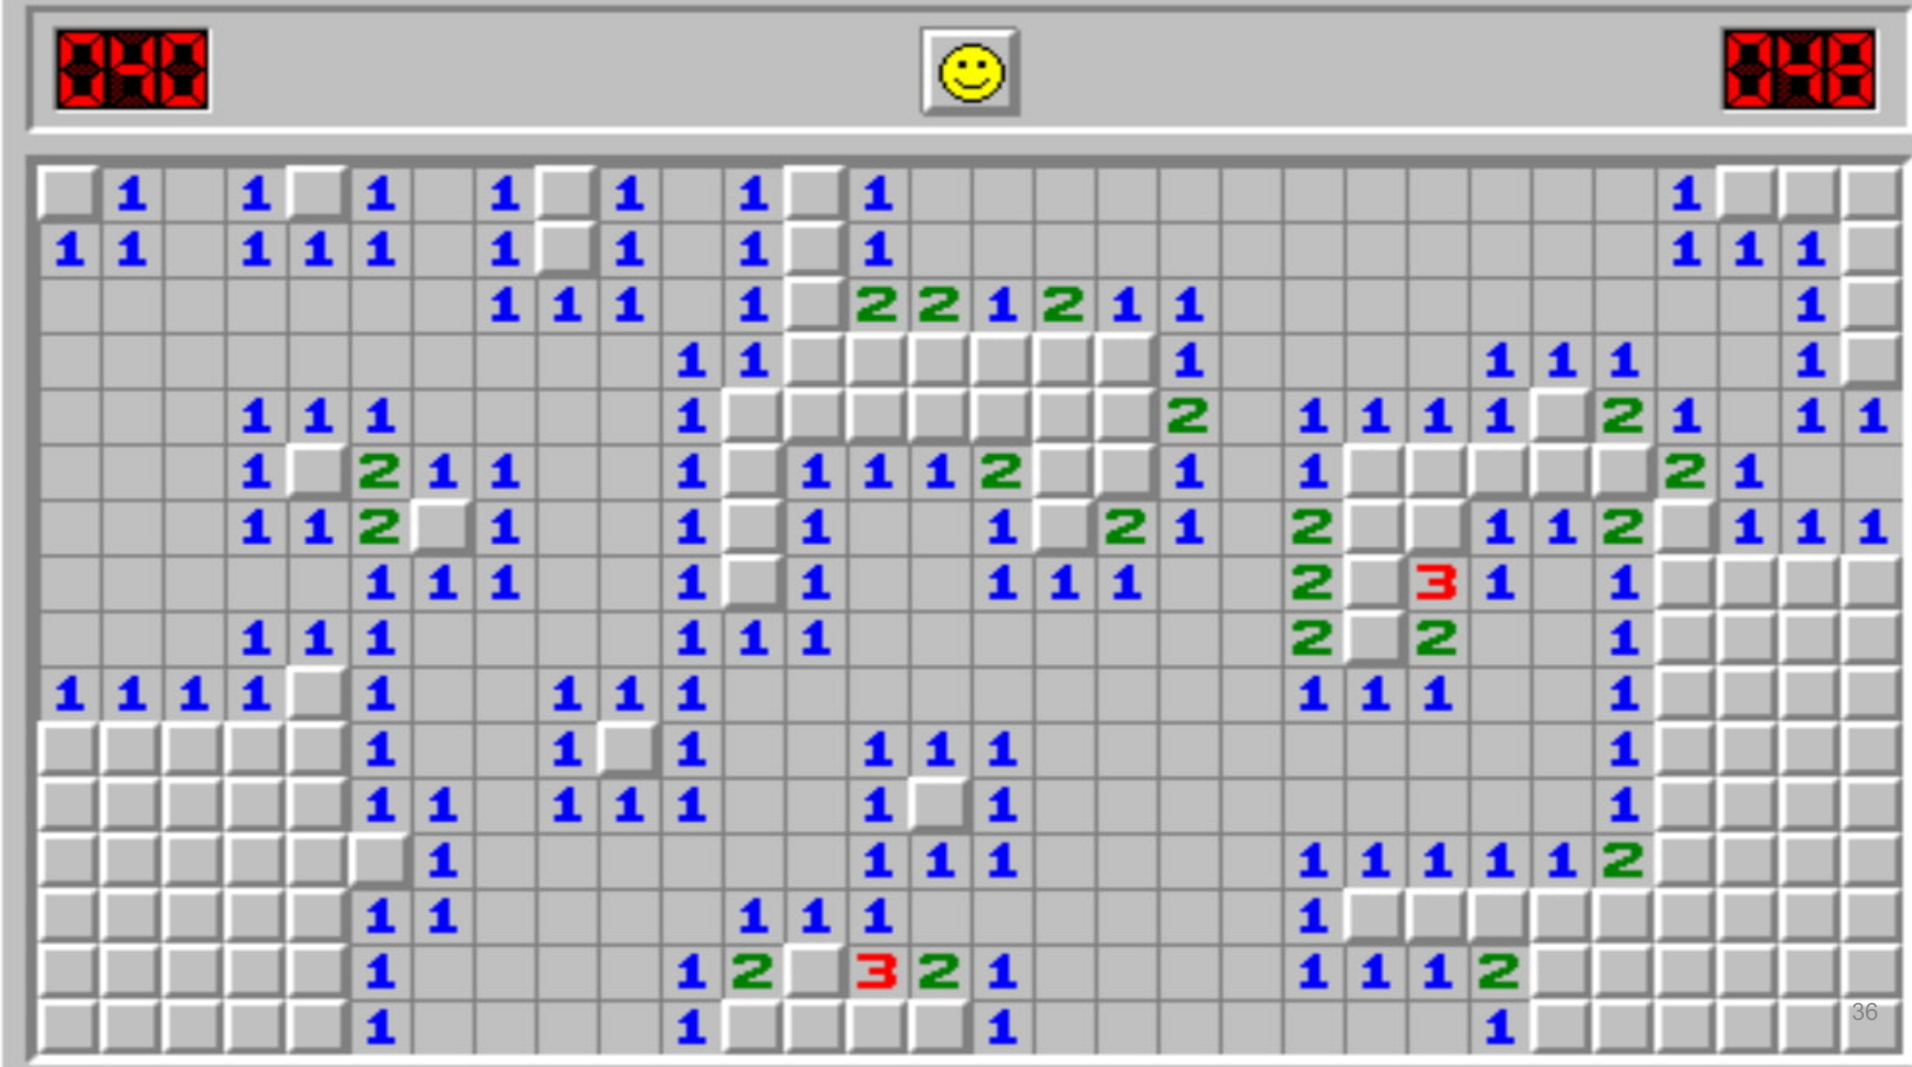

References

- Augustin et al., Avoiding traceroute anomalies with Paris Traceroute. *(IMC ‘06)*
- Veitch et al., Failure control in multipath route tracing. (*INFOCOM ’09*)
- Vermeulen et al,. Multilevel MDA-Lite Paris Traceroute (*IMC ’18)*
- Beverly. Yarrp'ing the internet: Randomized high-speed active topology discovery (*IMC ‘16*)
- Claffy et al., Internet mapping: from art to science (S*ecurity Applications & Technology Conference for Homeland Security ‘09)*

Laboratory Affiliations

- Kevin Vermeulen, Olivier Fourmaux, and Timur Friedman are associated with Sorbonne Université, CNRS, Laboratoire d'informatique de Paris 6, LIP6, F-75005 Paris, France.
- Kevin Vermeulen and Timur Friedman are associated with the Laboratory of Information, Networking and Communication Sciences, LINCS, F-75013 Paris, France.
- Justin P. Rohrer and Robert Beverly are associated with the department of Computer Science of Naval Postgraduate School.