Stroboscope: Declarative Network Monitoring on a Budget

Olivier Tilmans*, Tobias Bühler†, Ingmar Poese‡, Stefano Vissicchio§ and Laurent Vanbever†

* Université catholique de Louvain, † ETH Zürich, ‡ BENOCS, § University College London

stroboscope.ethz.ch

net-stroboscope

ISP network monitoring has unmet challenges

- Collecting statistics [1–3] on the routers themselves prevents to match observations across vantage points
- Frameworks to get packet-level visibility [4–8] are incompatible with ISP constraints:
  1. No control over end hosts
  2. Limited data-plane flexibility
  3. Constrained monitoring bandwidth

Stroboscope achieves deterministic sampling

- Stroboscope iteratively collects and analyzes traffic slices
  - Most vendors support mirroring and encapsulation
  - Slices can be as small as **23** ms on our routers (Cisco C7018)
- Operators specify *what* they want to observe
- Stroboscope decides *where* and *when* to mirror with provably correct algorithms guaranteeing measurements' accuracy
- Measurement analysis is decoupled from data collection

Stroboscope translates monitoring requirements to a schedule of mirroring rule activations

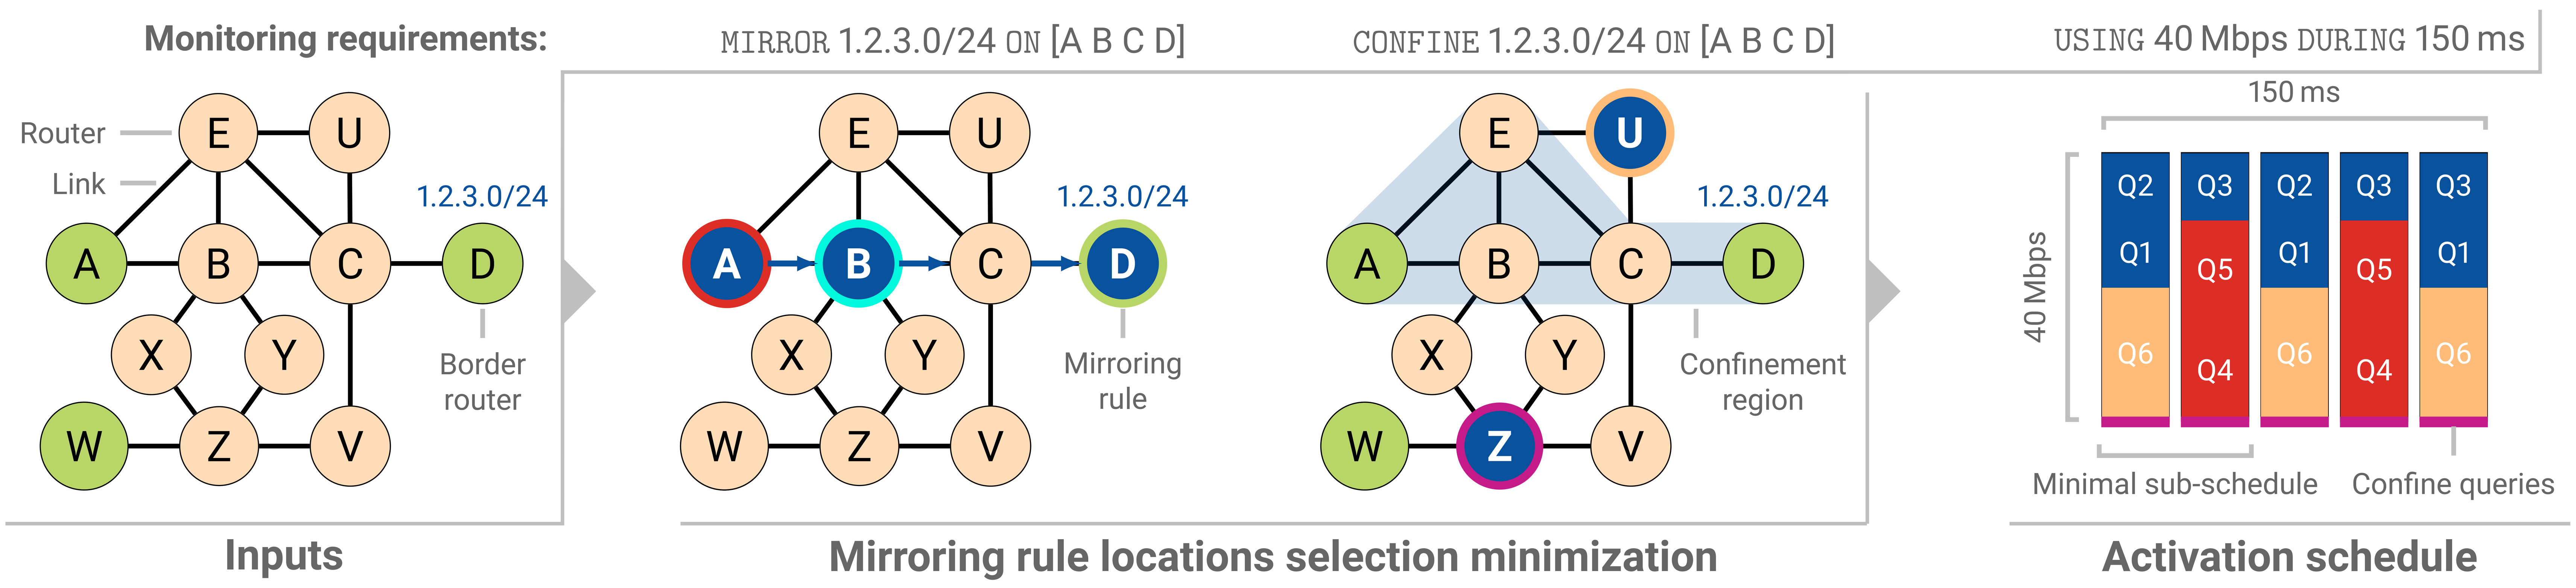

MIRROR queries reconstruct the path followed by packets
- The *Key-Points Sampling* algorithm minimizes their mirroring locations and guarantees non-ambiguous reconstructed paths
- The *KPS* algorithm takes ∼10ms on realistic inputs

CONFINE queries mirror any packet leaving a region
- The *Surrounding* algorithm minimizes their mirroring rules by computing a—possibly minimal—multi-terminal node cut
- The *Surrounding* algorithm runs in ms on realistic inputs

Scheduling queries is done in three steps:
1. Stroboscope conservatively estimates traffic demands
2. To scale to large inputs, a *minimal sub-schedule* is computed (variant of the bin-packing problem)
3. The budget usage is maximized by first replicating the sub-schedule then packing as many queries as possible

Computing measurement campaigns can be efficiently approximated in seconds on very large inputs (1000+ queries)

Traffic slices enable fine-grained measurements

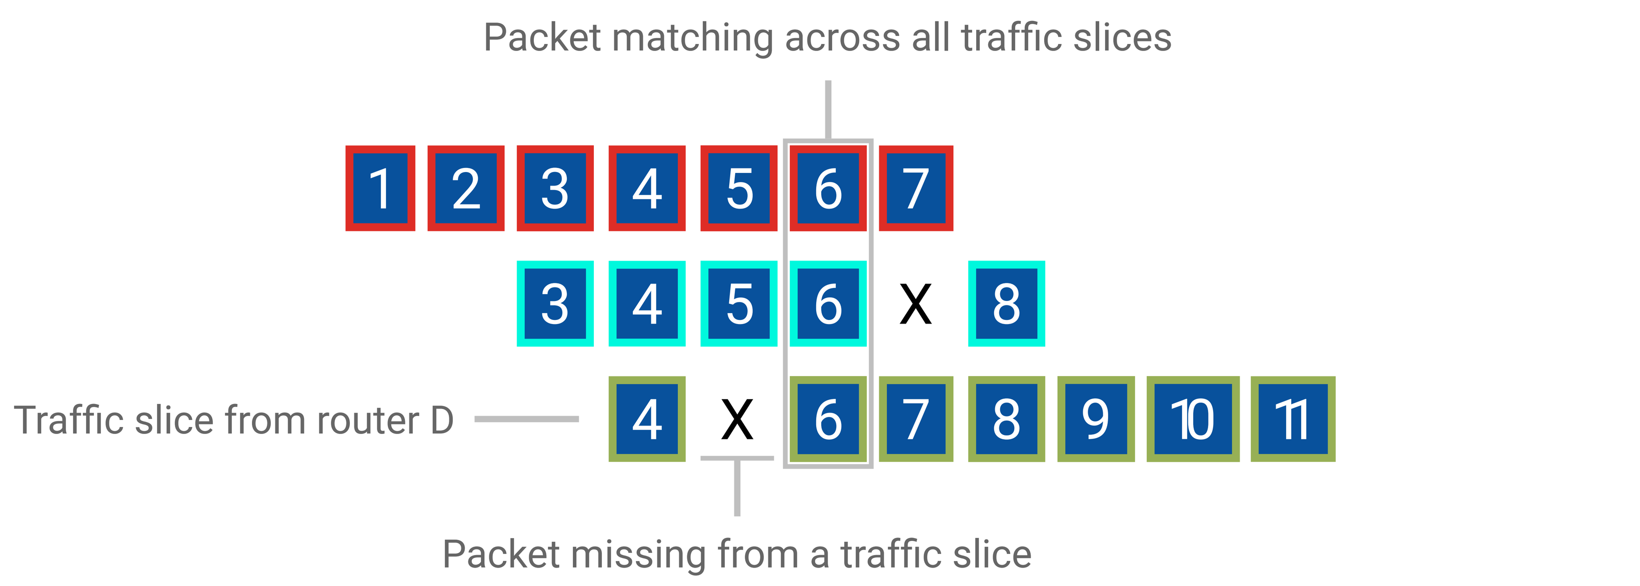

- The `CONFINE` query bounds the observation domain
- Packet Matches synchronize the traffic slices, removing the requirement to synchronize clocks across routers
- Packets missing in all last hops of a path signal losses
- Packets missing only within a path signal load-balancing

Stroboscope provides strong budget guarantees

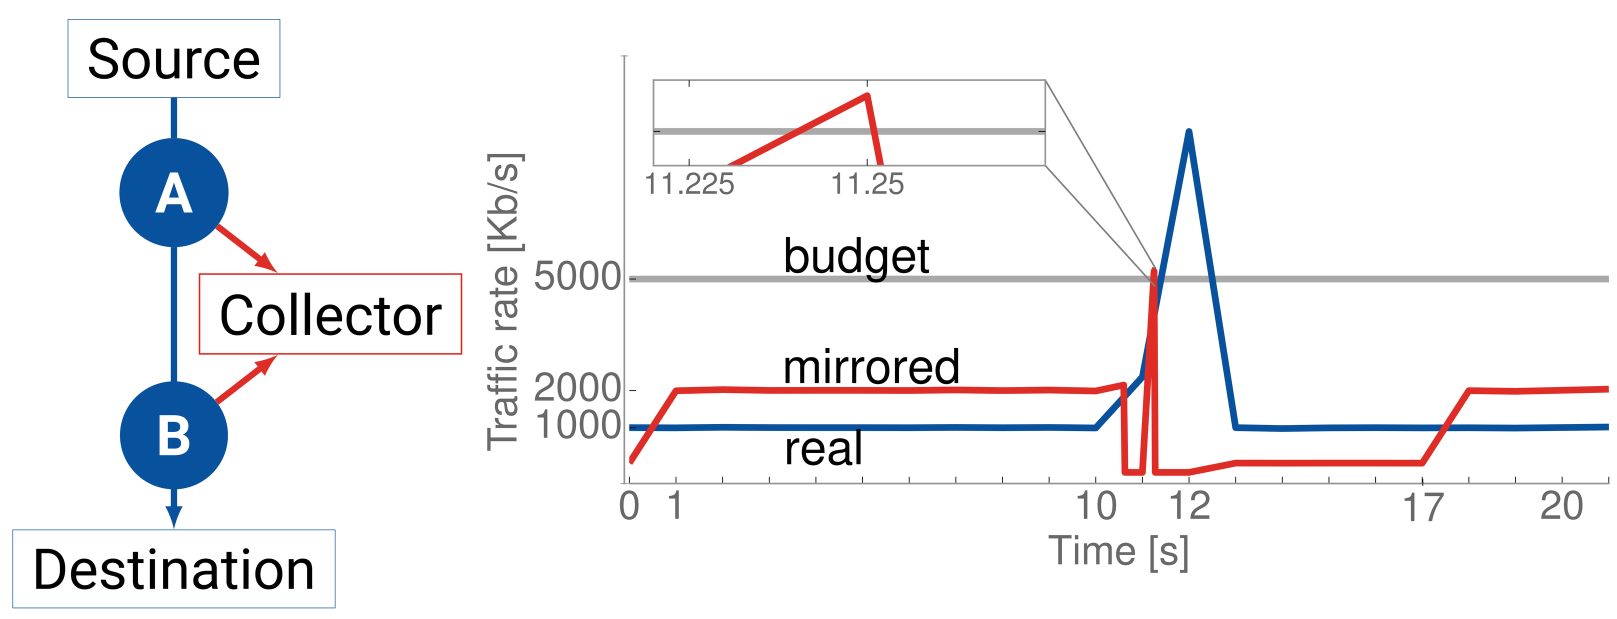

- Stroboscope tracks budget usage for every query
- Queries are not activated if their budget is exhausted
- Routers autonomously deactivate mirroring rules using timers
- `CONFINE` queries are heavily rate-limited
- Stroboscope exceeds the budget for at most one timeslot

[1] Benoit Claise. *Cisco Systems NetFlow Services Export Version 9*. RFC 3954 (Informational). `http://www.ietf.org/rfc/rfc3954.txt`. Internet Engineering Task Force, Oct. 2004.
[2] Chuck Cranor et al. "Gigascope: a stream database for network applications". In: *SIGMOD*. 2003, pp. 647–651.
[3] Qun Huang et al. "SketchVisor: Robust Network Measurement for Software Packet Processing". In: *SIGCOMM*. 2017, pp. 113–126.
[4] Masoud Moshref et al. "Trumpet: Timely and Precise Triggers in Data Centers". In: *SIGCOMM*. 2016, pp. 129–143.
[5] Srinivas Narayana et al. "Compiling Path Queries". In: *NSDI*. 2016, pp. 207–222.
[6] Srinivas Narayana et al. "Language-Directed Hardware Design for Network Performance Monitoring". In: *SIGCOMM*. 2017, pp. 85–98.
[7] Jeff Rasley et al. "Planck: Millisecond-scale Monitoring and Control for Commodity Networks". In: *SIGCOMM*. 2014, pp. 407–418.
[8] Yibo Zhu et al. "Packet-Level Telemetry in Large Datacenter Networks". In: *SIGCOMM*. 2015, pp. 479–491.

*O. Tilmans is supported by a grant from F.R.S.-FNRS FRIA.
†This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 688421, and was supported by the Swiss State Secretariat for Education, Research and Innovation (SERI) under contract number 15.0268.
The opinions expressed and arguments employed reflect only the authors' views. The European Commission is not responsible for any use that may be made of that information. Further, the opinions expressed and arguments employed herein do not necessarily reflect the official views of the Swiss Government.