Mille-Feuille: Putting ISP traffic under the scalpel

Olivier Tilmans
UCLouvain

HotNets-XV
Nov. 9, 2016

Joint work with
T. Bühler (ETH Zürich), S. Vissicchio (UCL) and L. Vanbever (ETH Zürich)

“What happens to the Skype traffic in my network?”

ISP operators only have access to poor and coarse-grained visibility over their network.

- Netflow, sFLOW, provide aggregated statistics over random packet sampling.
- Active probing scales poorly.
- Router Configuration/syslog analysis only covers a fraction of the control-plane.

ISP operators only have access to poor and coarse-grained visibility over their network.

- Netflow, sFLOW, provide aggregated statistics over random packet sampling.
- Active probing scales poorly.
- Router Configuration/syslog analysis only covers a fraction of the control-plane.

These techniques cannot provide real time information about the network state.

Research to provide complete traffic visibility in DC networks, leverages degrees of freedom unavailable in ISP networks.

ISP networks present unique challenges:

- No control on the end hosts.
- Geographically distributed.
- Wide-range of heterogeneous network equipments.

We aim to provide ISP operators a
fine-grained visibility over their networks.

Consider the following part of an ISP network.

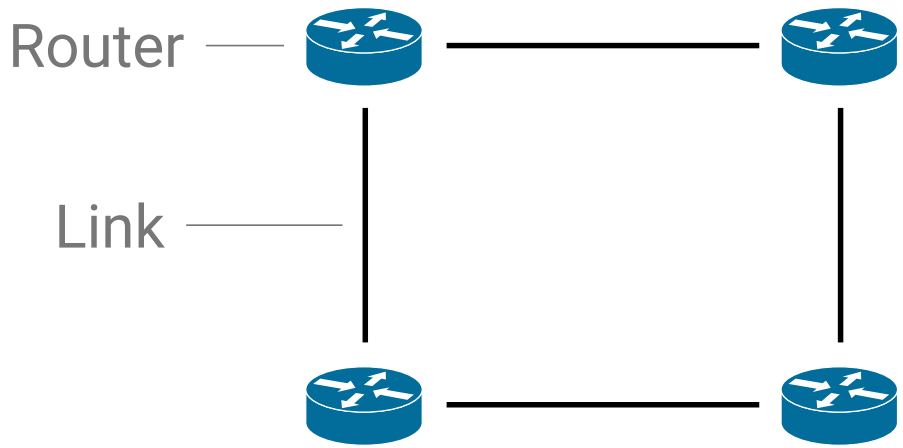

Consider the following part of an ISP network.

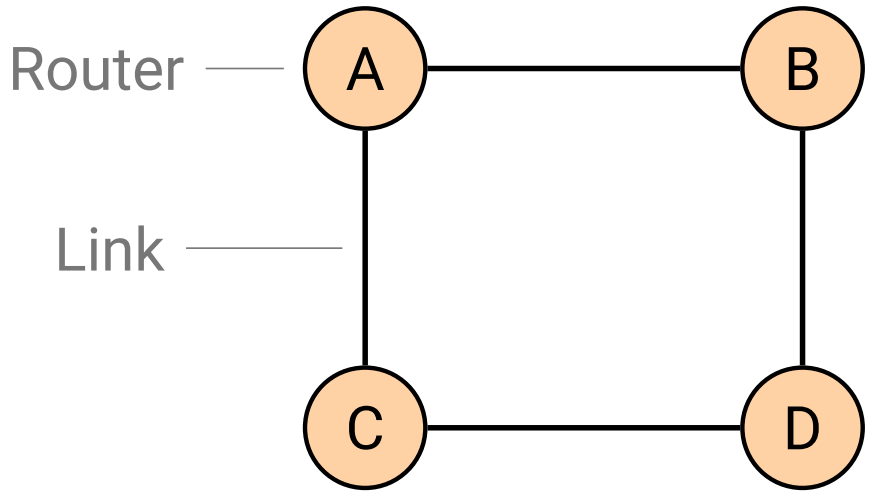

Consider the following part of an ISP network.

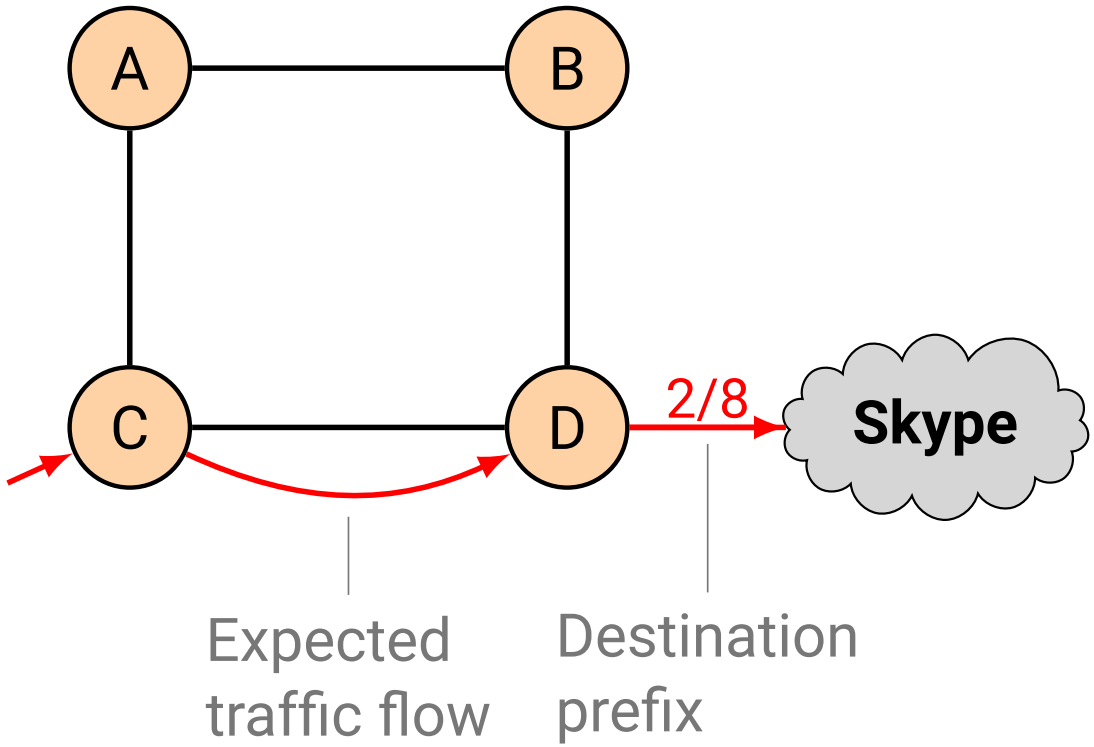

Mille-Feuille improves ISP monitoring with a traffic slicing primitive.

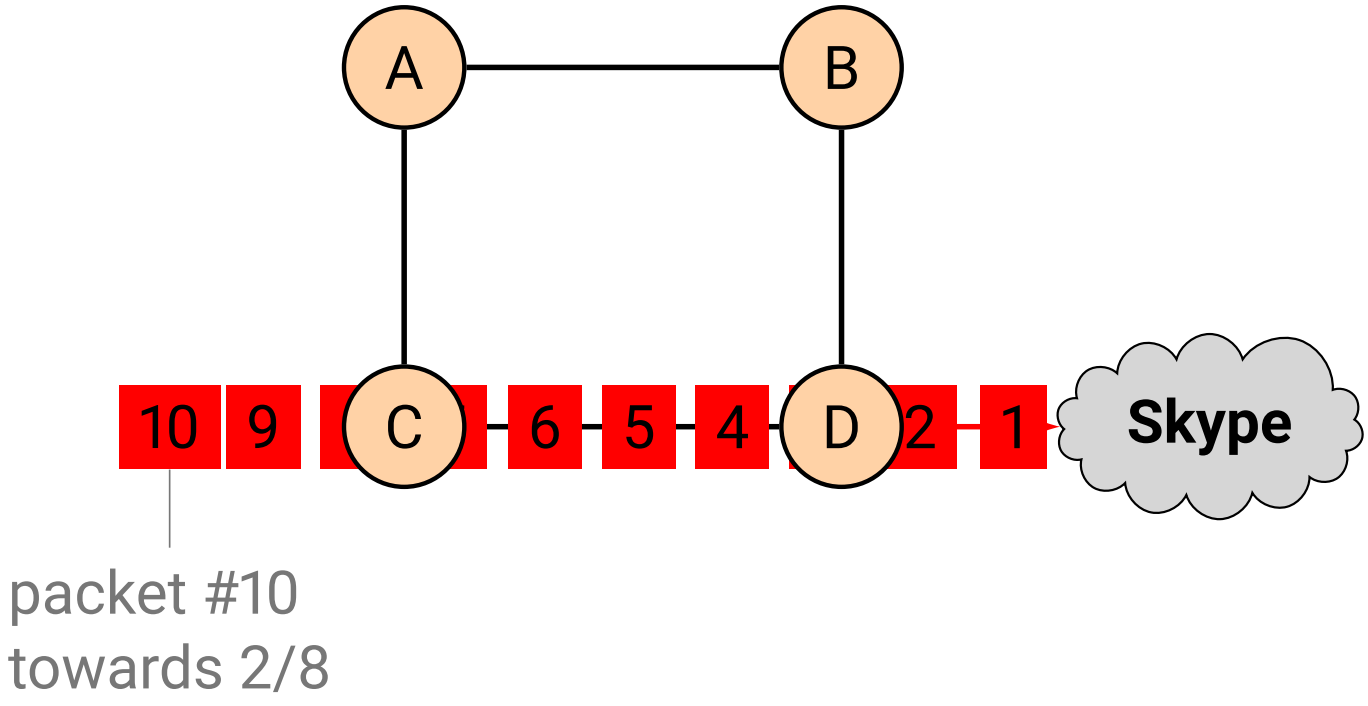

Mille-Feuille improves ISP monitoring with a traffic slicing primitive.

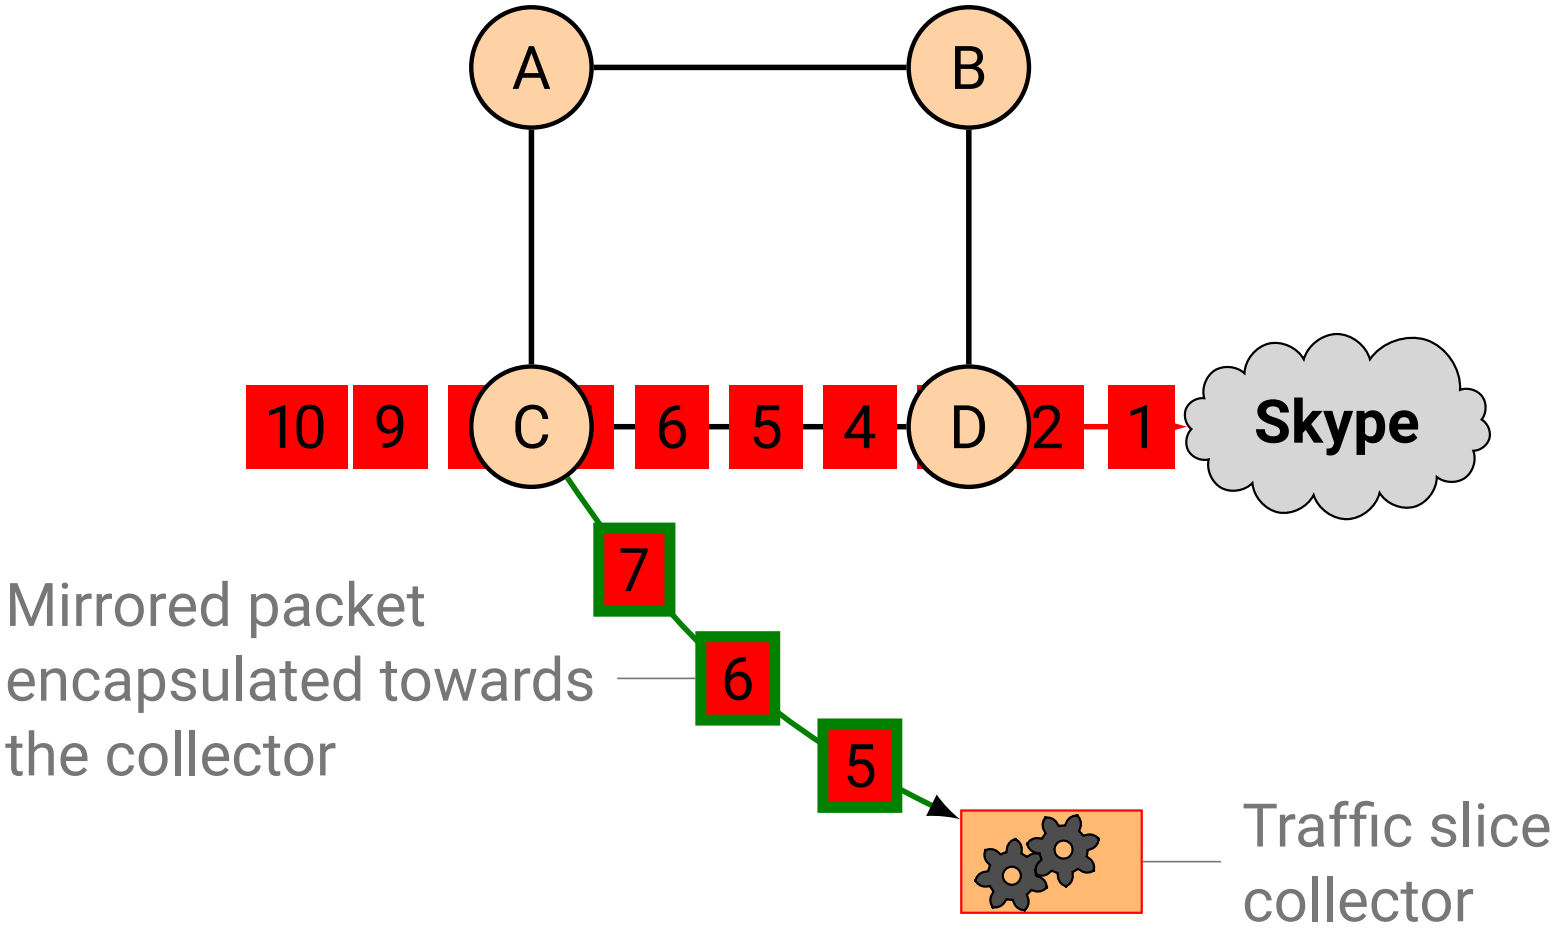

Mille-Feuille improves ISP monitoring with a traffic slicing primitive.

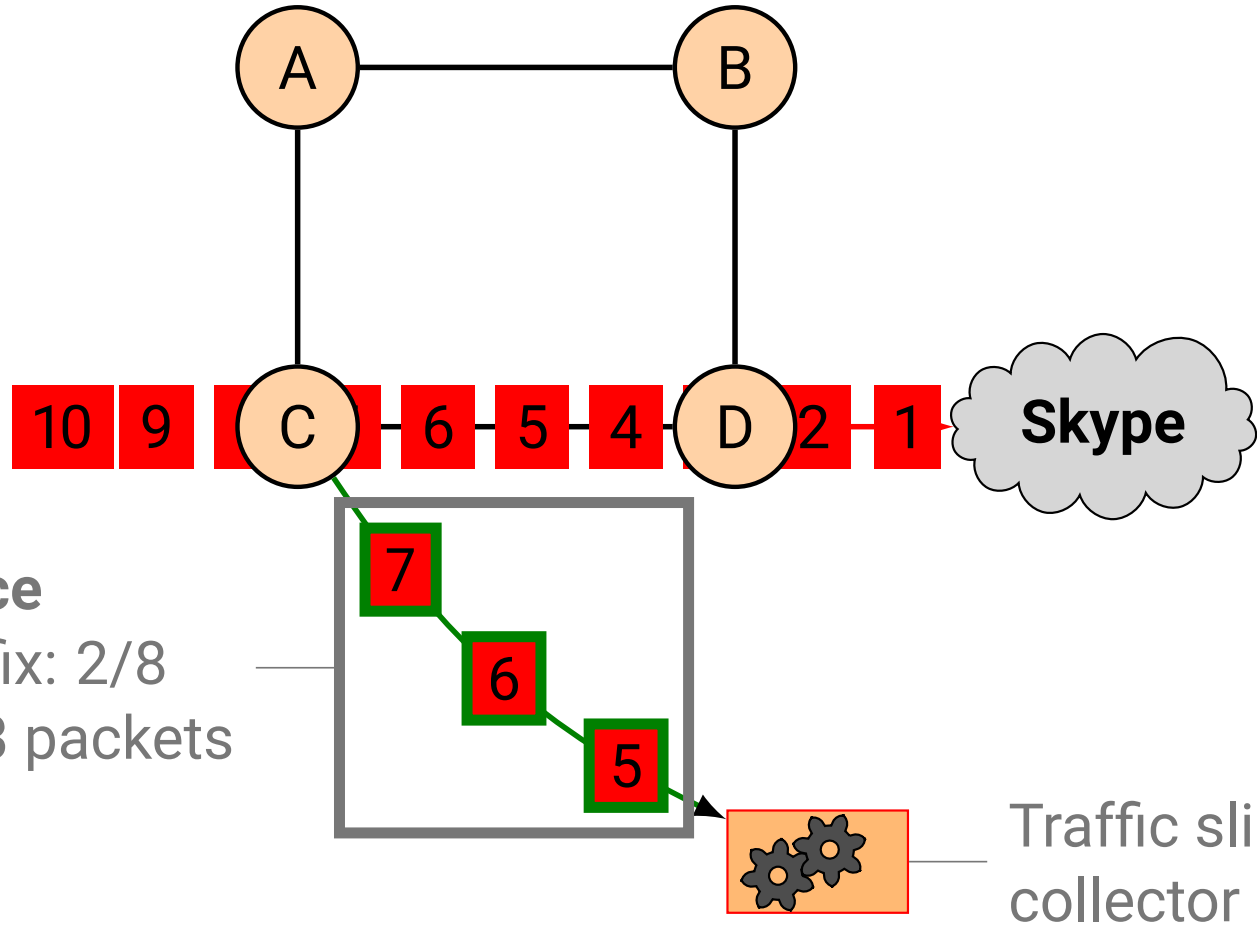

By concurrently capturing slices at different routers for the same prefix, Mille-Feuille can infer measurements about the traffic.

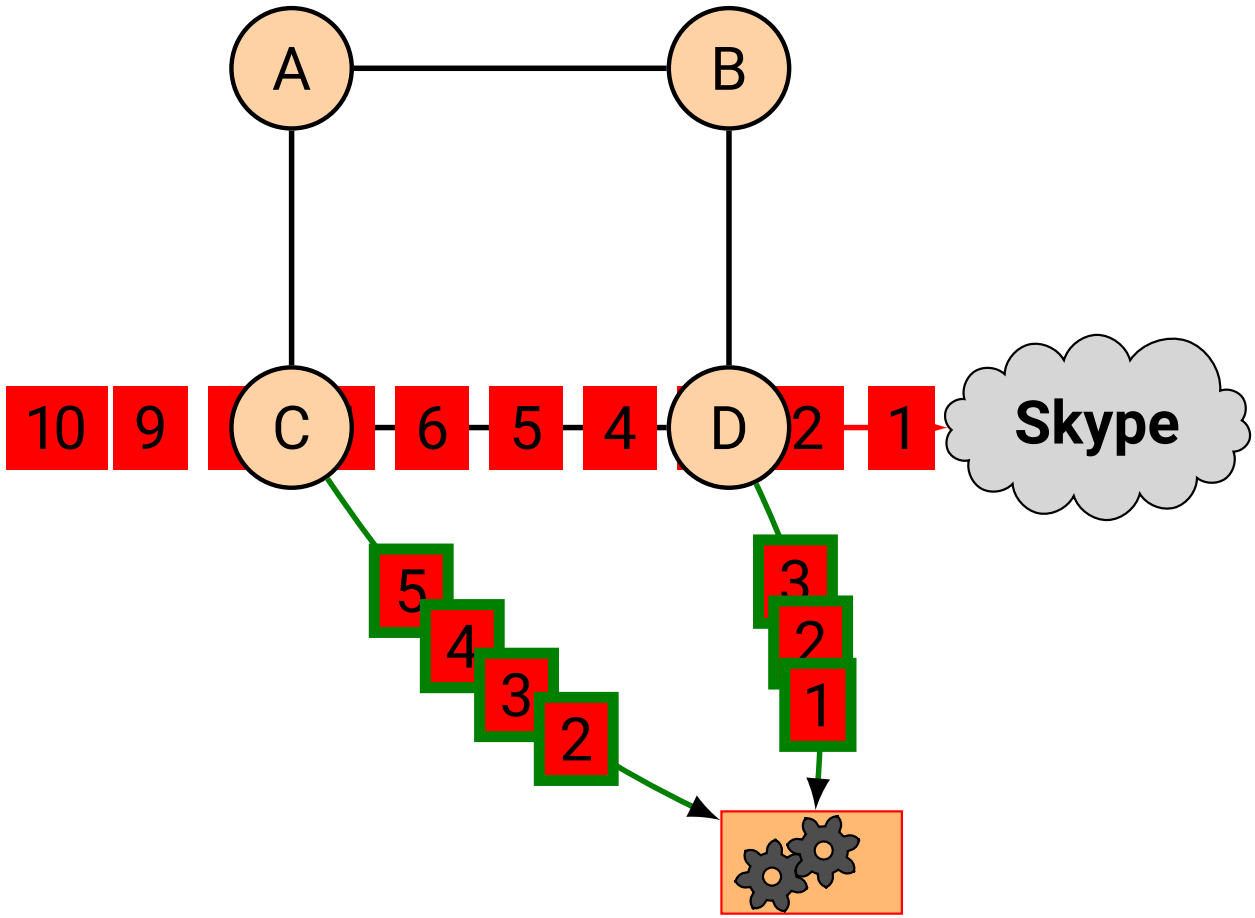

By concurrently capturing slices at different routers for the same prefix, Mille-Feuille can infer measurements about the traffic.

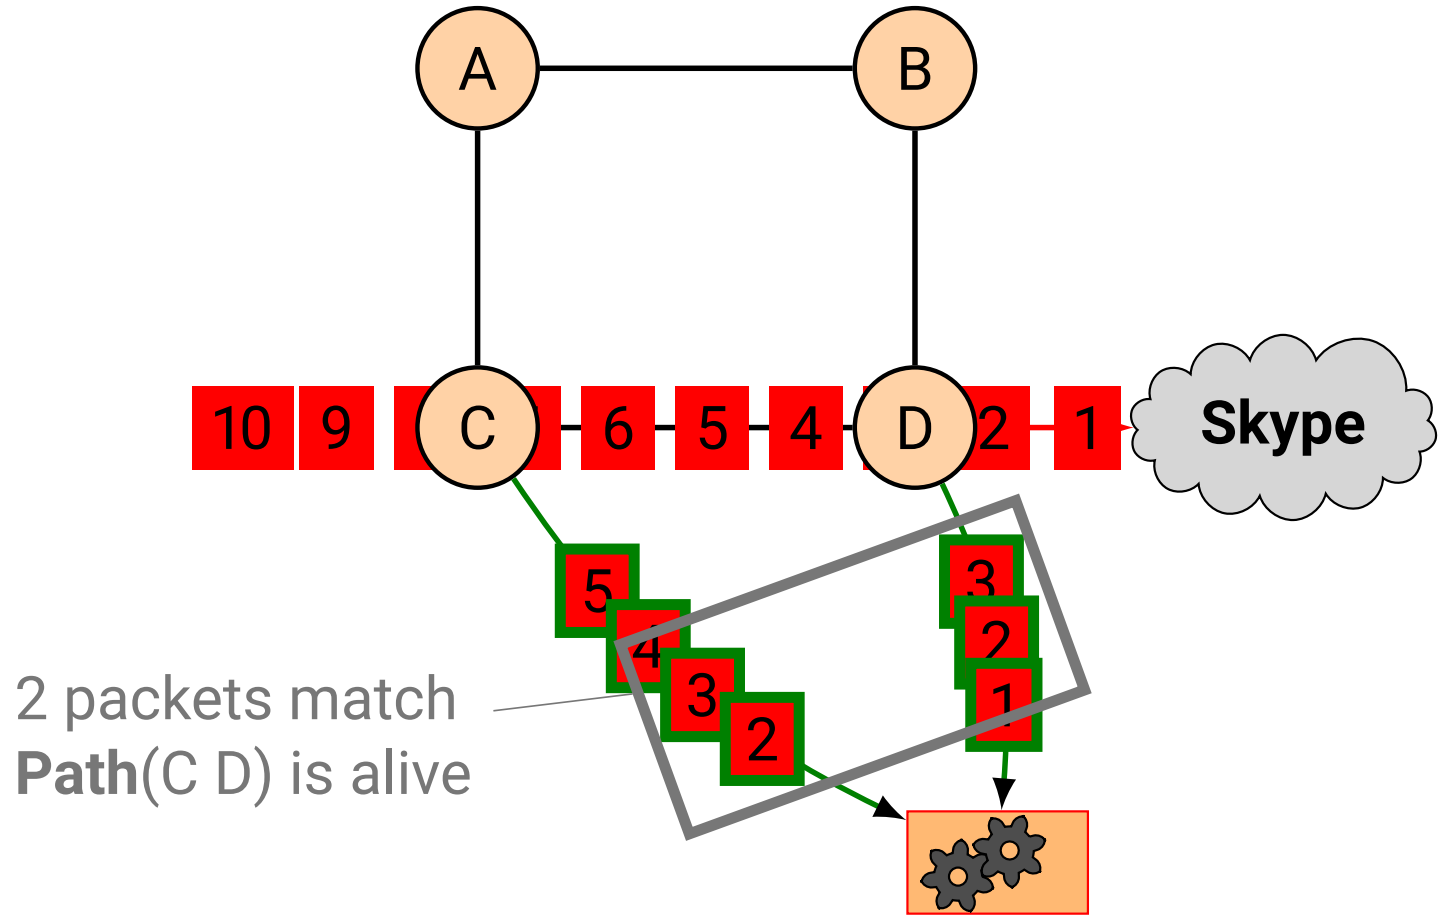

Capturing traffic slices is powerful.

- Slices contain the complete packet payload.
  Can remotely dissect traffic.

- Concurrent slices enable to trace a packet across the network and compute properties.
  e.g., proof of traversal, upper-bound on queuing delays.

- Fine-grained control on duration, point of capture and target prefix of slices.
  Explicit control on measurement overhead.

We implemented a collector prototype.

- Uses hardware-based mirroring features available in commercial routers.
  e.g., Cisco ERSPAN.

- Dynamically program the intra-domain routing protocol (OSPF) using Fibbing.
  can capture a traffic slice for *any* subprefix, network-wide.

We statically provision a mirroring VLAN on all links that must be monitored.

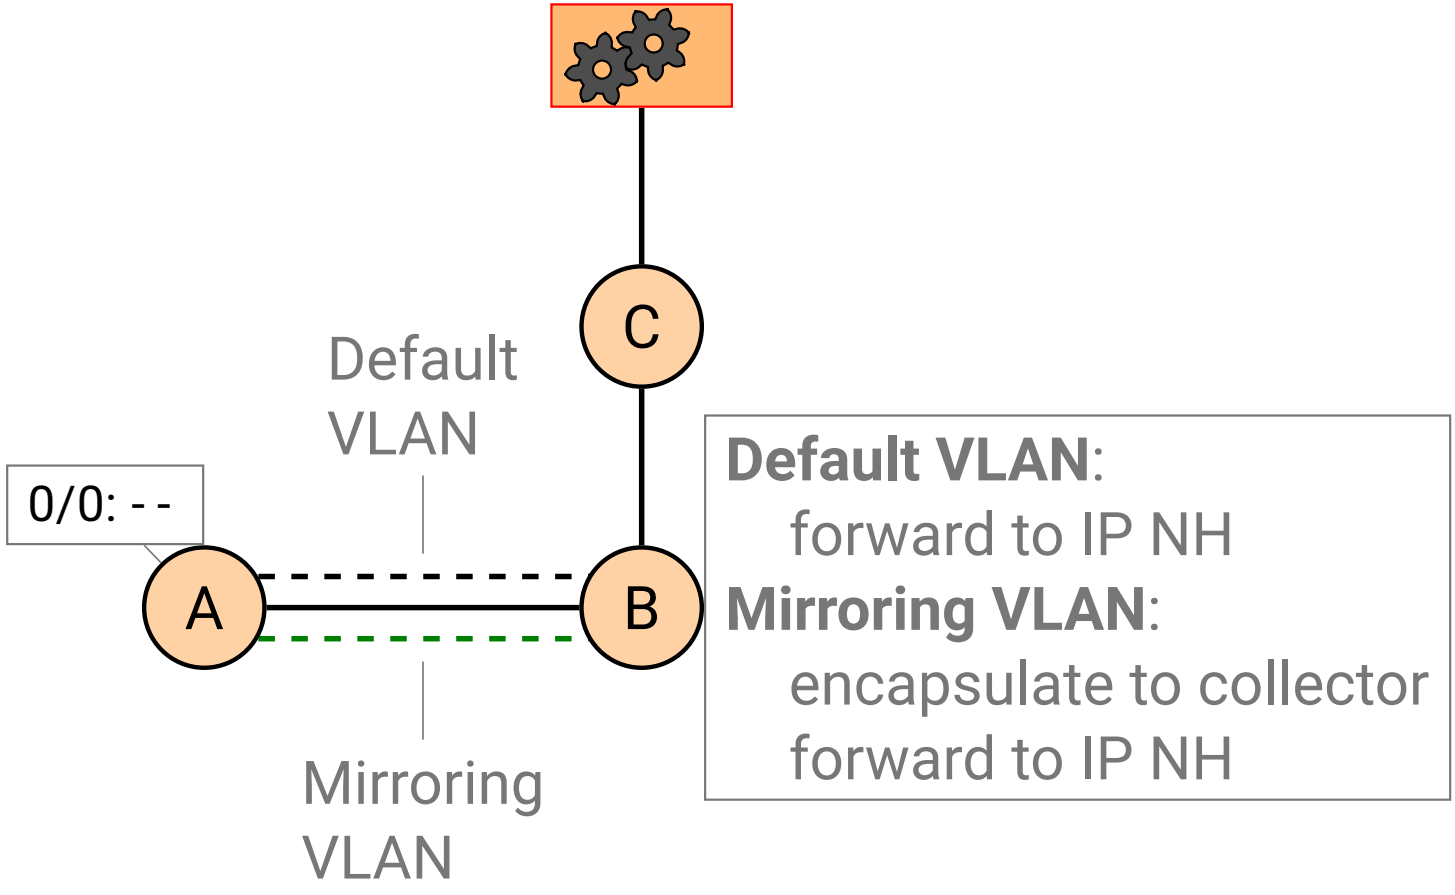

By default, all traffic is forwarded on the default VLAN.

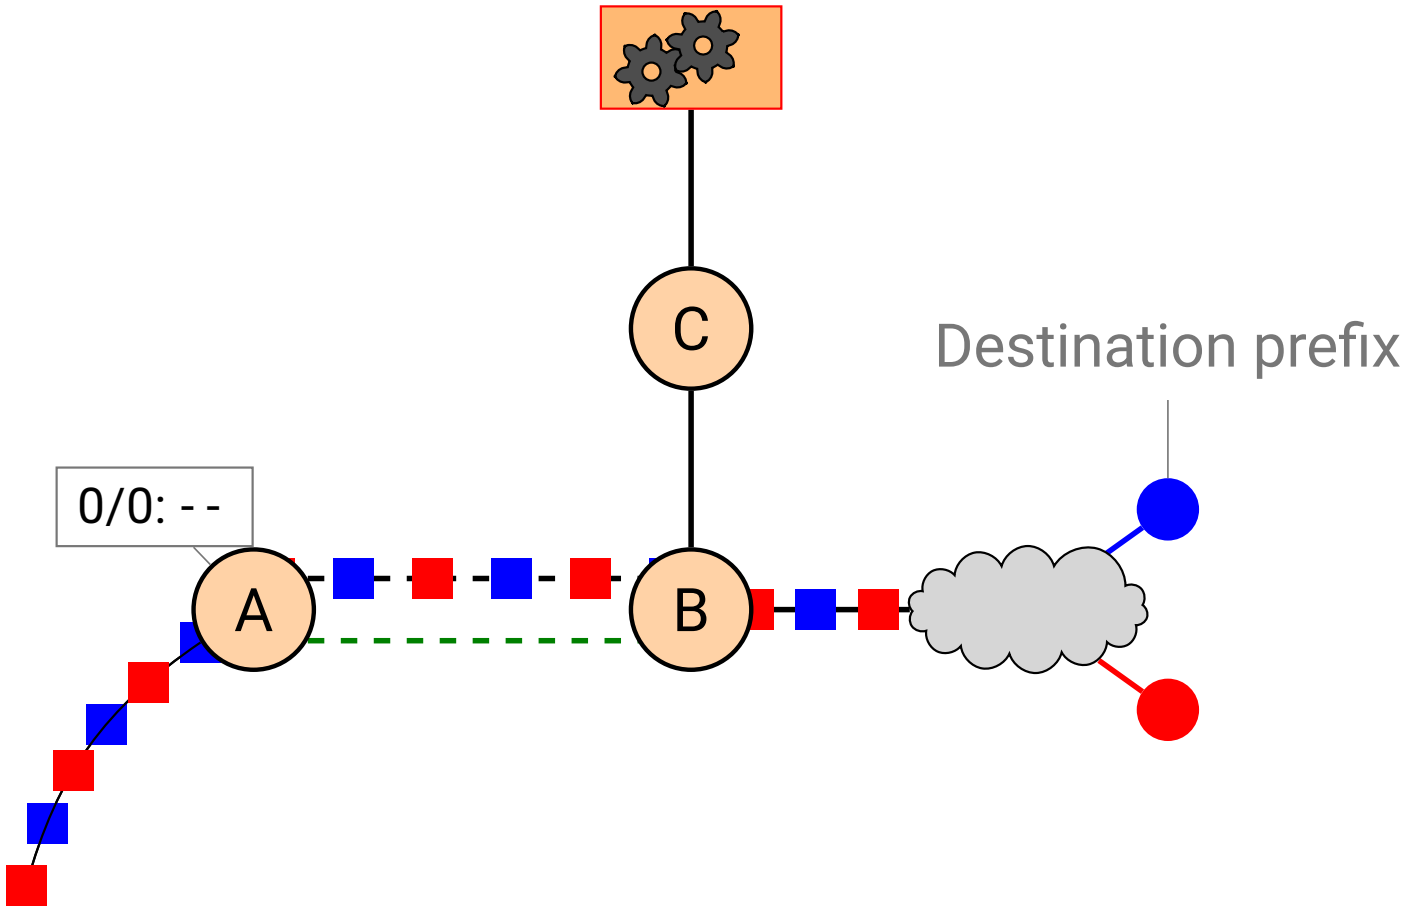

The collector sends an OSPF message to start a traffic slice.

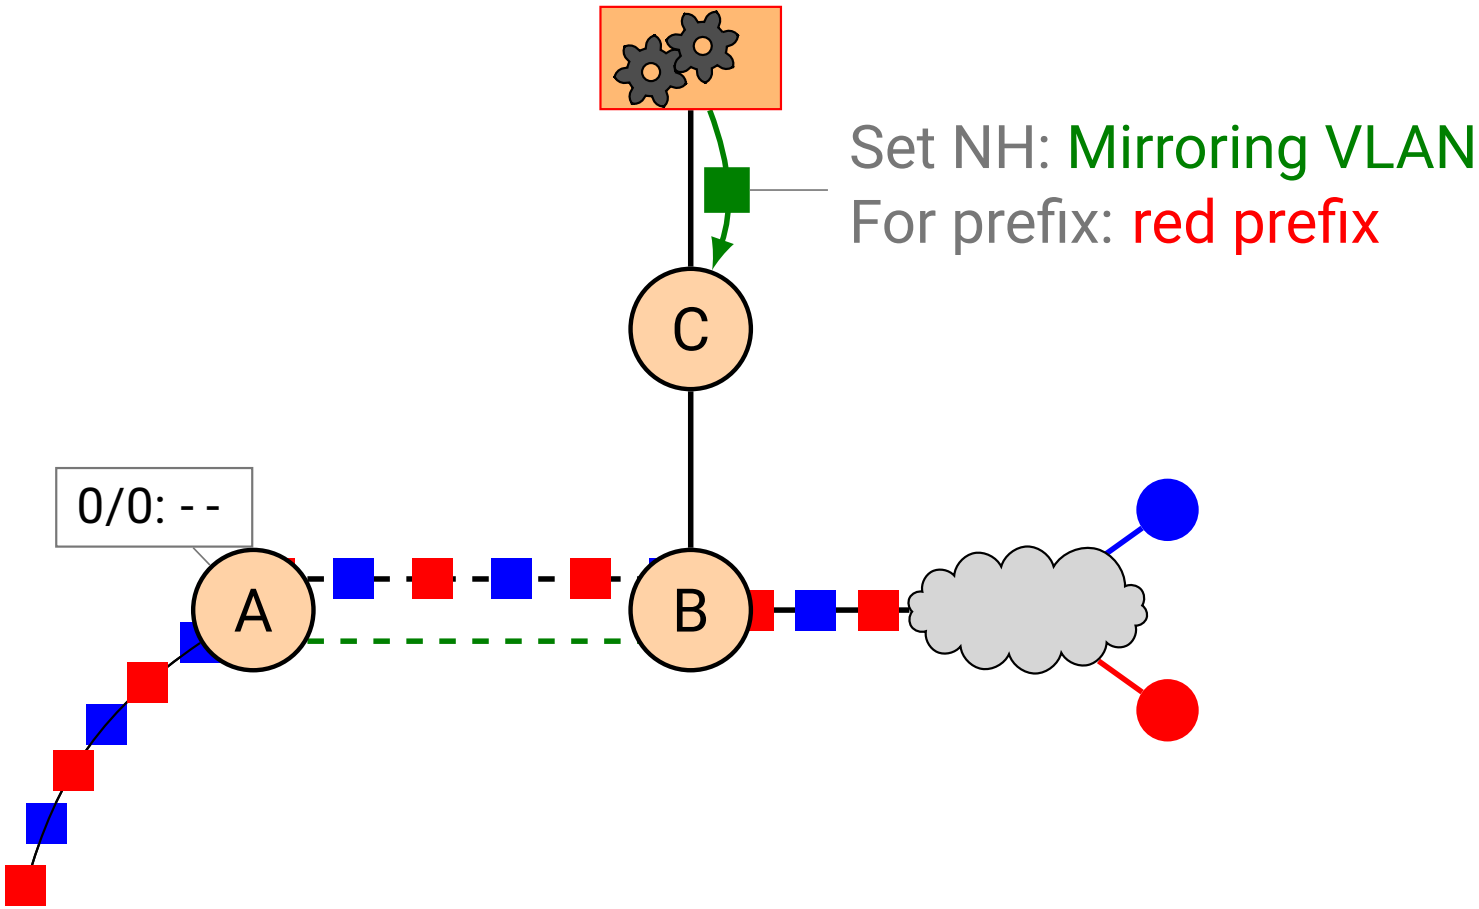

The OSPF message is flooded and reaches A, which then forwards traffic on the mirroring VLAN.

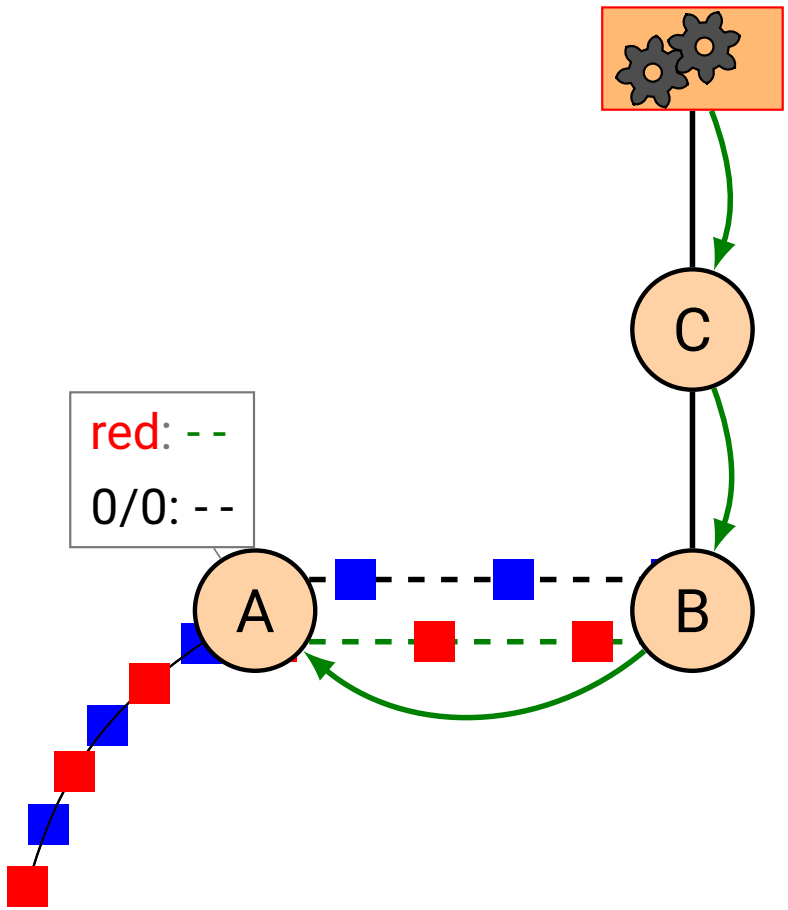

B then mirrors the packets towards the red prefix to the collector

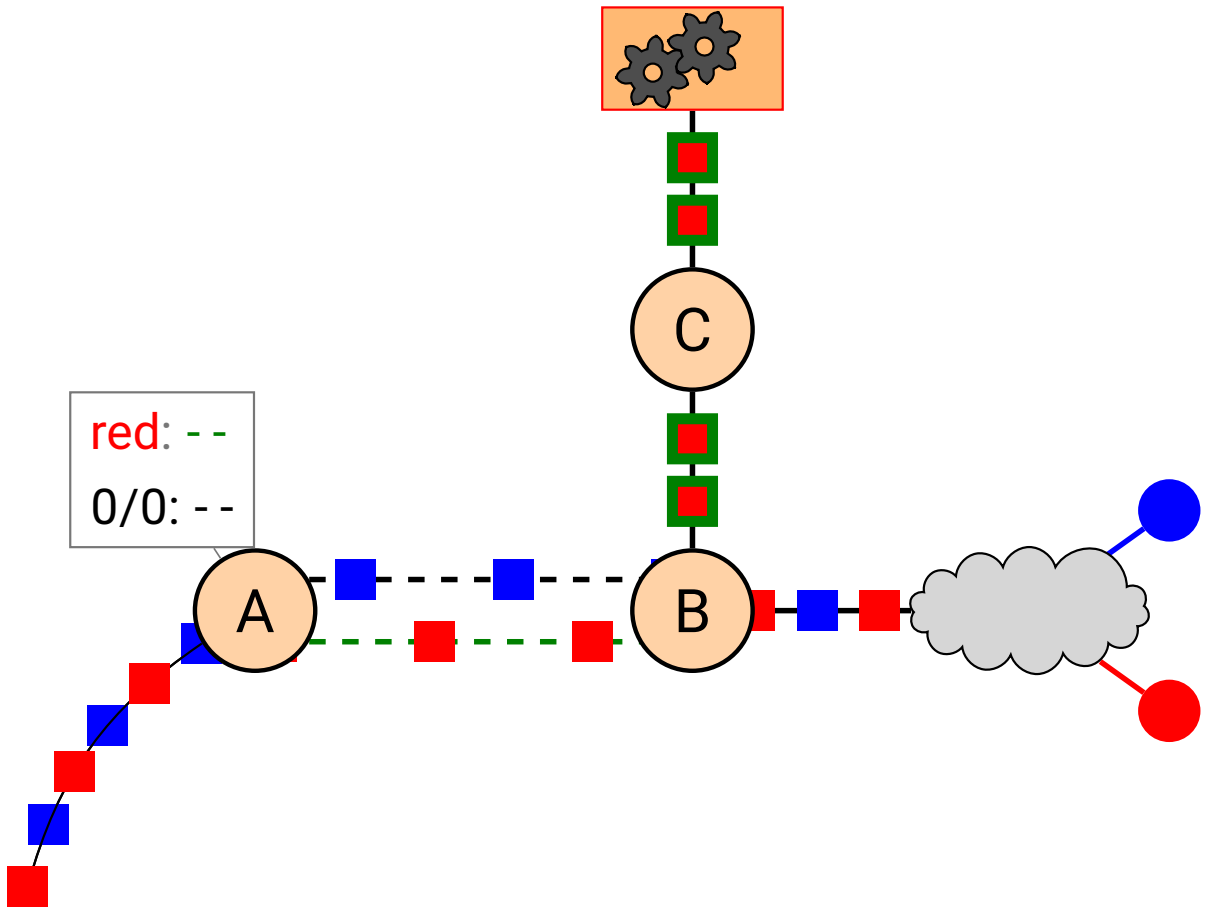

The collector stops the traffic slice similarly

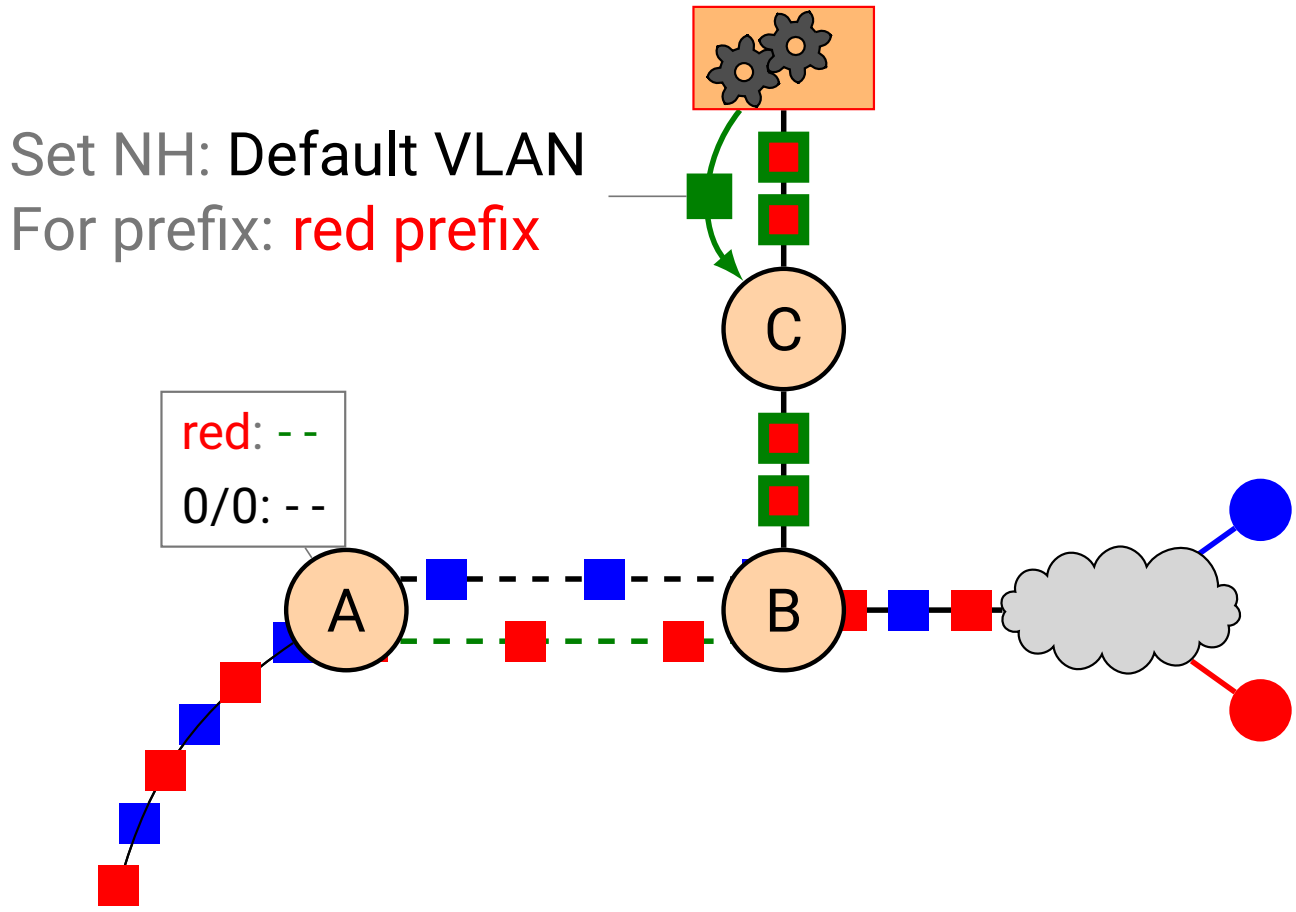

The collector stops the traffic slice similarly

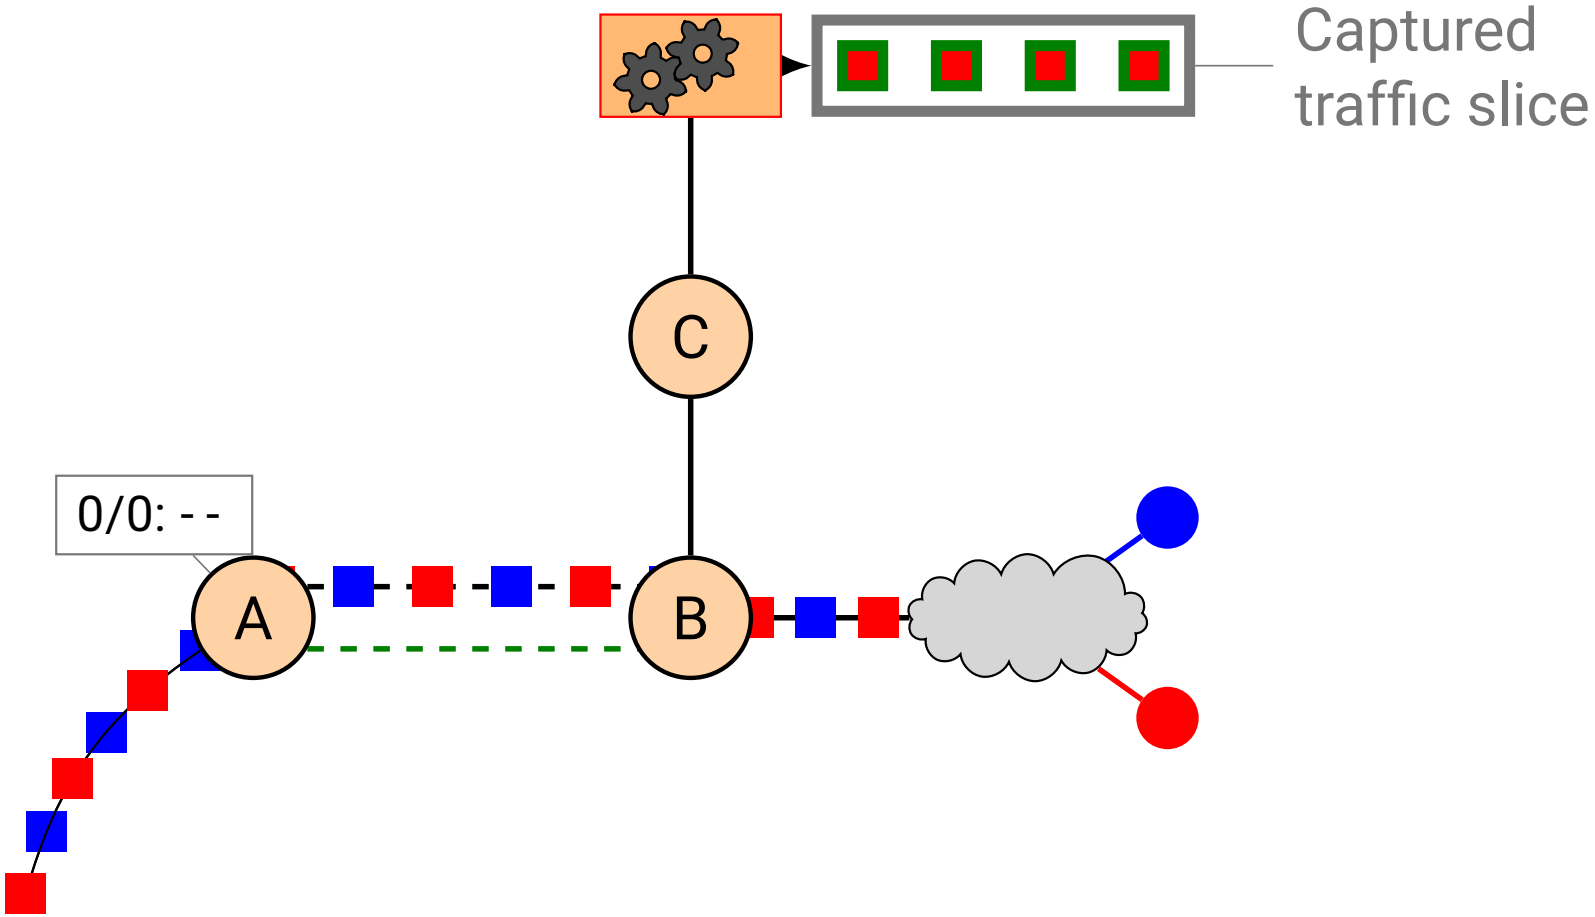

Our preliminary tests show that Mille-Feuille can work in practice.

- We were able to capture traffic **slices as thin as 14 ms**
- We **control the slice duration** through the delay between the activation and deactivation message.
- We were able to concurrently **(de)activate 1000 mirroring rules in 0.93 ms**, and 10 000 in 30 ms.

Mille-Feuille is a measurement framework realizing a deterministic sampling of the network in real time.

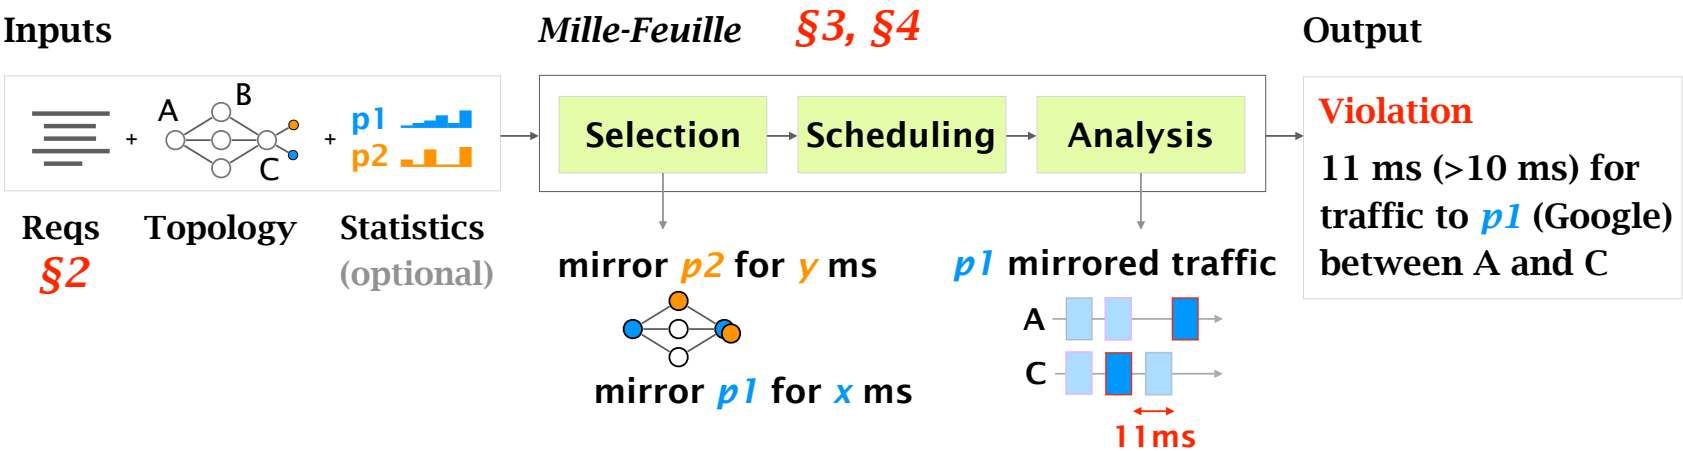

In Mille-Feuille, operators specify high-level measurement requirements and an associated measurement budget.

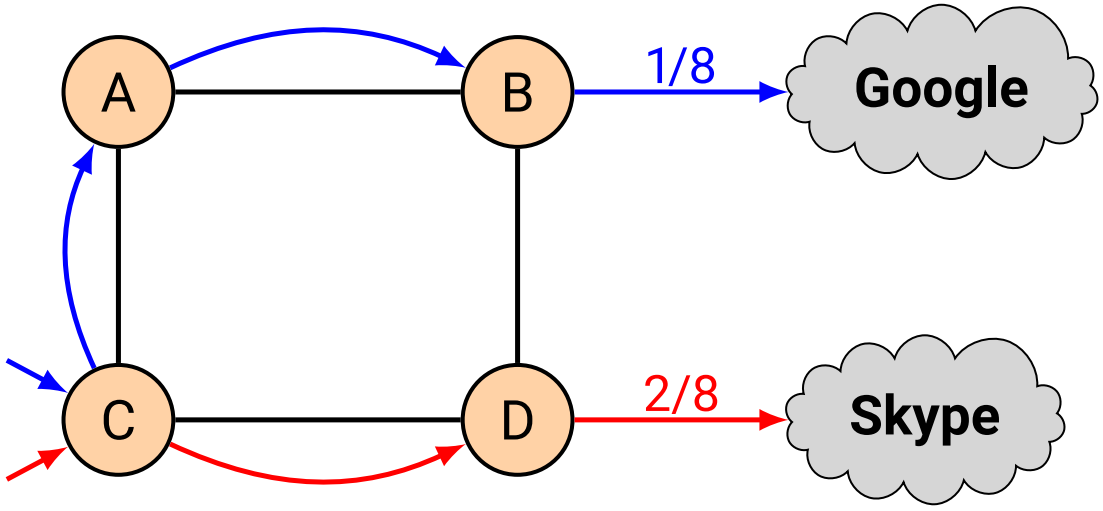

(

Path(C A B) **for** Google;

Path(*) **within**(20 ms) **for** Skype;

) **every**(1 s) **in**(30 ms) **using**(1 Gbps)

What? From traffic estimates, Mille-Feuille iteratively selects subprefixes to monitor.

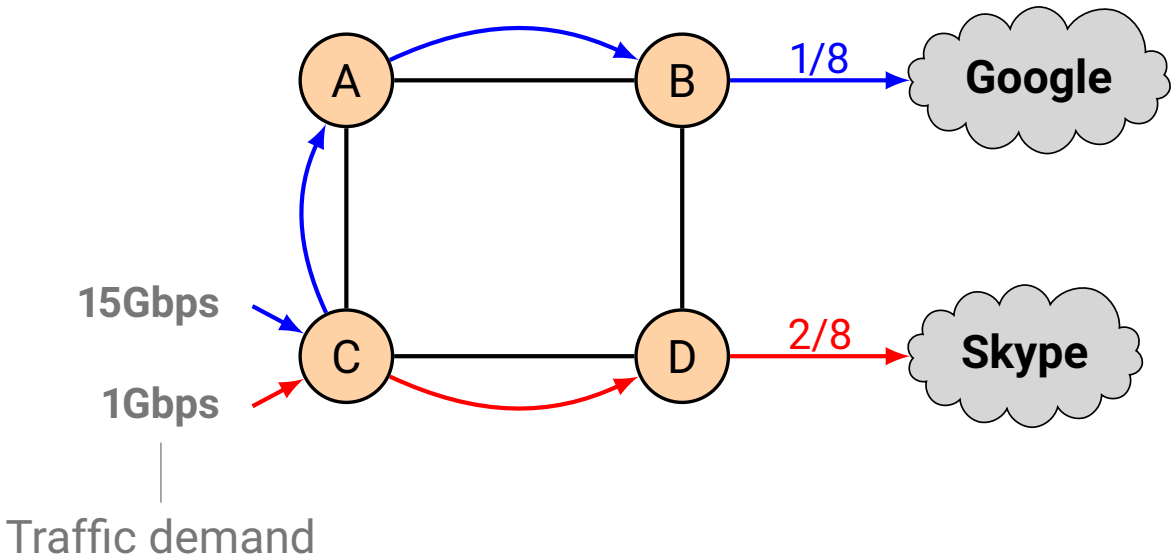

Traffic distribution

1/8	15	Gbps
2/8	1	Gbps

What? From traffic estimates, Mille-Feuille iteratively selects subprefixes to monitor.

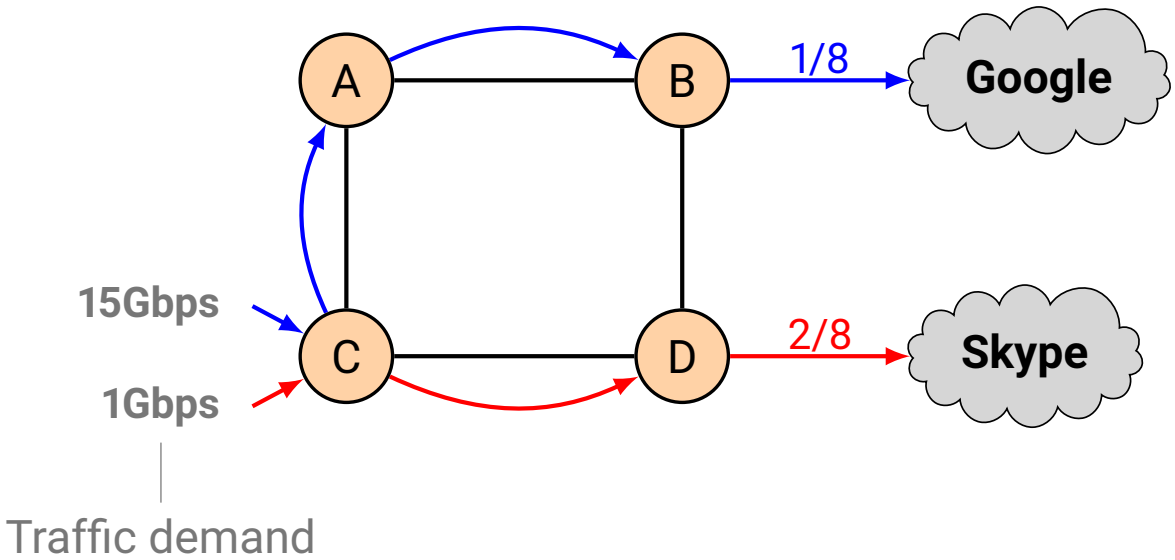

Traffic distribution

1/8	15	Gbps
1/24	.5	Gbps
2/8	1	Gbps
2/16	.1	Gbps

Target prefixes for schedule #1: 1.0.0.0/24, 2.0.0.0/16

What? From traffic estimates, Mille-Feuille iteratively selects subprefixes to monitor.

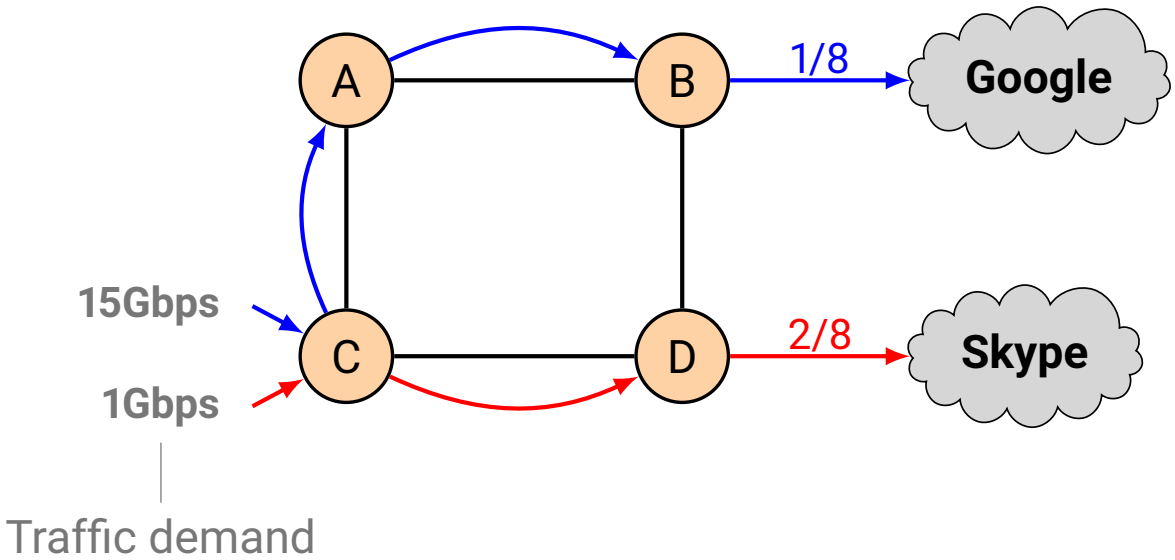

Traffic distribution

1/8	15	Gbps
1/24	.5	Gbps
2/8	1	Gbps
2/16	.1	Gbps

Target prefixes for schedule #1: 1.0.0.0/24, 2.0.0.0/16
Target prefixes for schedule #2: 1.0.1.0/24, 2.0.1.0/16
. . .

Where? Mille-Feuille creates mirroring rules and assigns them to one or more routers.

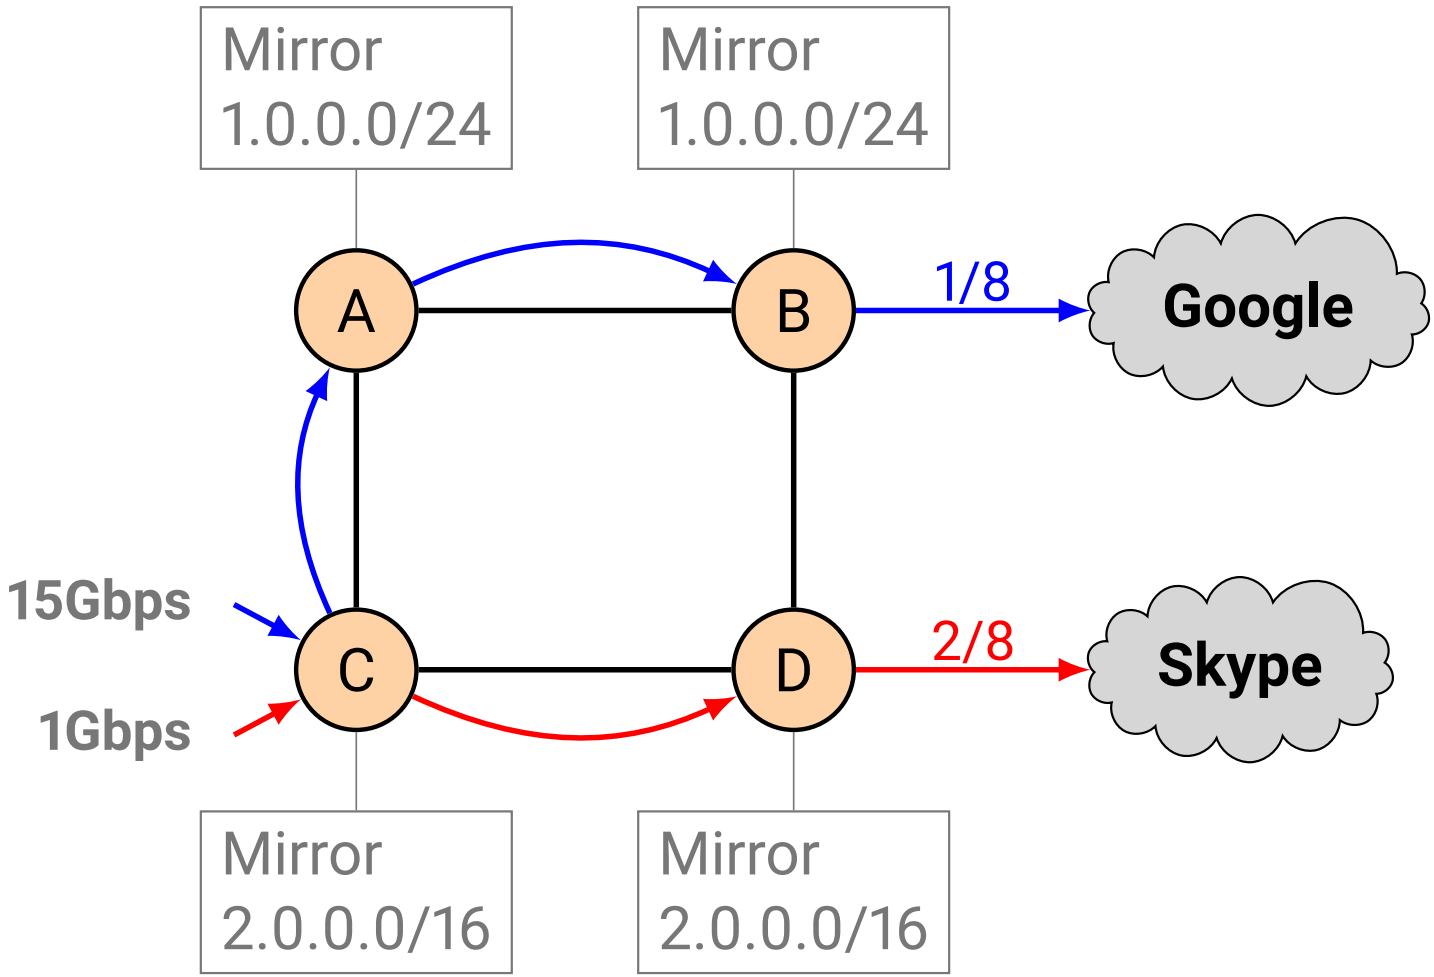

When? Mille-Feuille spreads the measurement campaign across time to meet the budget

$0\,ms \leqslant t < 15\,ms$

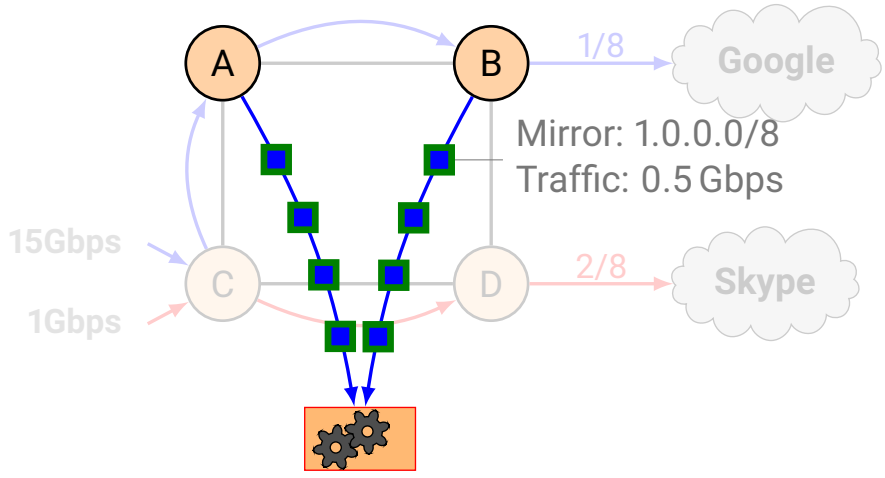

$15\,ms \leqslant t < 30\,ms$

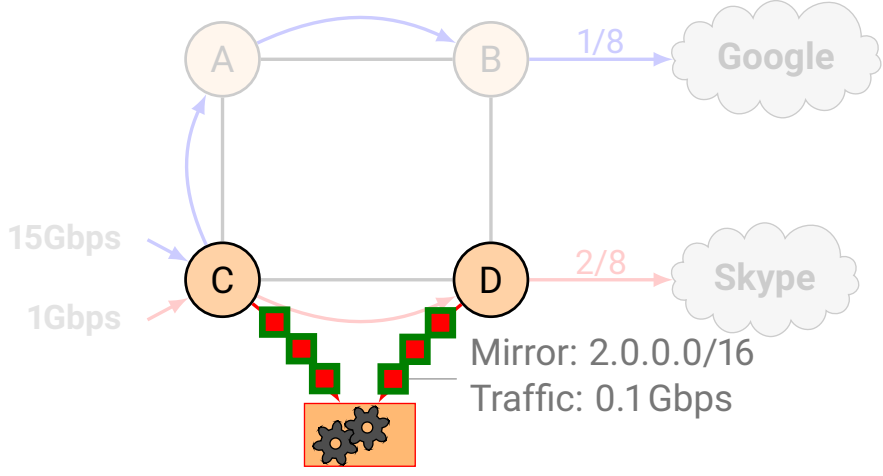

Mille-Feuille: Putting ISP traffic under the scalpel

- We collect **thin traffic slices** by programming the intra-domain routing protocol.
- We realize a **deterministic sampling** of the state of the network.
- We limit the **measurement** overhead according to a **budget**.