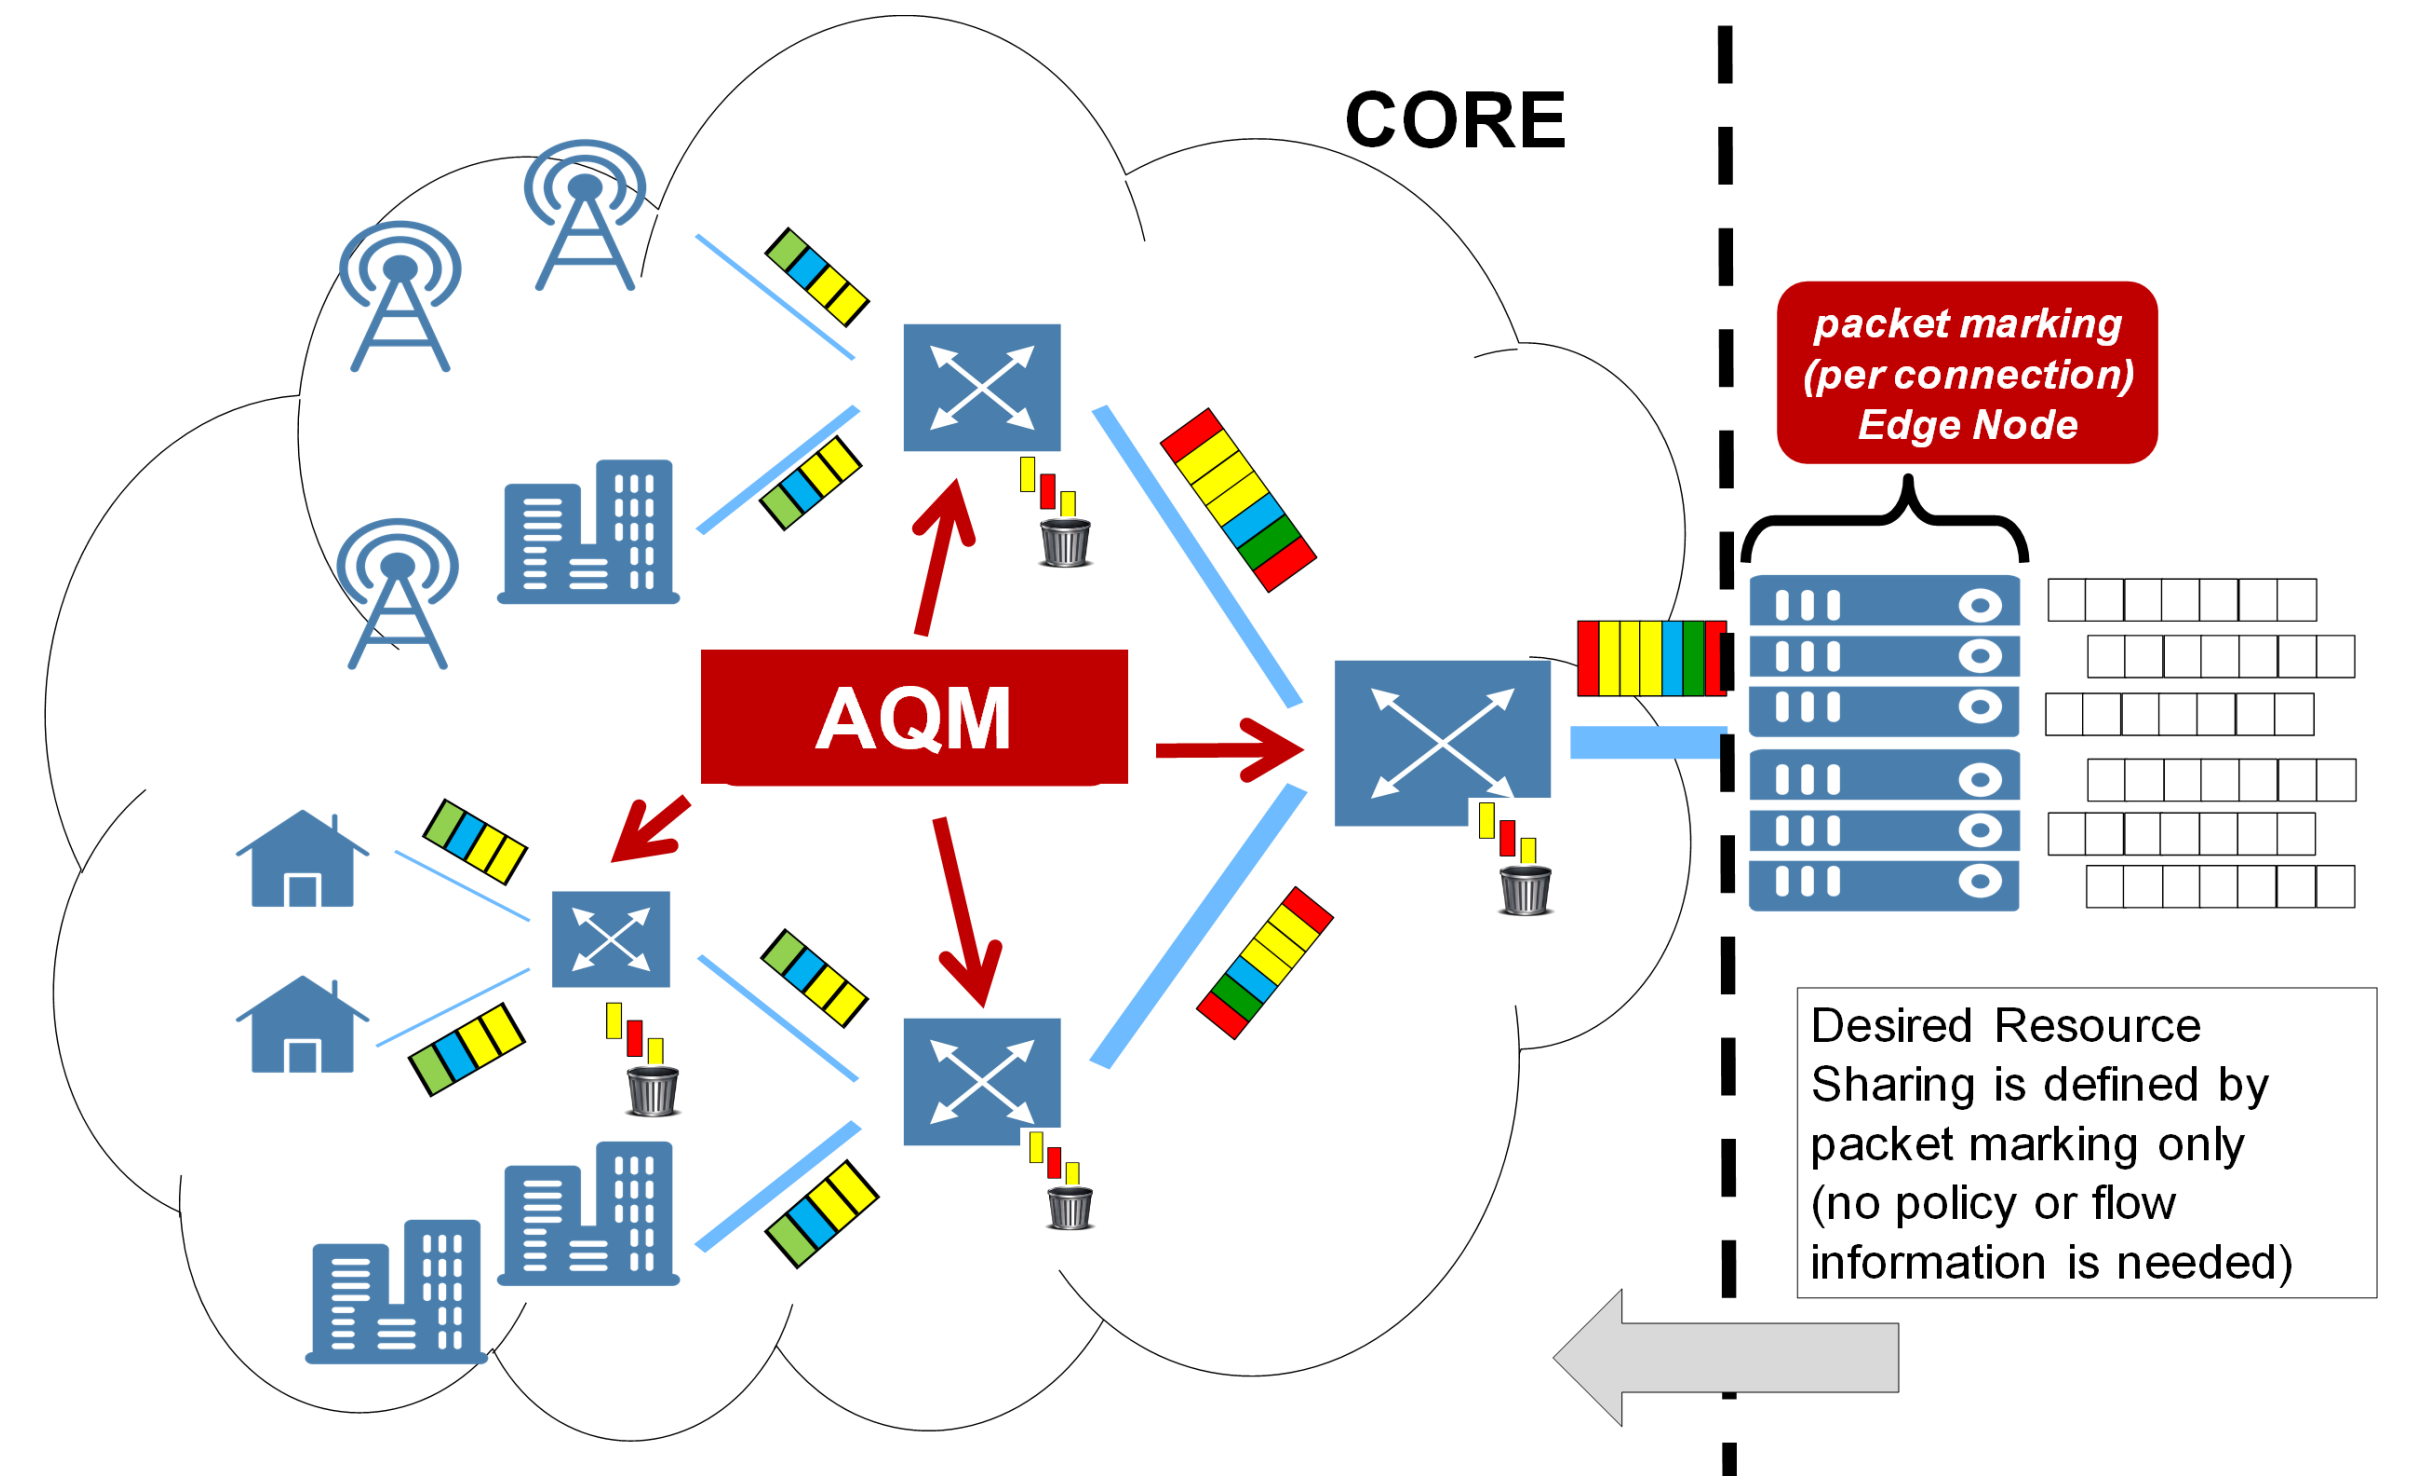


ABSTRACT

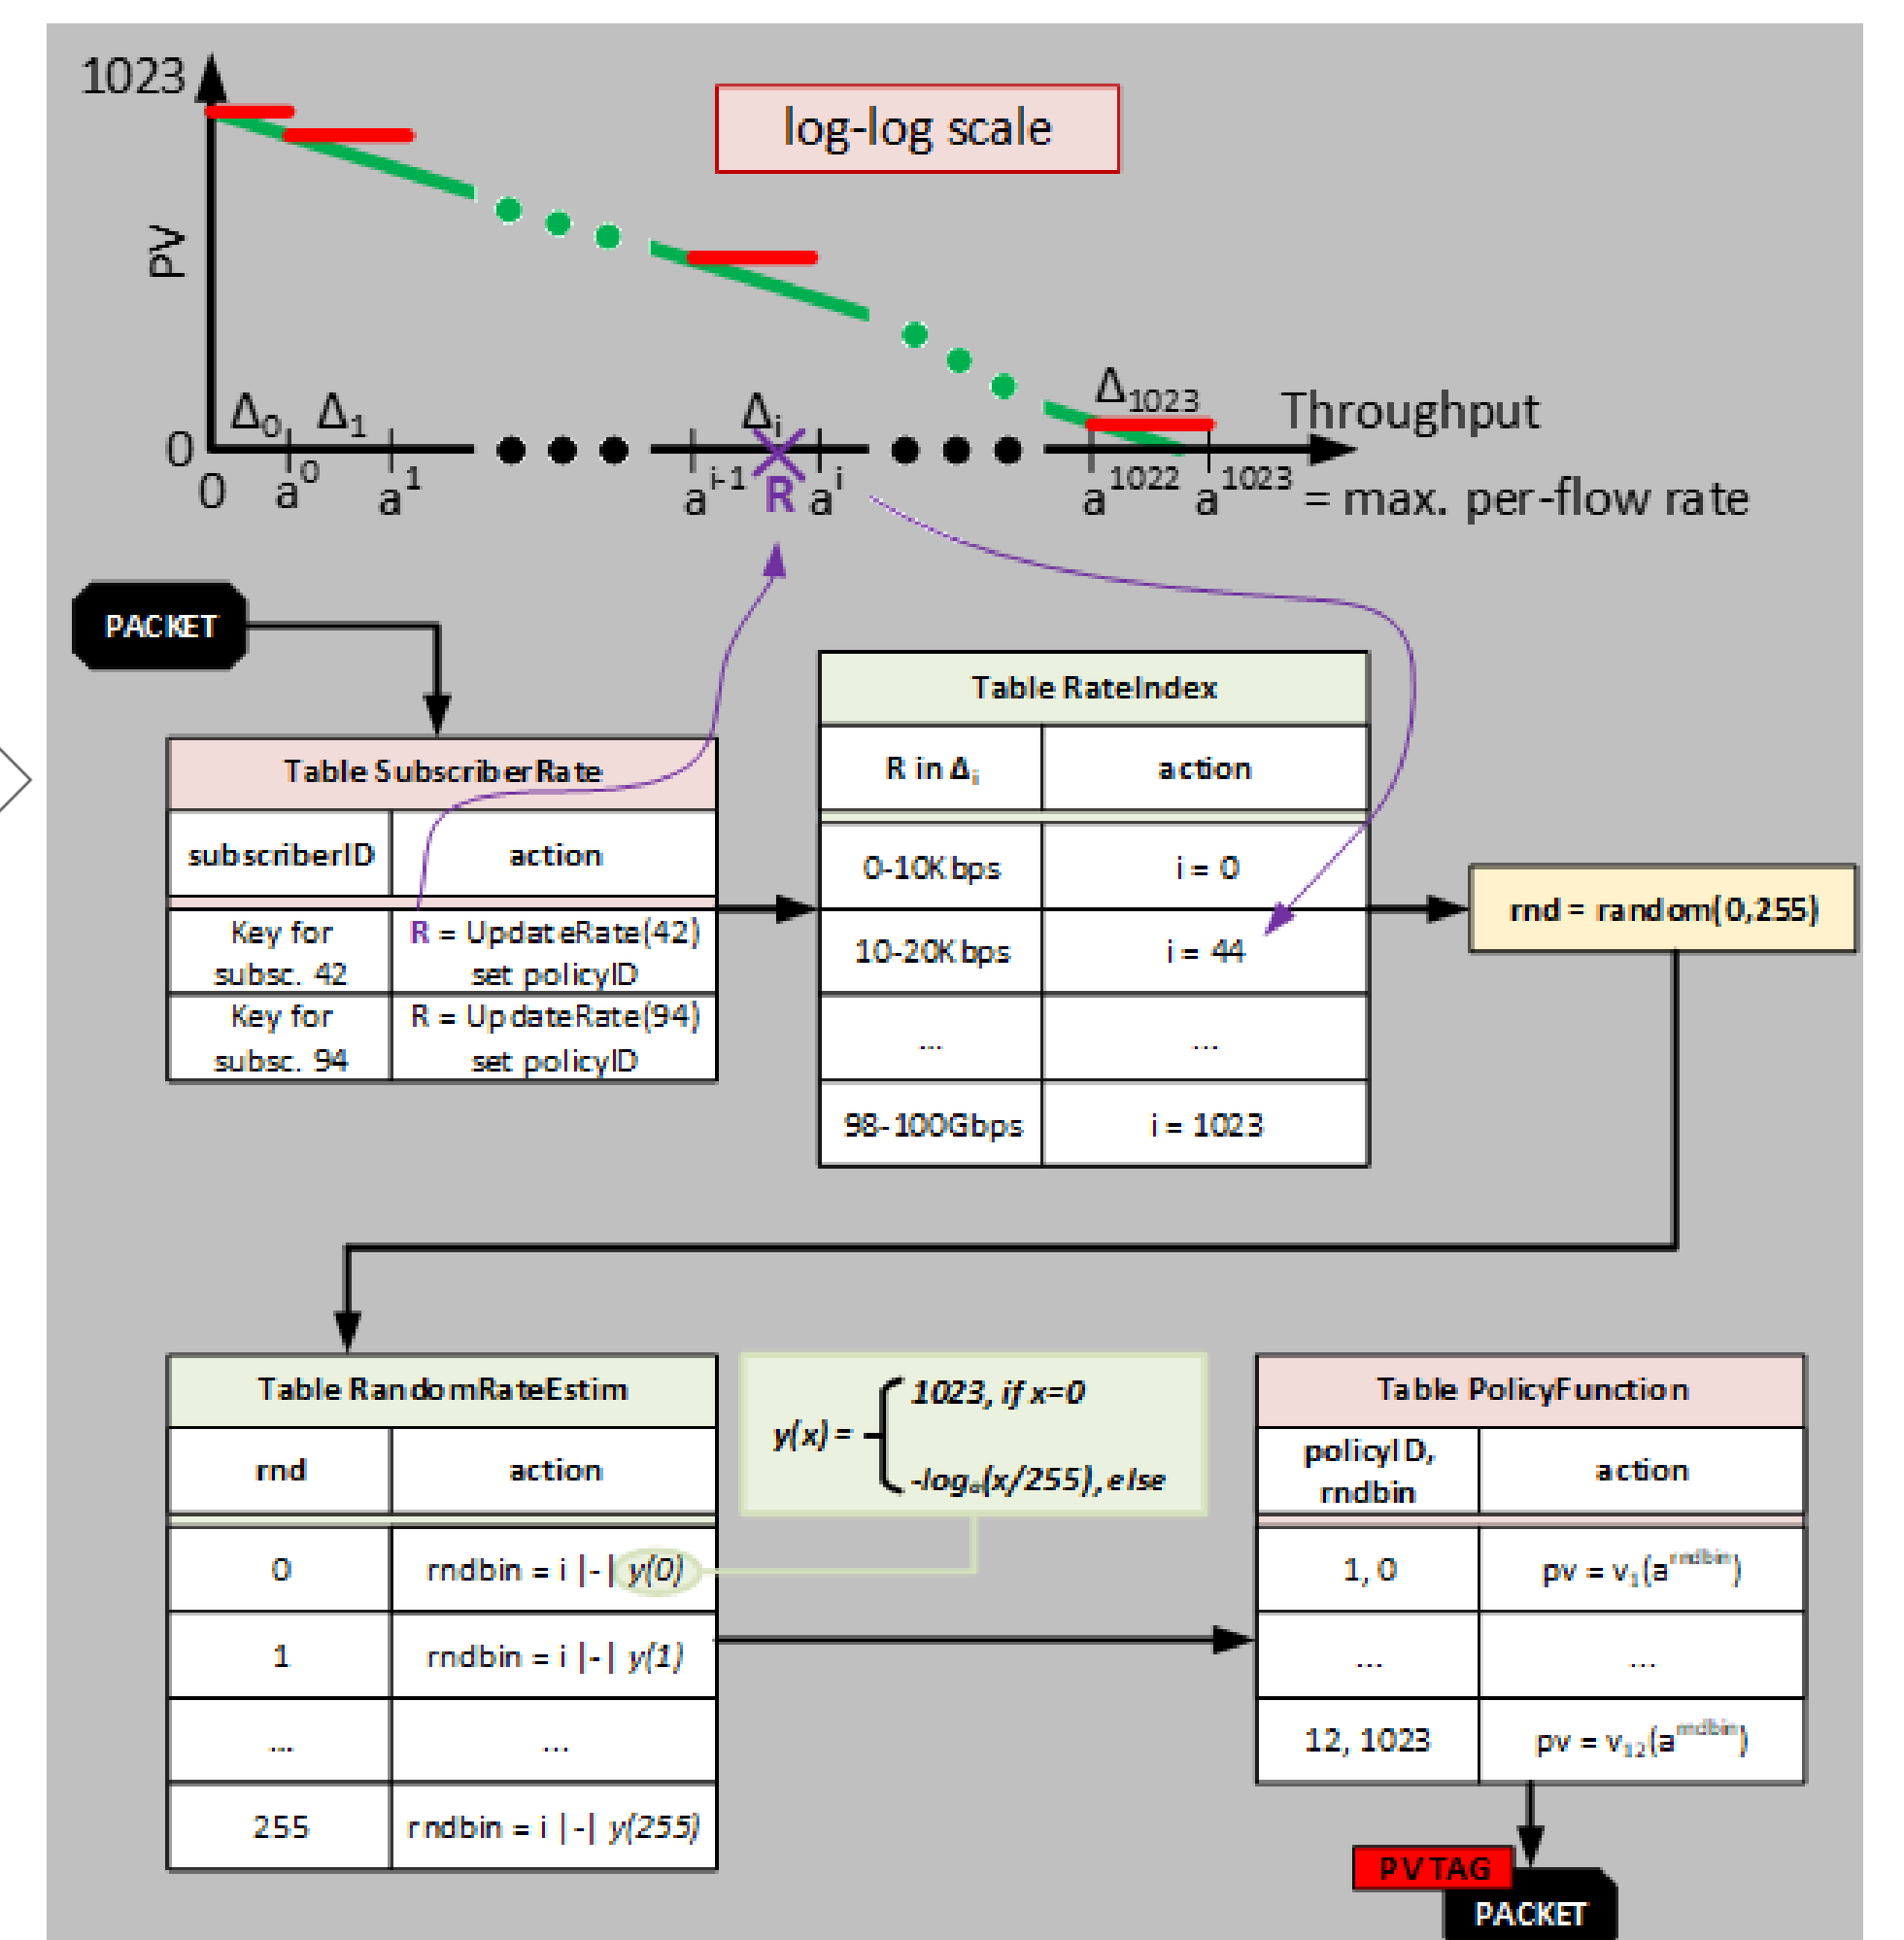
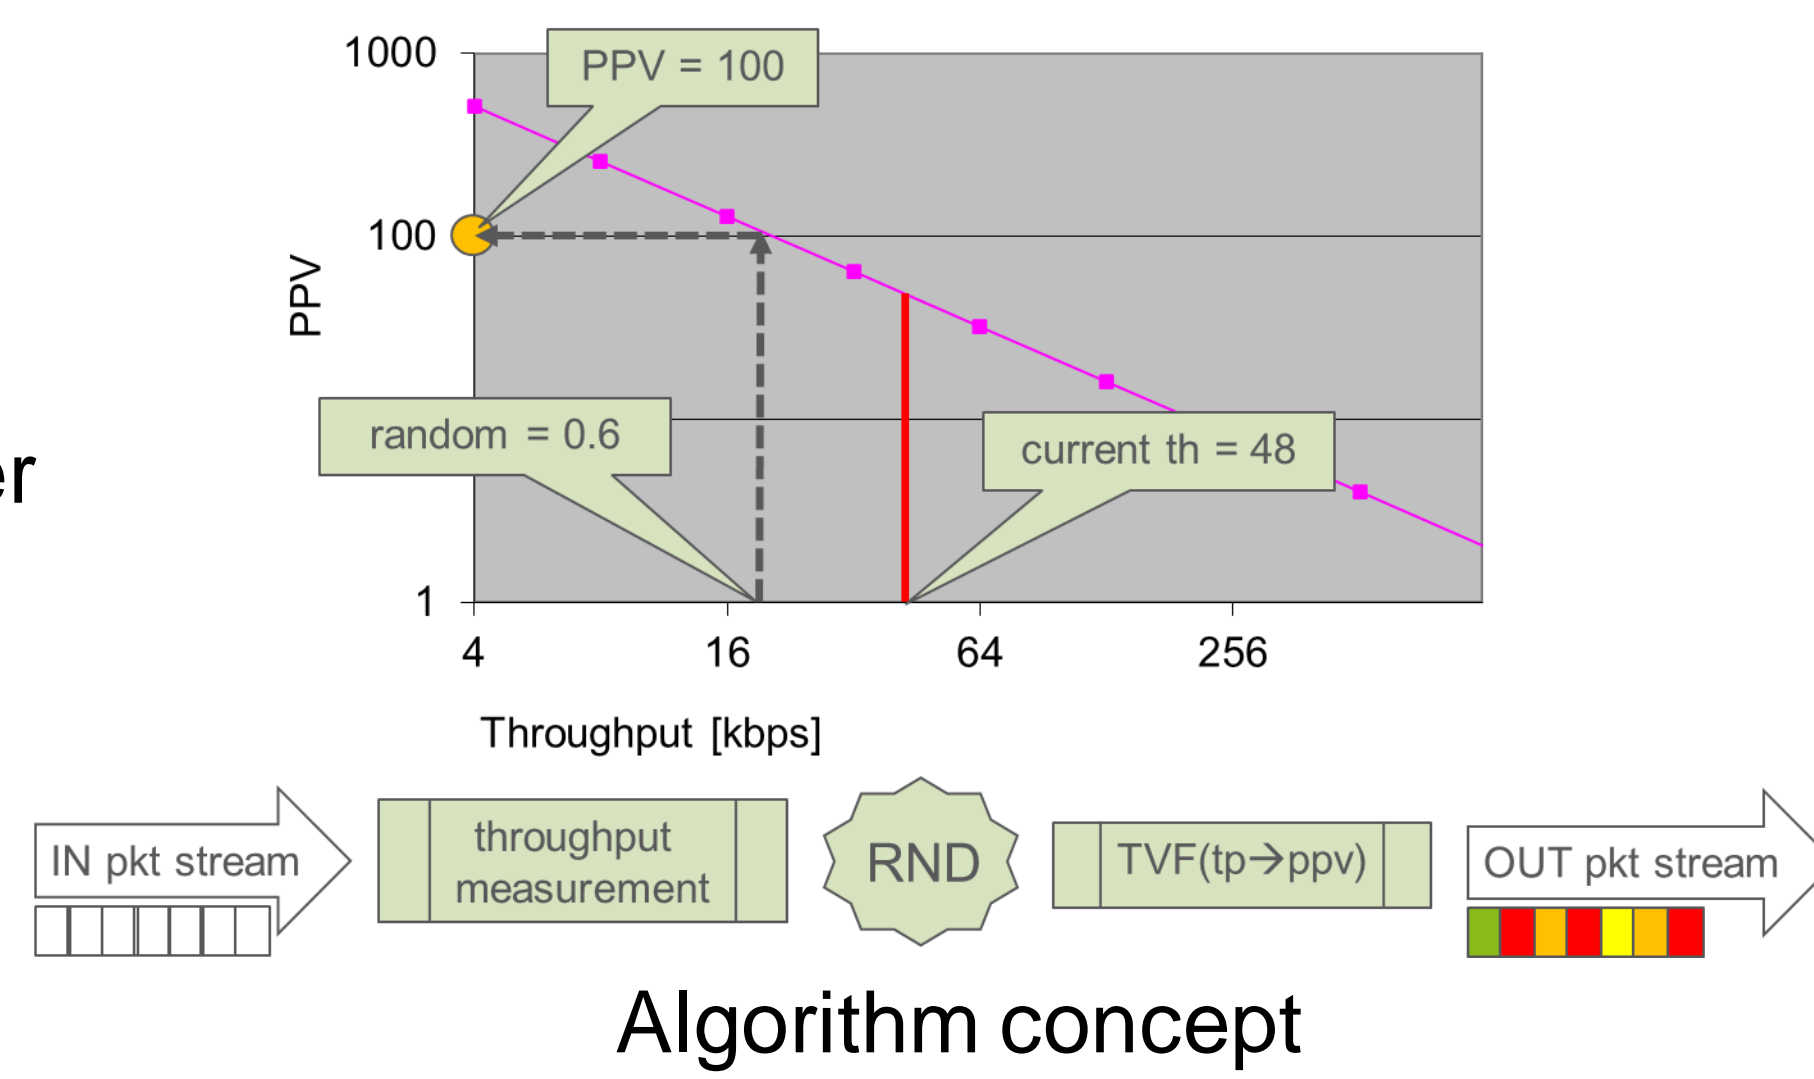
Core-stateless resource sharing solutions implemented in P4 hardware data planes have been proposed in the past few years. They share the idea of tagging packets with special values at the network edge that are then solely used for deciding how to handle packets in the network in case of congestion. Though the scheduler of our Core-Stateless Active Queue management (CSAQM) was implemented in P4 and was evaluated on Intel Tofino ASIC, the packet marker have only had a DPDK-based software implementation so far. In this demo:

- we present the **full data plane implementation of CSAQM**. Both packet marking and packet scheduling are executed by an **Intel Tofino ASIC**
- we demonstrate the scalability of our implementation by showing policy enforcement among up to 35000 subscribers at a 100~Gbps bottleneck using only a single queue
- In addition, we also present the resource sharing and isolation properties of CSAQM between flows with different rate control strategies, resulting in flow-specific congestion signals (drop probabilities) by design.




PACKET MARKING ON TOFINO


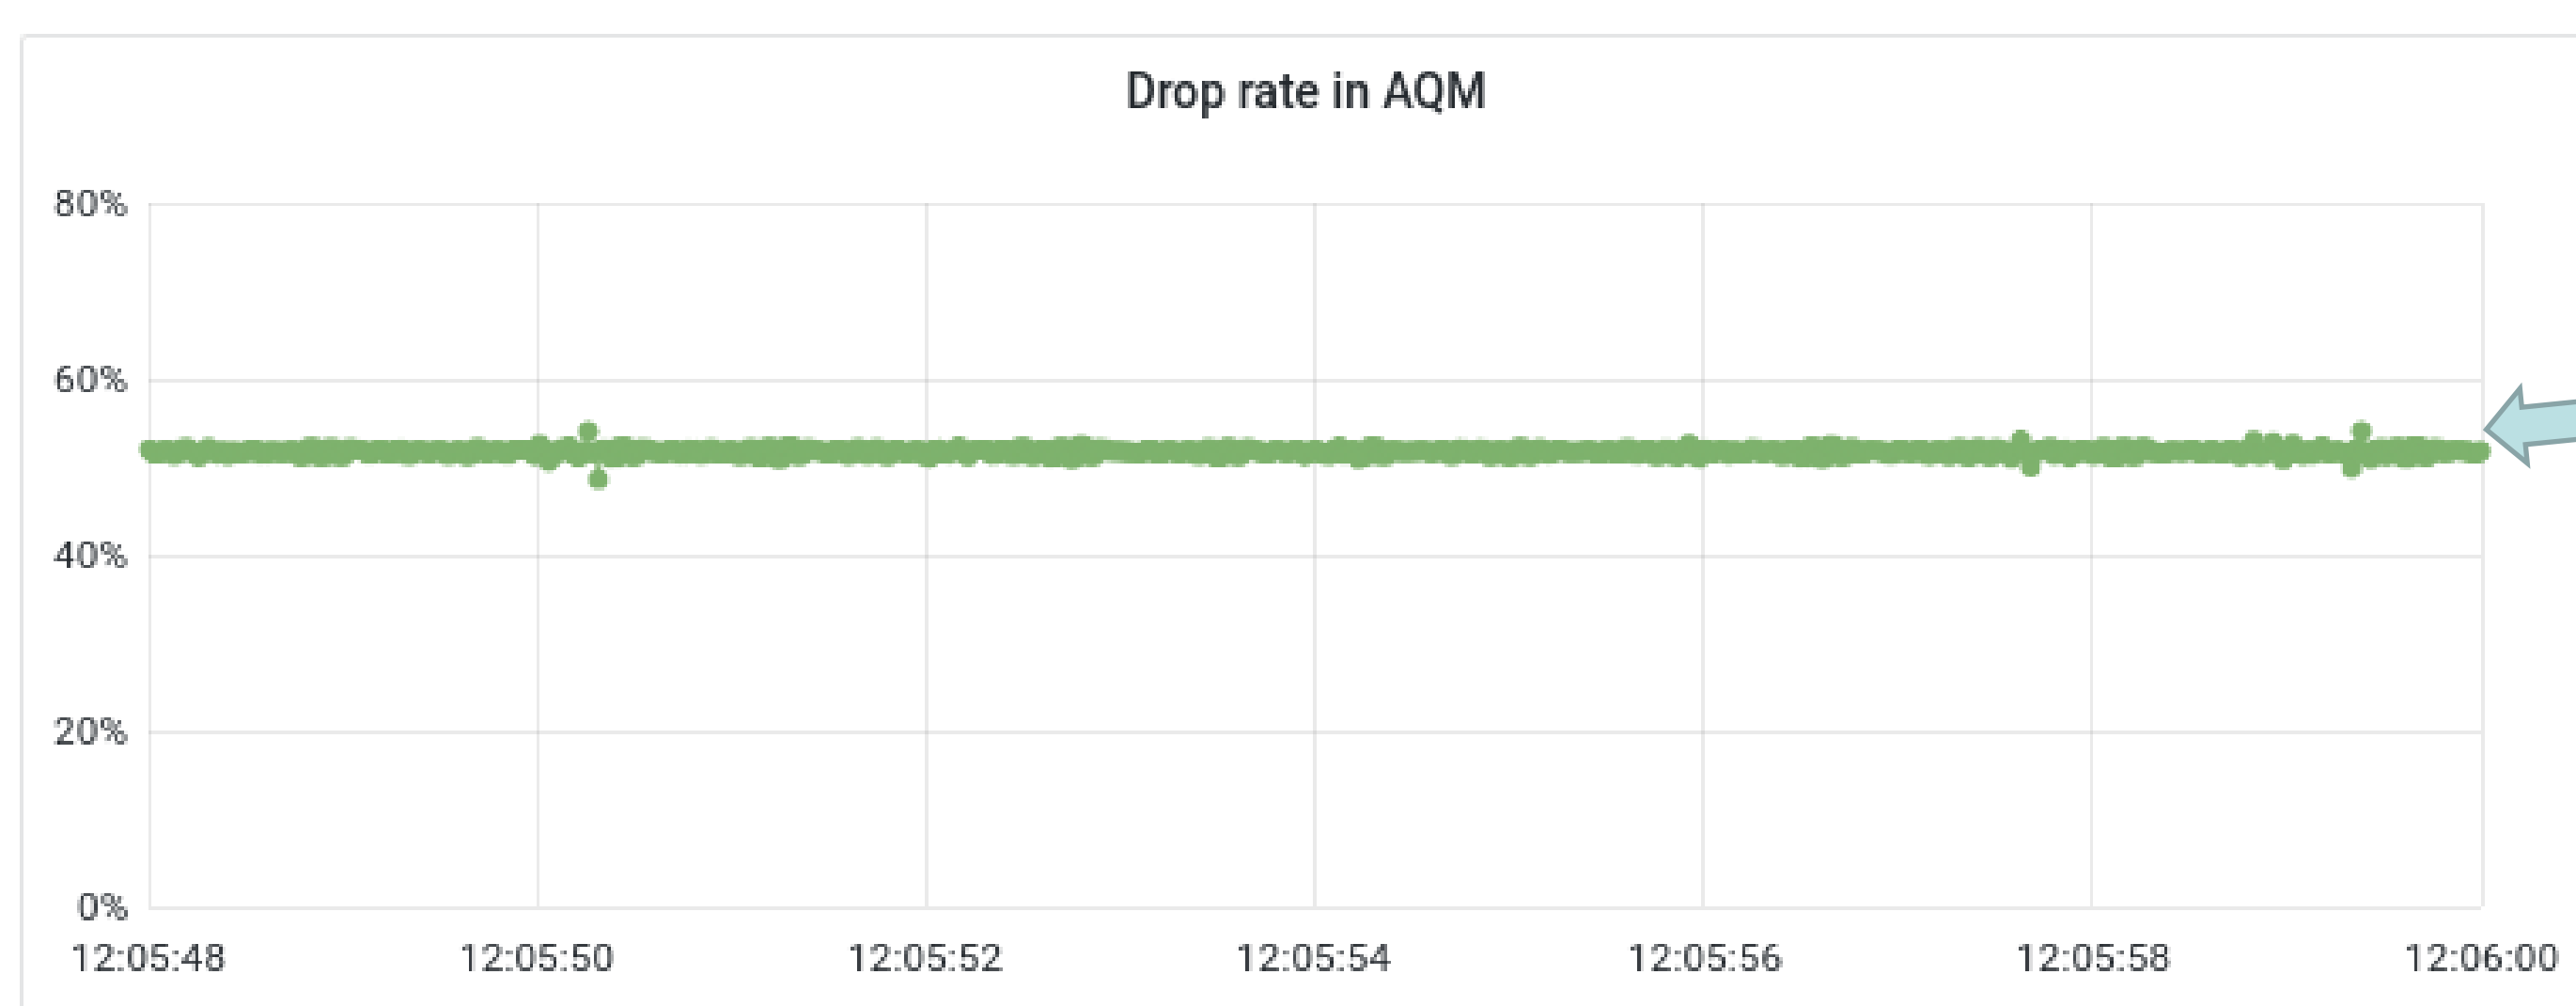
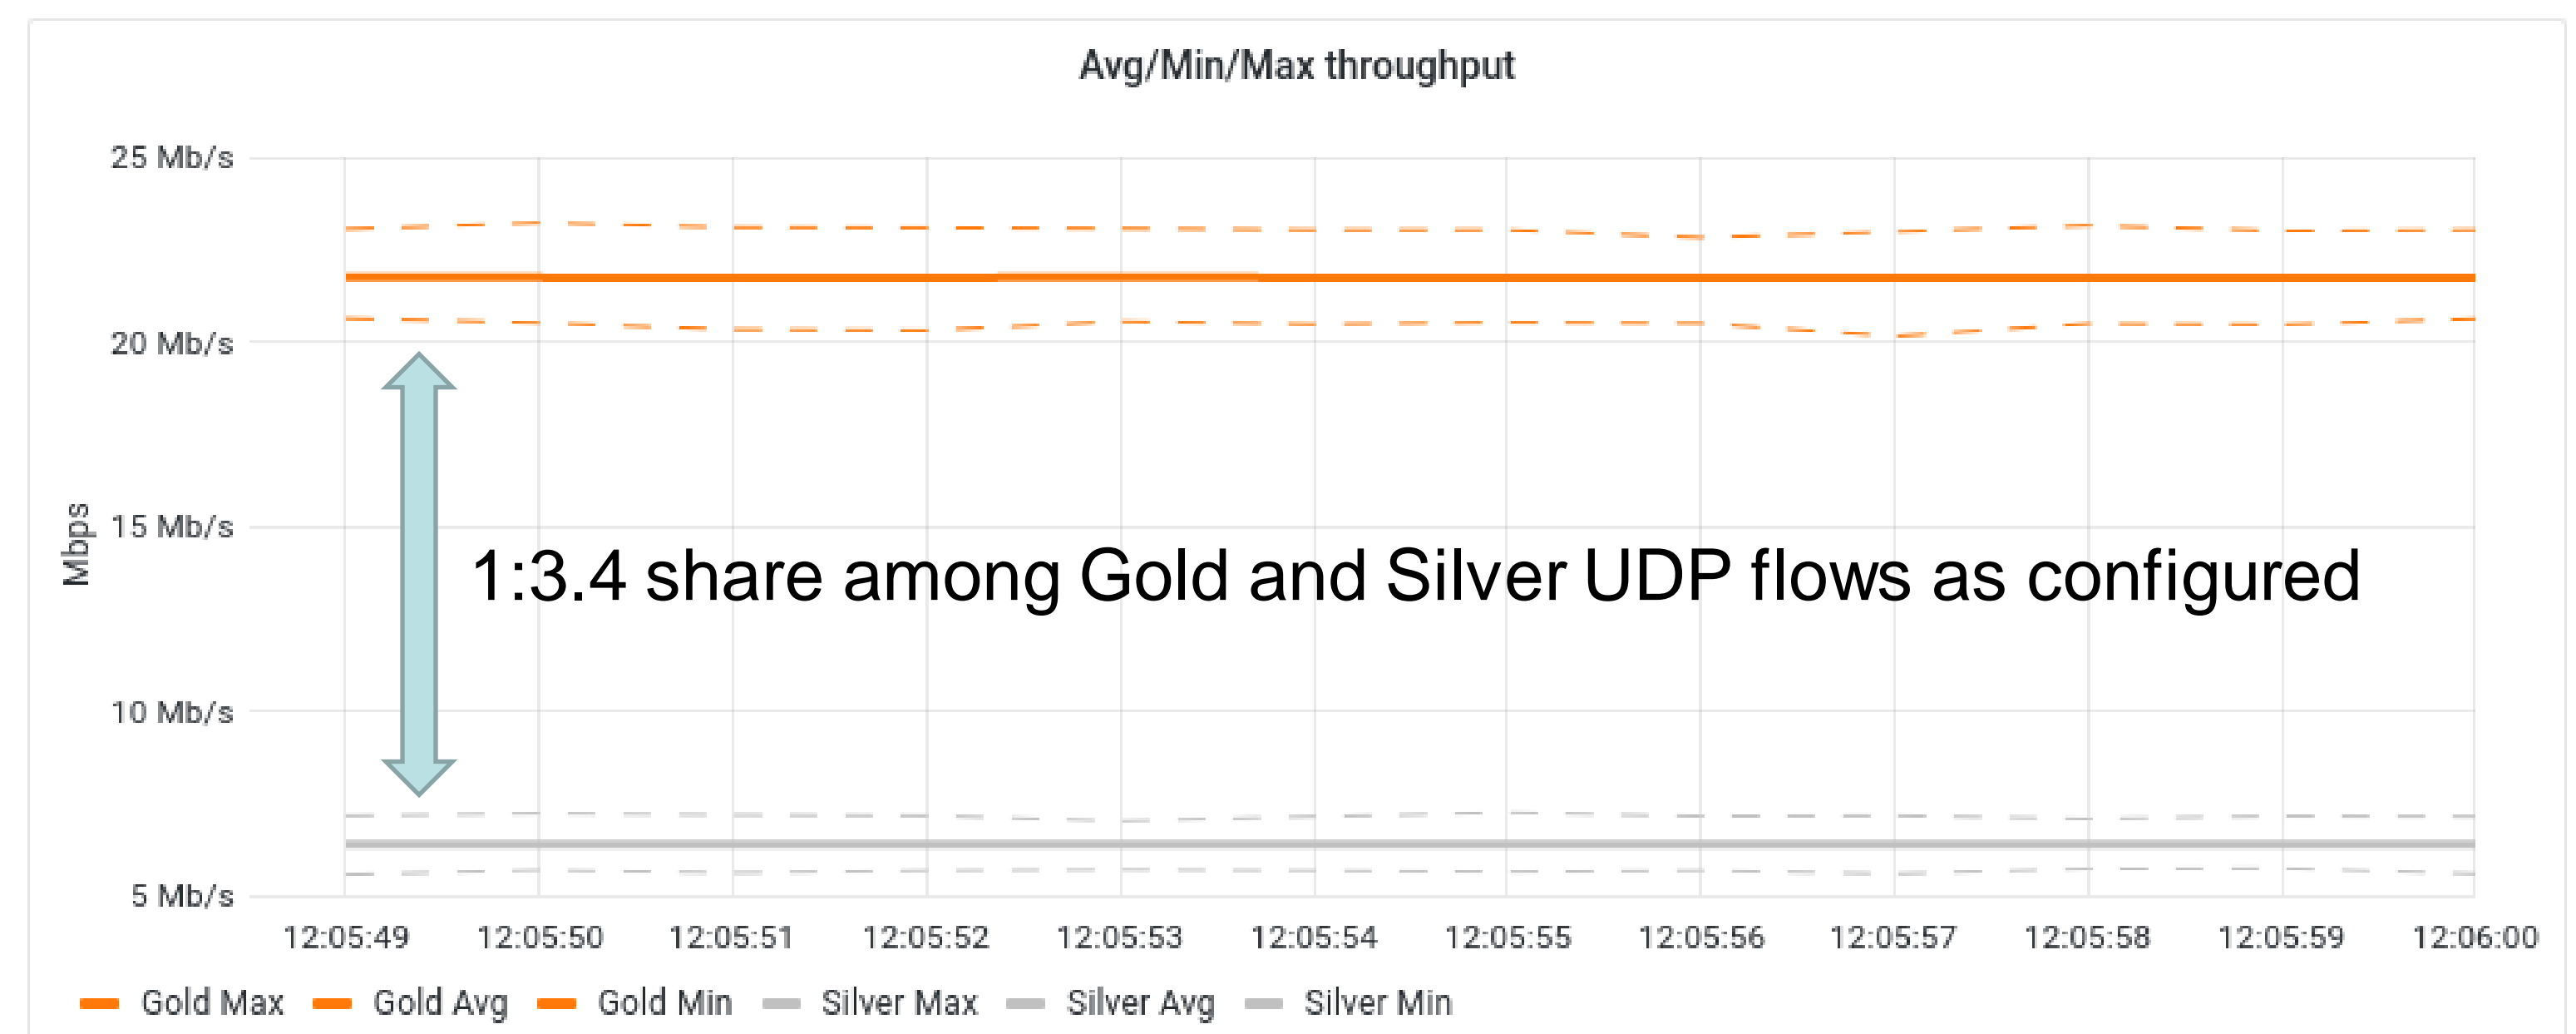

- The policy function is quantized logarithmically
- 1 marker instance marks the traffic of a single subscriber
- Logarithm tables to solve:
 - Calculate packet value from rate measurement
 - Random number generation between $[0, \text{rate}]$
 - Limited to $[0, 2^n - 1]$ $n \in \mathbb{N}$
- Data plane only implementation.
 - Policy functions are configured by the control plane
- We run 35000 marker instances in parallel



TESTBED SETUP



RESULTS



Contact: Szilveszter Nadas - szilveszter.nadas@ericsson.com, Sándor Laki – lakis@inf.elte.hu

ACM SIGCOMM 2022, August 22-26, 2022, Amsterdam, Netherlands

S. Laki and G. Gombos thank the support of the "Application Domain Specific Highly Reliable IT Solutions" project that has been implemented with the support provided from the National Research, Development and Innovation Fund of Hungary, financed under the Thematic Excellence Programme TKP2020-NKA-06 (National Challenges Subprogramme) funding scheme. The authors also thank the support of Ericsson.

