ERICSSON

STATELESS RESOURCE SHARING AND ATS

PUTTING TOGETHER THE BEST OF BREADS

IEEE P802.1Qcr, 2016-09-14, York, UK

Szilveszter Nádas, Balázs Varga (A), János Farkas

INTRODUCTION
SRS OVERVIEW AND PRESENTATION GOAL

- Stateless Resource Sharing (SRS) via per packet value (PPV) marking
  - Basic concept was presented in Budapest
  - [Link-to-contribution](http://www.ieee802.org/1/files/public/docs2016/cr-varga-srs-ppv-0526-v02.pdf)
    http://www.ieee802.org/1/files/public/docs2016/cr-varga-srs-ppv-0526-v02.pdf

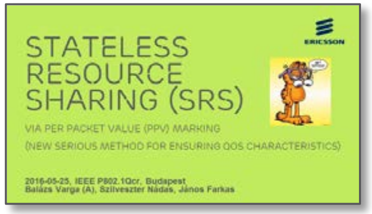

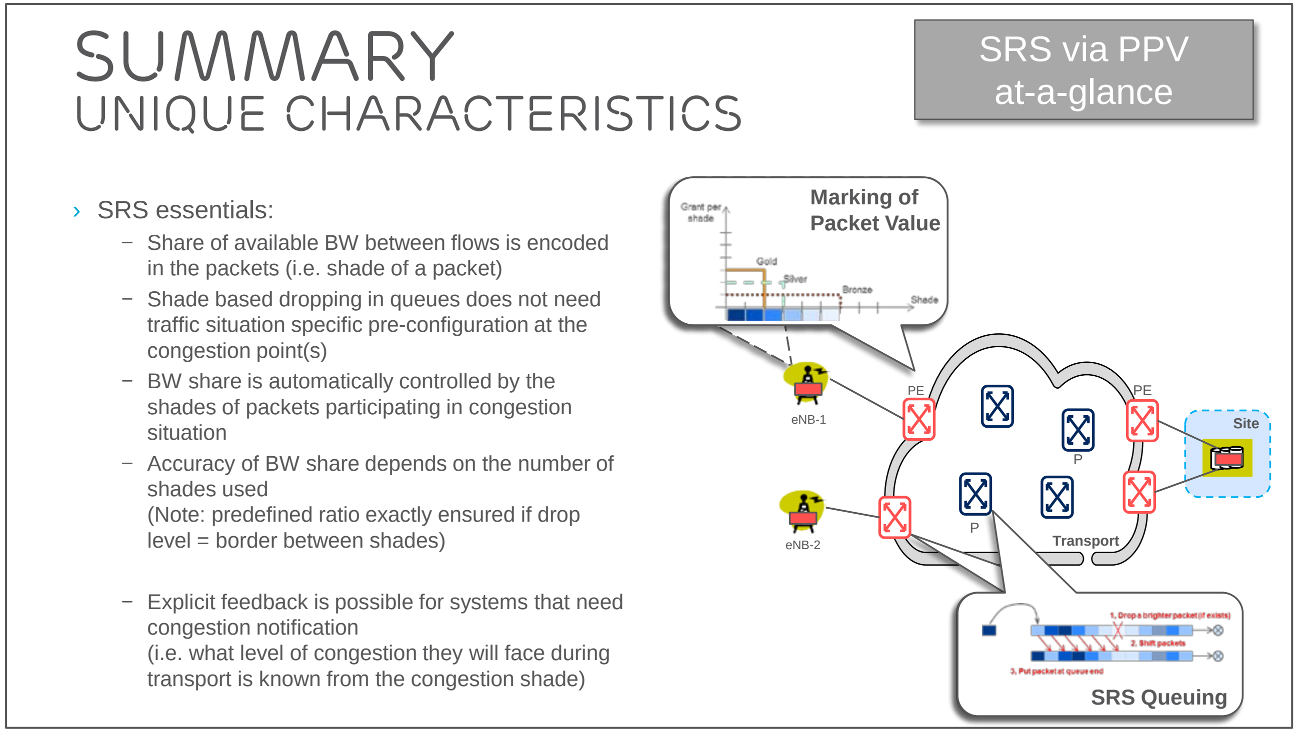

SUMMARY
UNIQUE CHARACTERISTICS

SRS via PPV at-a-glance

- SRS essentials:
  - Share of available BW between flows is encoded in the packets (i.e. shade of a packet)
  - Shade based dropping in queues does not need traffic situation specific pre-configuration at the congestion point(s)
  - BW share is automatically controlled by the shades of packets participating in congestion situation
  - Accuracy of BW share depends on the number of shades used (Note: predefined ratio exactly ensured if drop level = border between shades)
  - Explicit feedback is possible for systems that need congestion notification (i.e. what level of congestion they will face during transport is known from the congestion shade)

- Goal of this presentation to show
  - Combination of SRS and ATS
  - Achievable advantages by such a combination

SRS PROVIDES
ZERO CONGESTION LOSS AND DETERMINISTIC LATENCY

- SRS can achieve zero congestion loss and deterministic latency
  - Congestion results in packet drop:
    - Darker packets kick-off brighter ones
  - Congestion level:
    - Can be defined by a „congestion shade"
- **Note:** All traffic having darker shades than the "Congestion Threshold Value" will experience zero loss and deterministic delay

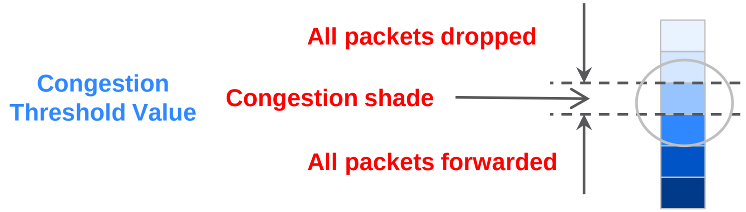

- SRS shortcomings
  - Congestion Threshold Value is not set in advance, it is rather the result of actual load and bottleneck capacity
  - SRS (itself) does not protect against bad behaving nodes

- **BUT:** Congestion Threshold Value can be dimensioned by resource allocation and worst case delay calculation (by e.g. ATS)

SRS ADD-ONS
DETERMINISTIC DELAY FOR LOSSLESS & LOSSY TRAFFIC

- It is possible to extend the ATS scenario with traffic that has the same delay guarantee as “guaranteed” TSN traffic, but some loss is allowed, i.e., it is lossy
- What does it mean in practice?
  - Allows reservation of less resources for ATS flows which have a loss tolerant component but need in-order-delivery for all packets of a data-flow
  - Additional loss tolerant flows that require the same deterministic delay can be served easily
  - Other scenarios may also exist …
- The positive effect
  - It can highly increase link utilization when some flows do not use their reservation all the time
  - Under the prerequisite that (some) lossy traffic has delay guarantee too
- This works hop-by-hop, not just for flows following the same path

SRS+ATS USE CASES

- Alarm traffic
  - There is a wide range of different importance from good to know to critical
    - Some alarms are loss critical, some also delay critical
  - It is hard to dimension for, because it is rare and bursty
  - Proposal: Control the priority of the alarm traffic by SRS. Allow less critical TSN traffic use this capacity when (typically) available.
- Layered video traffic
  - Key frames and other critical frames are guarantees
  - Other enhancement layers that are good to have are non-guaranted (can be lost, but delay is still bounded)

ATS AND ITS SUB-QUEUES
REMINDER

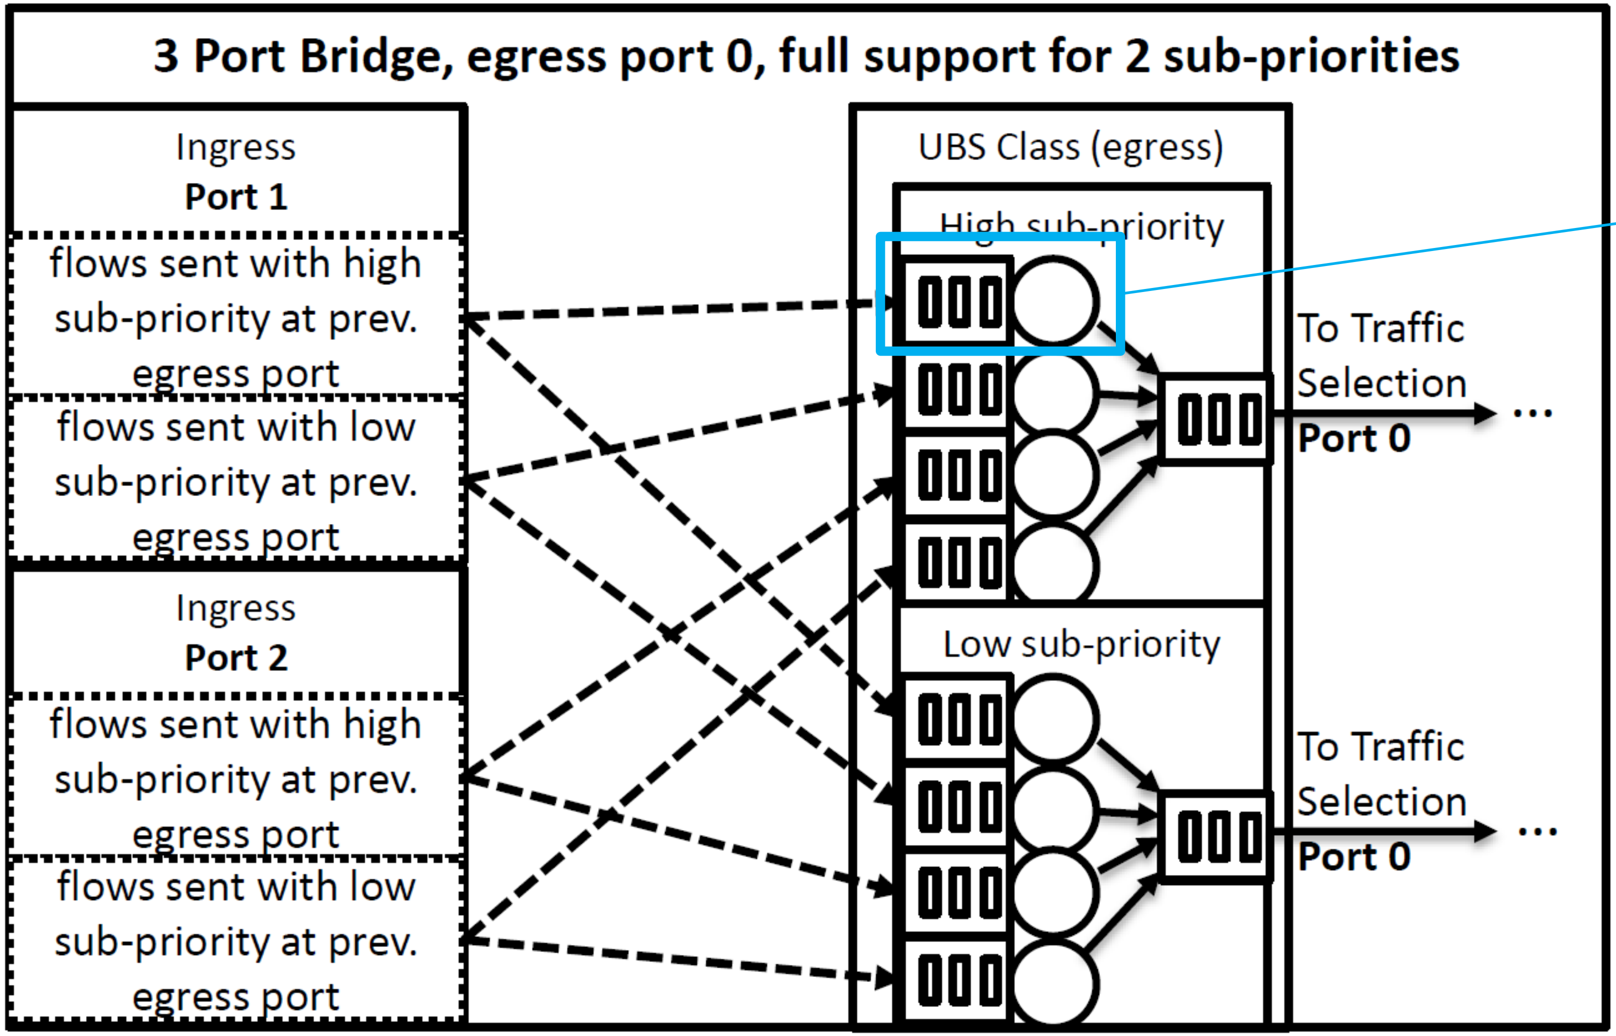

- We propose to extend the handling for sub-queues
- Sub-queue handling is detailed on following slides
- Number of the following is kept
  - Sub-priorities
  - Sub-queues
- State is the same as for ATS*

* Except tuning variables for *f* and *g* (few per sub-queues)

Taken form slide 10 of *http://www.ieee802.org/1/files/public/docs2015/new-tsn-specht-ubs-queues-0521-v0.pdf*

ARCHITECTURE
PUTTING IT ALL TOGETHER

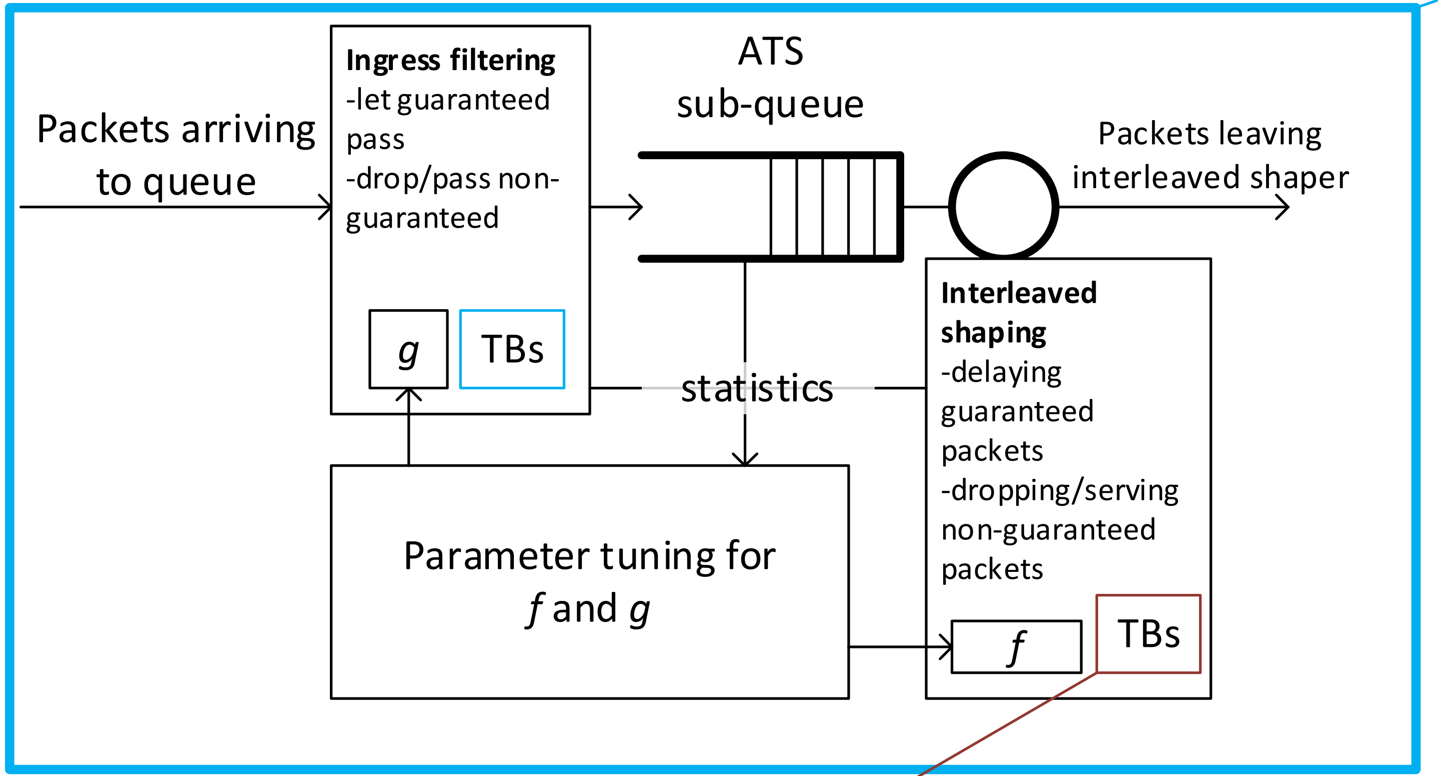

- l is the length of the packet
- v is the shade/value of the packet
- $l' \geq l$ is an effective packet length
  - calculated by *f* or *g* functions
- Non-guaranteed eligible if at least l' tokens in a nG TB space
  - E.g. $\hat{b}_i + l'$ tokens in a sub-shaper per flow bucket TB_i (till $\hat{b}_i$ reserved for guaranteed)

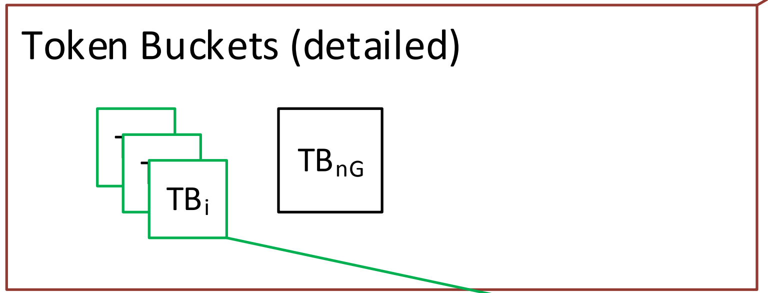

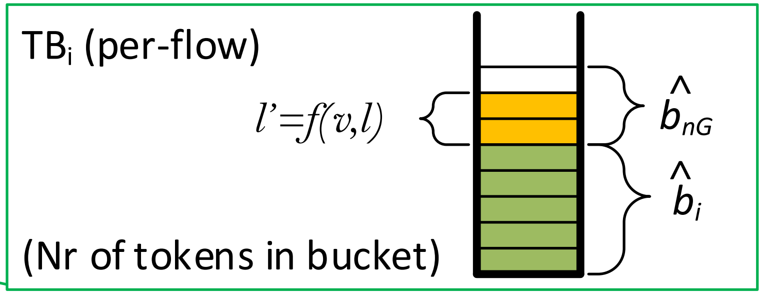

ARCHITECTURE ALT2
A SINGLE EXCESS BUCKET FOR NON-GUARANTEED

- A single excess bucket for non-Guaranteed traffic is also a possibility
  - All/selected sub-shaper per flow buckets overflow into this bucket

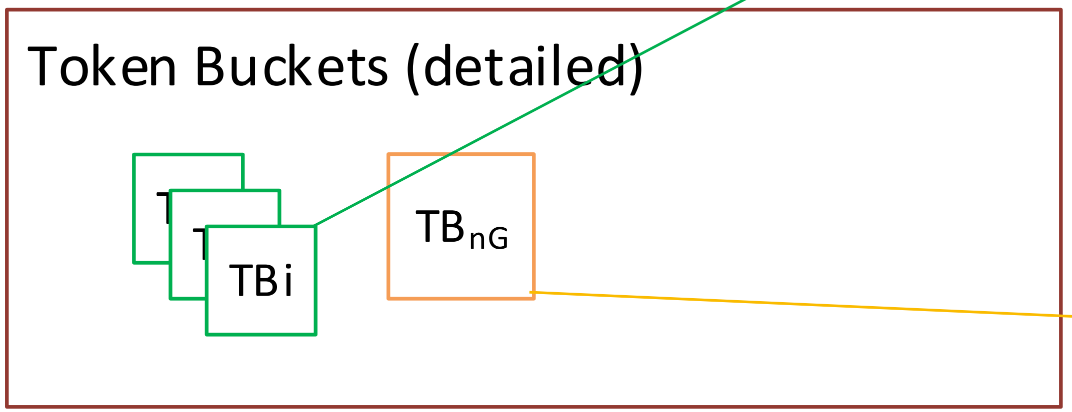

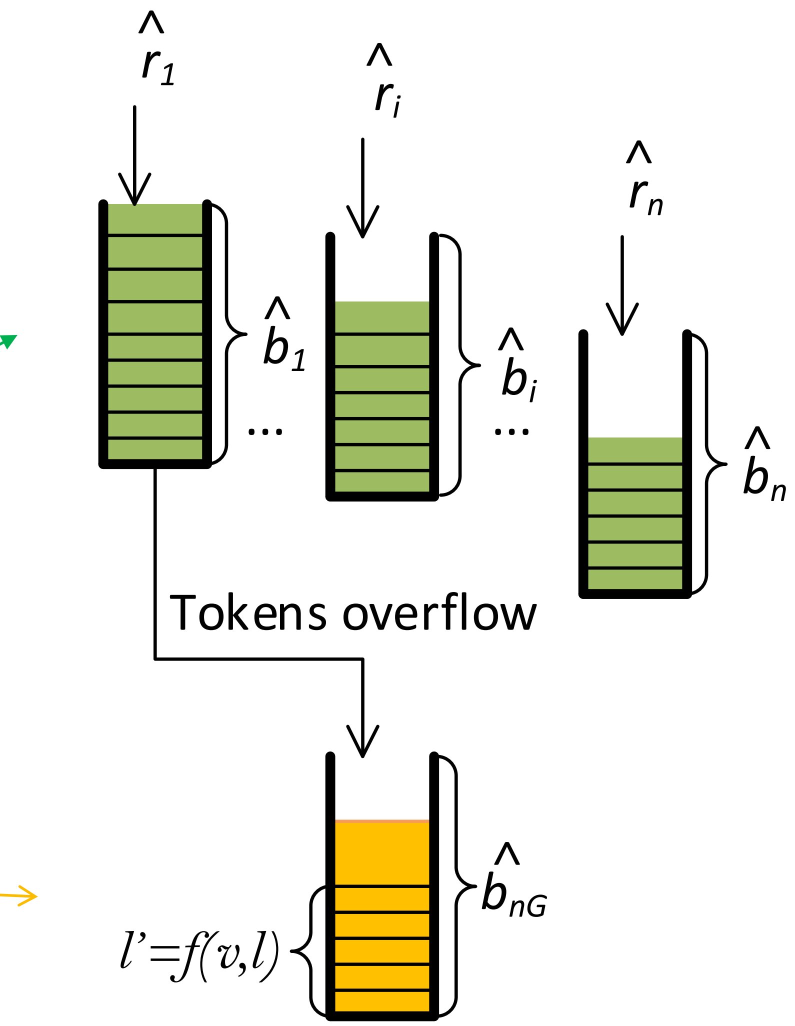

OUTCOME OF COMBINATION
LOSS VS. THROUGHPUT

- A slightly larger bucket size
- The possibility of forwarding non-guaranteed packets results in a larger achievable flow rate
- The packet value determines
  - Whether or not a packet is guaranteed
  - Whether a non-guaranteed packet is dropped or forwarded (note: there can be more important and less important non-guaranteed packets)
- The size of the larger bucket is chosen based on a compromise between (1) delay, (2) allowed total guaranteed rate, and (3) allowing excess traffic
  - The per hop delay remains bounded

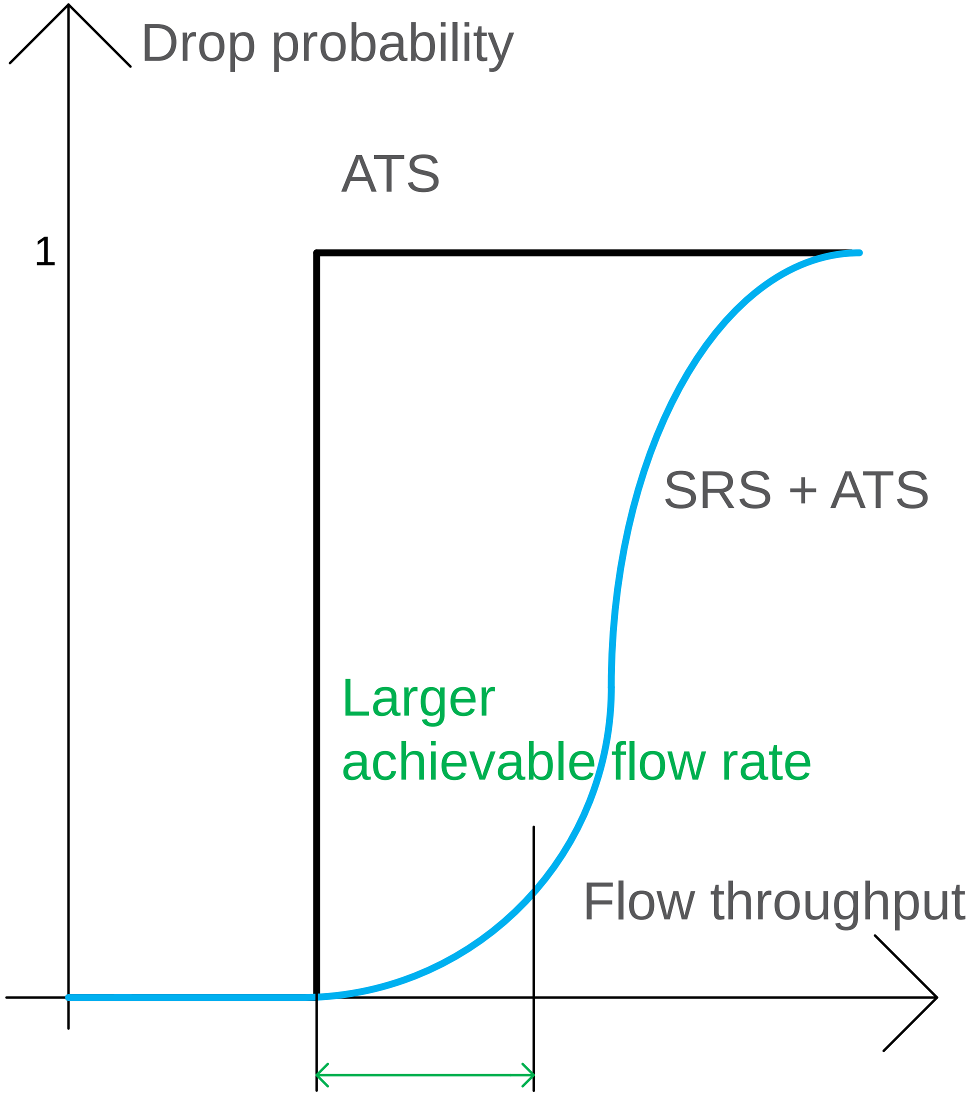

SUMMARY
FURTHER WORK

- This is an initial proposal to evolve ATS further
- Delay model and math to be discussed and verified in detail
- Further use cases

ERICSSON

CHANGED COMPONENTS
FOR COMBINING SRS AND ATS

- Need to be able to identify additional traffic (at all hops)

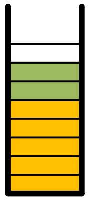

- Slightly larger bucket sizes are needed to allow excess traffic

- Interleaved shaping has to be able to
  - drop excess packets
  - read packet values, and based on that, influence whether a packet is dropped or not

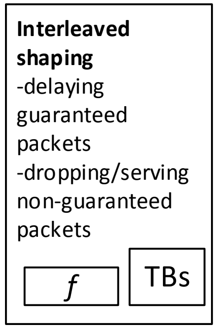

- Ingress filtering for excess packets is needed in order to avoid flooding the queues
  - Slightly larger buffers are needed

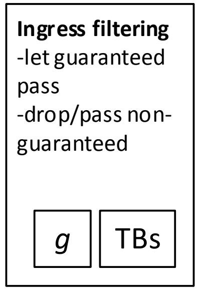

- Statistics from all boxes are needed to tune packet value aware dropping