Bridge-Local Guaranteed Latency with Strict Priority Scheduling

Alexej Grigorjew – March 02, 2020

comnet.informatik.uni-wuerzburg.de

Introduction – Distributed Admission Control

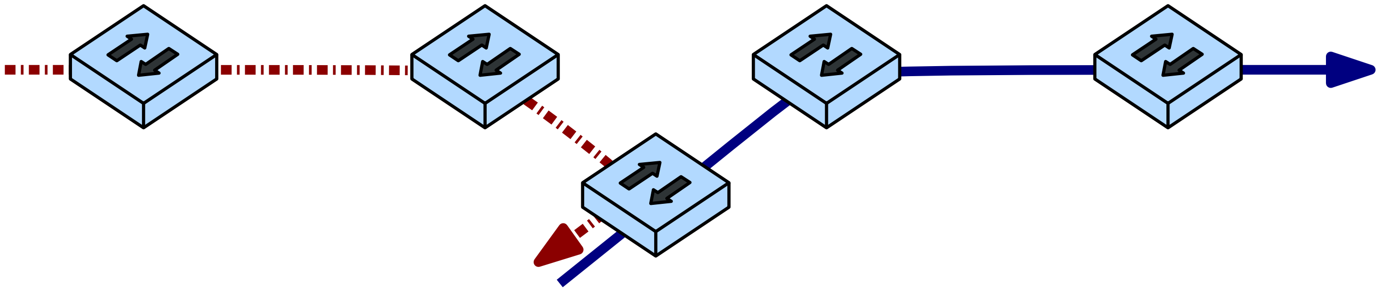

Introduction – Distributed Admission Control

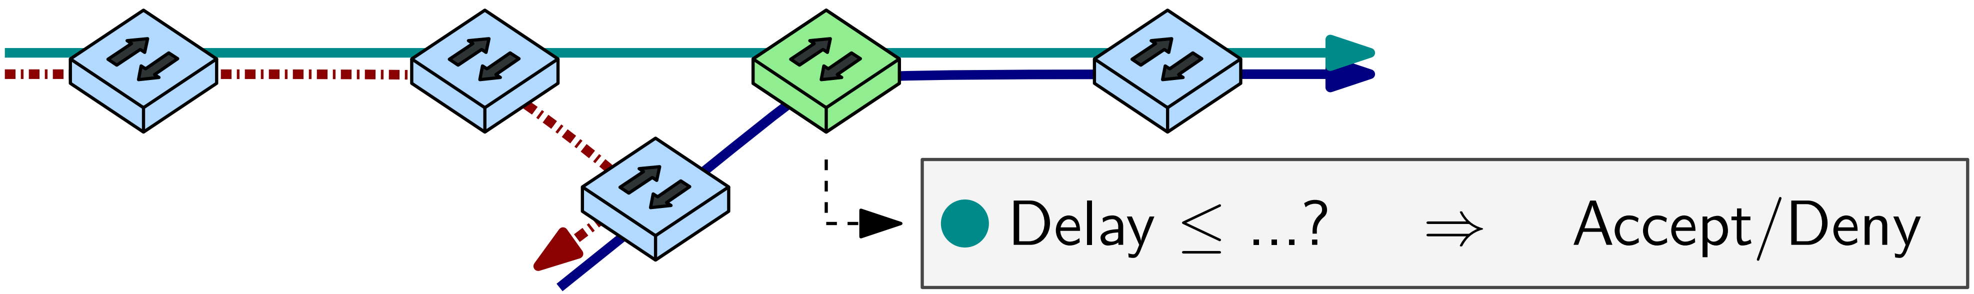

Introduction – Distributed Admission Control

Transmission Selection

SP CBS ...
ATS CQF

Traffic Specification (SRP, RAP)

MaxFrameSize Interval
MaxFramesPerInterval ...

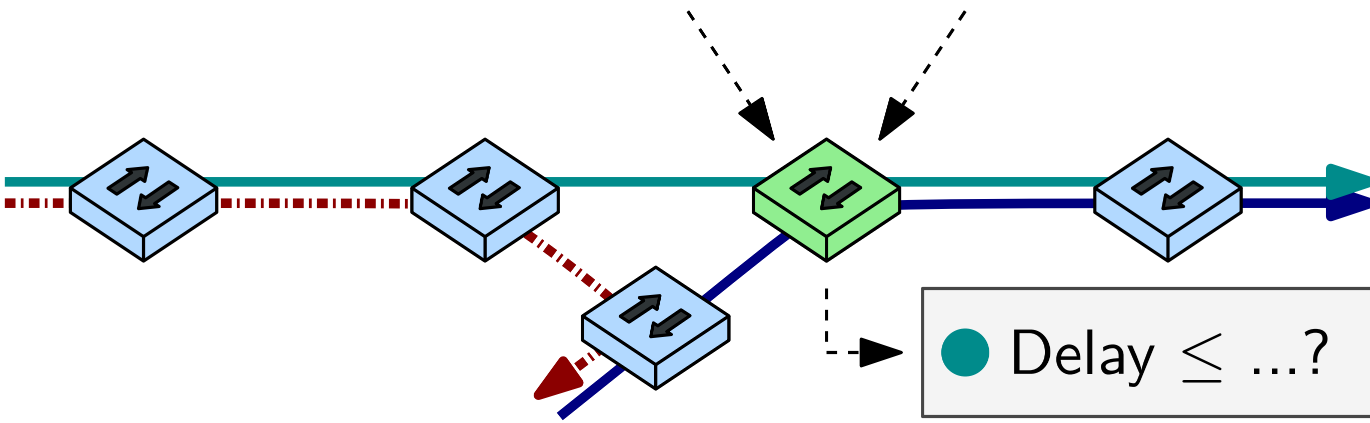

● Delay $\leq$ ...? $\Rightarrow$ Accept/Deny

Introduction – Distributed Admission Control

Transmission Selection

SP CBS ...
ATS CQF

Traffic Specification (SRP, RAP)

MaxFrameSize Interval
MaxFramesPerInterval ...

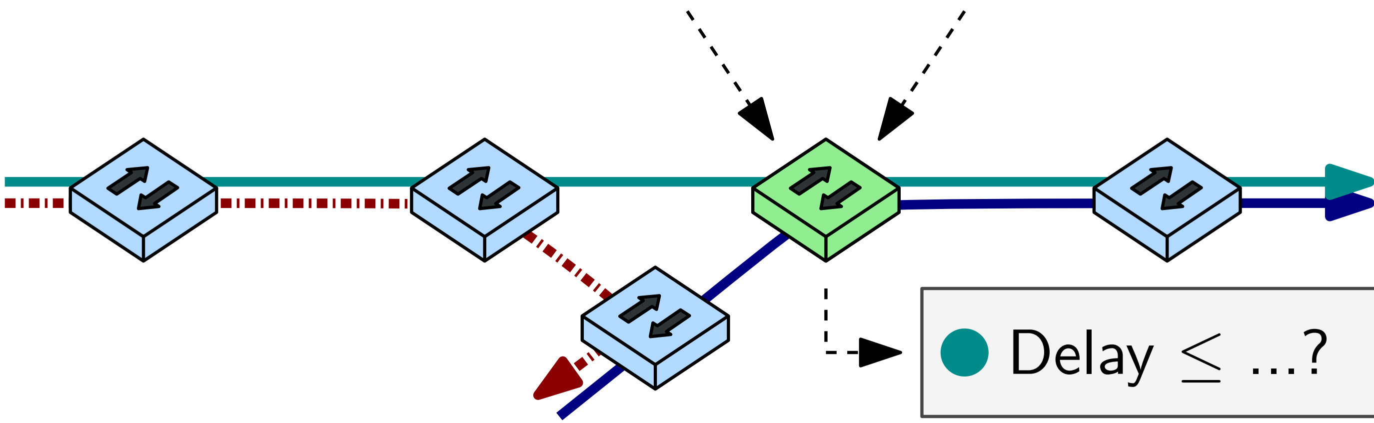

Delay $\leq$ ...? $\Rightarrow$ Accept/Deny

Desired Features:

- Computationally feasible
- Do not require global information (from ●)
- Support brownfield installations $\Rightarrow$ SP

Table of Contents

Preliminaries:

- Switch delay model
- Assumptions and constraints
  - Talker characteristics
  - Switch characteristics

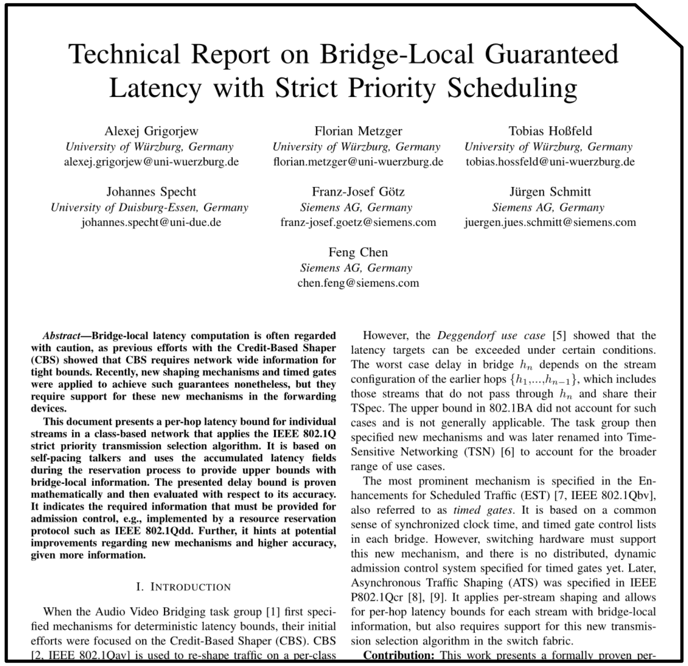

Technical Report on Bridge-Local Guaranteed Latency with Strict Priority Scheduling

Alexej Grigorjew, *University of Würzburg, Germany*, alexej.grigorjew@uni-wuerzburg.de
Florian Metzger, *University of Würzburg, Germany*, florian.metzger@uni-wuerzburg.de
Tobias Hoßfeld, *University of Würzburg, Germany*, tobias.hossfeld@uni-wuerzburg.de
Johannes Specht, *University of Duisburg-Essen, Germany*, johannes.specht@uni-due.de
Franz-Josef Götz, *Siemens AG, Germany*, franz-josef.goetz@siemens.com
Jürgen Schmitt, *Siemens AG, Germany*, juergen.jues.schmitt@siemens.com
Feng Chen, *Siemens AG, Germany*, chen.feng@siemens.com

***Abstract*—Bridge-local latency computation is often regarded with caution, as previous efforts with the Credit-Based Shaper (CBS) showed that CBS requires network wide information for tight bounds. Recently, new shaping mechanisms and timed gates were applied to achieve such guarantees nonetheless, but they require support for these new mechanisms in the forwarding devices.**

This document presents a per-hop latency bound for individual streams in a class-based network that applies the IEEE 802.1Q strict priority transmission selection algorithm. It is based on self-pacing talkers and uses the accumulated latency fields during the reservation process to provide upper bounds with bridge-local information. The presented delay bound is proven mathematically and then evaluated with respect to its accuracy. It indicates the required information that must be provided for admission control, e.g., implemented by a resource reservation protocol such as IEEE 802.1Qdd. Further, it hints at potential improvements regarding new mechanisms and higher accuracy, given more information.

I. Introduction

When the Audio Video Bridging task group [1] first specified mechanisms for deterministic latency bounds, their initial efforts were focused on the Credit-Based Shaper (CBS). CBS [2, IEEE 802.1Qav] is used to re-shape traffic on a per-class

However, the *Deggendorf use case* [5] showed that the latency targets can be exceeded under certain conditions. The worst case delay in bridge h_n depends on the stream configuration of the earlier hops $\{h_1,...,h_{n-1}\}$, which includes those streams that do not pass through h_n and share their TSpec. The upper bound in 802.1BA did not account for such cases and is not generally applicable. The task group then specified new mechanisms and was later renamed into Time-Sensitive Networking (TSN) [6] to account for the broader range of use cases.

The most prominent mechanism is specified in the Enhancements for Scheduled Traffic (EST) [7, IEEE 802.1Qbv], also referred to as *timed gates*. It is based on a common sense of synchronized clock time, and timed gate control lists in each bridge. However, switching hardware must support this new mechanism, and there is no distributed, dynamic admission control system specified for timed gates yet. Later, Asynchronous Traffic Shaping (ATS) was specified in IEEE P802.1Qcr [8], [9]. It applies per-stream shaping and allows for per-hop latency bounds for each stream with bridge-local information, but also requires support for this new transmission selection algorithm in the switch fabric.

Contribution: This work presents a formally proven per-

https://nbn-resolving.org/urn:nbn:de:bvb:20-opus-198310

Contribution:

- Overview of required information from the Resource Allocation Protocol (RAP)
- Proven per-hop latency bound for Strict Priority (SP) transmission selection with only bridge-local information
- Initial evaluation of network capacity for an admission control system using this bound

Preliminaries

Switch delay models, assumptions and constraints

Switch Delay Model

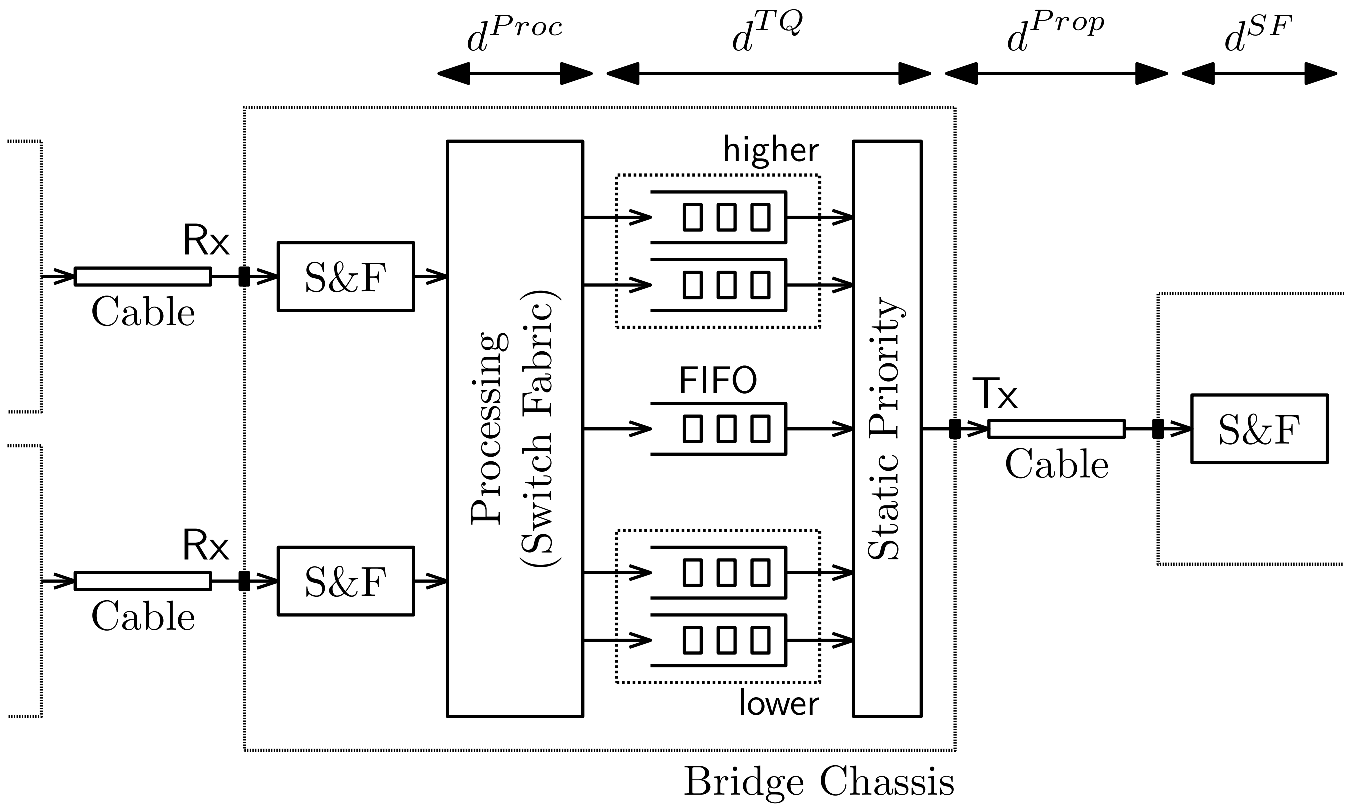

Switch Delay Model

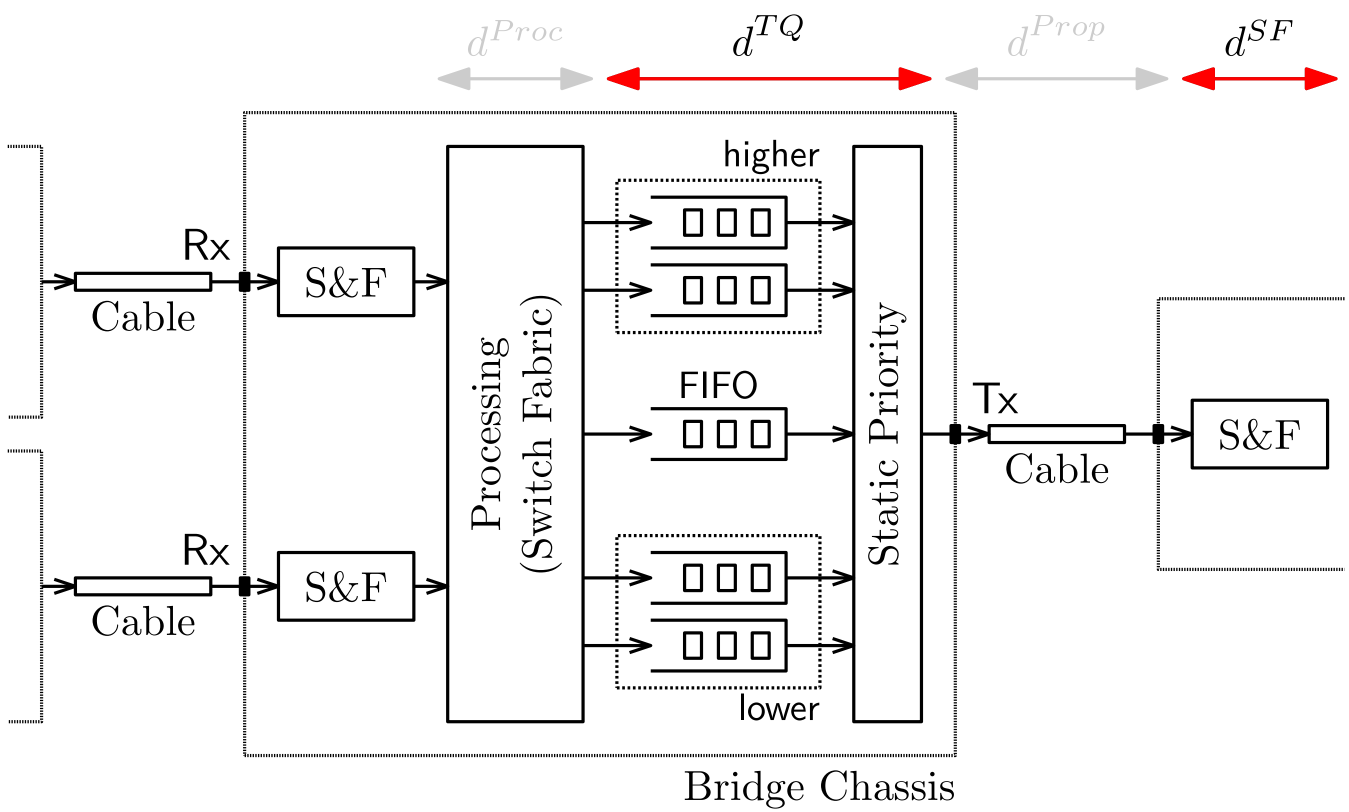

- Processing delay d^{Proc} is device specific and not considered
- Propagation delay d^{Prop} is bounded by max cable length
- Upper bound for $d^{TQ} + d^{SF}$ desired (queuing and transmission delay)

Assumptions and Constraints – Talkers

1. Frames of stream i do not exceed their **max frame size** $\hat{\ell}_i$ and **min frame size** $\check{\ell}_i$.

Assumptions and Constraints – Talkers

1. Frames of stream i do not exceed their **max frame size** $\hat{\ell}_i$ and **min frame size** $\check{\ell}_i$.

2. Talkers pace their traffic according to a **burst size** b_i and a **burst interval** τ_i. For any point t in time, the traffic sent by stream i in the time interval $[t, t + \tau_i]$ may not exceed b_i.

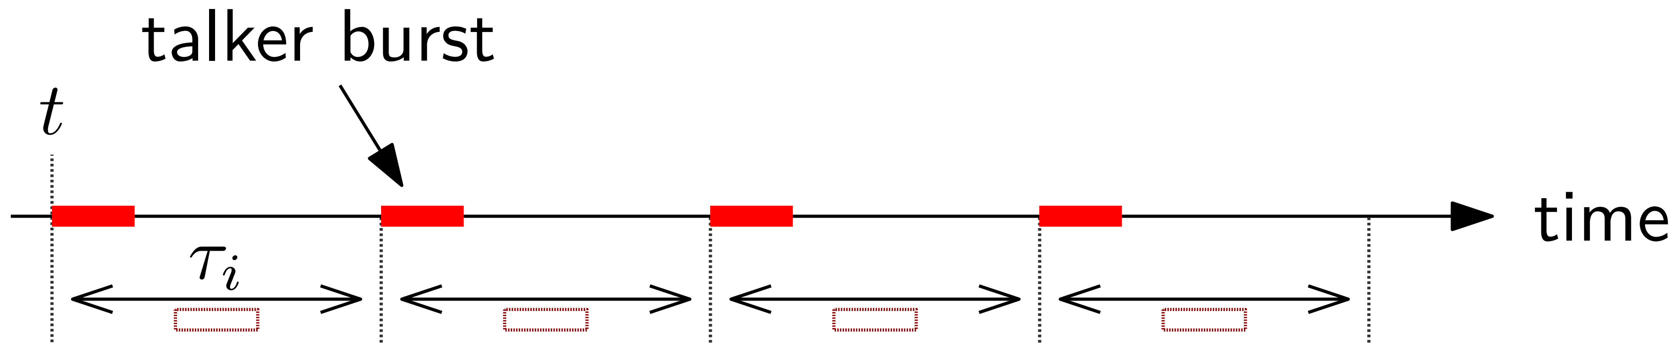

Assumptions and Constraints – Talkers

1. Frames of stream i do not exceed their **max frame size** $\hat{\ell}_i$ and **min frame size** $\check{\ell}_i$.

2. Talkers pace their traffic according to a **burst size** b_i and a **burst interval** τ_i. For any point t in time, the traffic sent by stream i in the time interval $[t, t + \tau_i]$ may not exceed b_i.

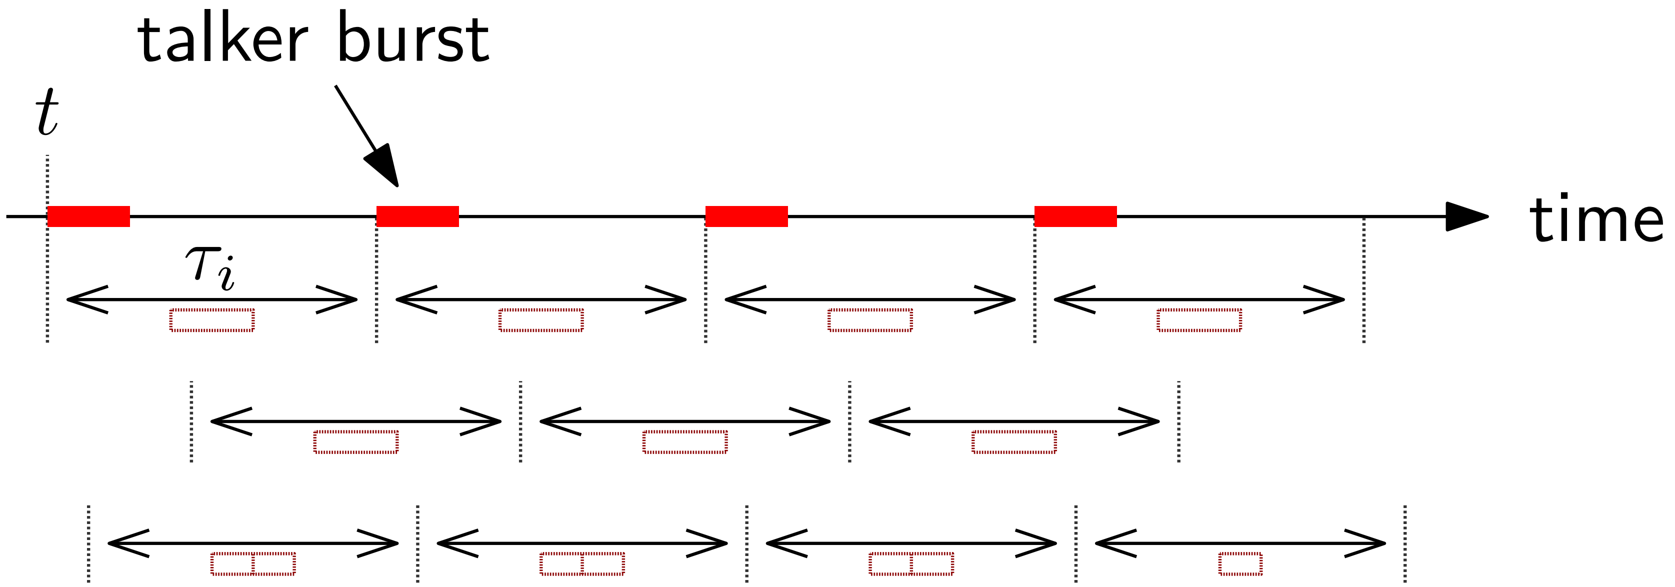

Assumptions and Constraints – Talkers

1. Frames of stream i do not exceed their **max frame size** $\hat{\ell}_i$ and **min frame size** $\check{\ell}_i$.

2. Talkers pace their traffic according to a **burst size** b_i and a **burst interval** τ_i. For any point t in time, the traffic sent by stream i in the time interval $[t, t+\tau_i]$ may not exceed b_i.

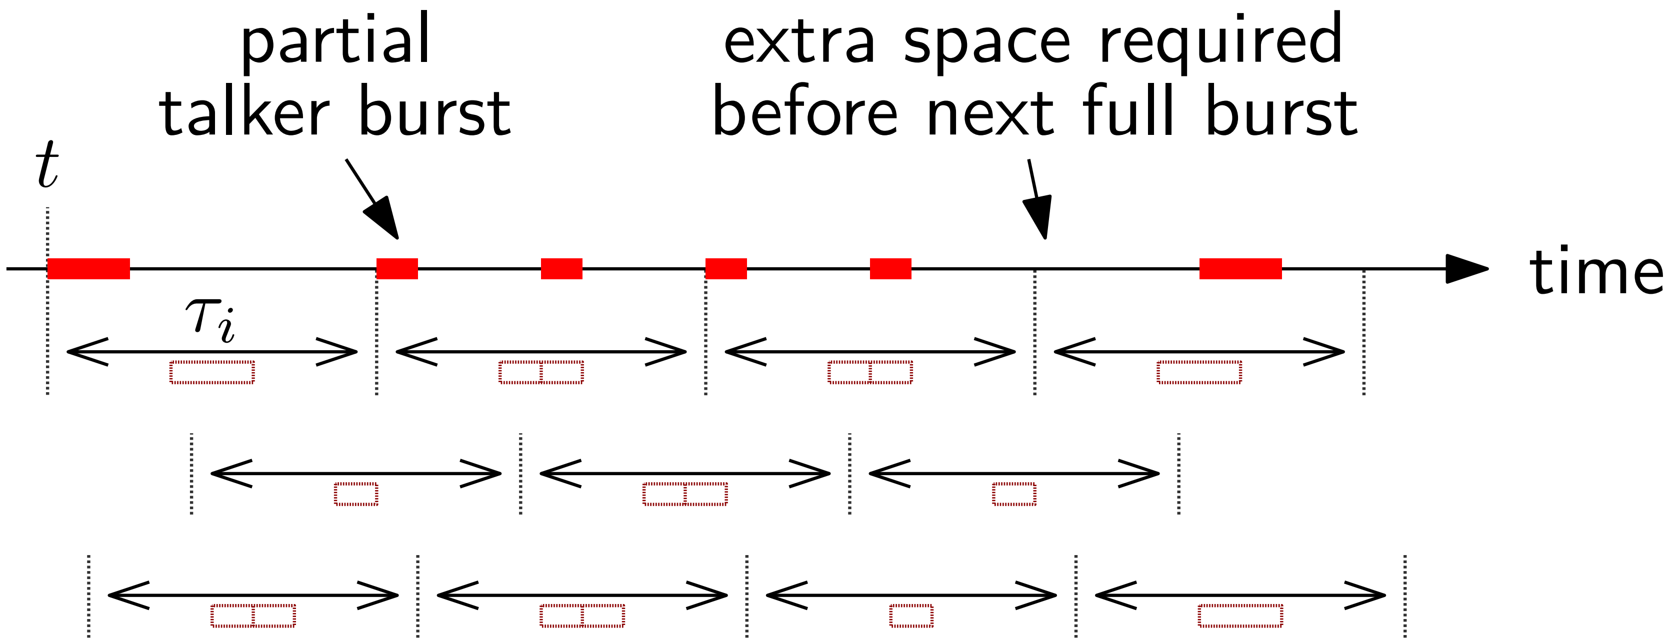

Assumptions and Constraints – Talkers

1. Frames of stream i do not exceed their **max frame size** $\hat{\ell}_i$ and **min frame size** $\check{\ell}_i$.

2. Talkers pace their traffic according to a **burst size** b_i and a **burst interval** τ_i. For any point t in time, the traffic sent by stream i in the time interval $[t, t + \tau_i]$ may not exceed b_i.

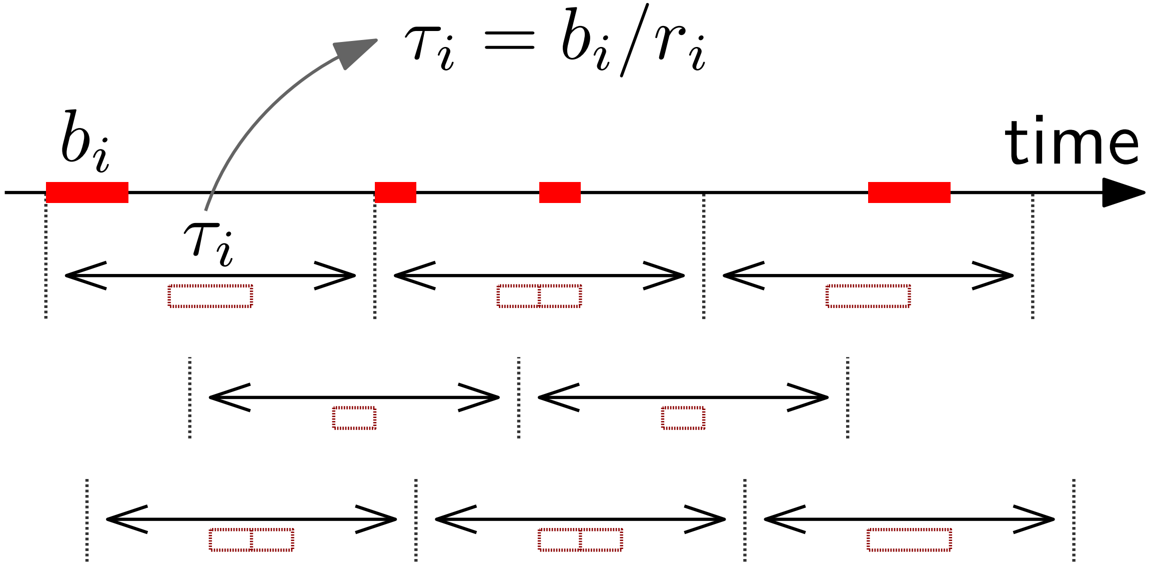

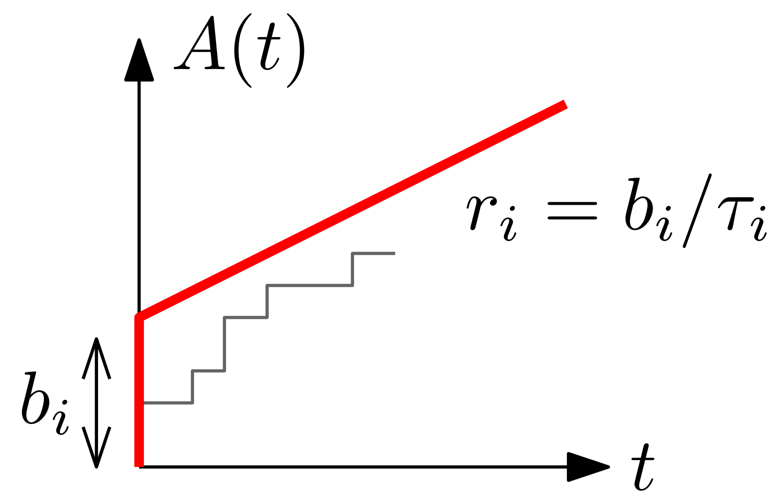

Assumptions and Constraints – Bridges

3. Bridges use IEEE 802.1Q **priority transmission selection**, i.e., frames with a higher traffic class are always selected for transmission before frames with lower traffic classes.

   (a) Within each traffic class, **FIFO** transmission selection is used.

   (b) **No shaping** mechanisms are used in any considered traffic class. The earliest frame of each class is always regarded eligible for transmission.

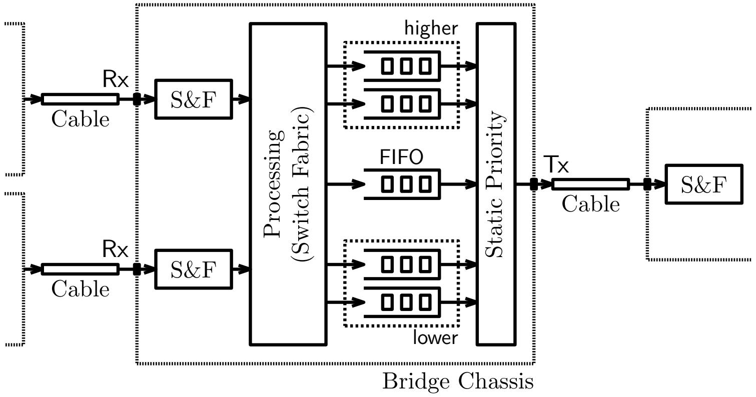

4. Each bridge h has a pre-configured maximum per-hop **delay guarantee** δ_p^h for each traffic class p.

   (a) **Admission control** prevents the deployment of new streams that would cause delay violations for any deployed stream.

Latency Bound

Required information, formula and reasoning

Required Information $\rightarrow$ TSpec

TSpec should include for stream i:

- Traffic class p_i
- Max frame size $\hat{\ell}_i$ (e.g., 1542 B)
- Min frame size $\check{\ell}_i$ (e.g., 84 B)
- Committed burst size b_i
- Burst interval τ_i

including preamble and IPG

Required Information → TSpec

TSpec should include for stream i:

- Traffic class p_i
- Max frame size $\hat{\ell}_i$ (e.g., 1542 B)
- Min frame size $\check{\ell}_i$ (e.g., 84 B)
- Committed burst size b_i
- Burst interval τ_i
- Accumulated max latency $accMaxD_i^{h_k}$
- Accumulated min latency $accMinD_i^{h_k}$

including preamble and IPG

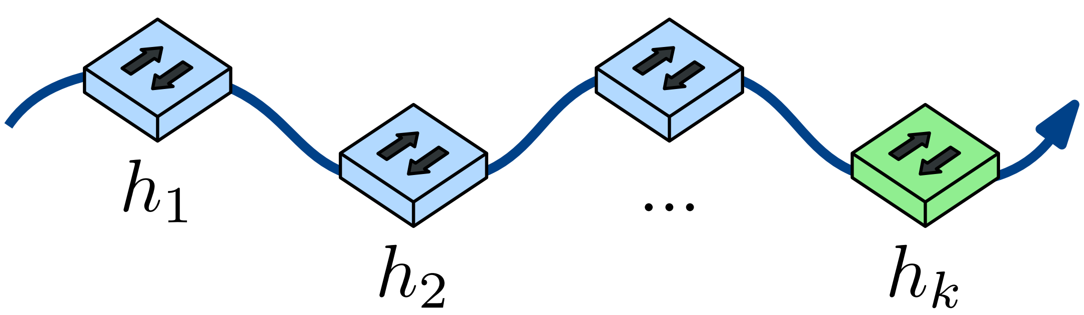

$$accMaxD_i^{h_k} = \sum_{j=1}^{k} \delta_{p_i}^{h_j}$$

$$accMinD_i^{h_k} = \sum_{j=1}^{k} \frac{\check{\ell}_i}{link\ speed_{h_j}}$$

Latency Bound

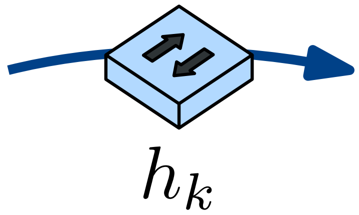

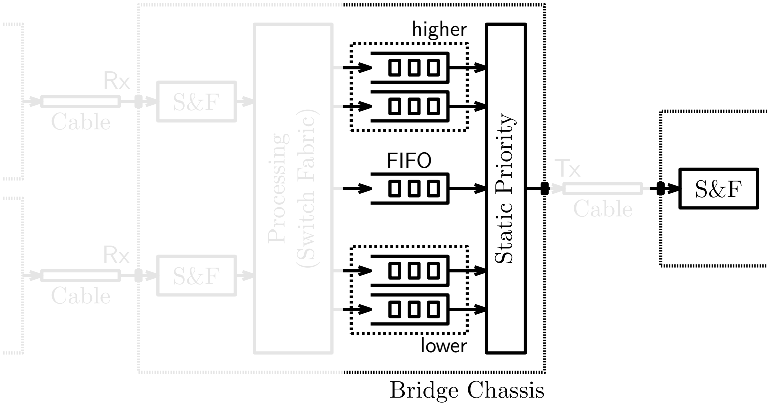

p_i	traffic class
$\hat{\ell}_i$	max frame size
$\check{\ell}_i$	min frame size
b_i	burst size
τ_i	burst interval
$\delta_{p_i}^{h_k}$	delay guarantee
r	link speed
$\mathcal{S}$	set of all streams

- Worst case latency of stream i at bridge h_k is bounded by:

$$d_i^{TQ,SF} \leq \sum_{\{x \in \mathcal{S} | p_x > p_i\}} y_{i,x} b_x / r \; + \sum_{\{x \in \mathcal{S} | p_x = p_i\}} z_x b_x / r \; + \max_{\{x \in \mathcal{S} | p_x < p_i\}} \hat{\ell}_x / r$$

Latency Bound

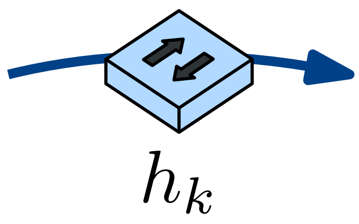

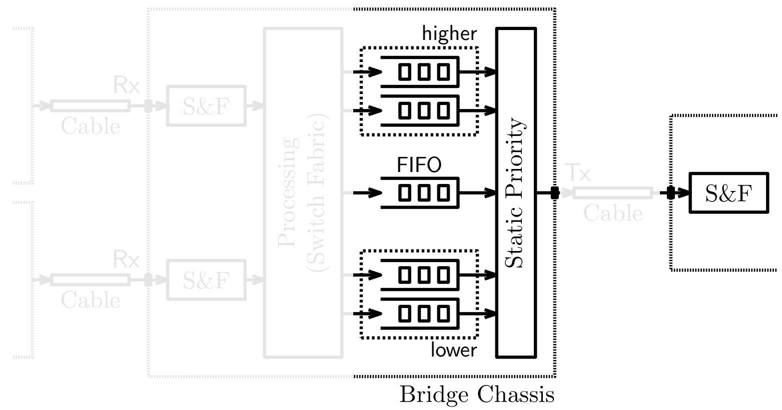

p_i	traffic class
$\hat{\ell}_i$	max frame size
$\check{\ell}_i$	min frame size
b_i	burst size
τ_i	burst interval
$\delta_{p_i}^{h_k}$	delay guarantee
r	link speed
$\mathcal{S}$	set of all streams

- Worst case latency of stream i at bridge h_k is bounded by:

$$d_i^{TQ,SF} \leq \sum_{\{x \in \mathcal{S} | p_x > p_i\}} y_{i,x} b_x / r \; + \sum_{\{x \in \mathcal{S} | p_x = p_i\}} z_x b_x / r \; + \max_{\{x \in \mathcal{S} | p_x < p_i\}} \hat{\ell}_x / r$$

number of bursts from interfering streams

$$y_{i,x} \geq \left\lceil \frac{accMaxD_x^{h_k} - accMinD_x^{h_{k-1}} + \delta_{p_i}^{h_k}}{\tau_x} \right\rceil$$

$$z_x \geq \left\lceil \frac{accMaxD_x^{h_k} - accMinD_x^{h_{k-1}}}{\tau_x} \right\rceil$$

Reasoning – Residence Times of Frames in TQ

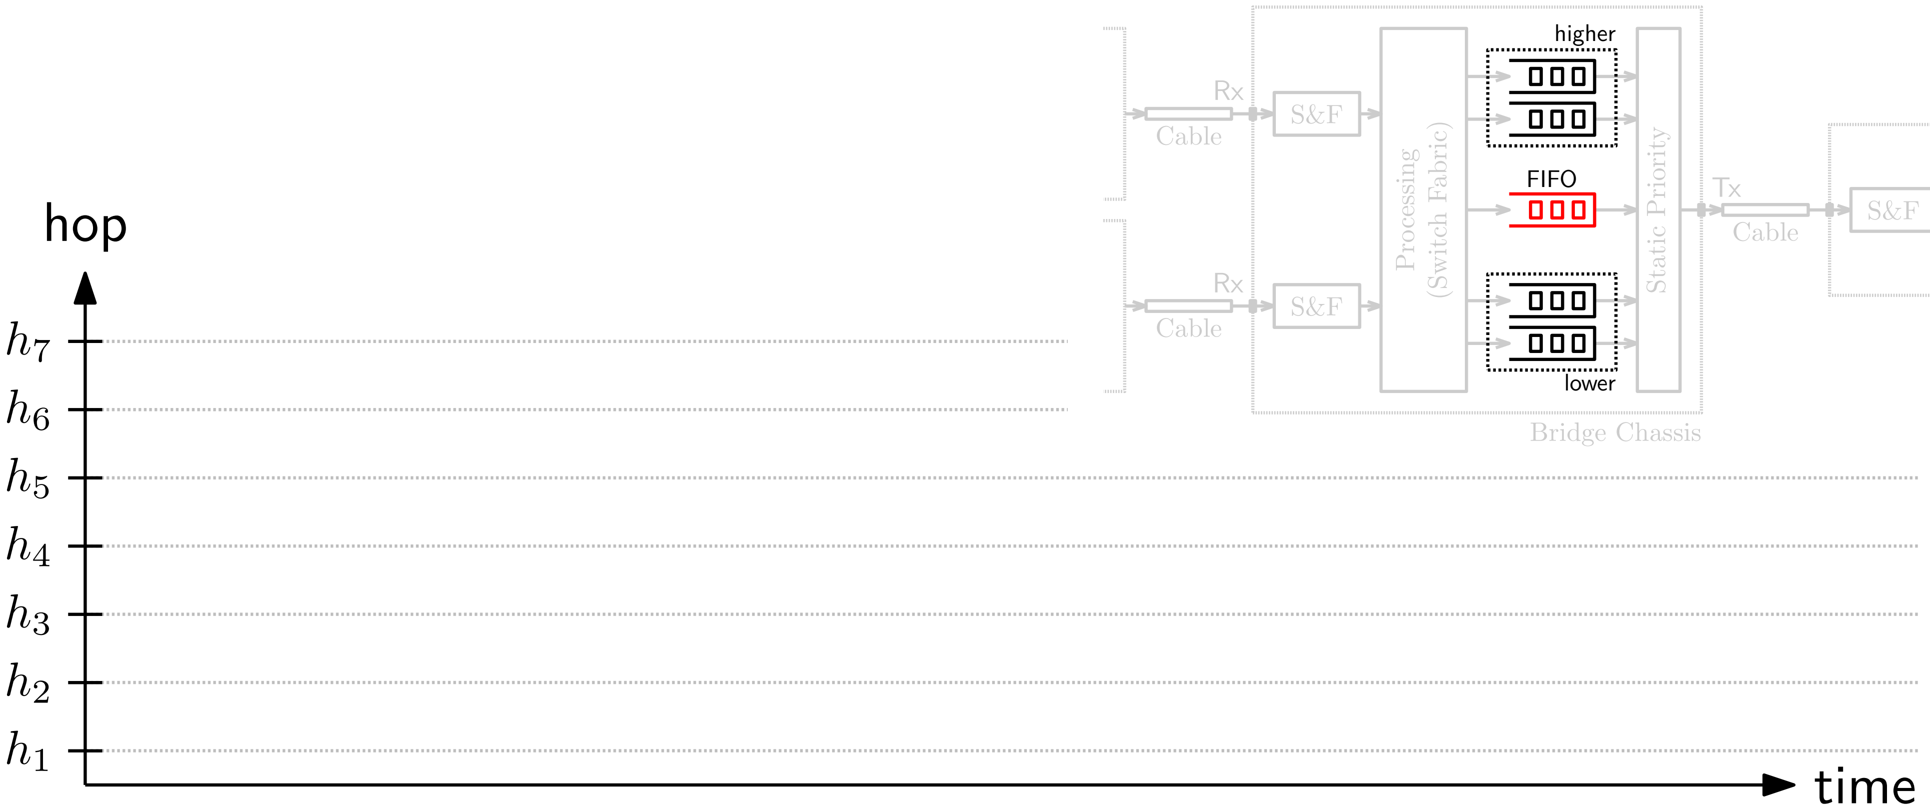

Reasoning – Residence Times of Frames in TQ

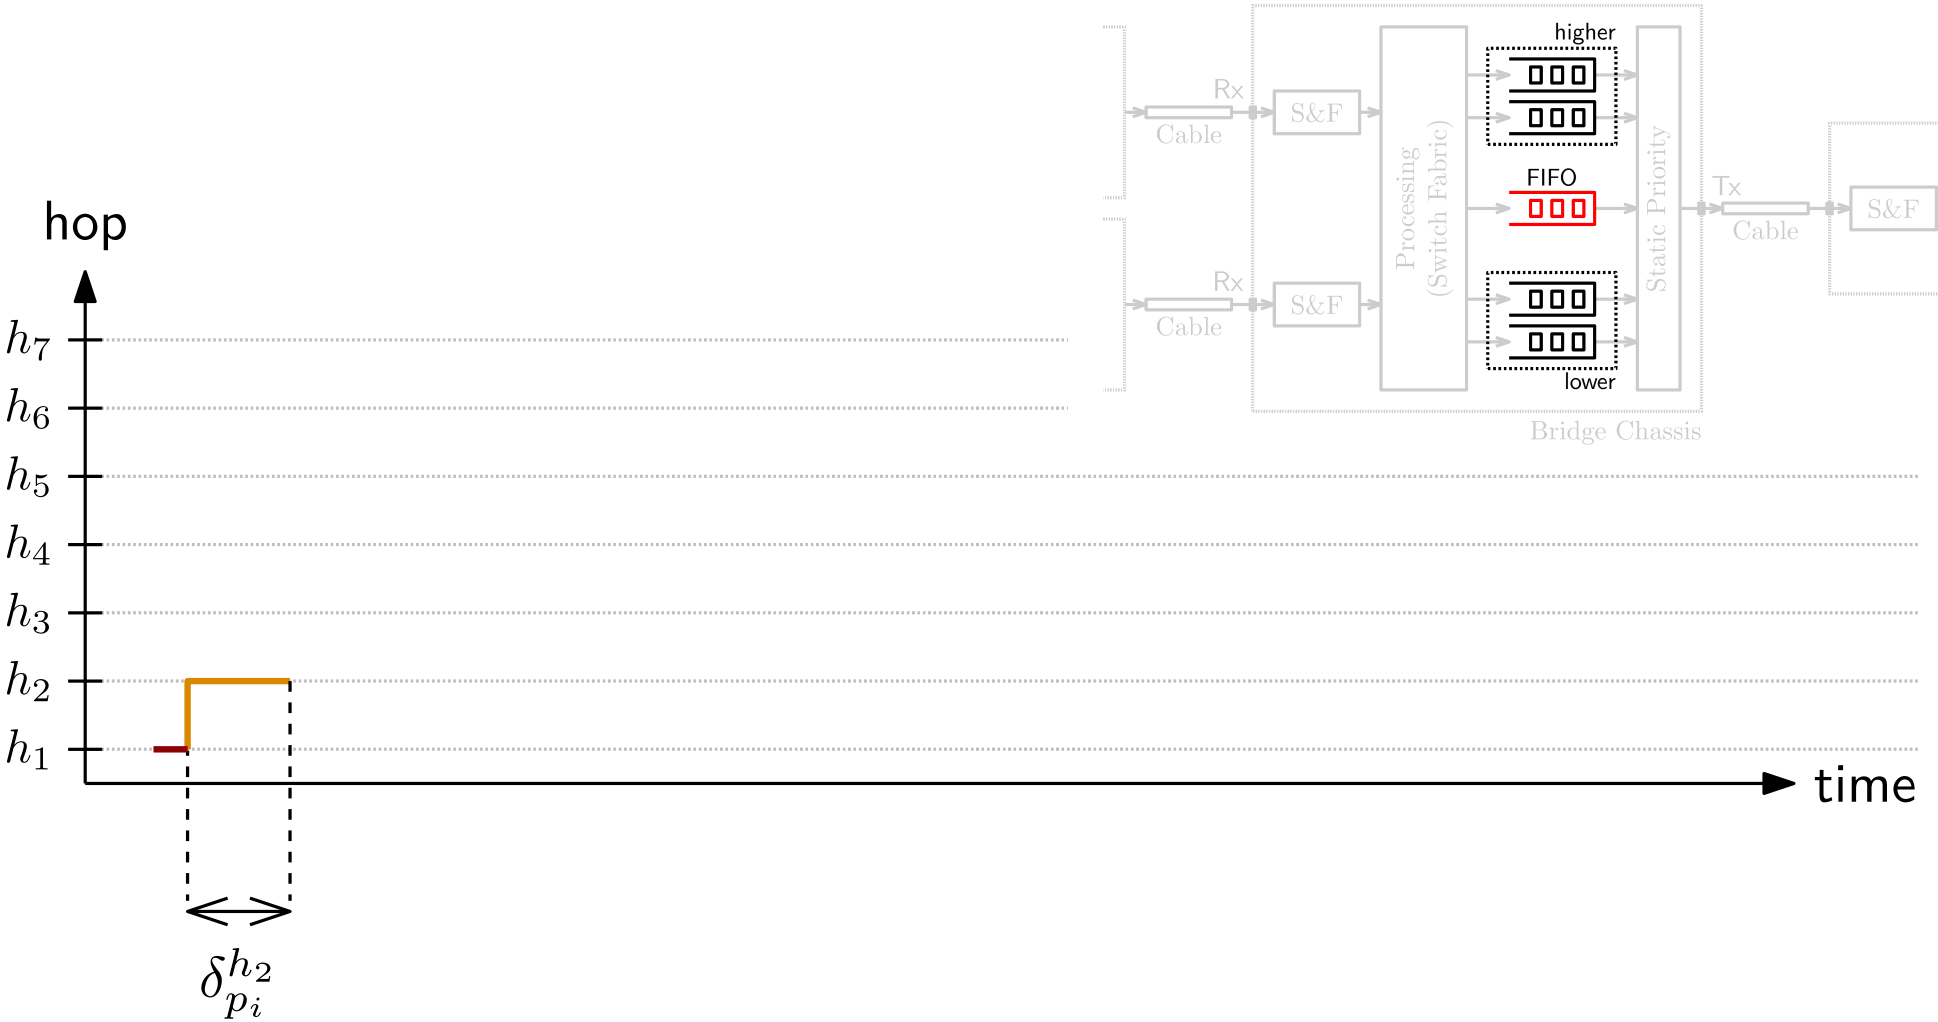

Reasoning – Residence Times of Frames in TQ

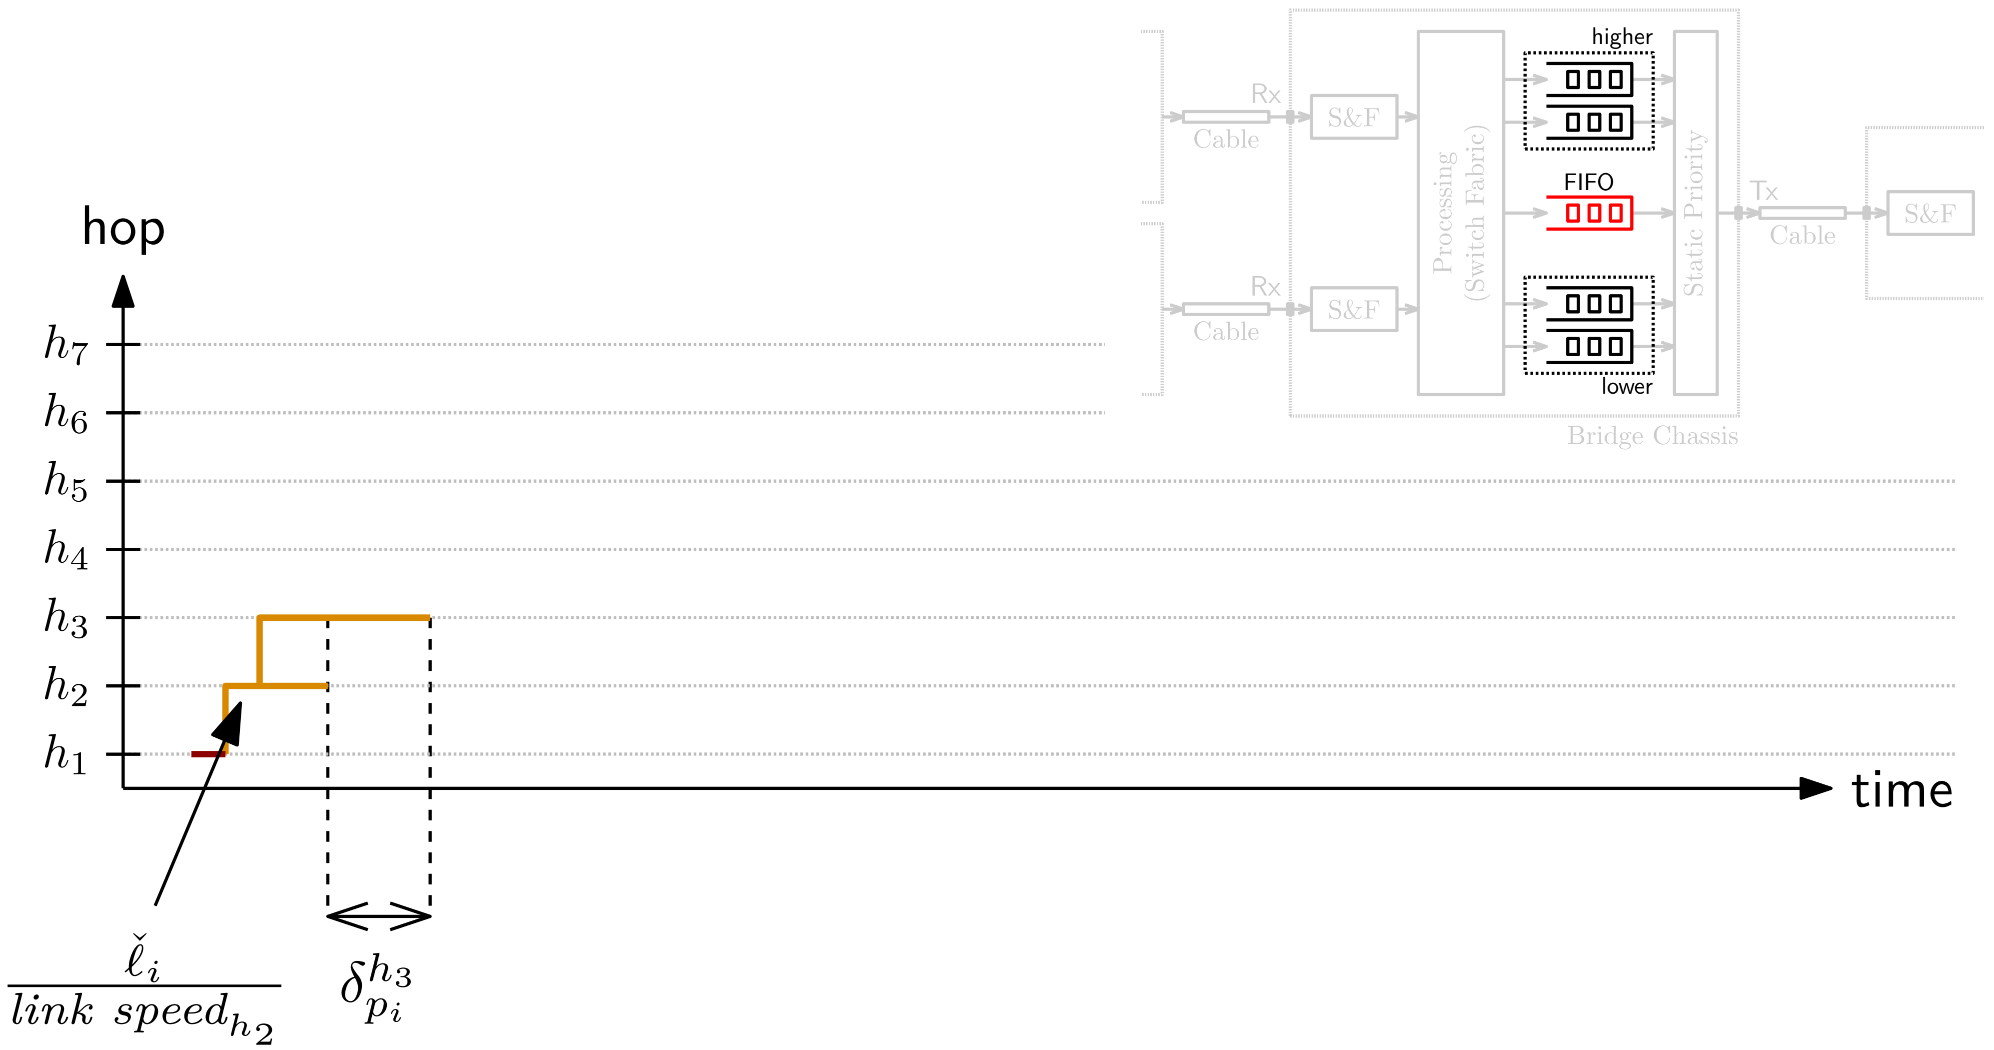

Reasoning – Residence Times of Frames in TQ

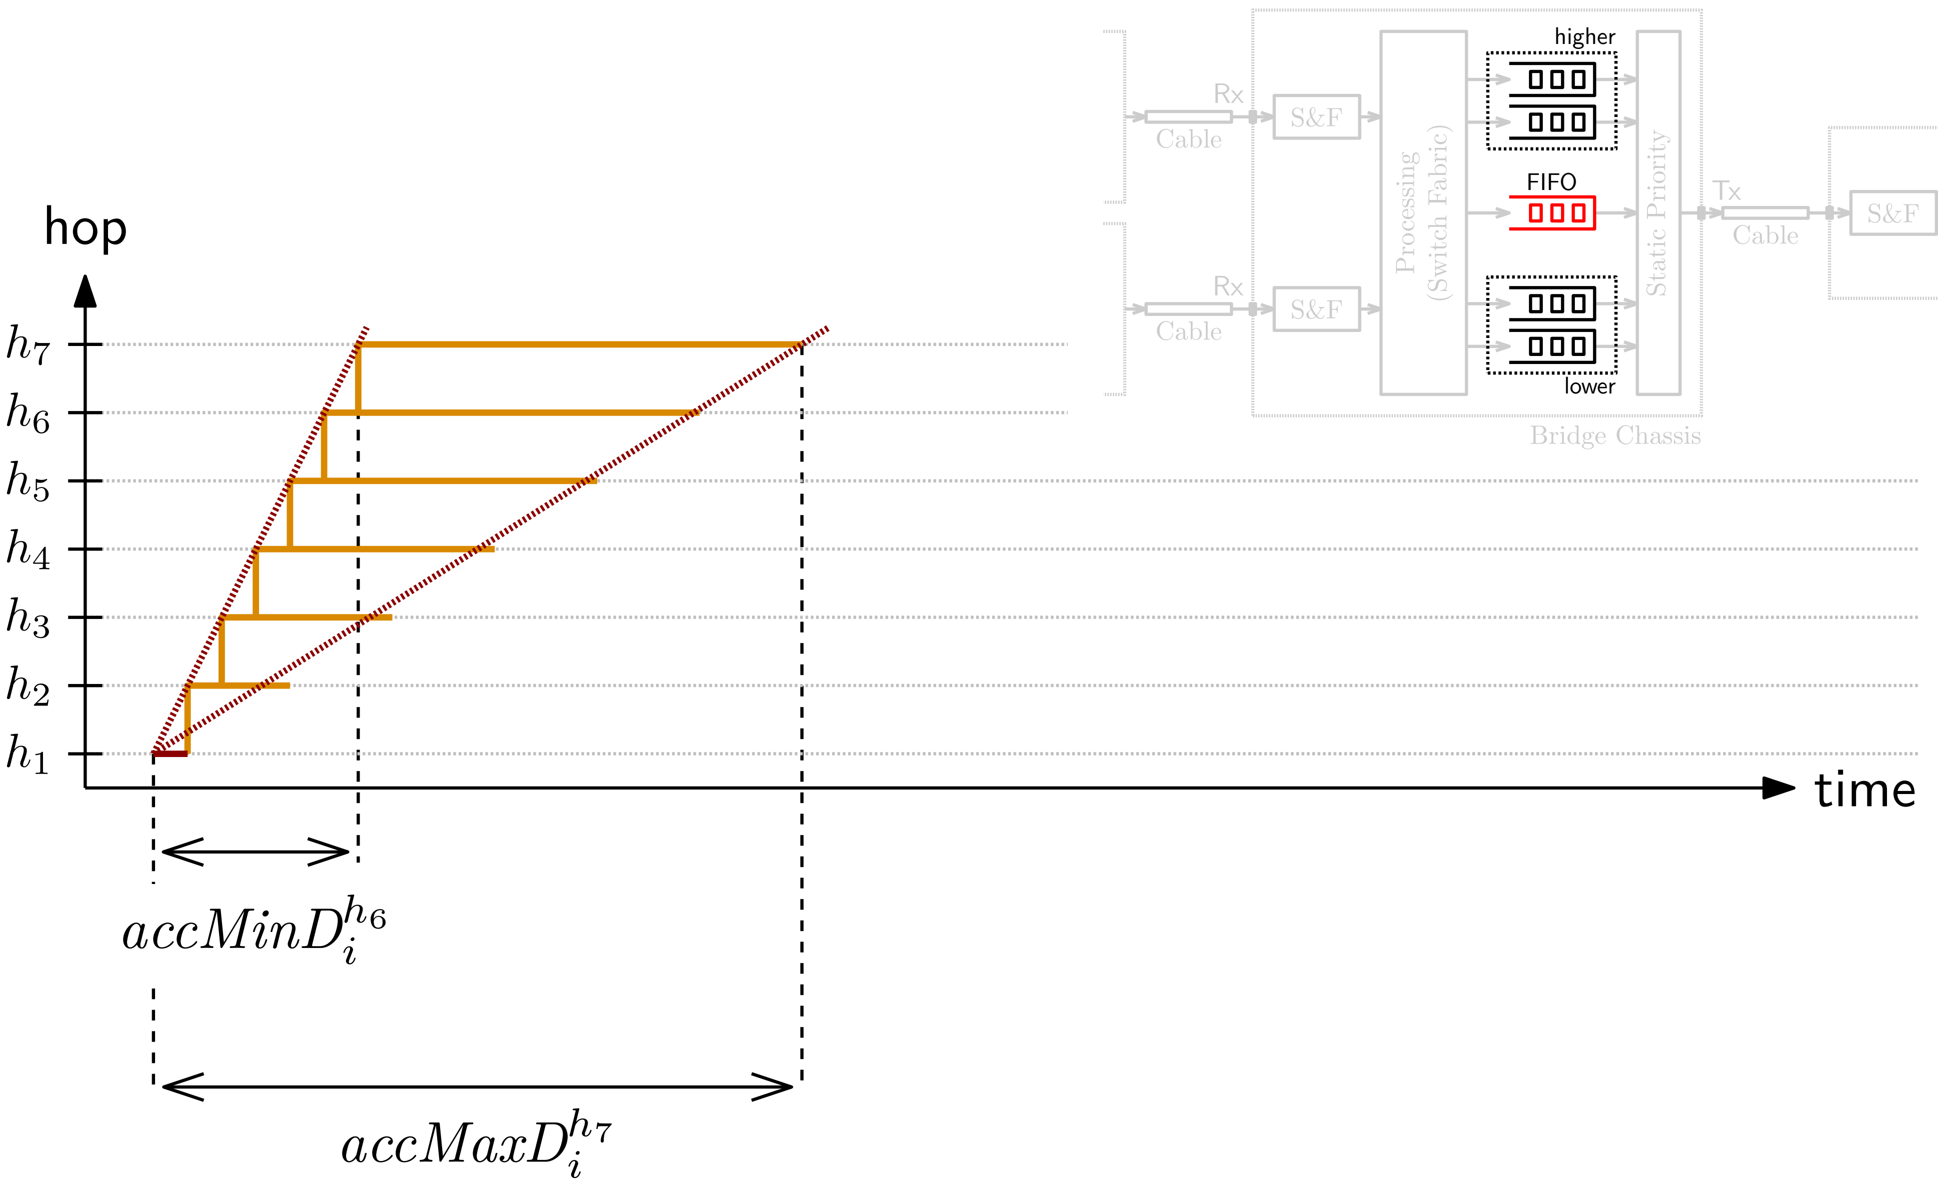

Reasoning – Residence Times of Frames in TQ

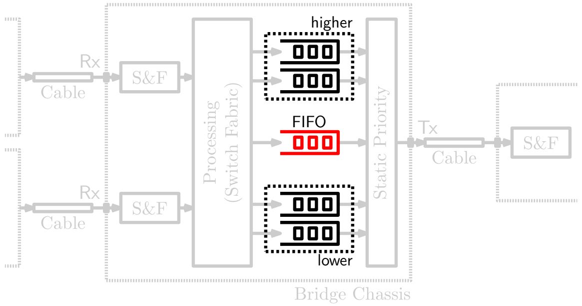

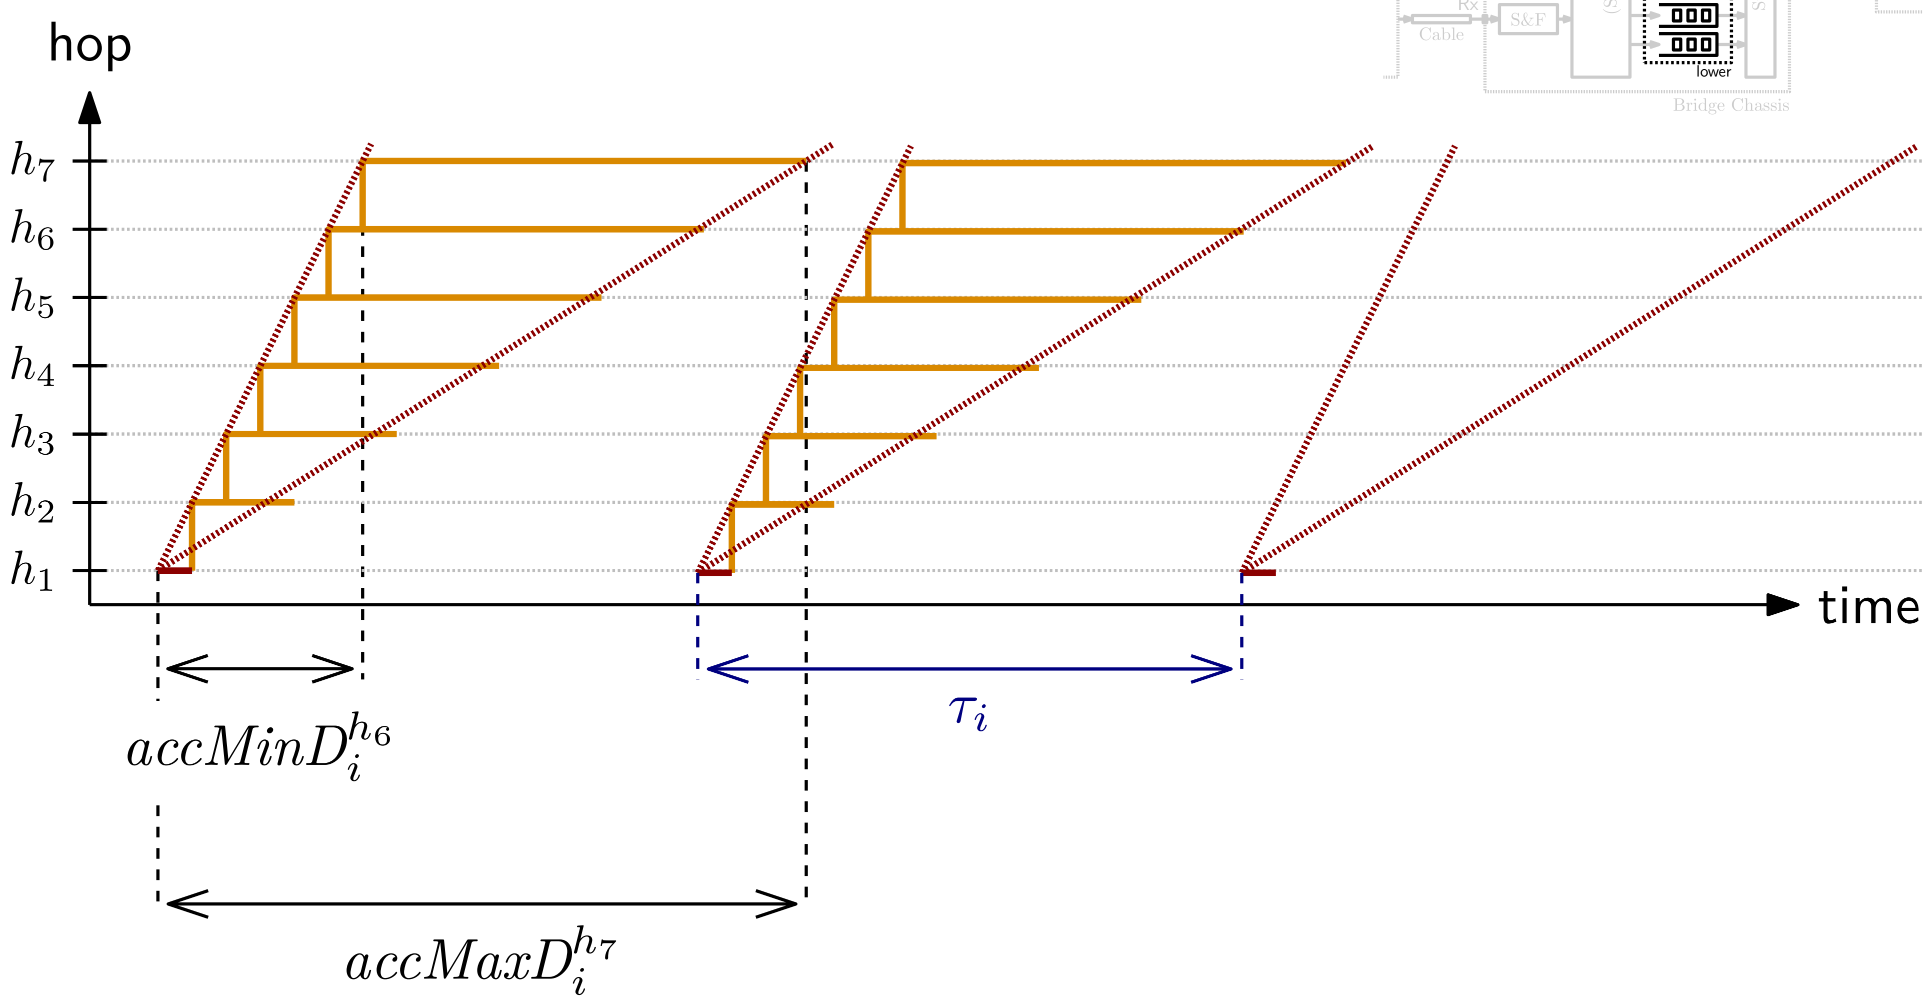

Reasoning – $\mathbf{z_x}$

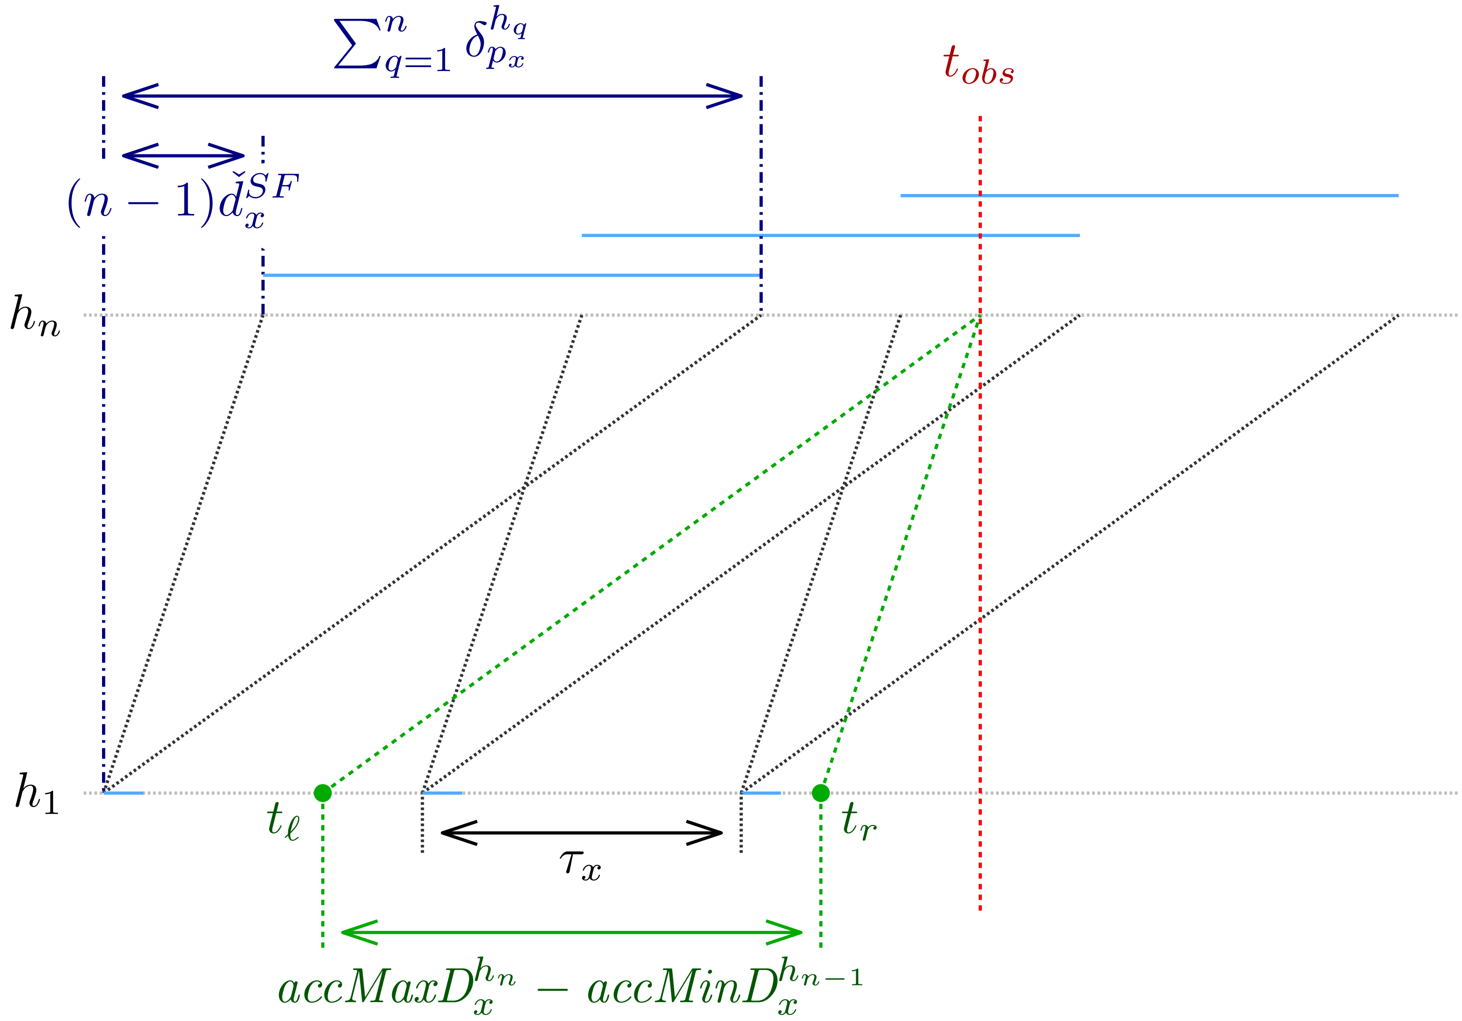

Reasoning – $\mathbf{y_{i,x}}$

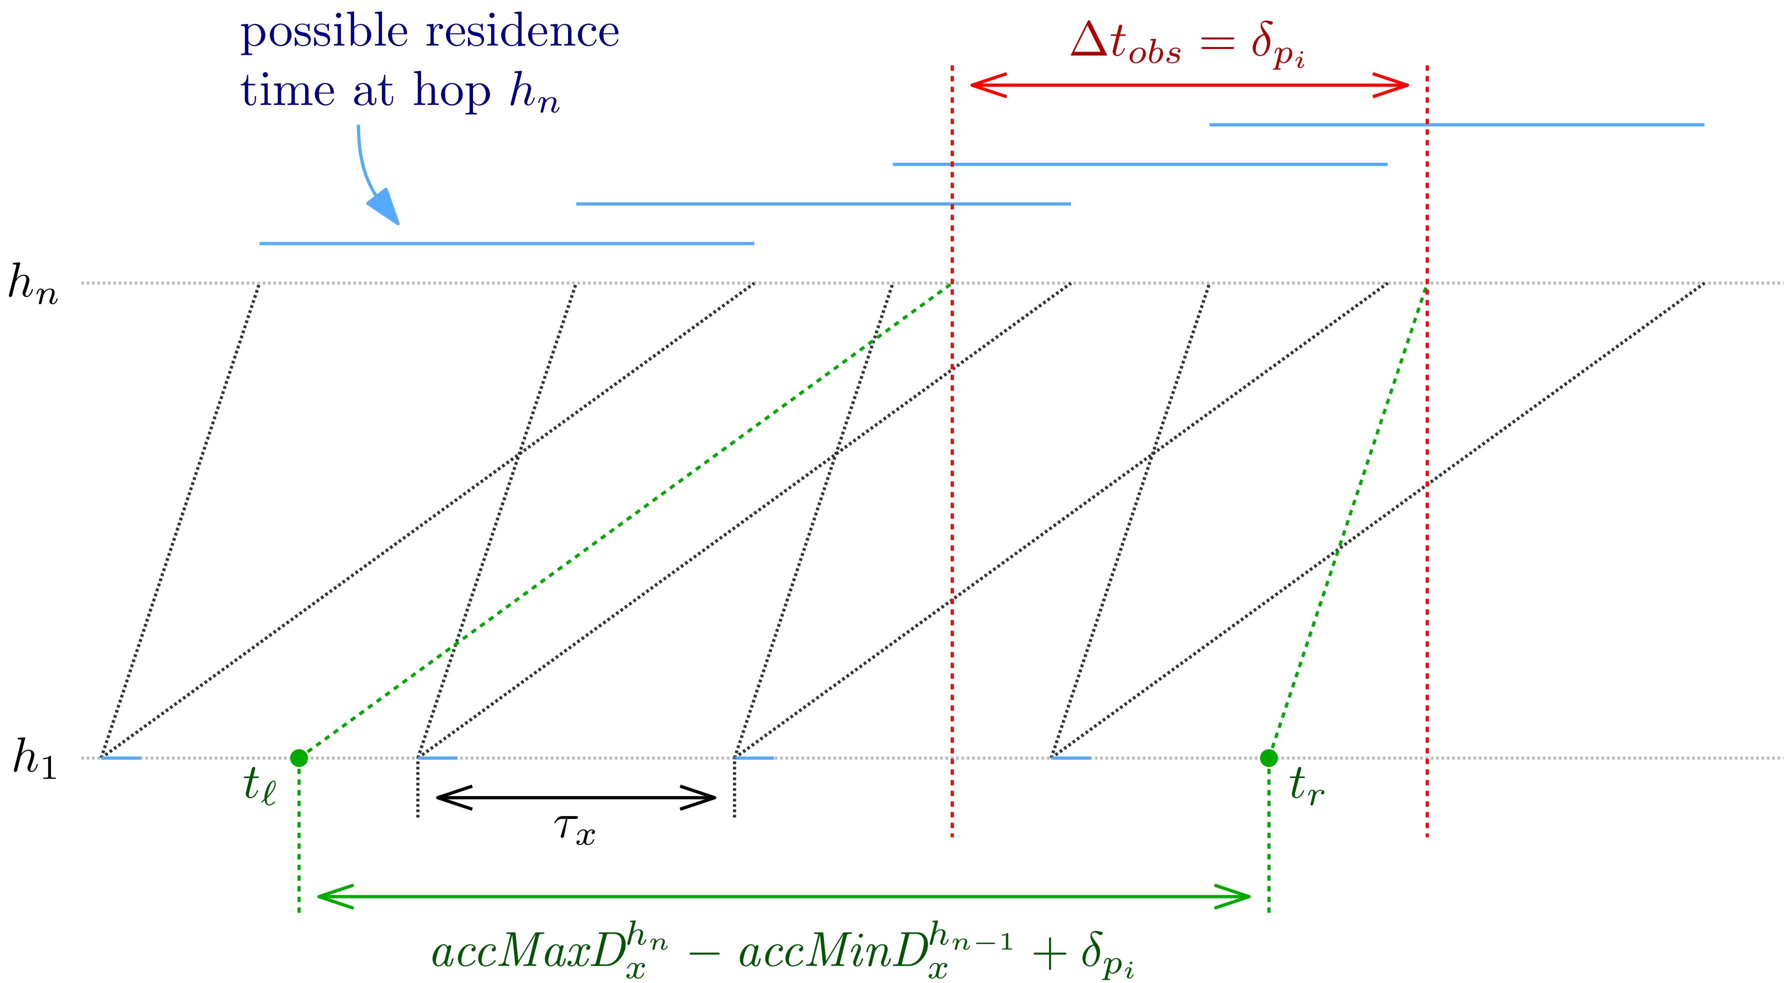

Evaluation

Inaccuracies and comparison to ATS

Worst Case Scenario Simulation

Table II
BURST AND INTERVAL PARAMETERS FOR THE THREE CONSIDERED CLASSES IN THE WORST-CASE SCENARIO.

Traffic class p_x	**Burst** b_x	**Interval** τ_x	**Per-hop delay** δ_{p_x}
lower ($< p_i$)	1500 B	100 ms	100 ms
same (p_i)	256 B	1 ms	1 ms
higher ($> p_i$)	64 B	250 µs	250 µs

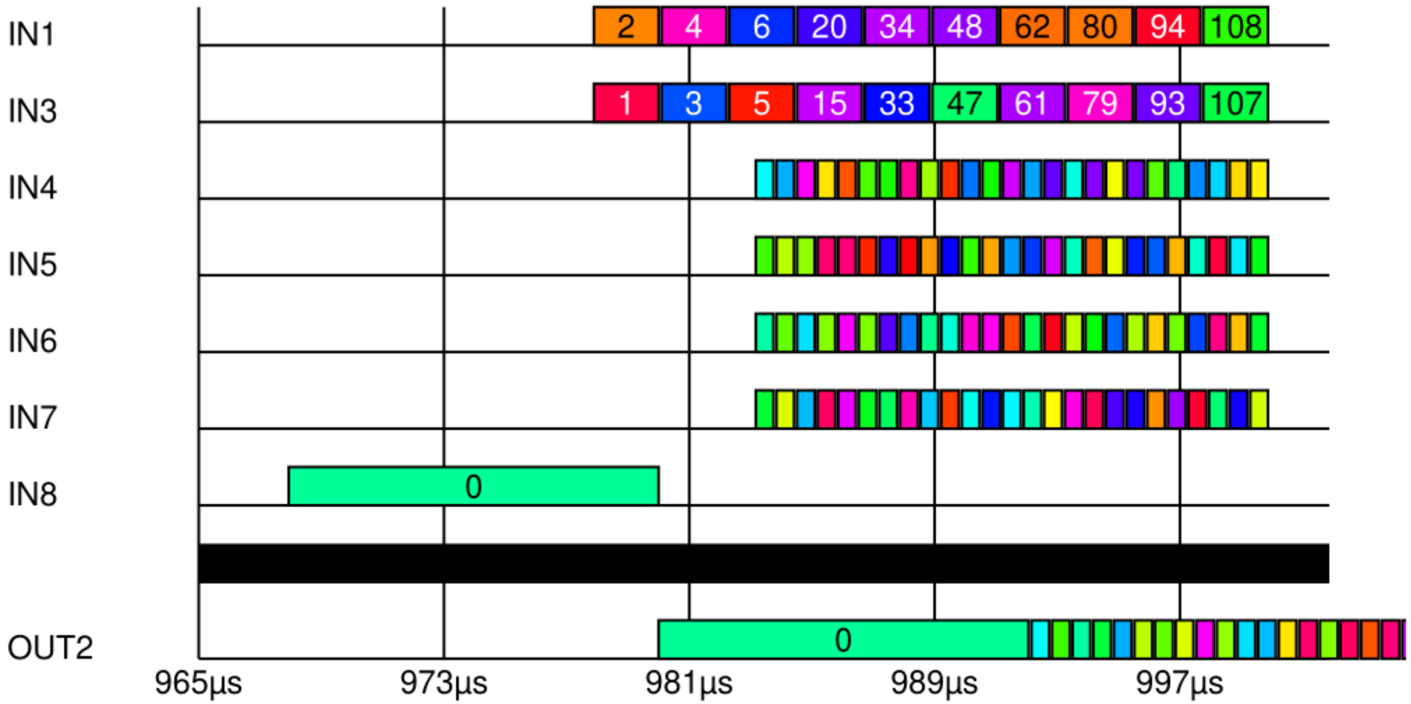

Sources of Inaccuracy

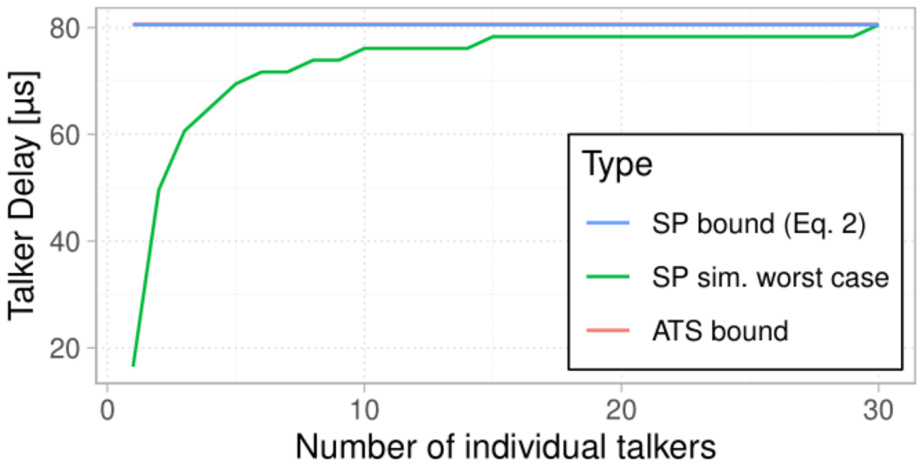

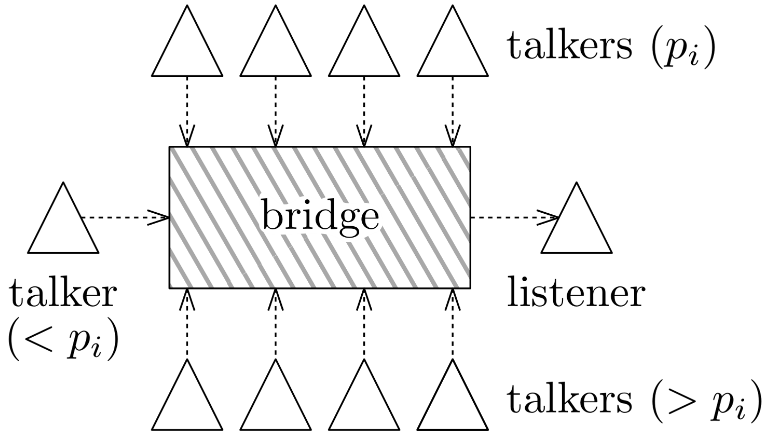

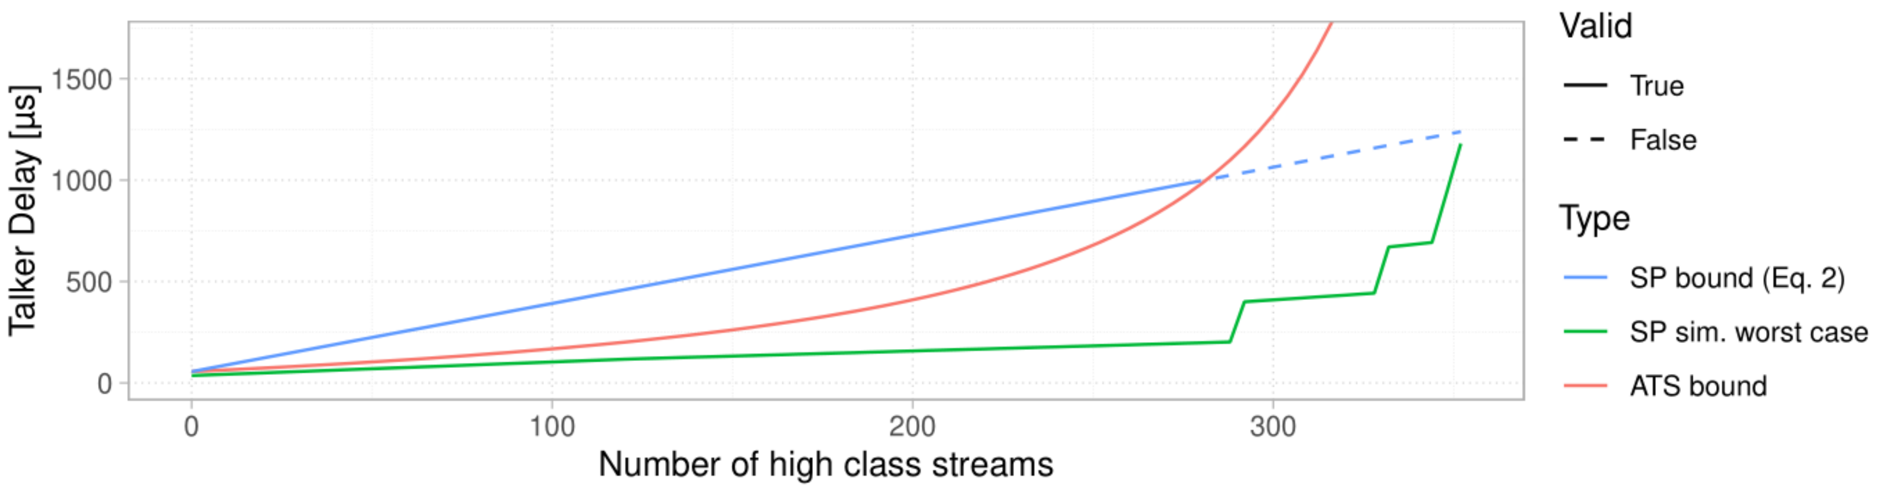

Comparison of Network Capacities

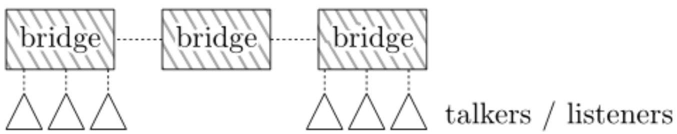

Figure 9. Evaluation topology for capacity comparisons. Talker and listener of each stream is chosen at random. Every link has a data rate of 1 Gbit/s.

Table III
STREAM PARAMETERS FOR THE CAPACITY COMPARISON.

Traffic class p_x	**Burst** $b_x = \hat{\ell}_x$	**Burst interval** τ_x
3 (high)	128 B	250 µs
3 (high)	256 B	500 µs
3 (high)	512 B	1000 µs
2 (low)	1024 B	2000 µs
2 (low)	1522 B	4000 µs

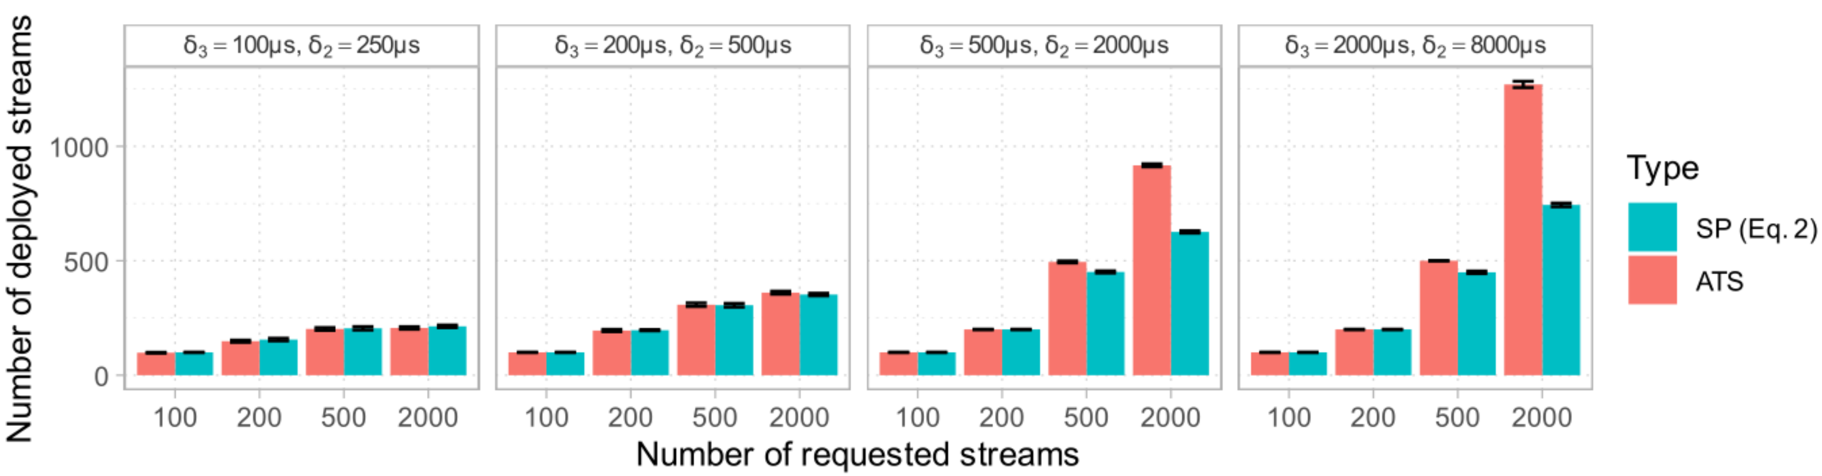

Conclusion

- Bridge-local bounded latency with SP possible
  - Bound only applicable in admission control scenarios
  - Streams whose latency exceeds their guarantee **must** be denied
- Feasible, but not as efficient as per-hop reshaping
- Requirements (self pacing talkers) similar to other mechanisms
- Most required information is already contained in TSpec fields of Qcc
  - + accMinLatency
  - + minFrameSize