Technische Universität Braunschweig

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Decentralized Resource Reservation for Real-Time Communication in TSN

Lisa Maile, September 10, 2024

Decentralized Reservation Protocols

General Procedure

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Decentralized Admission Control

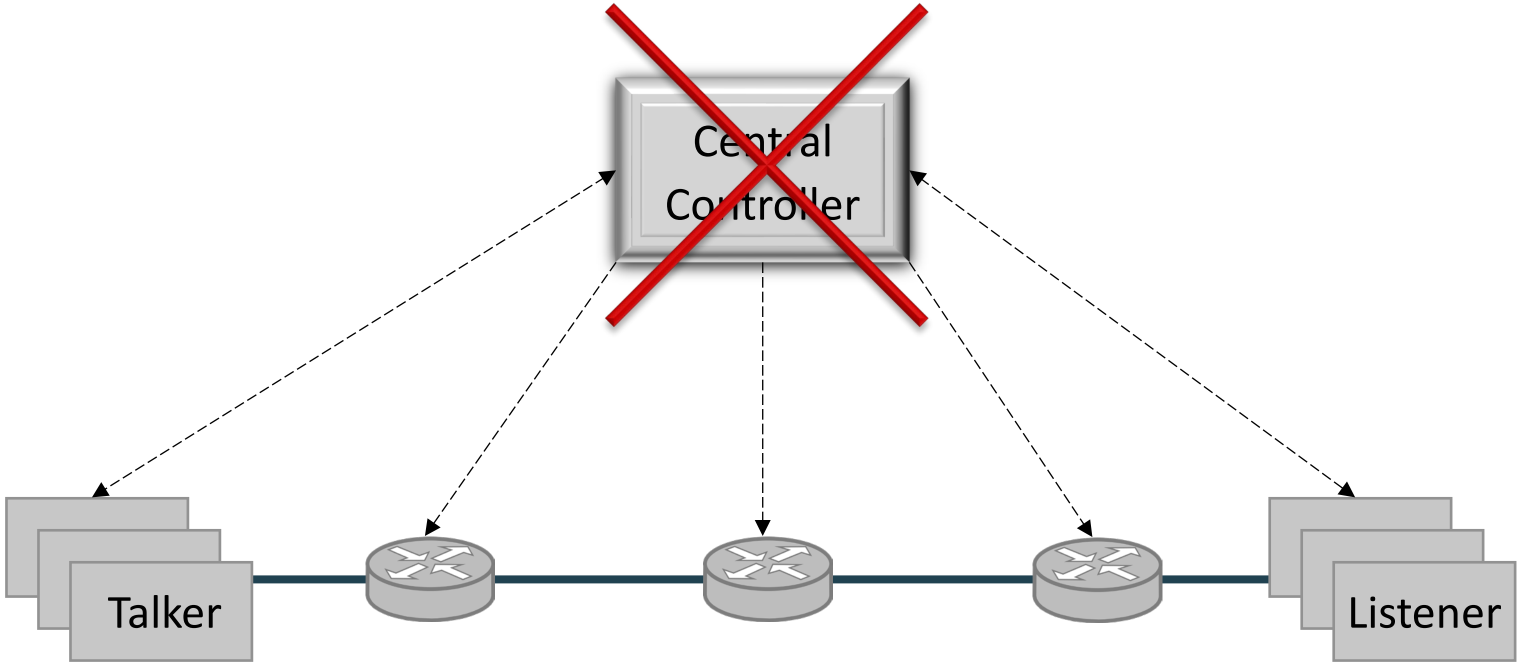

Decentralized Reservation Protocols

General Procedure

Decentralized Admission Control

1) Talker advertises the availability of data

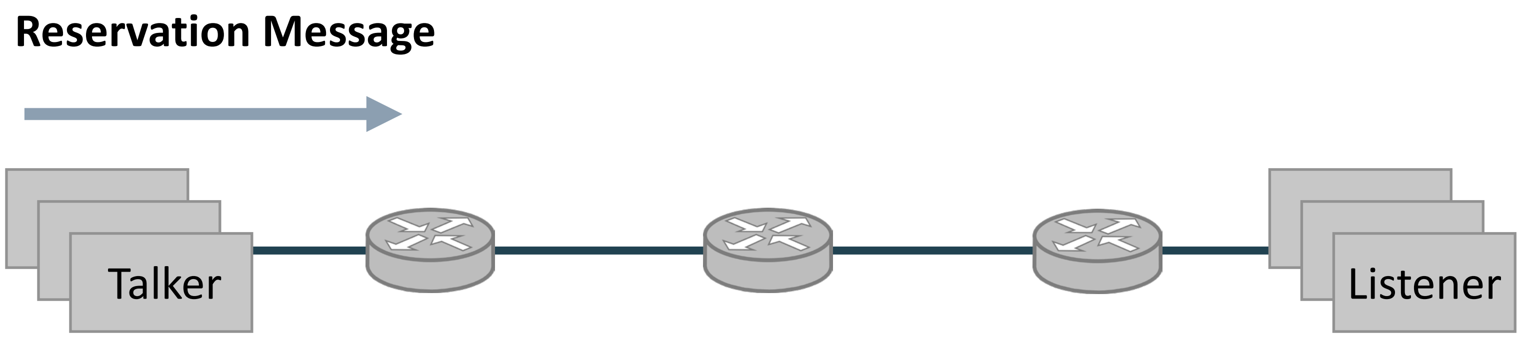

Decentralized Reservation Protocols

General Procedure

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Decentralized Admission Control

1) Talker advertises the availability of data

2) Bridges check for enough resources

What about Paths?

assumed to be given, e.g., by

- broadcasting (pub-sub principle)
- pre-established:
  - static / manual
  - additional central control unit [IEEE Std 802.1Qcc]
  - other protocols:
    - Rapid Spanning Tree Protocol (RSTP) [IEEE Std 802.1w]
    - Shortest Path Bridging (SPB) [IEEE Std 802.1aq]
      → extended: Intermediate System to Intermediate System (IS-IS) for Path Control and Reservation [IEEE Std 802.1Qca]

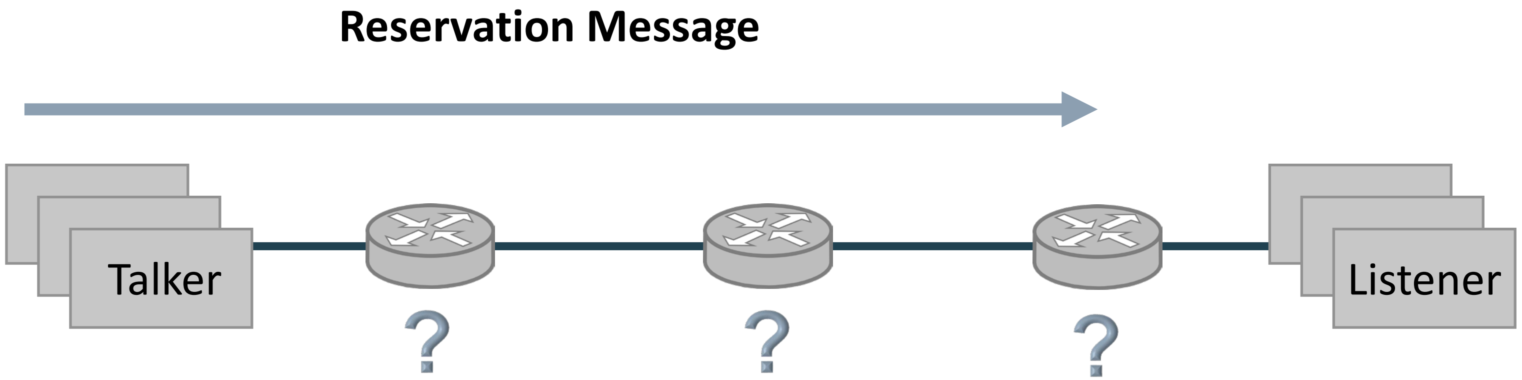

Decentralized Reservation Protocols

General Procedure

What about Paths?

assumed to be given, e.g., by
- broadcasting (pub-sub principle)
- pre-established:
  - static / manual
  - additional central control unit [IEEE Std 802.1Qcc]
  - other protocols:
    - Rapid Spanning Tree Protocol (RSTP) [IEEE Std 802.1w]
    - Shortest Path Bridging (SPB) [IEEE Std 802.1aq]
      → extended: Intermediate System to Intermediate System (IS-IS) for Path Control and Reservation [IEEE Std 802.1Qca]

Decentralized Admission Control

1) Talker advertises the availability of data
2) Bridges check for enough resources
3) Listener respond if they are interested

Confirmation Message

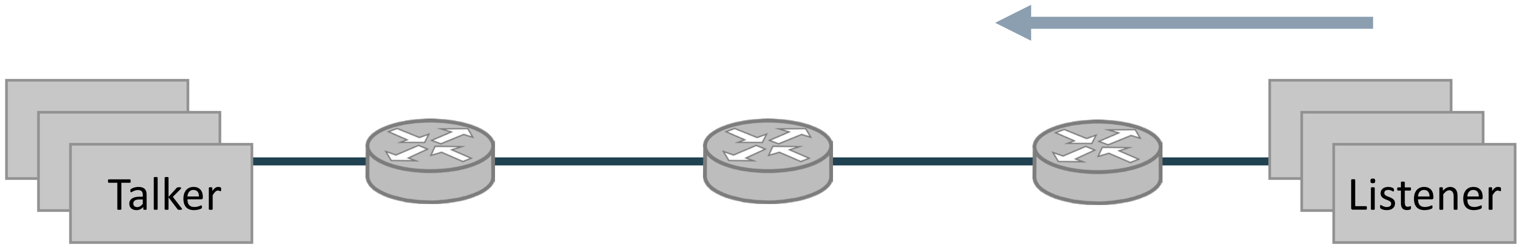

Decentralized Reservation Protocols

General Procedure

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Decentralized Admission Control

1) Talker advertises the availability of data
2) Bridges check for enough resources
3) Listener respond if they are interested
4) Bridges reserve resources

What about Paths?

assumed to be given, e.g., by

- broadcasting (pub-sub principle)
- pre-established:
  - static / manual
  - additional central control unit [IEEE Std 802.1Qcc]
  - other protocols:
    - Rapid Spanning Tree Protocol (RSTP) [IEEE Std 802.1w]
    - Shortest Path Bridging (SPB) [IEEE Std 802.1aq]
      → extended: Intermediate System to Intermediate System (IS-IS) for Path Control and Reservation [IEEE Std 802.1Qca]

Confirmation Message

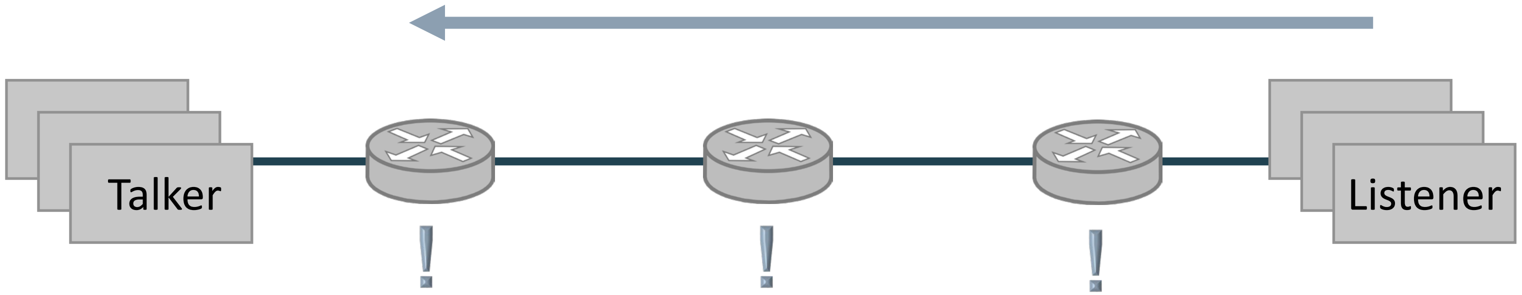

Decentralized Reservation Protocols

General Procedure

Decentralized Admission Control

1) Talker advertises the availability of data
2) Bridges check for enough resources
3) Listener respond if they are interested
4) Bridges reserve resources
5) Start of Transmission

What about Paths?

assumed to be given, e.g., by

- broadcasting (pub-sub principle)
- pre-established:
  - static / manual
  - additional central control unit [IEEE Std 802.1Qcc]
  - other protocols:
    - Rapid Spanning Tree Protocol (RSTP) [IEEE Std 802.1w]
    - Shortest Path Bridging (SPB) [IEEE Std 802.1aq]
      → extended: Intermediate System to Intermediate System (IS-IS) for Path Control and Reservation [IEEE Std 802.1Qca]

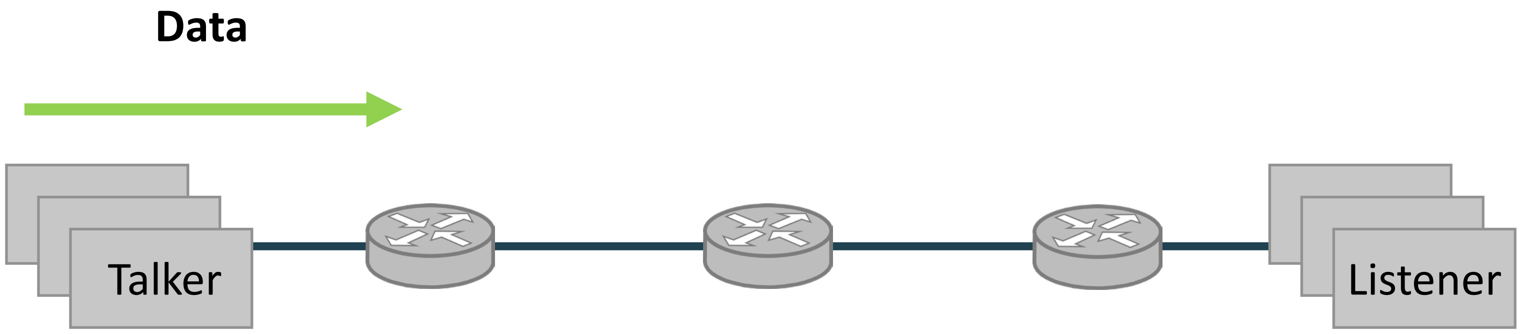

Decentralized Reservation Protocols

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Protocols for Decentralized Resource Reservations

Wide Area Networks (WAN)
Layer 4

Time-Sensitive Networking (TSN) / Local Area Networks (LAN)
Layer 2

1997
Resource Reservation Protocol (RSVP)

mostly relied on IntServ
with per-flow shaping

2010
Stream Reservation Protocol (SRP)

for Credit-Based Shaper
with per-class shaping
IEEE Std 802.1Qat

2018
Resource Allocation Protocol (RAP)

support planned for various schedulers
IEEE Std P802.1Qdd **(draft)**

interoperability

How “it” works…

where
“it” ≔ decentralized reservation protocols in TSN

Decentralized Reservation in TSN

simplified

Stream Reservation Protocol (SRP)

Resource Allocation Protocol (RAP)

differences in message format and bridge logic

Talker Advertise
StreamID: "D4:85:64:12:A3:09:00:01"
Data Frame Parameter:
- Destination: "01:00:5E:AB:CD:EF"
- Priority: 7
- Rank: 0

TSpec:
- Interval: 125 µs
- MaxFramesPerInterval: 1
- MaxFrameSize: 1500B

AccumulatedLatency: 0 µs

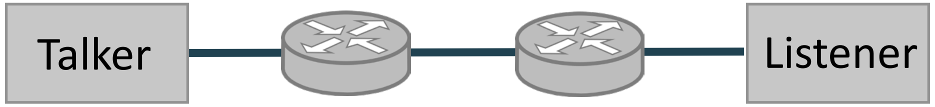

Talker Announce
StreamID: "D4:85:64:12:A3:09:00:01"
Data Frame Parameter:
- Destination: "01:00:5E:AB:CD:EF"
- Priority: 7
- Rank: 0

TSpec:
- MaxTransmittedFrameLength: 1500 B
- MinTransmittedFrameLength: 100 B
- CommittedBurstSize: 1520 B
- CommittedInformationRate: 12.16 MB/s

MaxAccumulatedLatency: 0 µs
MinAccumulatedLatency: 0 µs

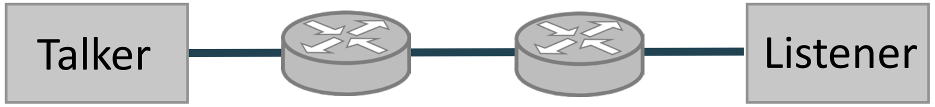

Decentralized Reservation in TSN

simplified

Stream Reservation Protocol (SRP)

Resource Allocation Protocol (RAP)

differences in message format and bridge logic

Talker Advertise

StreamID: "D4:85:64:12:A3:09:00:01"
Data Frame Parameter:
    Destination: "01:00:5E:AB:CD:EF"
    Priority: 7
    Rank: 0
TSpec:
    Interval: 125 µs
    MaxFramesPerInterval: 1
    MaxFrameSize: 1500B

AccumulatedLatency: **500 µs**

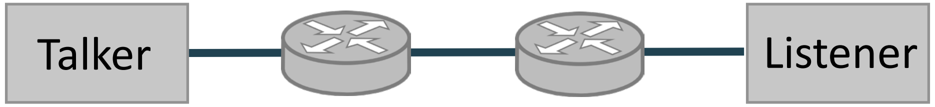

Talker Announce

StreamID: "D4:85:64:12:A3:09:00:01"
Data Frame Parameter:
    Destination: "01:00:5E:AB:CD:EF"
    Priority: 7
    Rank: 0
TSpec:
    MaxTransmittedFrameLength: 1500 B
    MinTransmittedFrameLength: 100 B
    CommittedBurstSize: 1520 B
    CommittedInformationRate: 12.16 MB/s
MaxAccumulatedLatency: **500 µs**
MinAccumulatedLatency: **1 µs**

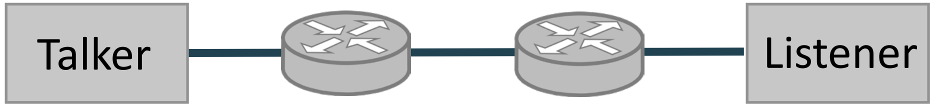

Decentralized Reservation in TSN

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

simplified

Stream Reservation Protocol (SRP)

Talker Advertise

StreamID: "D4:85:64:12:A3:09:00:01"
Data Frame Parameter:
- Destination: "01:00:5E:AB:CD:EF"
- Priority: 7
- Rank: 0

TSpec:
- Interval: 125 µs
- MaxFramesPerInterval: 1
- MaxFrameSize: 1500B

AccumulatedLatency: **1000 µs**

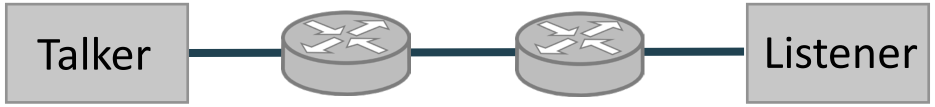

Resource Allocation Protocol (RAP)

Talker Announce

StreamID: "D4:85:64:12:A3:09:00:01"
Data Frame Parameter:
- Destination: "01:00:5E:AB:CD:EF"
- Priority: 7
- Rank: 0

TSpec:
- MaxTransmittedFrameLength: 1500 B
- MinTransmittedFrameLength: 100 B
- CommittedBurstSize: 1520 B
- CommittedInformationRate: 12.16 MB/s

MaxAccumulatedLatency: **1000 µs**
MinAccumulatedLatency: **2 µs**

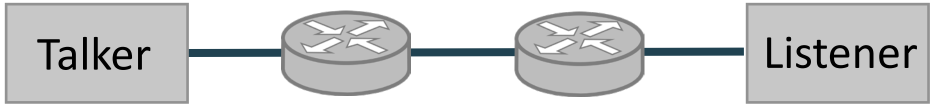

differences in message format and bridge logic

Decentralized Reservation in TSN

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

simplified

Stream Reservation Protocol (SRP) | **Resource Allocation Protocol (RAP)**

differences in message format and bridge logic

Listener Ready
StreamID: "D4:85:64:12:A3:09:00:01"

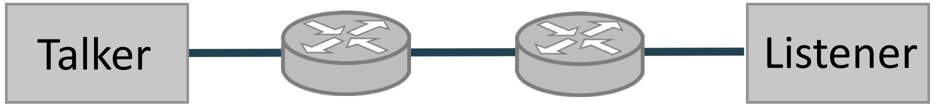

Listener Attach
StreamID: "D4:85:64:12:A3:09:00:01“
Status: Attach Ready

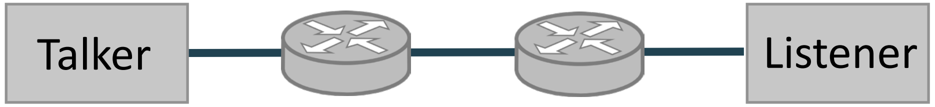

Decentralized Reservation in TSN

simplified

Stream Reservation Protocol (SRP)

Resource Allocation Protocol (RAP)

differences in message format and bridge logic

Bridge Logic
Reserve, if checks successful for:
- Bandwidth?
- Queue Size?
- Internal Resources?

Listener Ready
StreamID: "D4:85:64:12:A3:09:00:01"

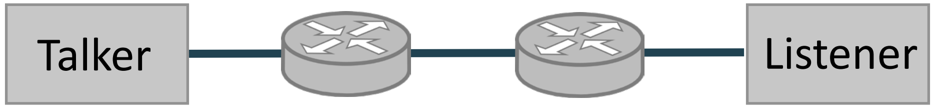

Bridge Logic
Reserve, if checks successful for:
- Bandwidth?
- Queue Size?
- Internal Resources?
- Delay?

Listener Attach
StreamID: "D4:85:64:12:A3:09:00:01“
Status: Attach Ready

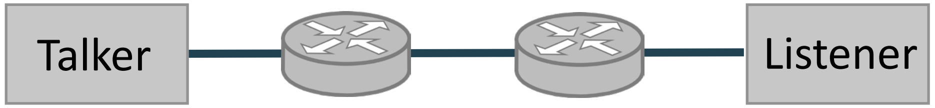

Decentralized Reservation in TSN

simplified

Stream Reservation Protocol (SRP)

Resource Allocation Protocol (RAP)

differences in message format and bridge logic

Link-Local Registration Protocol (LRP) [IEEE Std 802.1CS-2020]:
→ small data bases for each port synchronized

SRP

Bridge Logic
Reserve, if checks successful for:
- Bandwidth?
- Queue Size?
- Internal Resources?

Listener Ready
StreamID: "D4:85:64:12:A3:09:00:01"

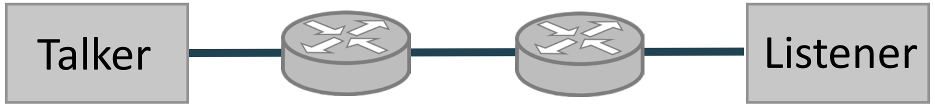

RAP

Bridge Logic
Reserve, if checks successful for:
- Bandwidth?
- Queue Size?
- Internal Resources?
- Delay?

Listener Attach
StreamID: "D4:85:64:12:A3:09:00:01“
Status: Attach Ready

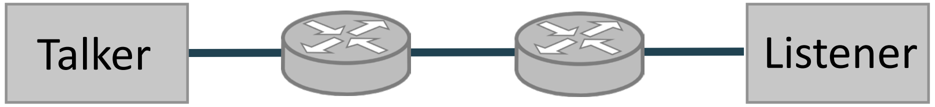

The Problem

Latency Calculation

Looking back: RSVP/IntServ

Flow latency does not depend on other flows – service guaranteed at all times [Zhang&Ferrari 1993, Frangioni 2017]

But now?

Service offered by TSN schedulers depends on scheduled flows
→ flow burstiness changes, which changes the delay in the subsequent network

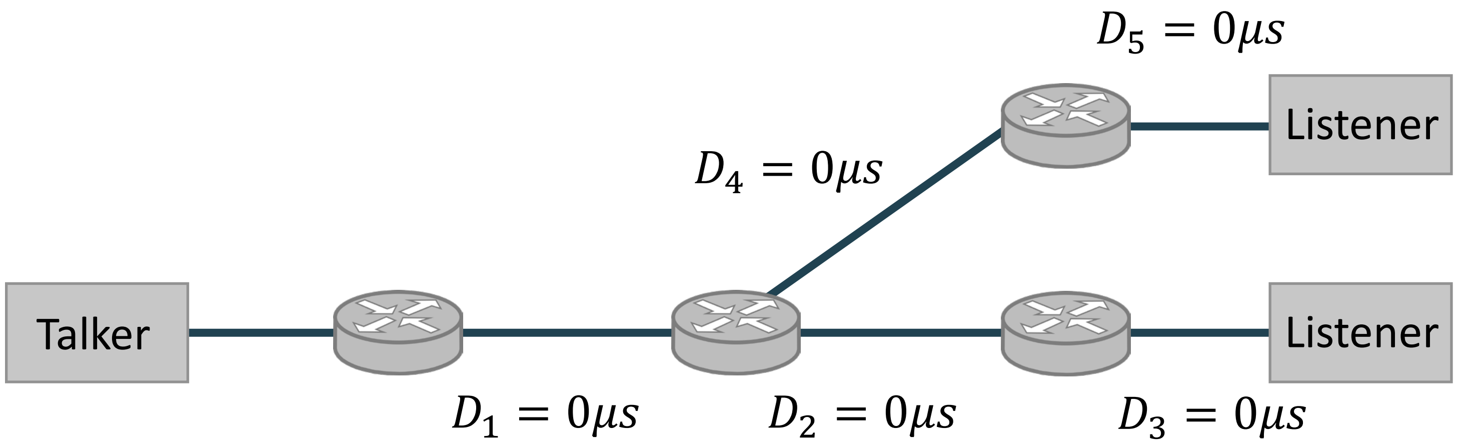

Latency Calculation

Looking back: RSVP/IntServ

Flow latency does not depend on other flows – service guaranteed at all times [Zhang&Ferrari 1993, Frangioni 2017]

But now?

Service offered by TSN schedulers depends on scheduled flows
→ flow burstiness changes, which changes the delay in the subsequent network

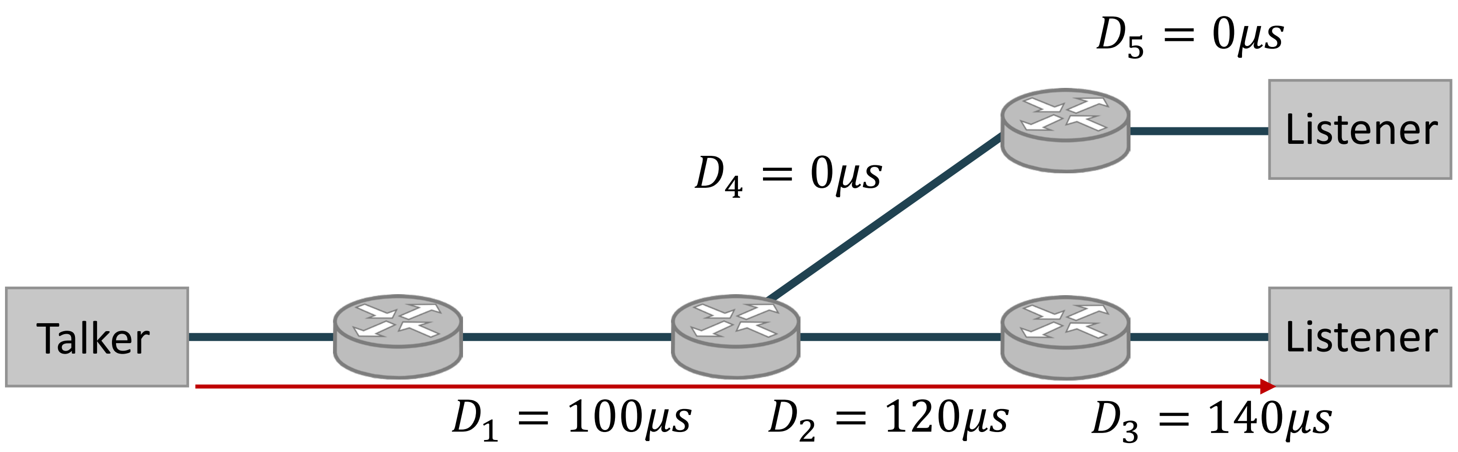

Latency Calculation

Looking back: RSVP/IntServ

Flow latency does not depend on other flows – service guaranteed at all times [Zhang&Ferrari 1993, Frangioni 2017]

But now?

Service offered by TSN schedulers depends on scheduled flows
→ flow burstiness changes, which changes the delay in the subsequent network

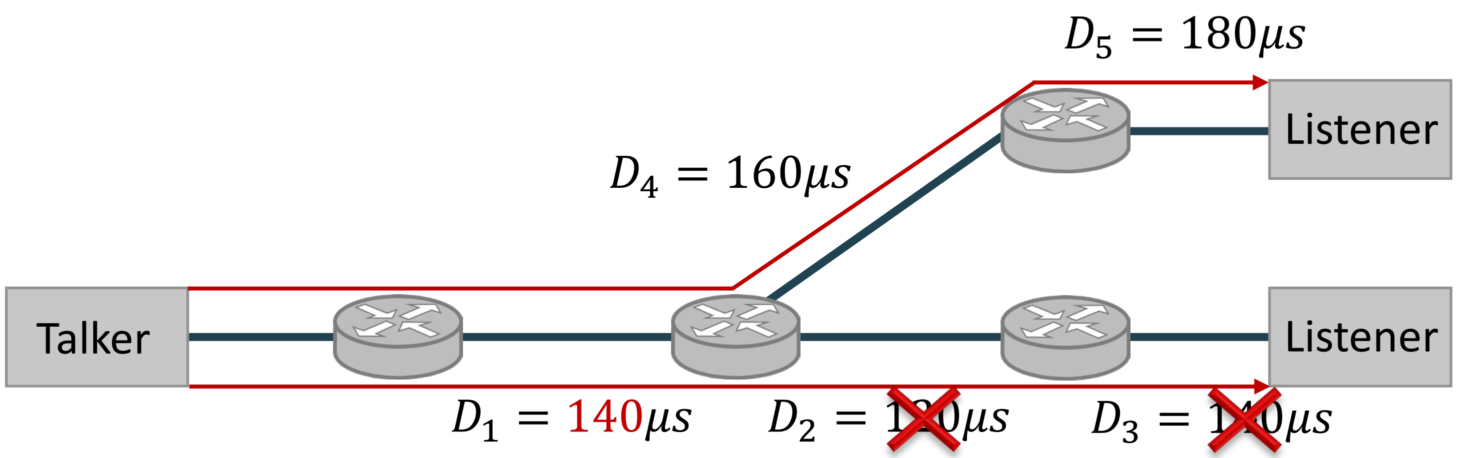

Latency Calculation

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Looking back: RSVP/IntServ

Flow latency does not depend on other flows – service guaranteed at all times

But now?

Service offered by TSN schedulers depends on scheduled flows
→ flow burstiness changes, which changes the delay in the subsequent network

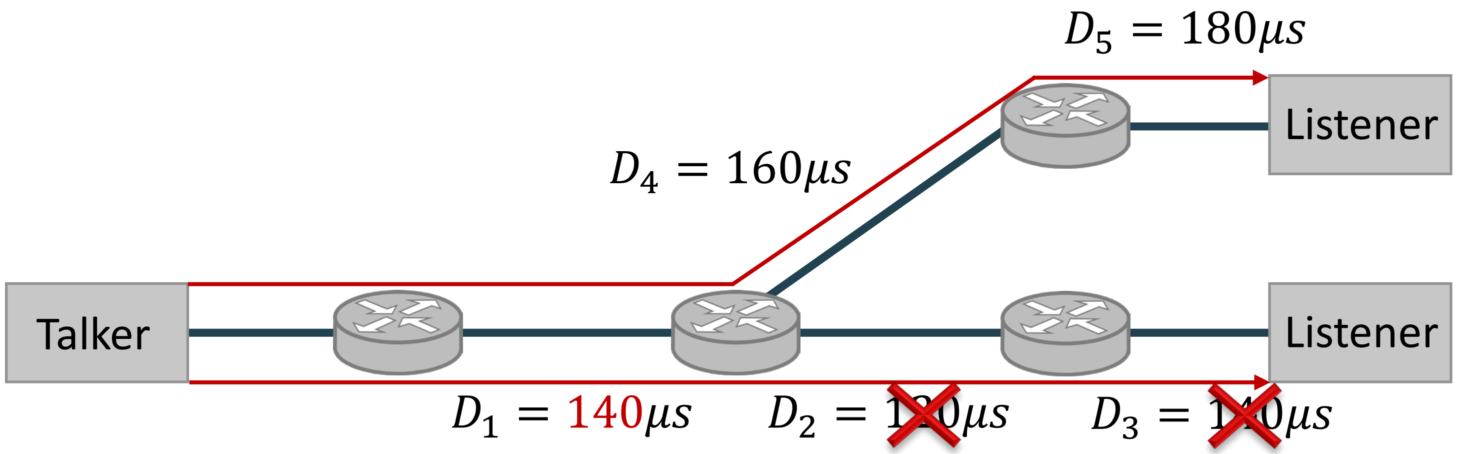

Required:

Topology-Independent
Per-Hop Latency Calculation

- latency/backlog bounds of flow not affected by other flows
- using only local information (no topology overview, no information of interference before hop)

Independent Delay Bounds

Independent Per-Queue Latency

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Notation:

Current Worst-Case Per-Queue Latency d_q

Per-Queue Delay Bound $\overline{D_q}$

<u>Bridge Logic (SRP)</u>
Check bounds for:
- Bandwidth?
- Queue Size?
- Internal Resources?

SRP

assumption: there exists a maximum d_q

RAP

added delay bound check: $d_q \leq \overline{D_q}$

<u>Bridge Logic (RAP)</u>
Check bounds for:
- Bandwidth?
- Queue Size?
- Internal Resources?
- Delay?

Independent Per-Queue Latency

Stream Reservation Protocol (SRP)

Worst-case d in the TSN standards for Credit-Based Shaper (e.g., for AccumulatedLatency in SRP):

1) IEEE 802.1BA

$$\overline{D} = t_{proc} + \underbrace{t_{L_{max}}}_{\text{other pr.}} + \underbrace{\left(\frac{idSl}{C}\cdot CMI - t_{L_{FoI}}\right)\cdot\frac{C}{idSl}}_{\text{same priority}} + t_{L_{FoI}-IPG}$$

2) IEEE 802.1Q

$$\overline{D} = t_{inQueue} + t_{int} + t_{L} + t_{pr}\ldots$$

$$t_{queue} = \begin{cases} L_{\ldots} & \text{for prio. 7} \\ (L_{max} + L^{(7)})/(C - idSl^{(7)}) & \text{for prio. 6} \end{cases}$$

$$\frac{idSl}{C}\cdot CMI \cdot \frac{C}{idSl}$$

% of time traffic arrives

CMI = sending interval of all prio. 7 flows

3) Plenary 100 Mbit/s

$$\overline{D} = \left(\underbrace{L_{max}}_{\text{other pr.}} + \underbrace{2(R_{max} - L_{FoI}) - \left\lceil\frac{R_{max} - L_{FoI}}{N}\right\rceil + L_{FoI}}_{\text{same priority}}\right)\cdot t_{oct}$$

$$R_{max} = \left\lceil\frac{CMI}{t_{oct}}\cdot\frac{idSl}{C}\right\rceil,\ N = \min\left(|\mathcal{L}^-|,\ \left\lfloor\frac{R_{max} - L_{FoI}}{L_{min}}\right\rfloor\right)$$

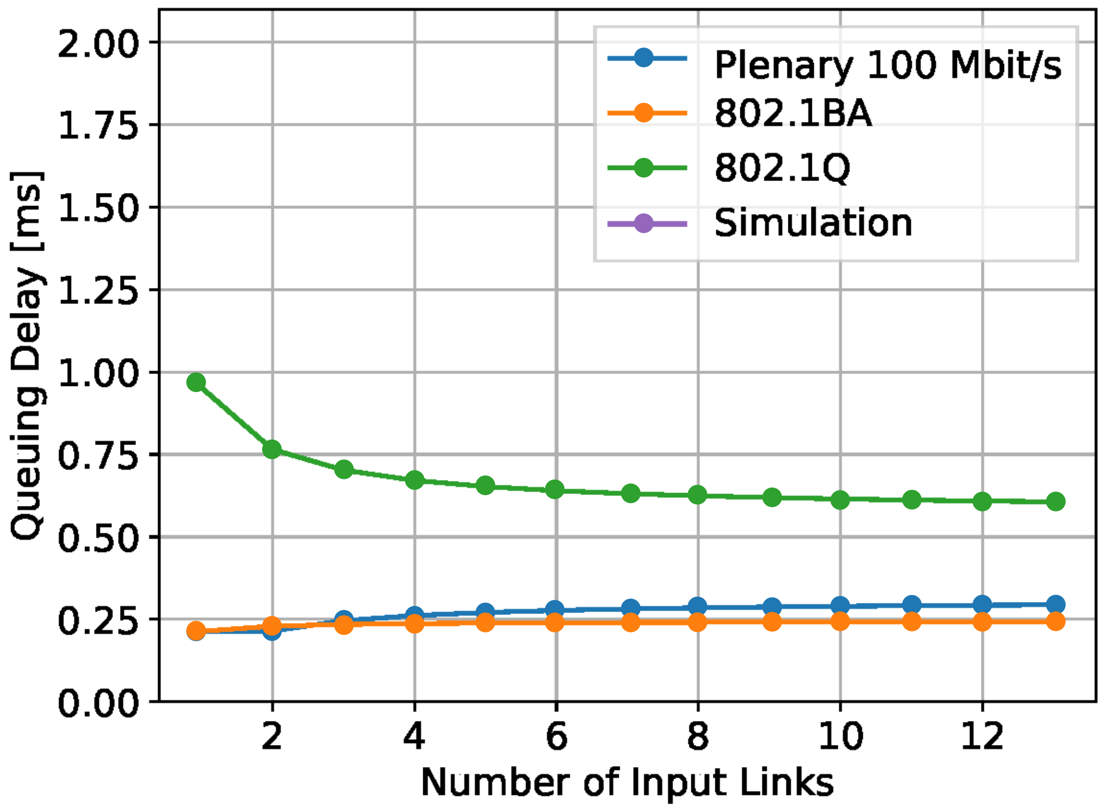

results and analysis from [Maile 2023]

Independent Per-Queue Latency

Stream Reservation Protocol (SRP)

Worst-case d in the TSN standards for Credit-Based Shaper (e.g., for AccumulatedLatency in SRP):

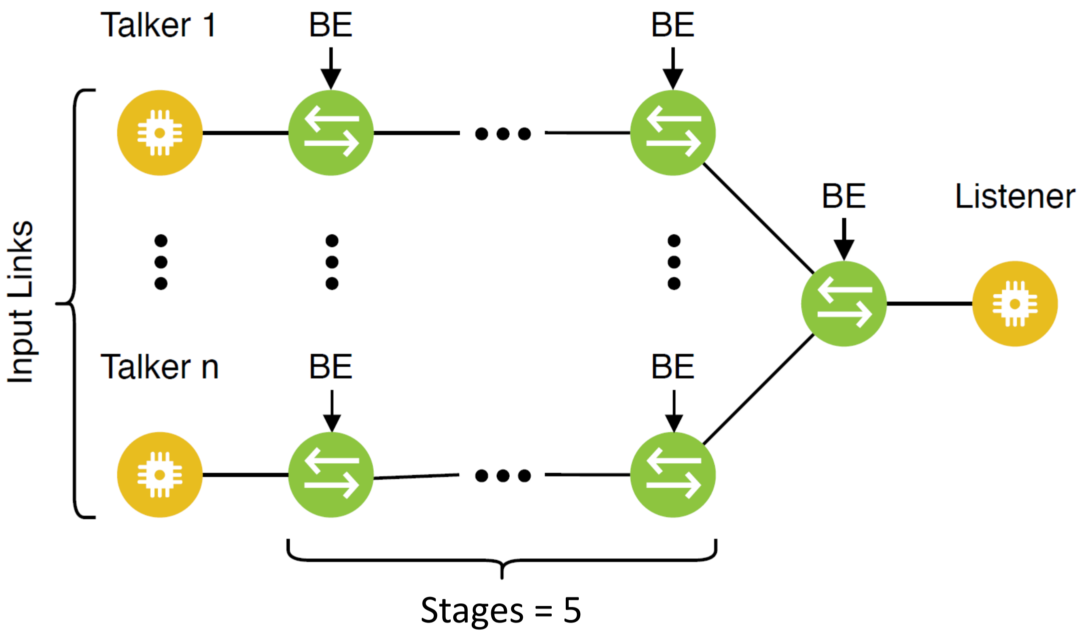

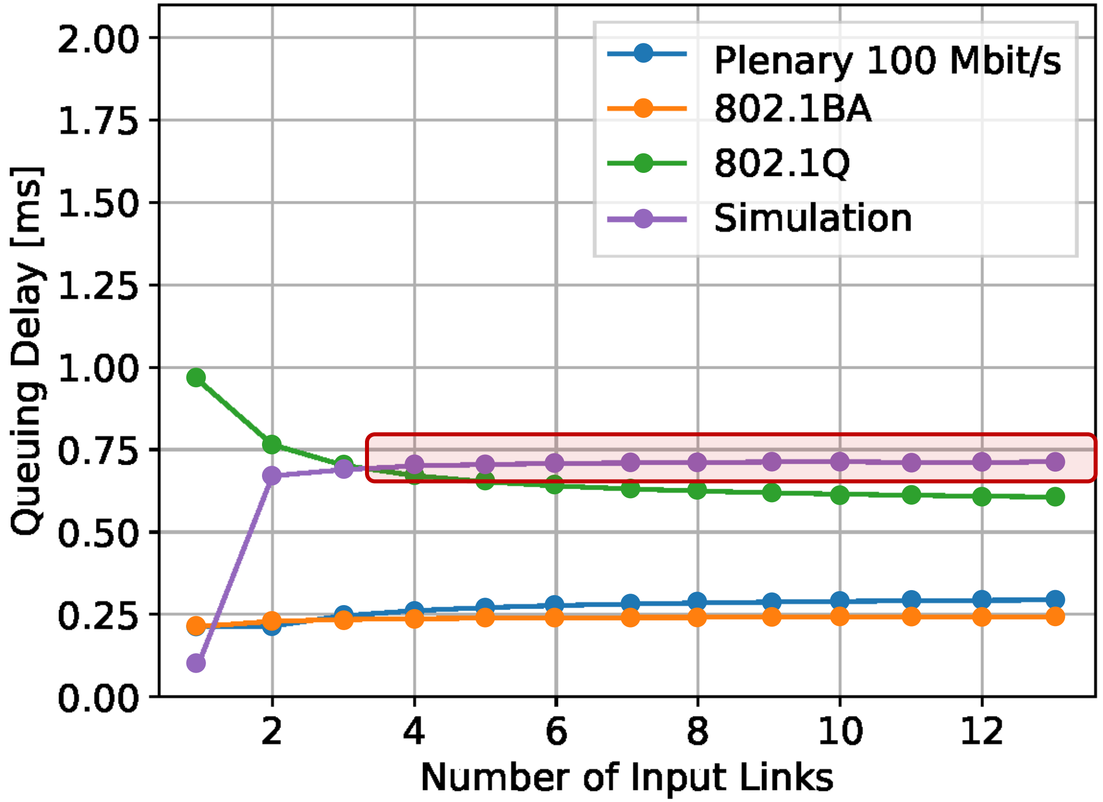

results and analysis from [Maile 2023]

The New Approach

Independent Per-Queue Latency

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

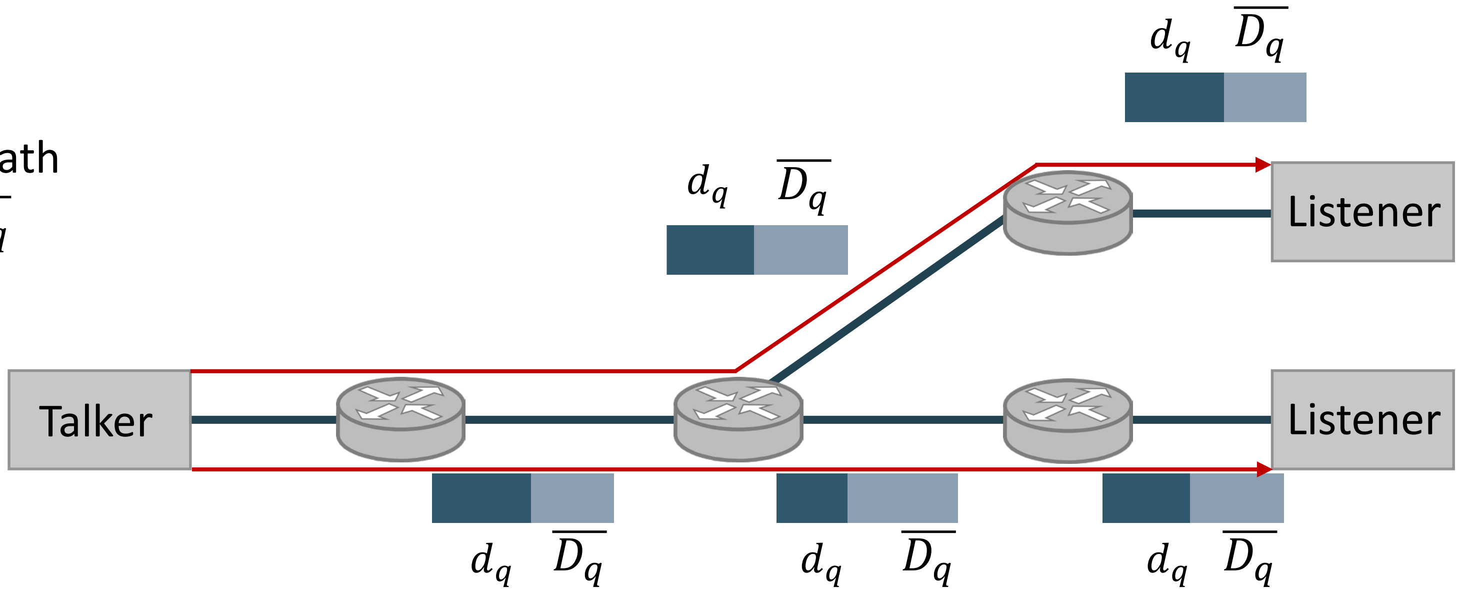

Independent Per-Queue Latency

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

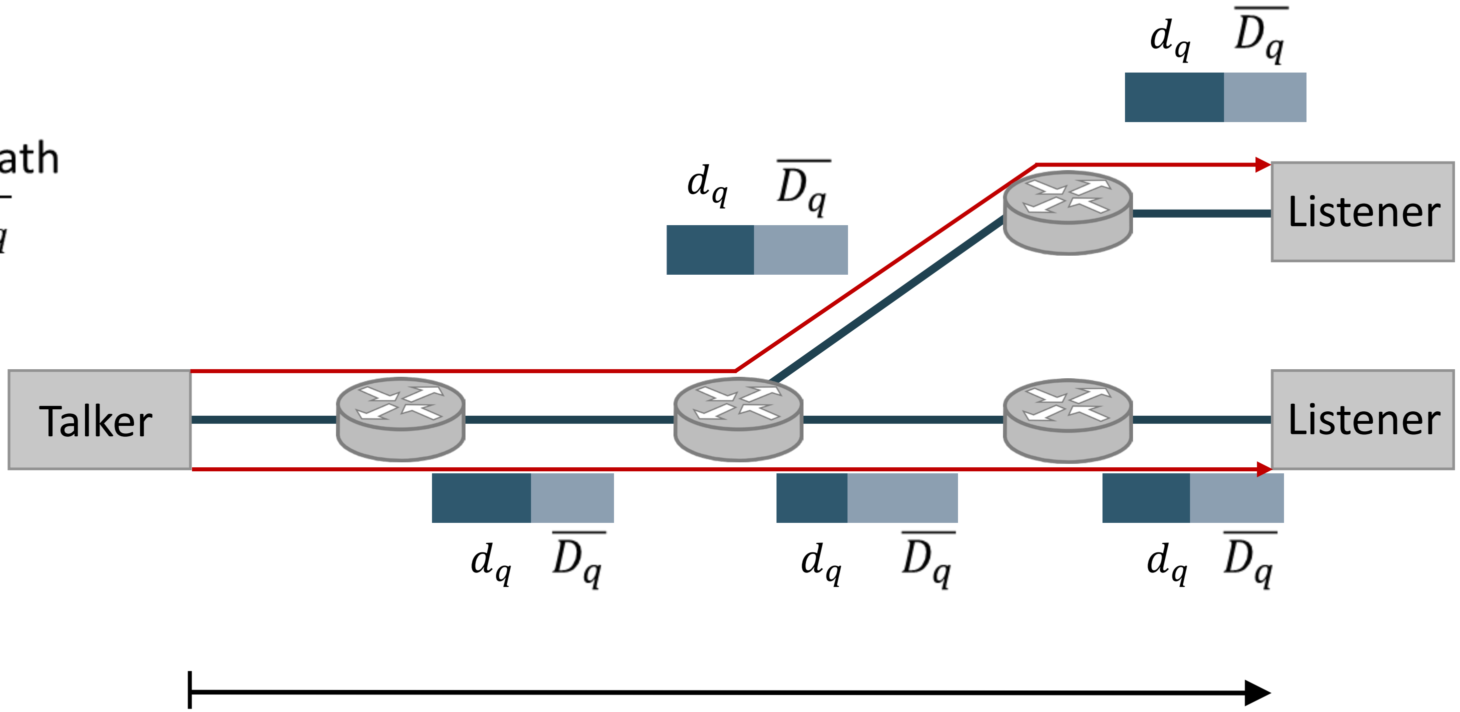

Independent Per-Queue Latency

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

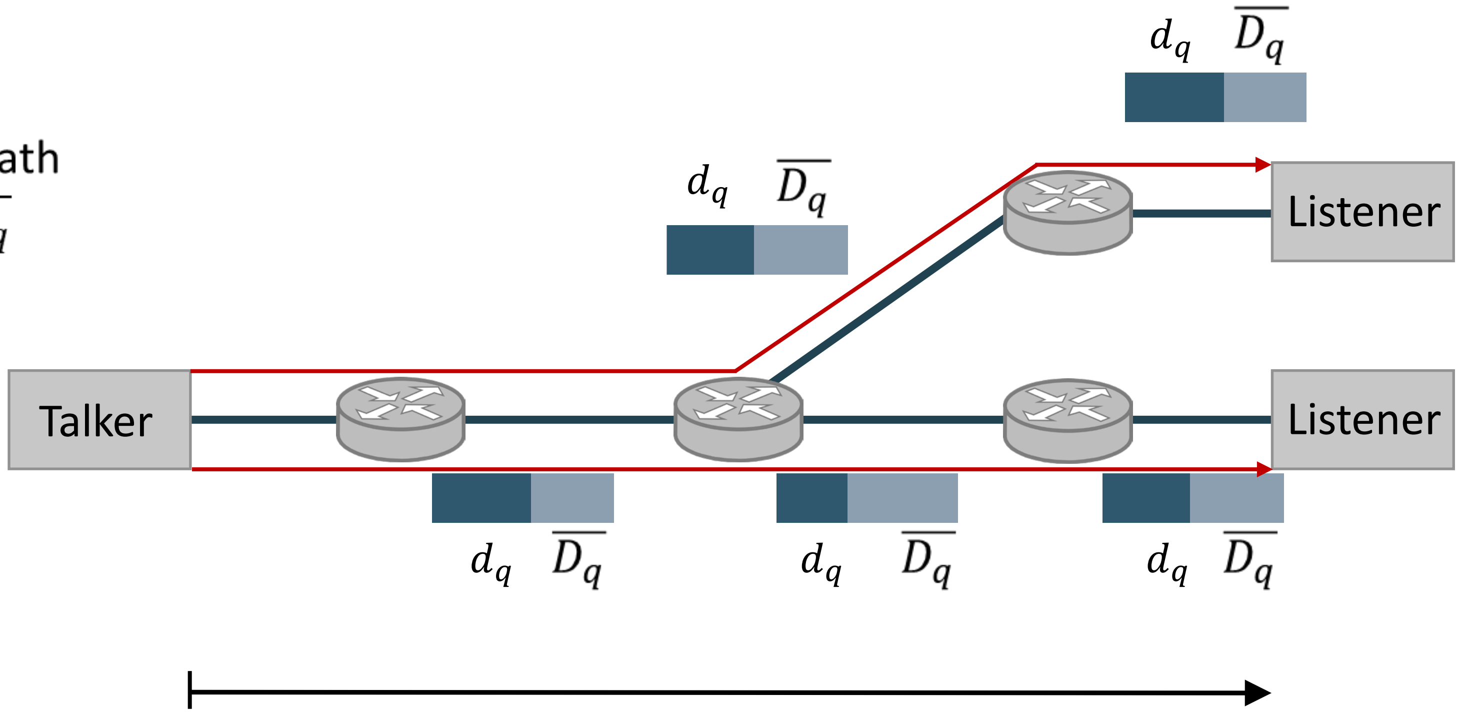

Independent Per-Queue Latency

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

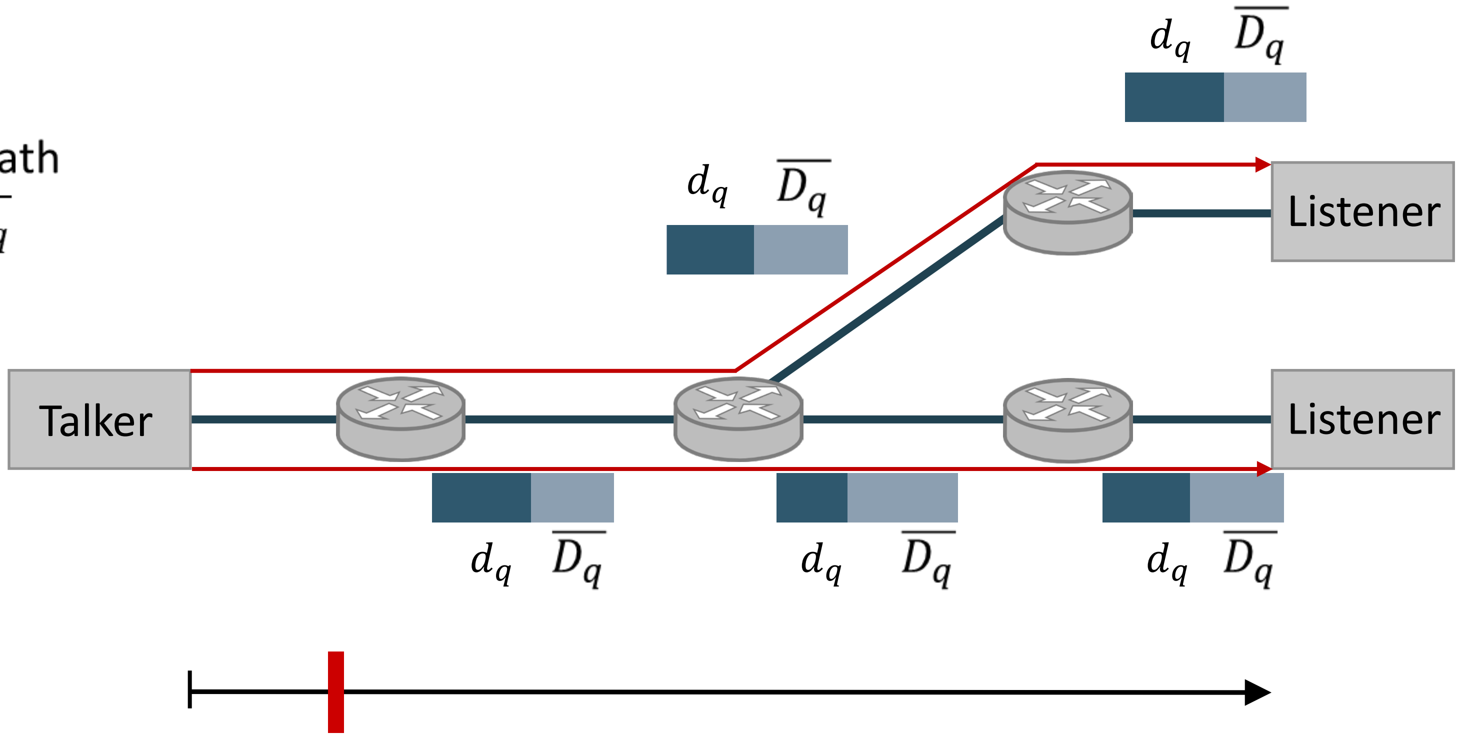

Independent Per-Queue Latency

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

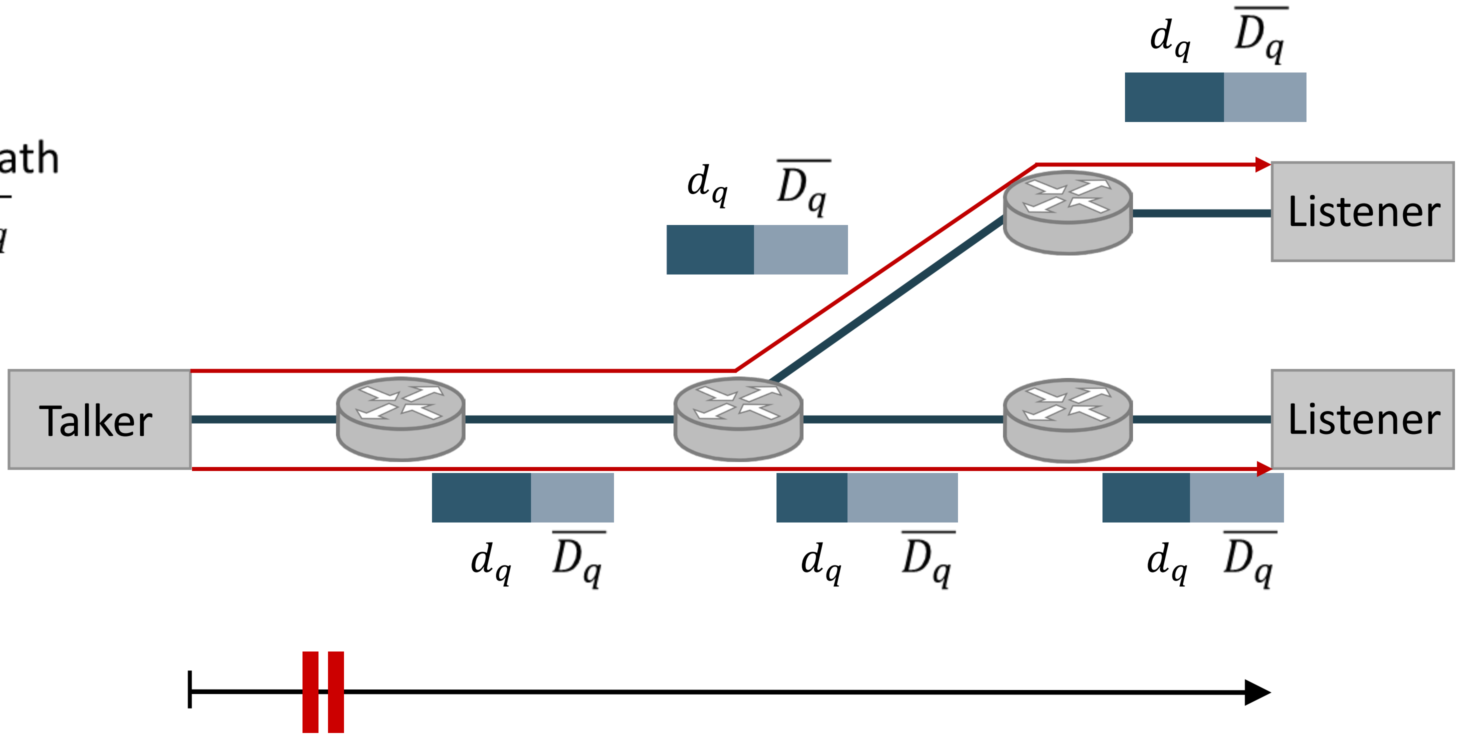

Independent Per-Queue Latency

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

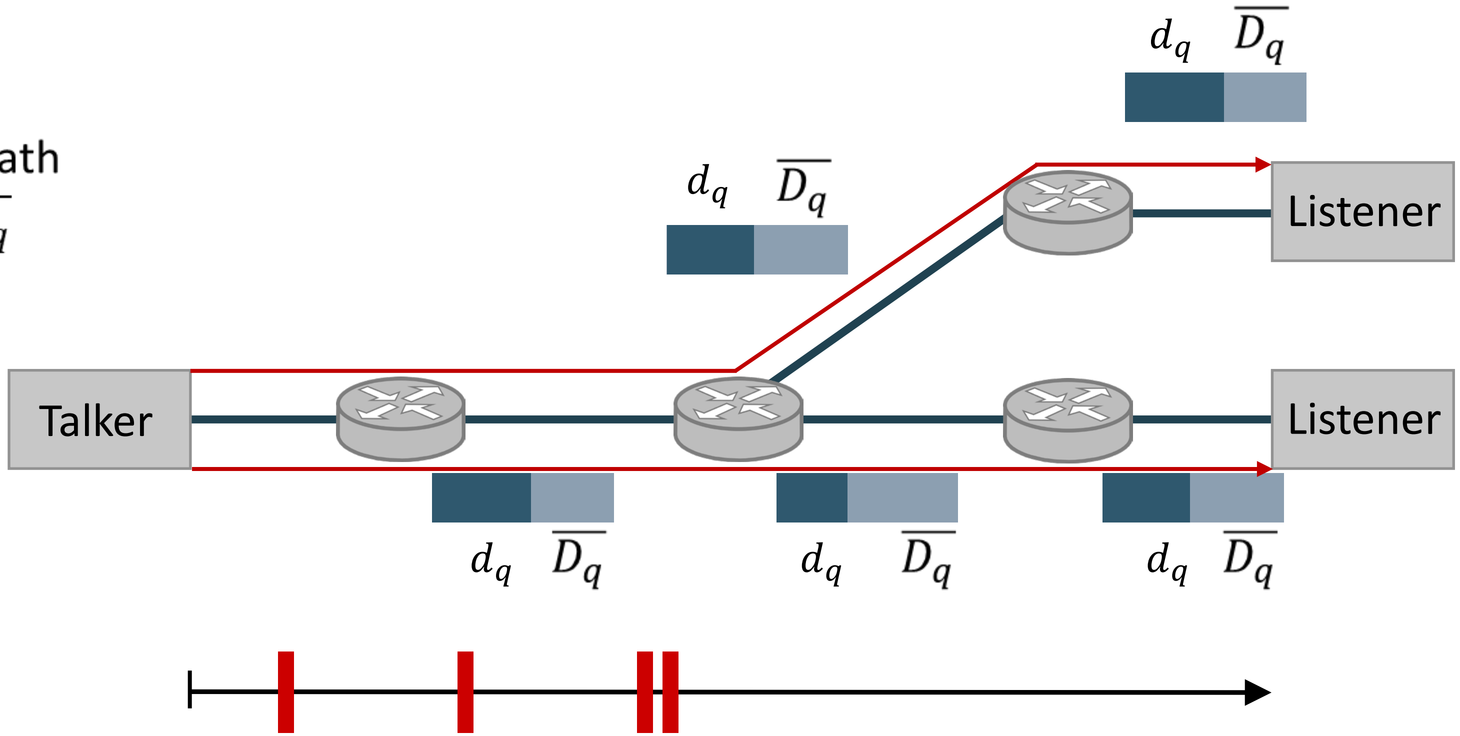

Independent Per-Queue Latency

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

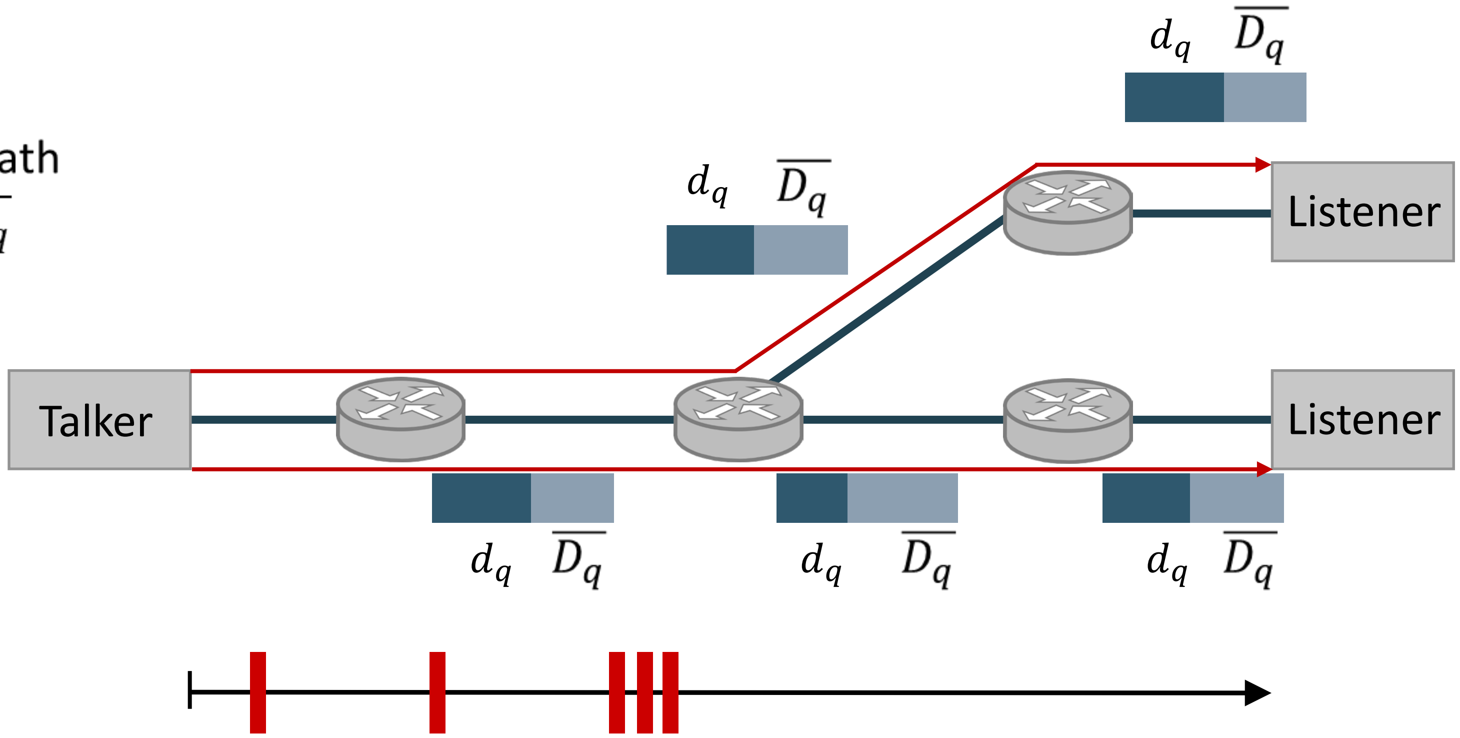

Independent Per-Queue Latency

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

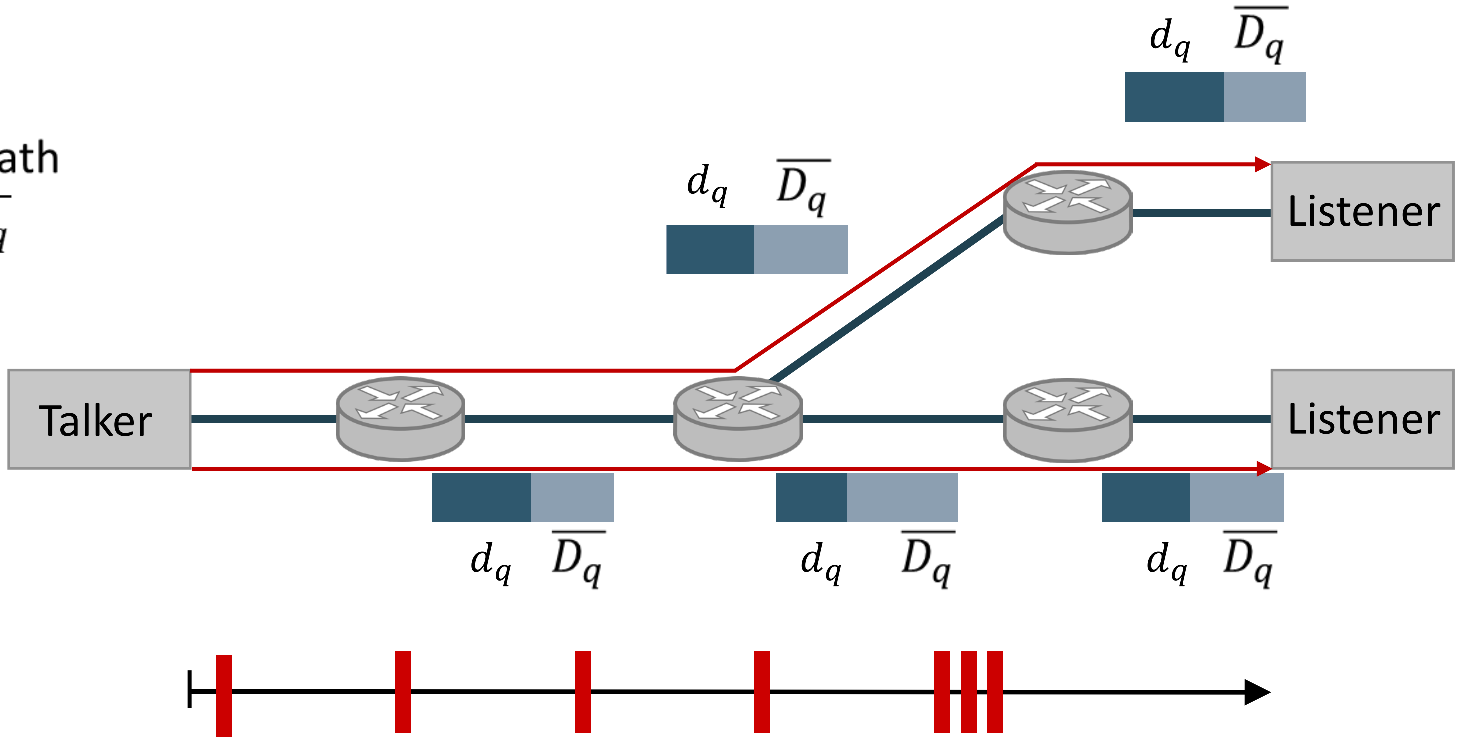

Independent Per-Queue Latency

simplified

Admission Control

- Current Worst-Case Per-Queue Latency d_q
- Per-Queue Delay Bound $\overline{D_q}$

- Allow flow only if: $d_q \leq \overline{D_q}, \forall q$ on path
- Latency bound of flow: Determined by $\overline{D_q}$

Supported Mechanisms

- Strict Priority [Grigorjew 2020]
- Credit-Based Shaper [Maile 2023]
- Asynchronous Traffic Shaper [Grigorjew 2022]
- …

using analytical models (e.g., Network Calculus) or measurement points

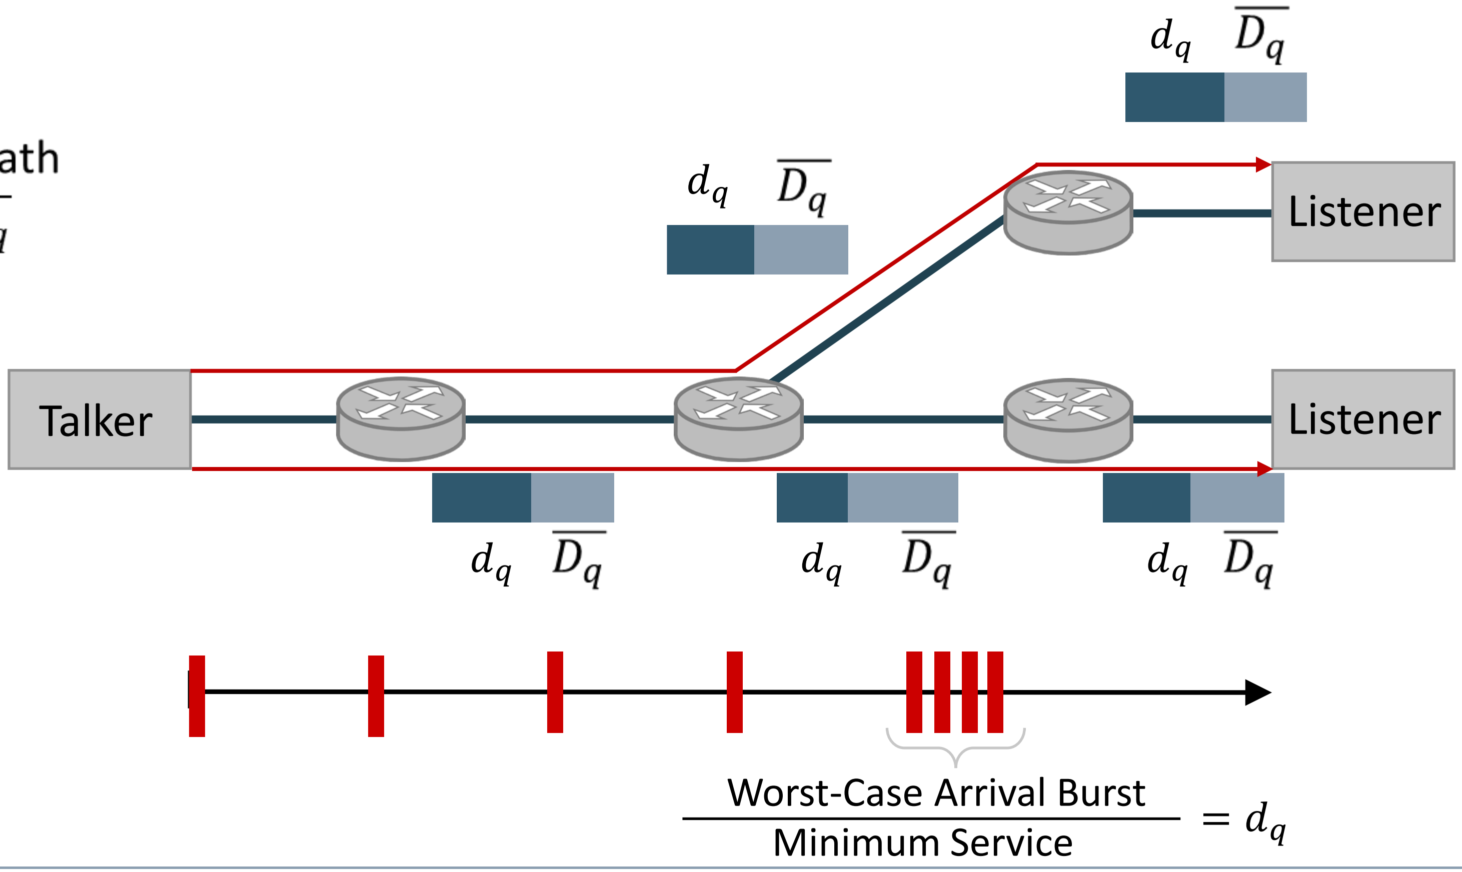

Delay Budgets

Importance

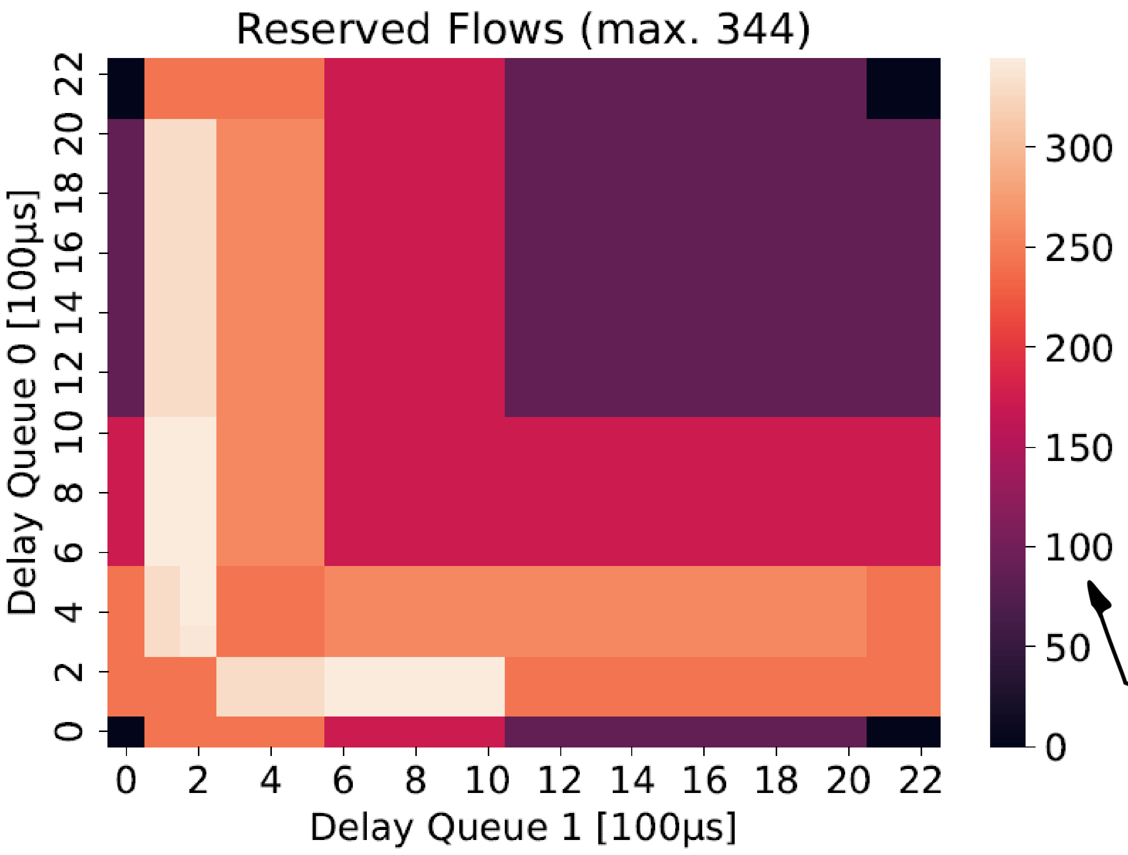

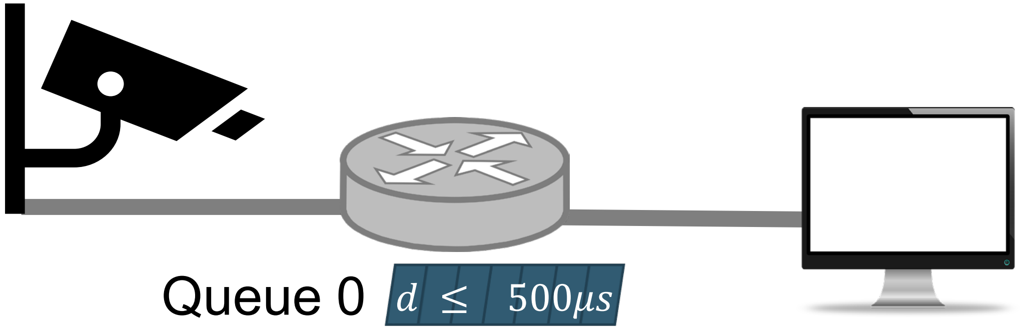

Queue 0 $d \leq 500\mu s$

Queue 1 $d \leq 1200\mu s$

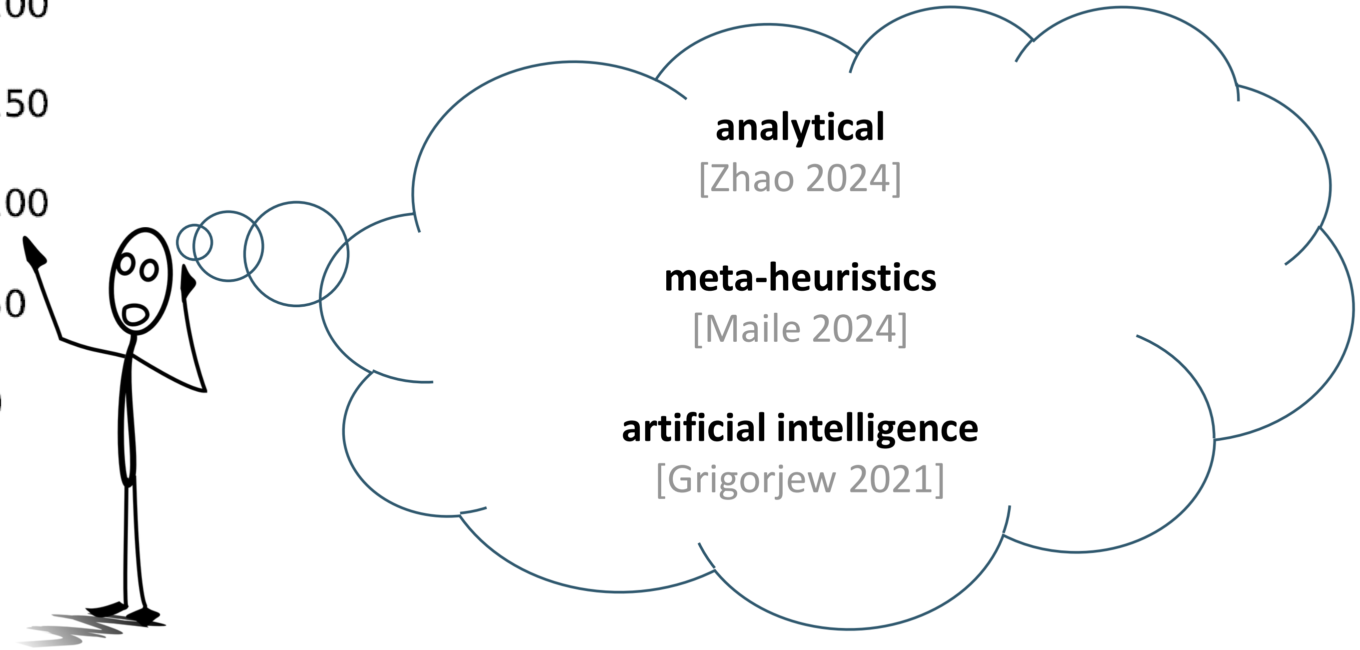

Delay Budgets

Some Insights

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

Benchmark Algorithms

Genetic Algorithm (GA)
Particle Swarm Optimization (PSO)
new solution → to be evaluated

Exhaustive Search (ES)
all solutions → optimum

Intuitive Approach (IA)
deadline of static flows uniformly distributed over path
→ “educated guess”

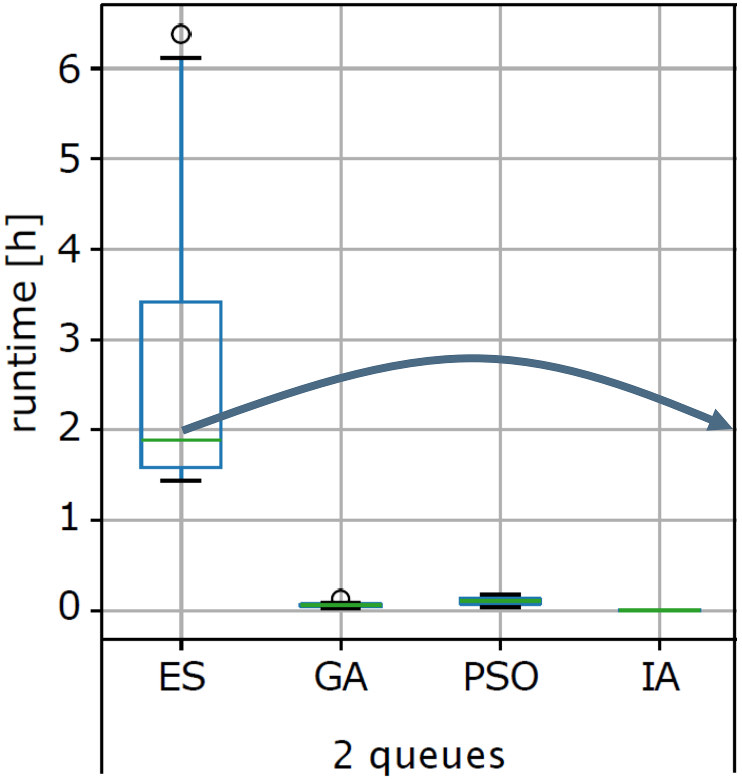

results and analysis from [Maile 2024]

Delay Budgets

Some Insights

Benchmark Algorithms

Genetic Algorithm (GA)
Particle Swarm Optimization (PSO)
new solution → to be evaluated

Exhaustive Search (ES)
all solutions → optimum

Intuitive Approach (IA)
deadline of static flows uniformly distributed over path
→ “educated guess”

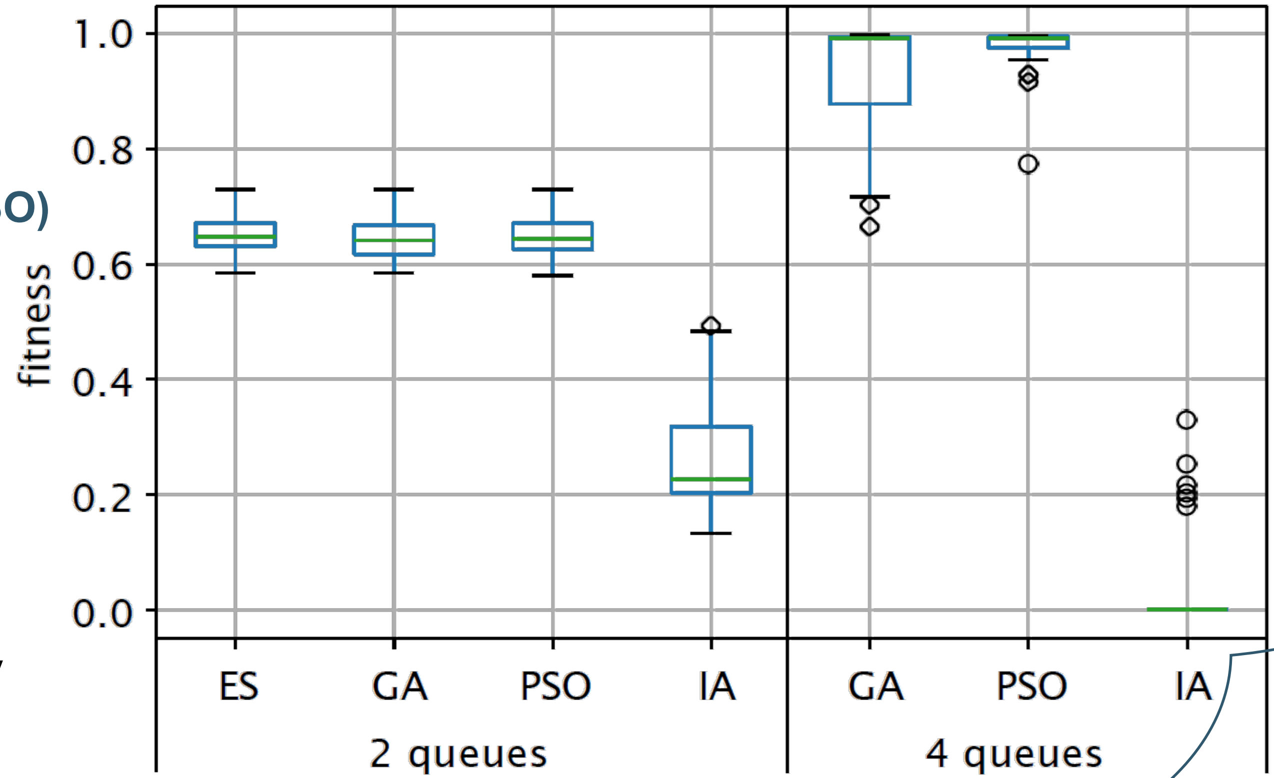

results and analysis from [Maile 2024]

Discussion & Conclusion

Conclusion

INSTITUTE OF COMPUTER AND NETWORK ENGINEERING

Technische Universität Braunschweig

decentralized reservation protocols have a long history and are still ongoing

current problems focus on new hardware **without per flow shaping**

introducing the concept of per-queue delay budgets

future research: choice of per-queue delay budgets

standardization still ongoing

GitHub

Open Source Implementation

https://github.com/Kathess/DYRECTsn

References

[Zhang&Ferrari 1993] H. Zhang and D. Ferrari, “Rate-controlled static-priority queueing,” in IEEE INFOCOM ’93 The Conference on Computer Communications, Proceedings, San Francisco, CA, USA: IEEE Comnput. Soc. Press, 1993, pp. 227–236. doi: 10.1109/INFCOM.1993.253355.

[Frangioni 2017] A. Frangioni, L. Galli, and G. Stea, “QoS routing with worst-case delay constraints: Models, algorithms and performance analysis,” *Computer Communications*, vol. 103, pp. 104–115, May 2017, doi: 10.1016/j.comcom.2016.09.006.

[Grigorjew 2020] A. Grigorjew et al., "Bounded Latency with Bridge-Local Stream Reservation and Strict Priority Queuing," 2020 11th International Conference on Network of the Future (NoF), Bordeaux, France, 2020, pp. 55-63, doi: 10.1109/NoF50125.2020.9249224.

[Grigorjew 2021] A. Grigorjew, M. Seufert, N. Wehner, J. Hofmann, and T. Hoßfeld, “ML-Assisted Latency Assignments in Time-Sensitive Networking,” in *2021 IFIP/IEEE International Symposium on Integrated Network Management (IM)*, May 2021, pp. 116–124.

[Grigorjew 2022] A. Grigorjew et al., "Constant Delay Switching: Asynchronous Traffic Shaping with Jitter Control," 2022 IFIP Networking Conference (IFIP Networking), Catania, Italy, 2022, pp. 1-9, doi: 10.23919/IFIPNetworking55013.2022.9829777.

[Maile 2023] L. Maile, D. Voitlein, A. Grigorjew, K.-S. J. Hielscher, and R. German, “On the Validity of Credit-Based Shaper Delay Guarantees in Decentralized Reservation Protocols,” in *Proceedings of the 31st International Conference on Real-Time Networks and Systems*, in RTNS ’23. New York, NY, USA: Association for Computing Machinery, Jun. 2023, pp. 108–118. doi: 10.1145/3575757.3593644.

[Maile 2024] L. Maile, K.-S. Hielscher, and R. German, “Combining Static and Dynamic Traffic with Delay Guarantees in Time-Sensitive Networking,” in *Performance Evaluation Methodologies and Tools*, vol. 539, E. Kalyvianaki and M. Paolieri, Eds., in Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering, vol. 539. , Cham: Springer Nature Switzerland, 2024, pp. 117–132. doi: 10.1007/978-3-031-48885-6_8.

[Zhao 2024] L. Zhao, Y. Yan, and X. Zhou, “Minimum Bandwidth Reservation for CBS in TSN With Real-Time QoS Guarantees,” IEEE Trans. Ind. Inf., vol. 20, no. 4, pp. 6187–6198, Apr. 2024, doi: 10.1109/TII.2023.3342466.

Thank you!

More information?

Lisa Maile, lisa.maile@tu-bs.de
Institute of Computer and Network Engineering
TU Braunschweig, Germany