University of Stuttgart
Institute of Parallel and Distributed Systems[1]
Institute of Systems Theory and Automatic Control[2]

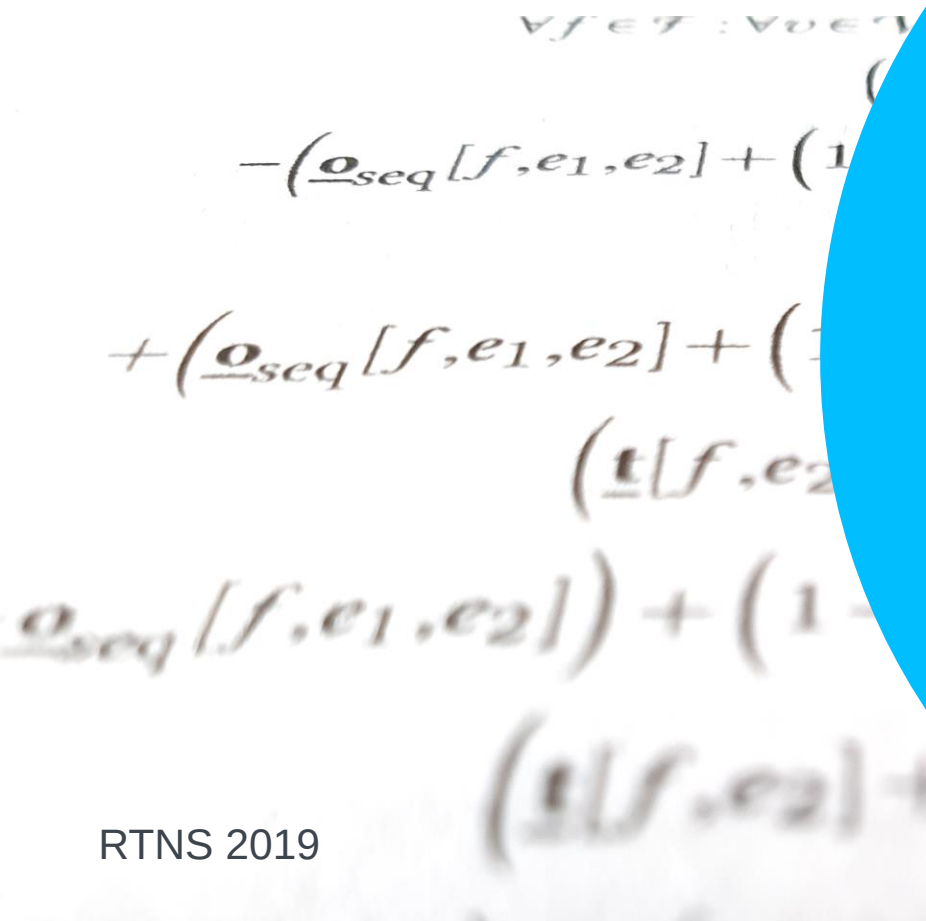

Optimal Routing and Scheduling of Complemental Flows in Converged Networks

Jonathan Falk[1], Frank Dürr[1], Steffen Linsenmayer[2], Stefan Wildhagen[2], Ben Carabelli[1], Kurt Rothermel[1]

RTNS 2019

The age of the cyber-physical machine.

More distributed systems interfacing with the physical world

- Manipulate the physical world via computers
- “Smart” {city, factory, home}
- Industry 4.0
- Autonomous Driving
- IEEE Time-sensitive Networking (TSN) Workgroup

The age of the cyber-physical machine.

More distributed systems interfacing with the physical world

- Manipulate the physical world via computers
- “Smart” {city, factory, home}
- Industry 4.0
- Autonomous Driving
- IEEE Time-sensitive Networking (TSN) Workgroup

requires real-time communication

But! There’s a catch.

Ubiquitous quality requirements

- message delivery with bounded delay & jitter
- introduced by “coupling” with the physical world
- lack of quality potentially disastrous
- often implemented with time-triggered paradigm with static schedule

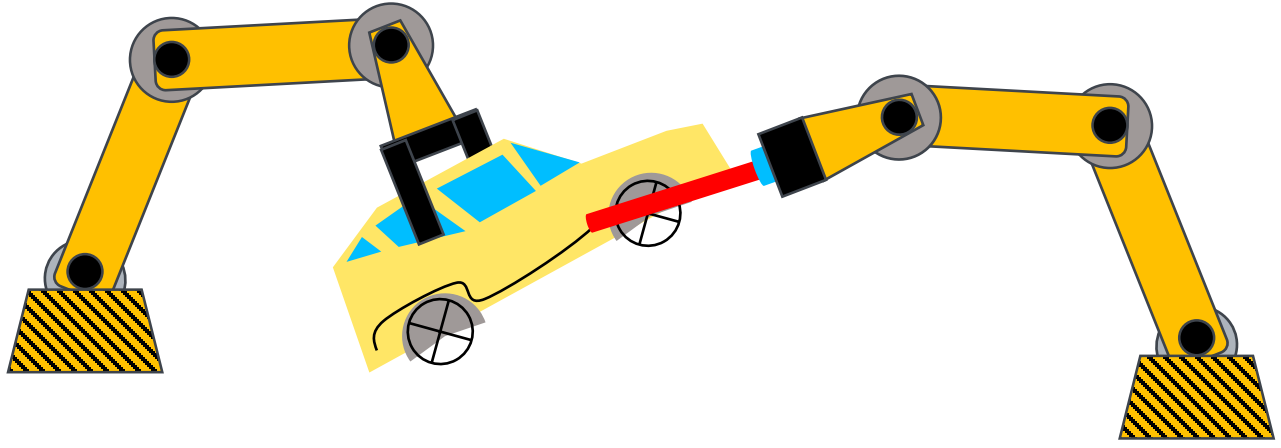

But! There’s a catch.

Applications can benefit from dynamic communication

- cf. in “stable” position

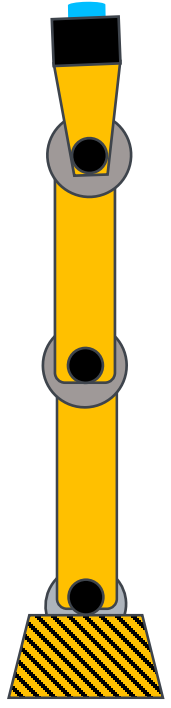

But! There’s a catch.

Applications can benefit from dynamic communication

- cf. in “stable” position

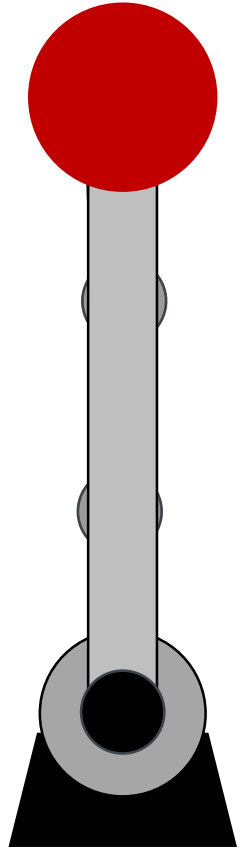

But! There's a catch.

Applications can benefit from dynamic communication

- cf. in "stable" position

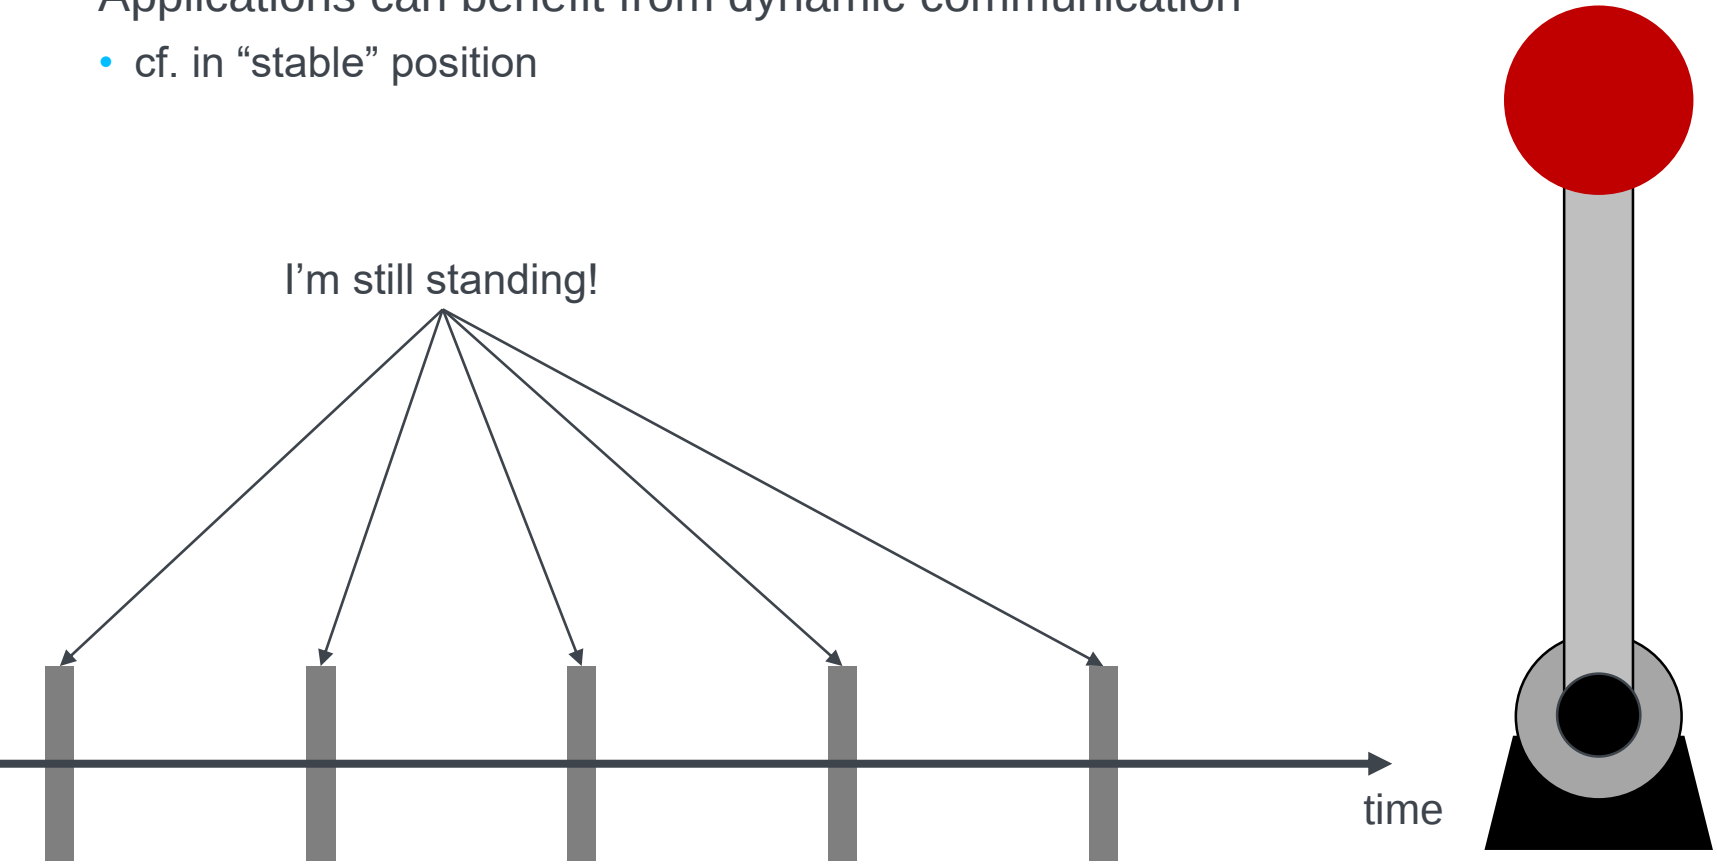

But! There’s a catch.

Applications can benefit from dynamic communication

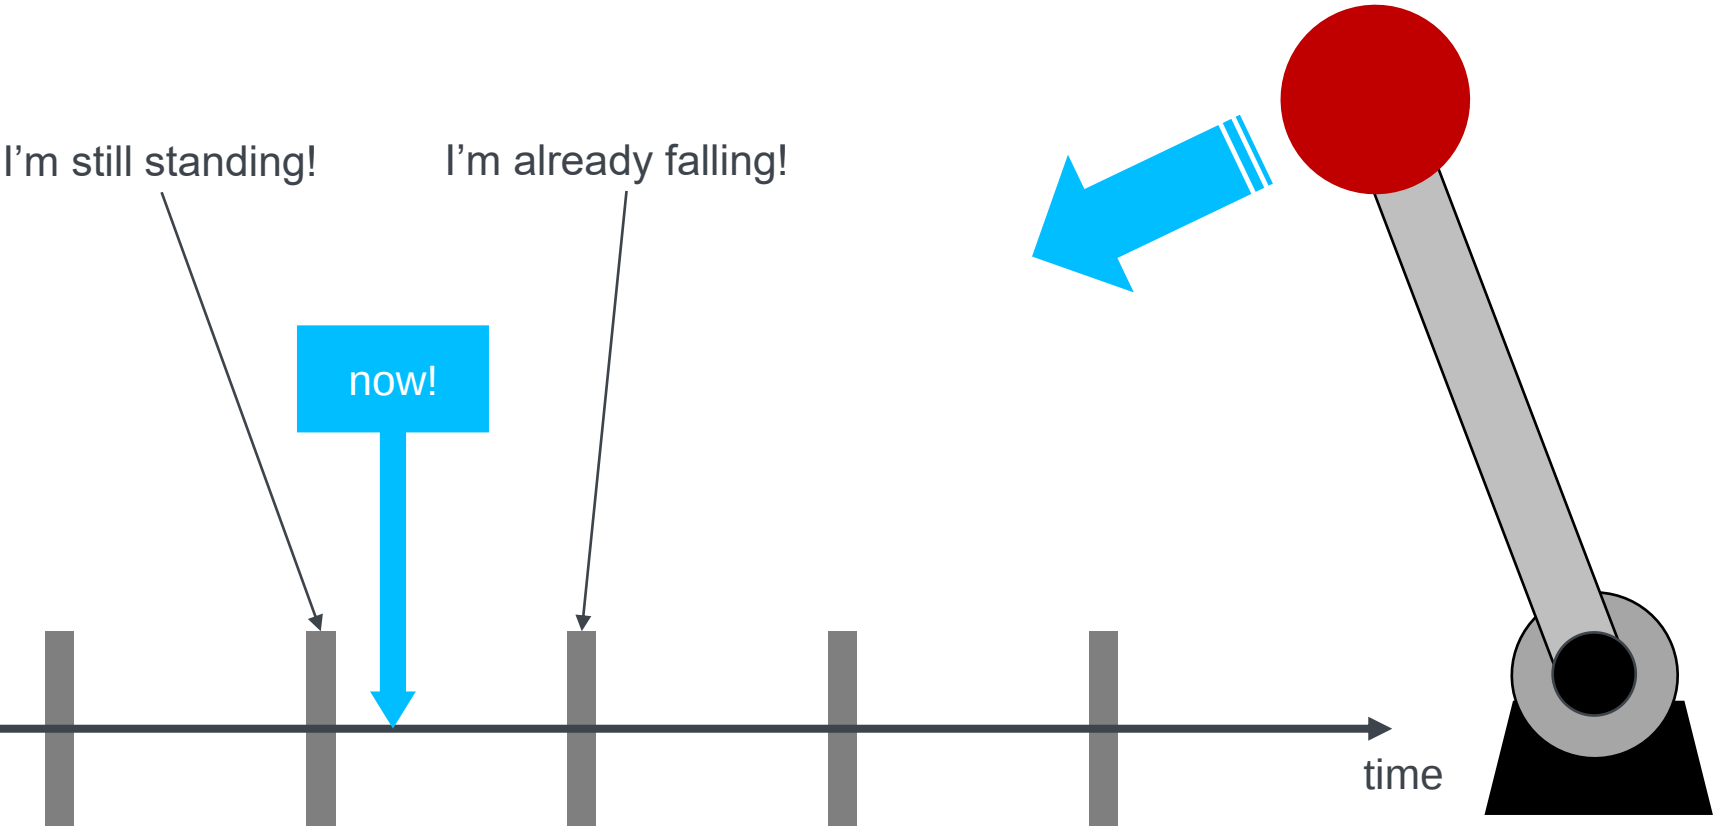

But! There's a catch.

Applications can benefit from dynamic communication

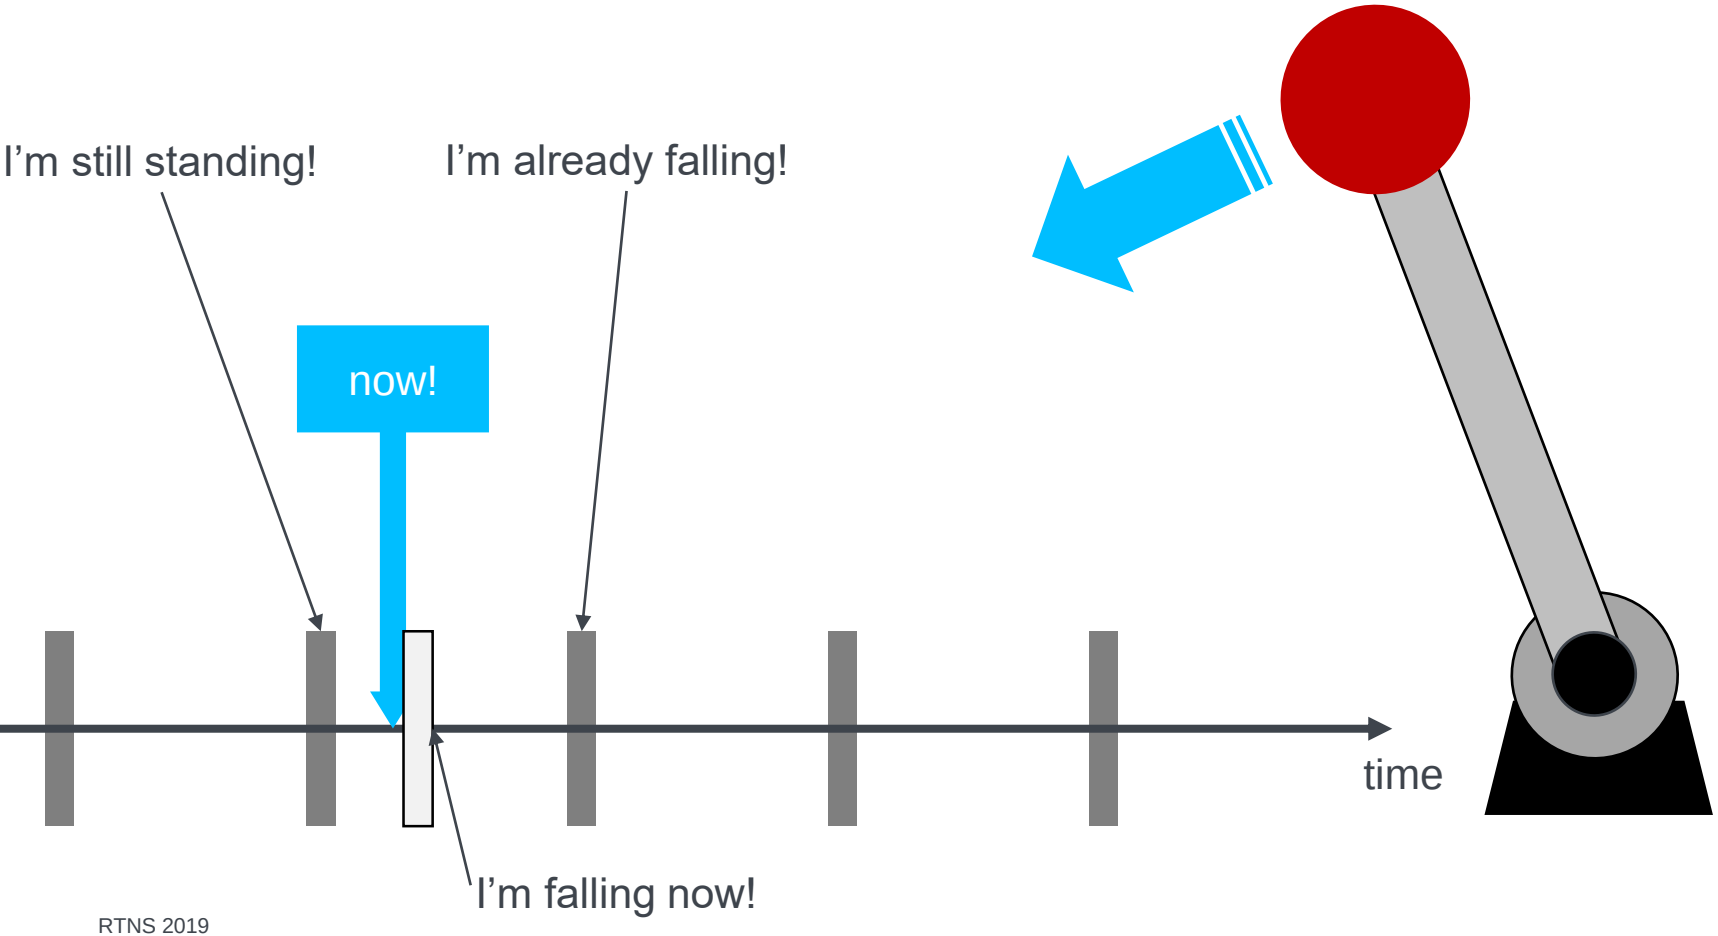

But! There’s a catch.

Messages can have different “importance”

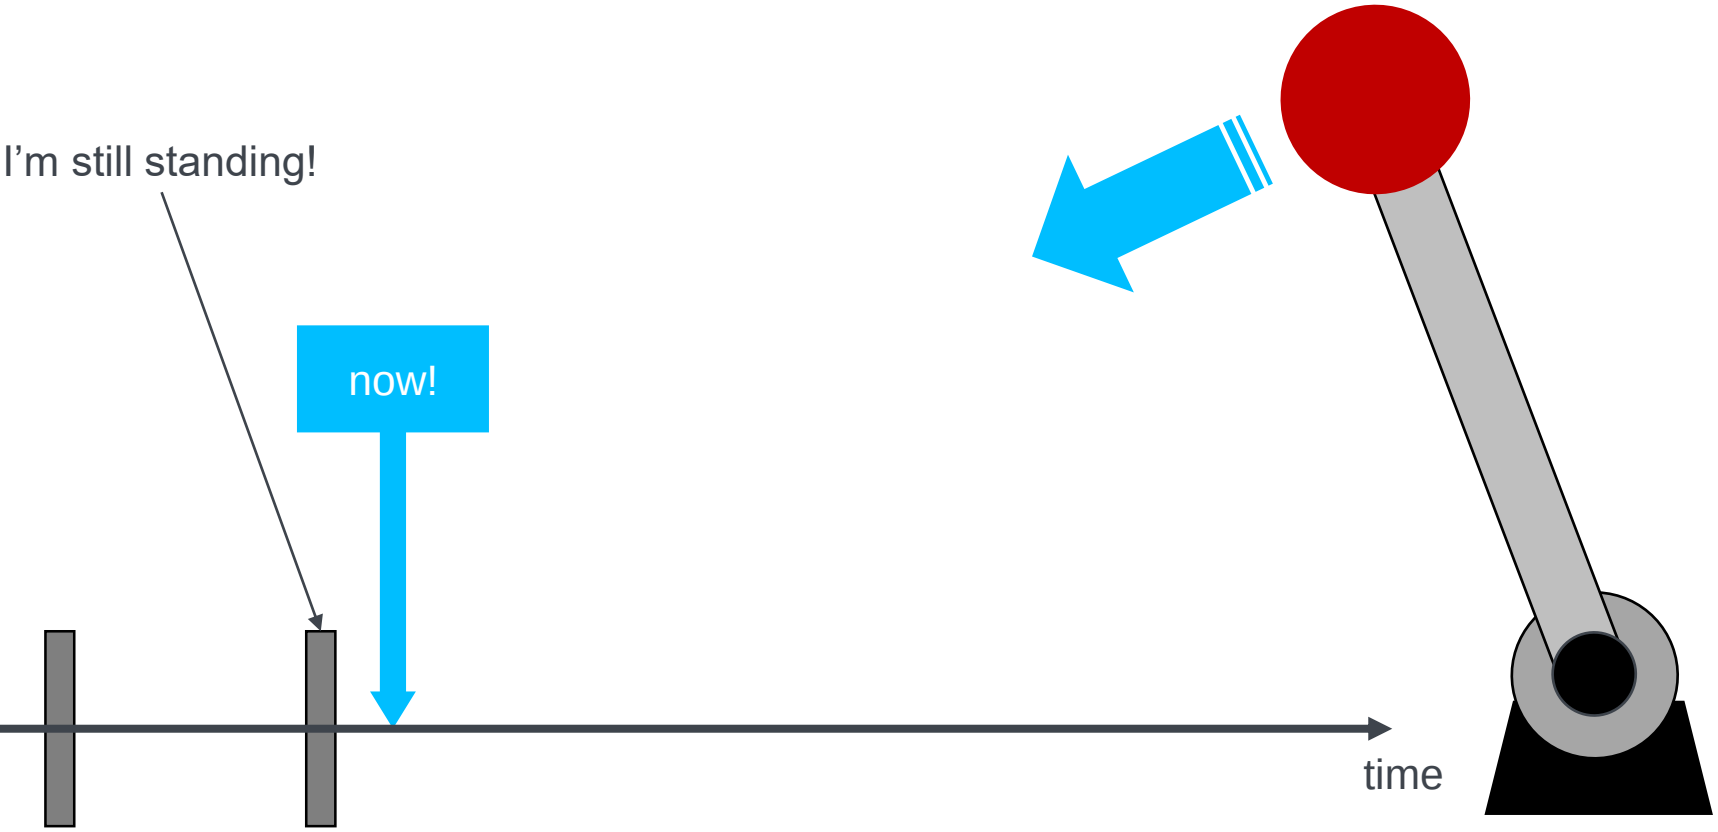

But! There's a catch.

Messages can have different "importance"

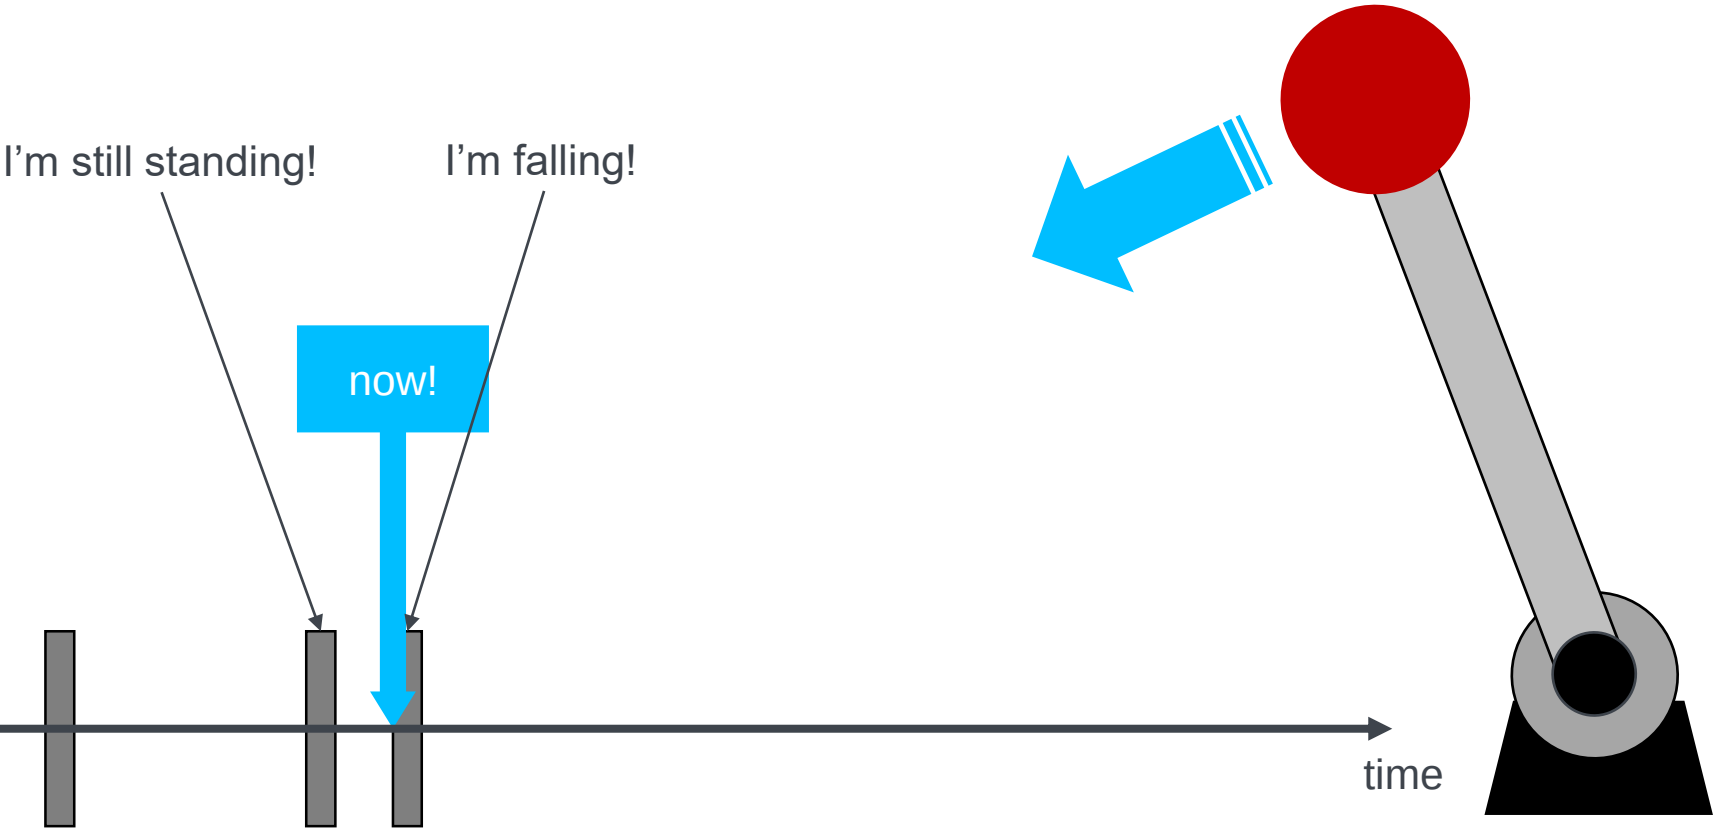

But! There's a catch.

Messages can have different "importance"

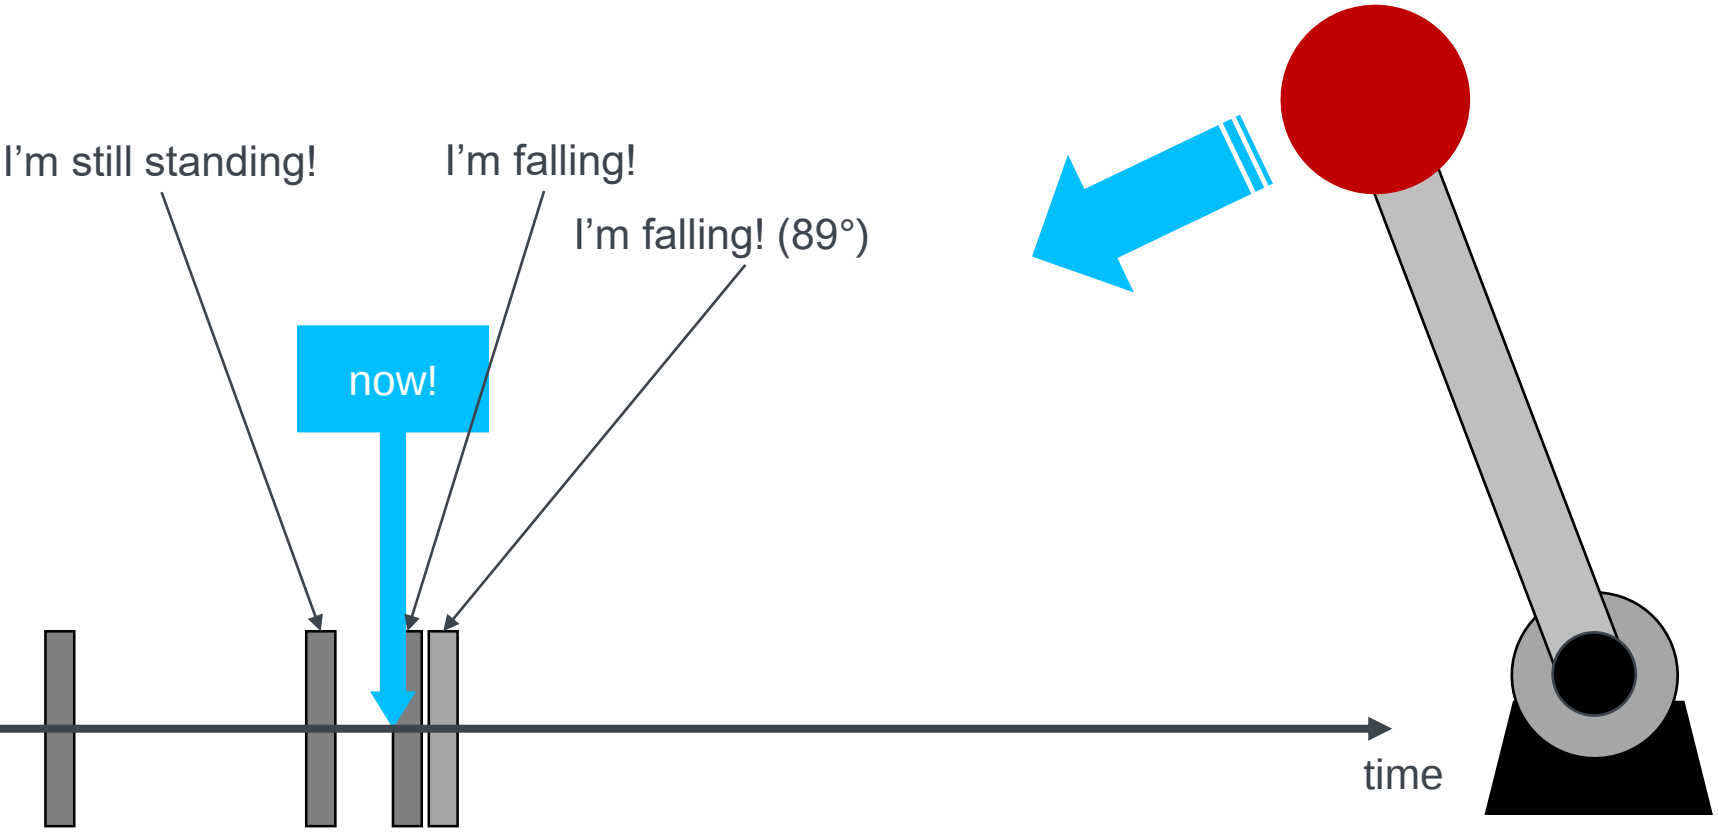

But! There’s a catch.

Messages can have different “importance”

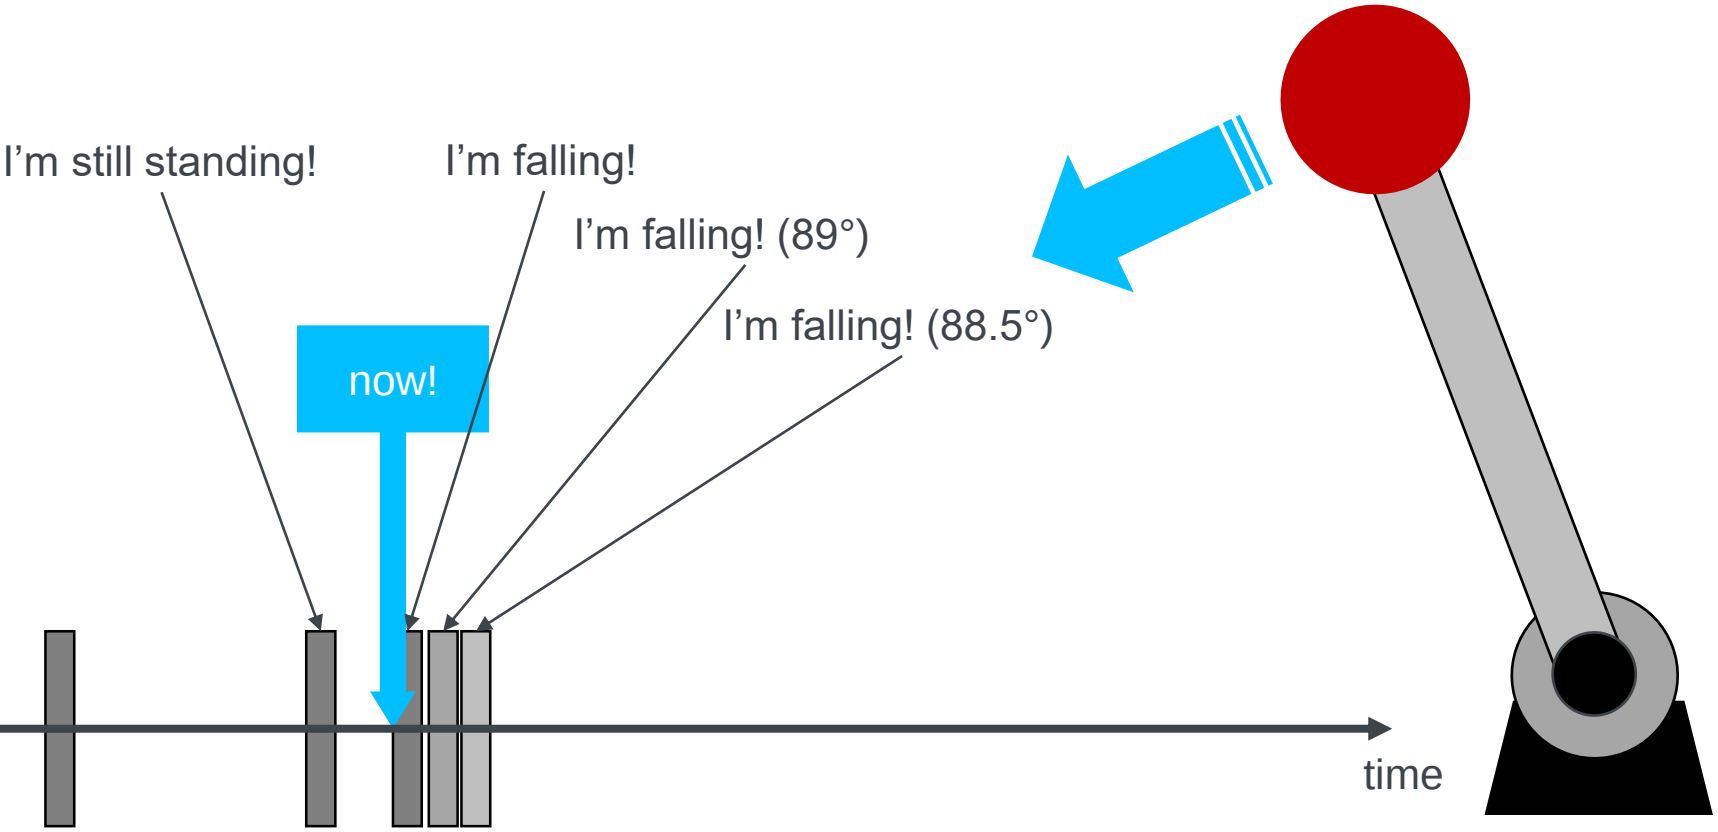

But! There's a catch.

Messages can have different "importance"

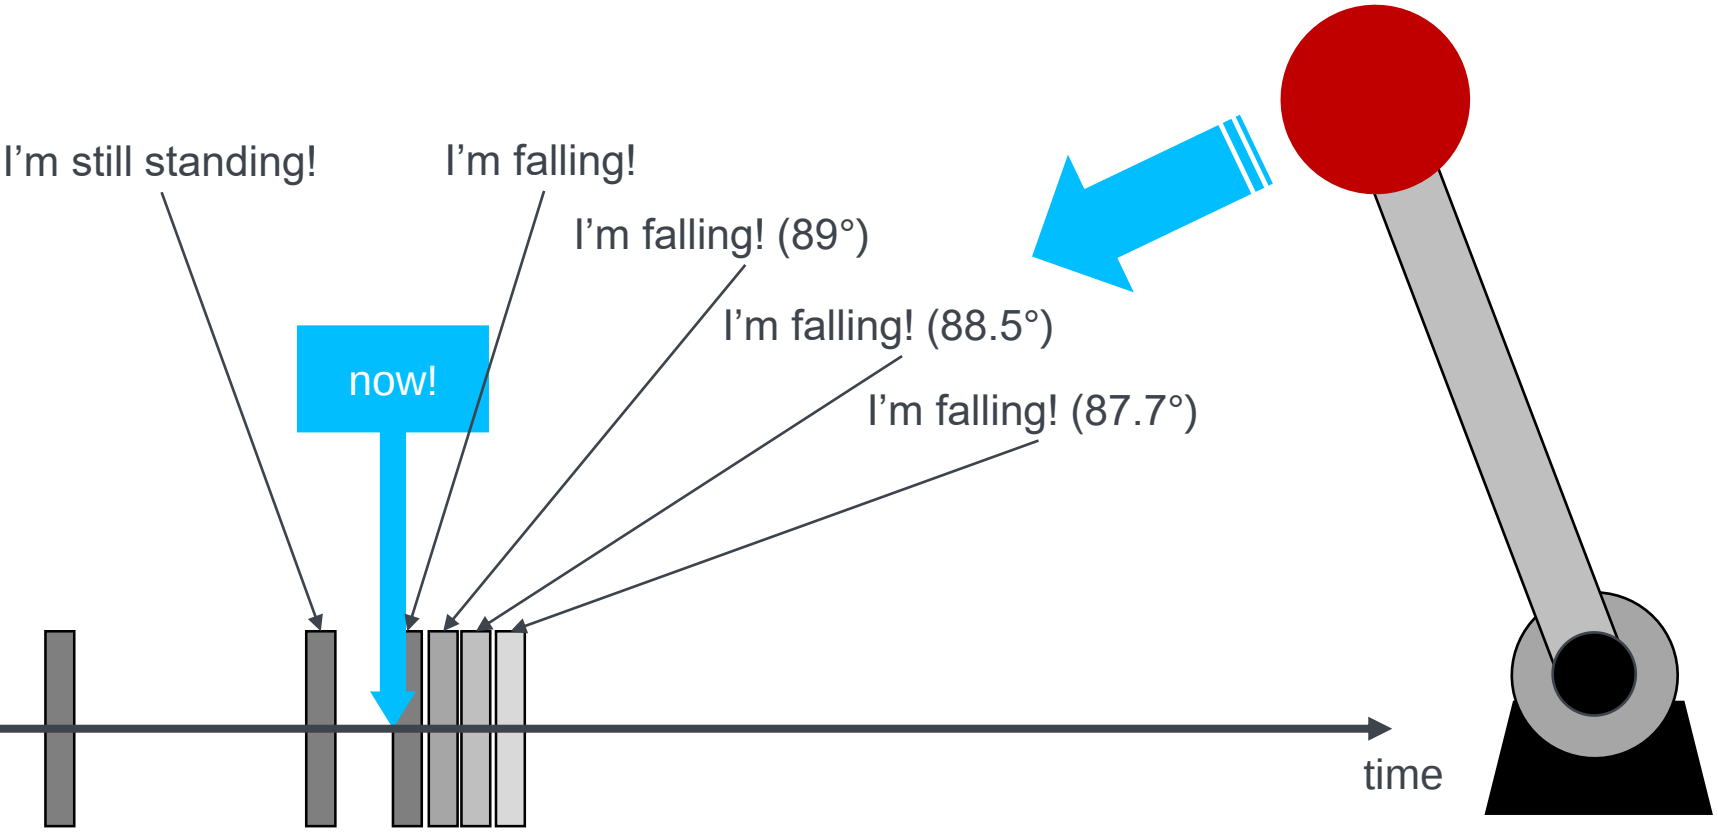

But! There's a catch.

Messages can have different "importance"

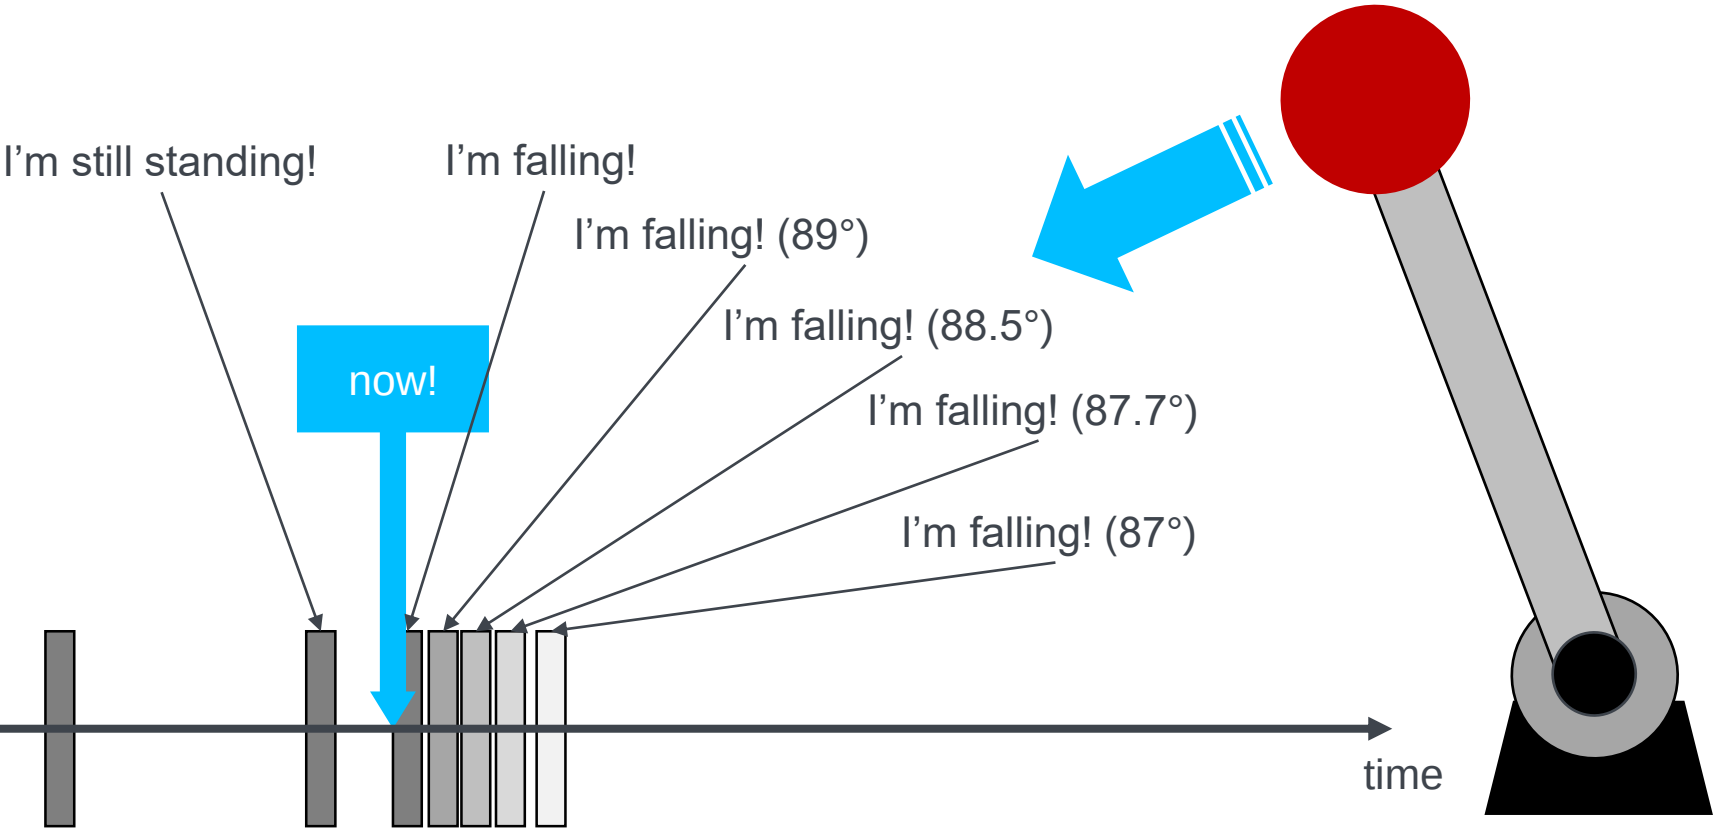

But! There's a catch.

Messages can have different "importance"

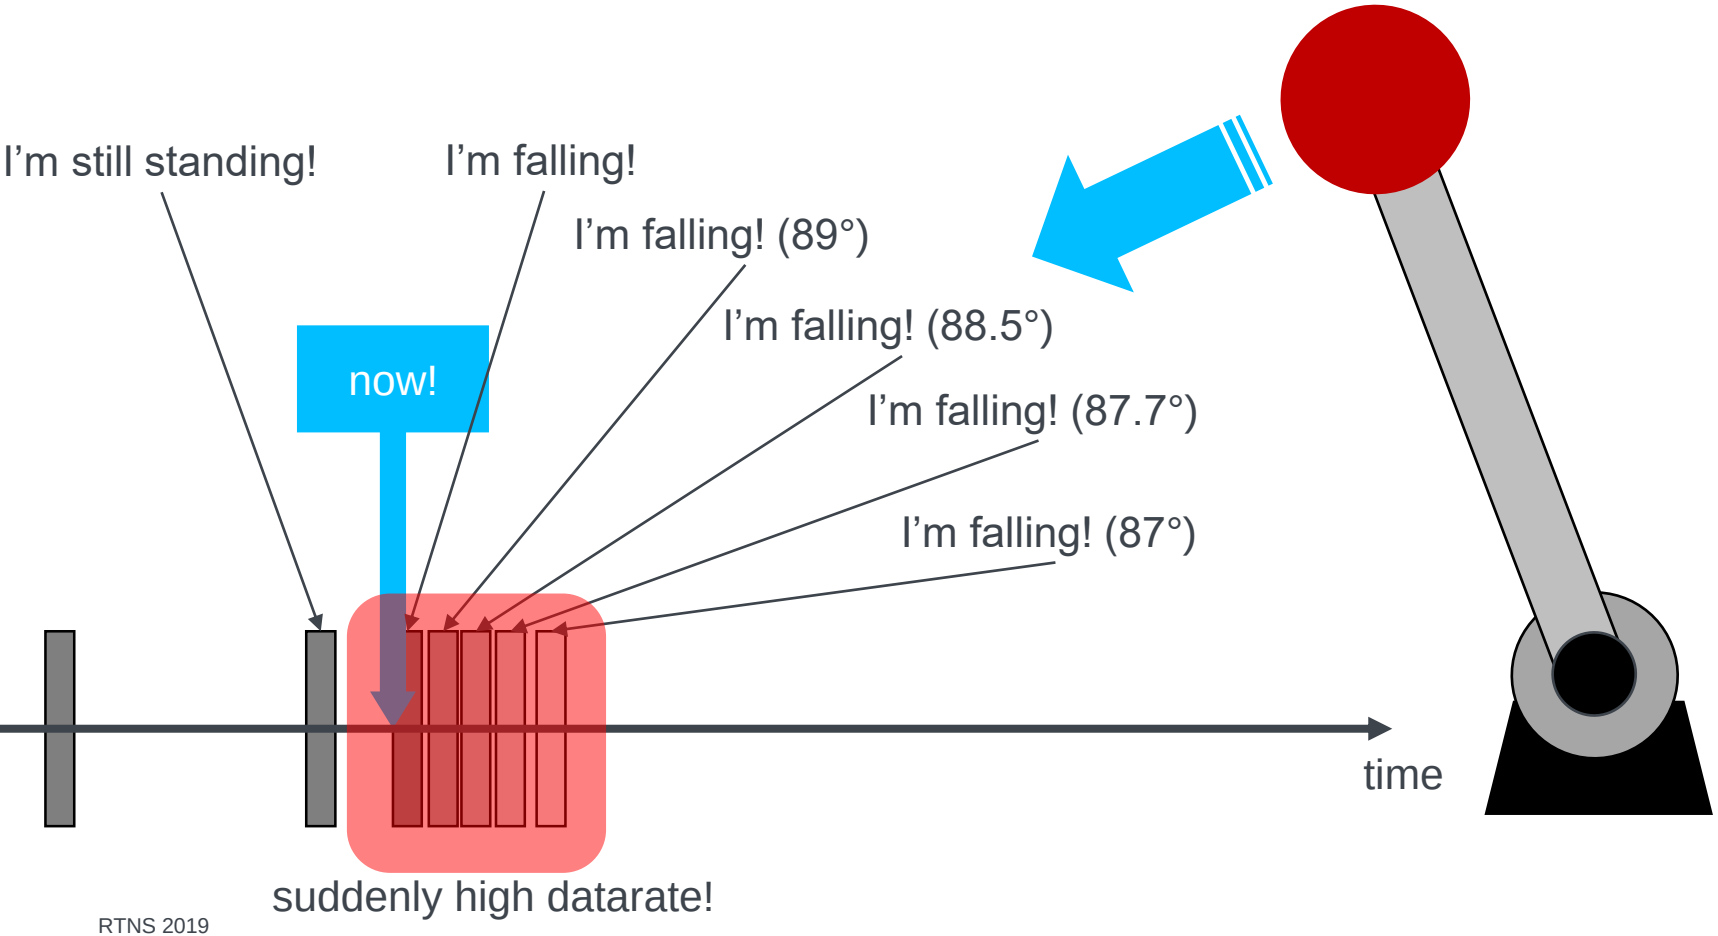

But! There’s a catch.

Conflicts for real-time communication in shared networks

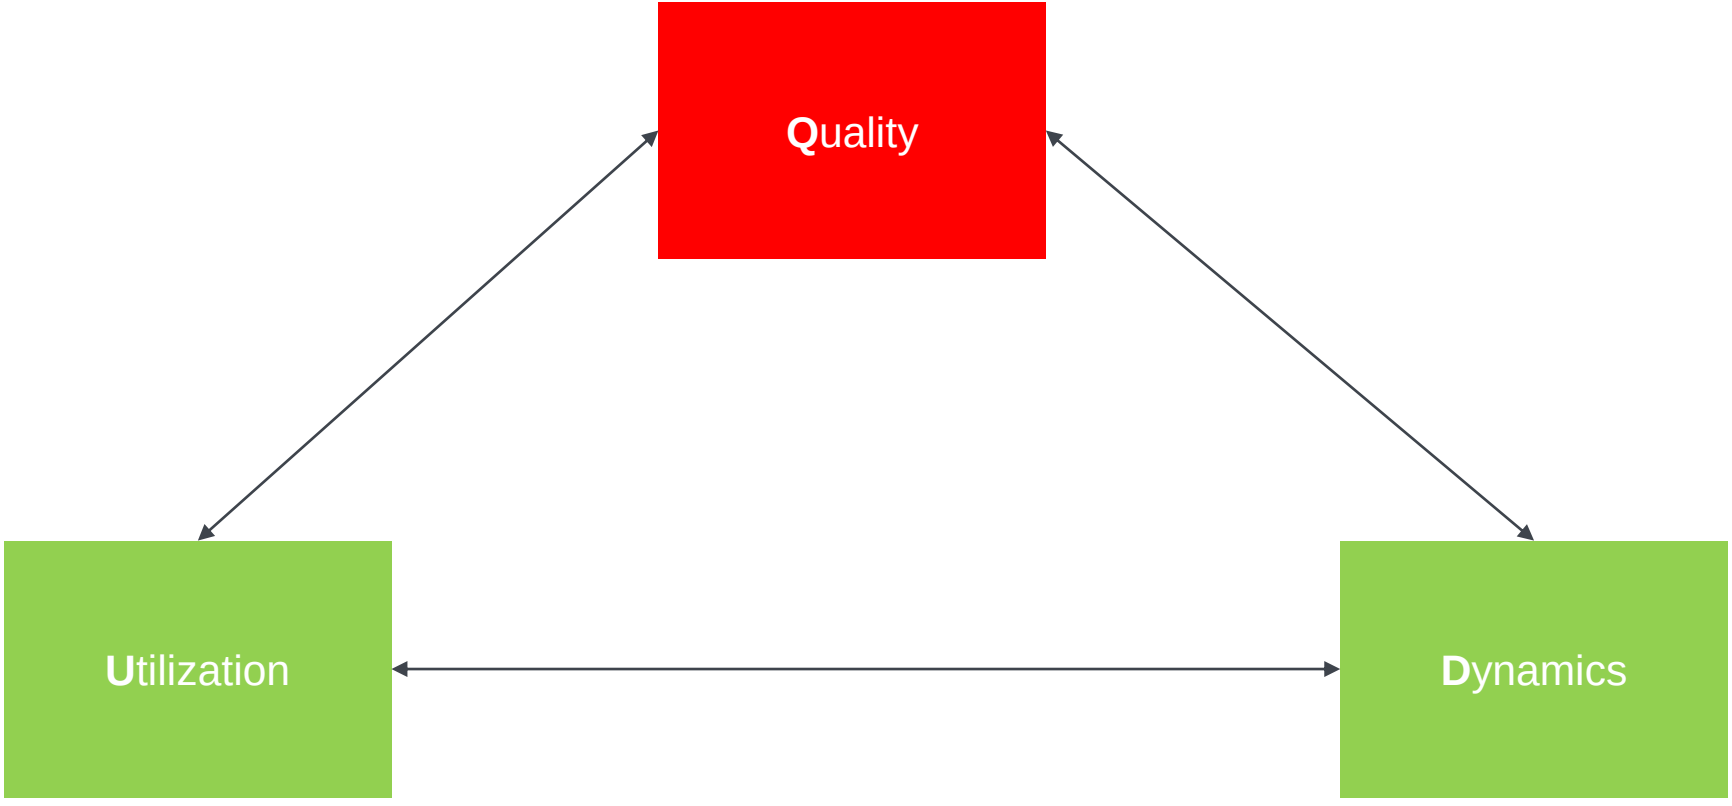

But! There’s a catch.

Conflicts for real-time communication in shared networks

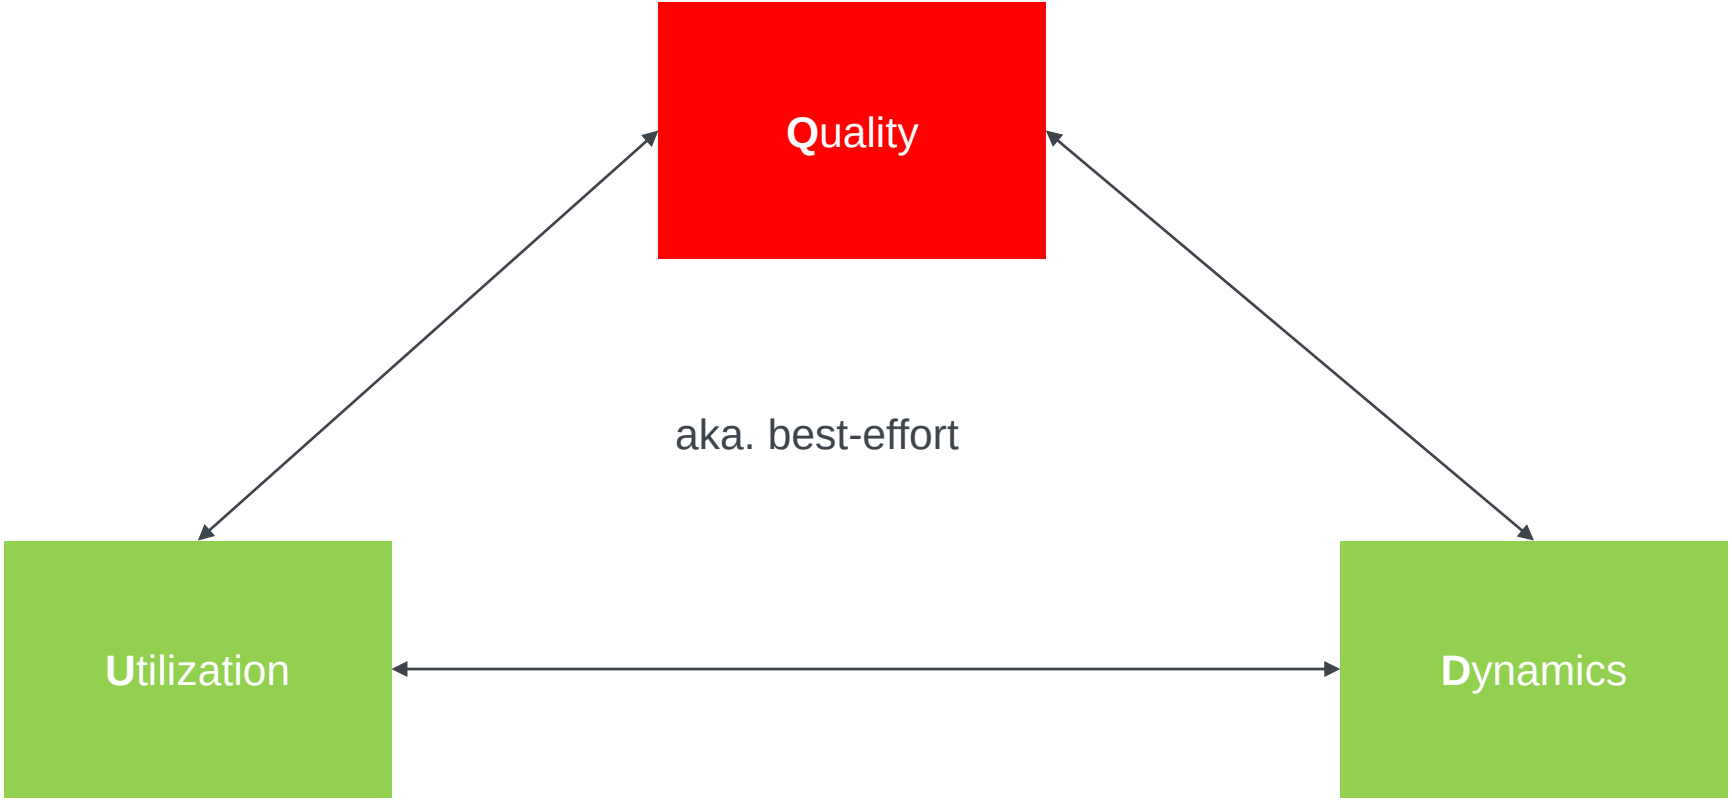

But! There’s a catch.

Conflicts for real-time communication in shared networks

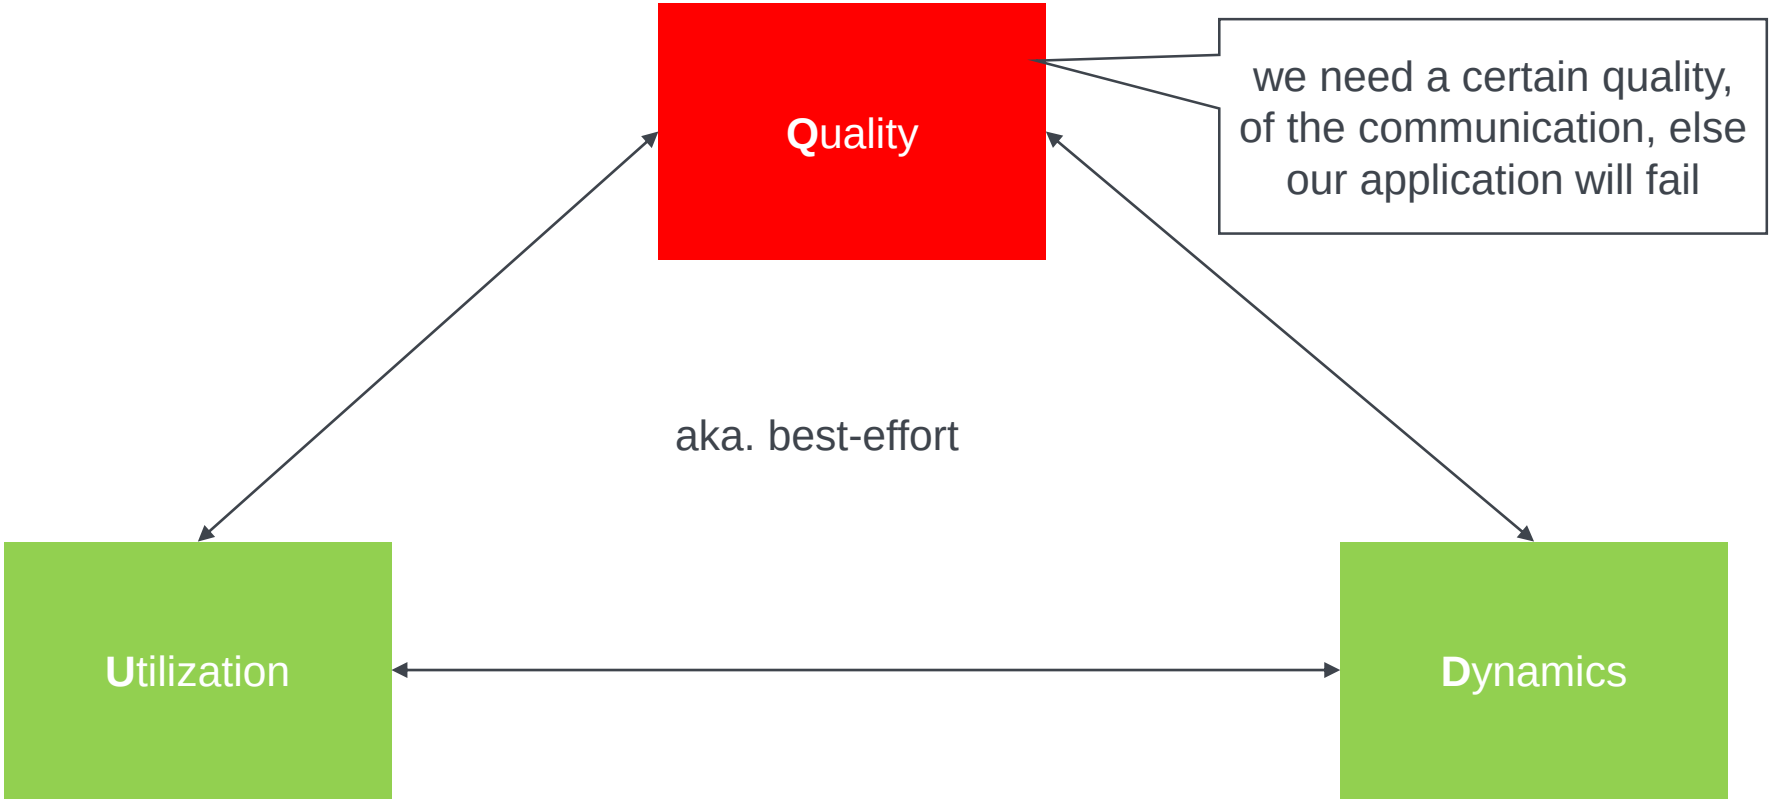

But! There’s a catch.

Conflicts for real-time communication in shared networks

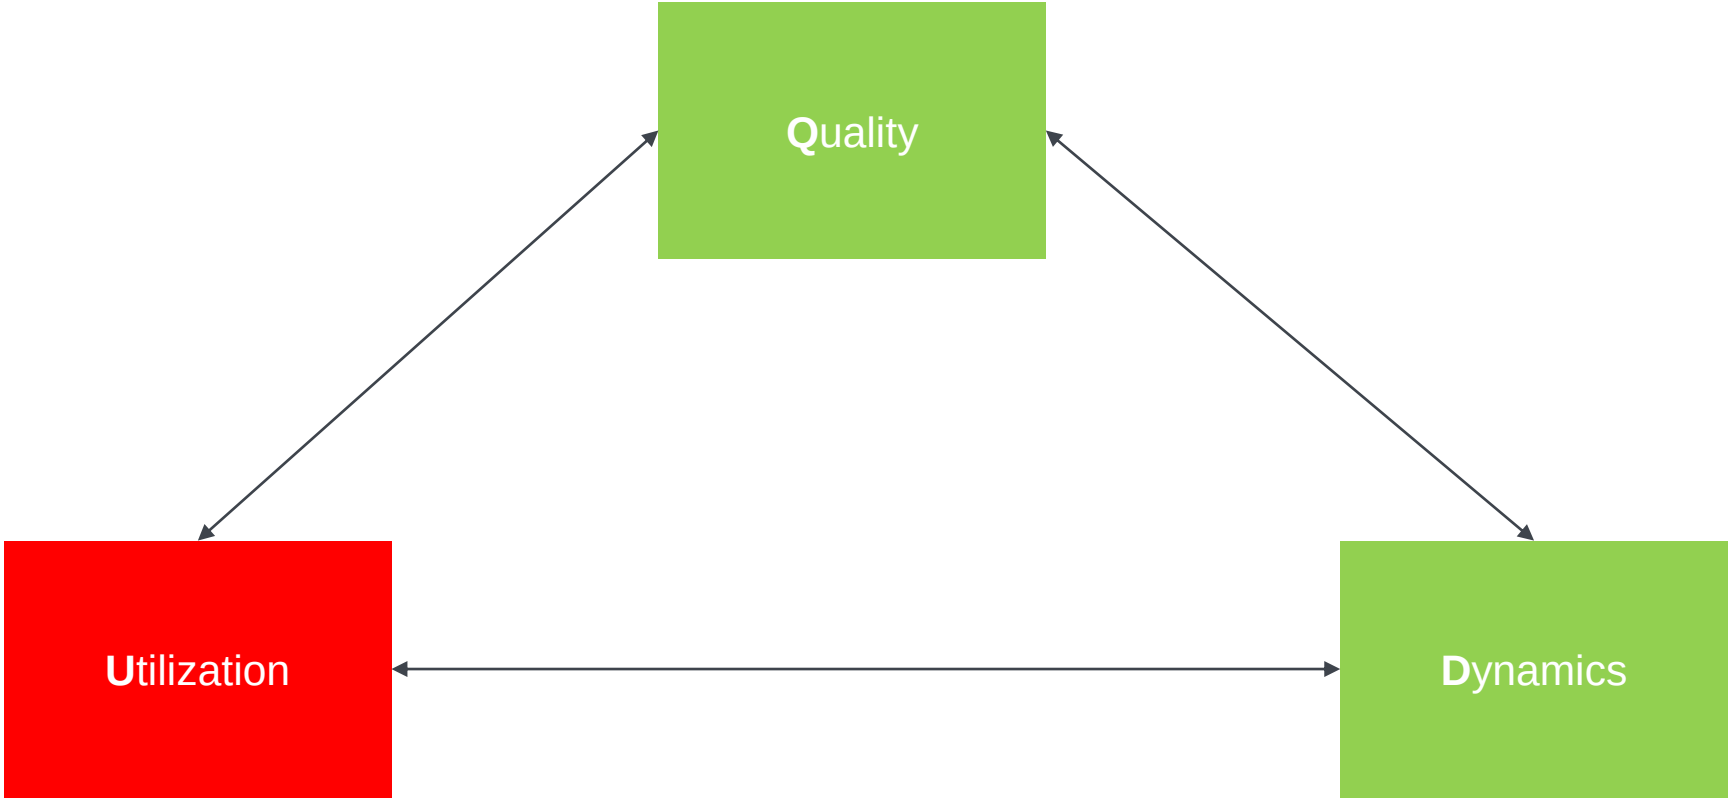

But! There’s a catch.

Conflicts for real-time communication in shared networks

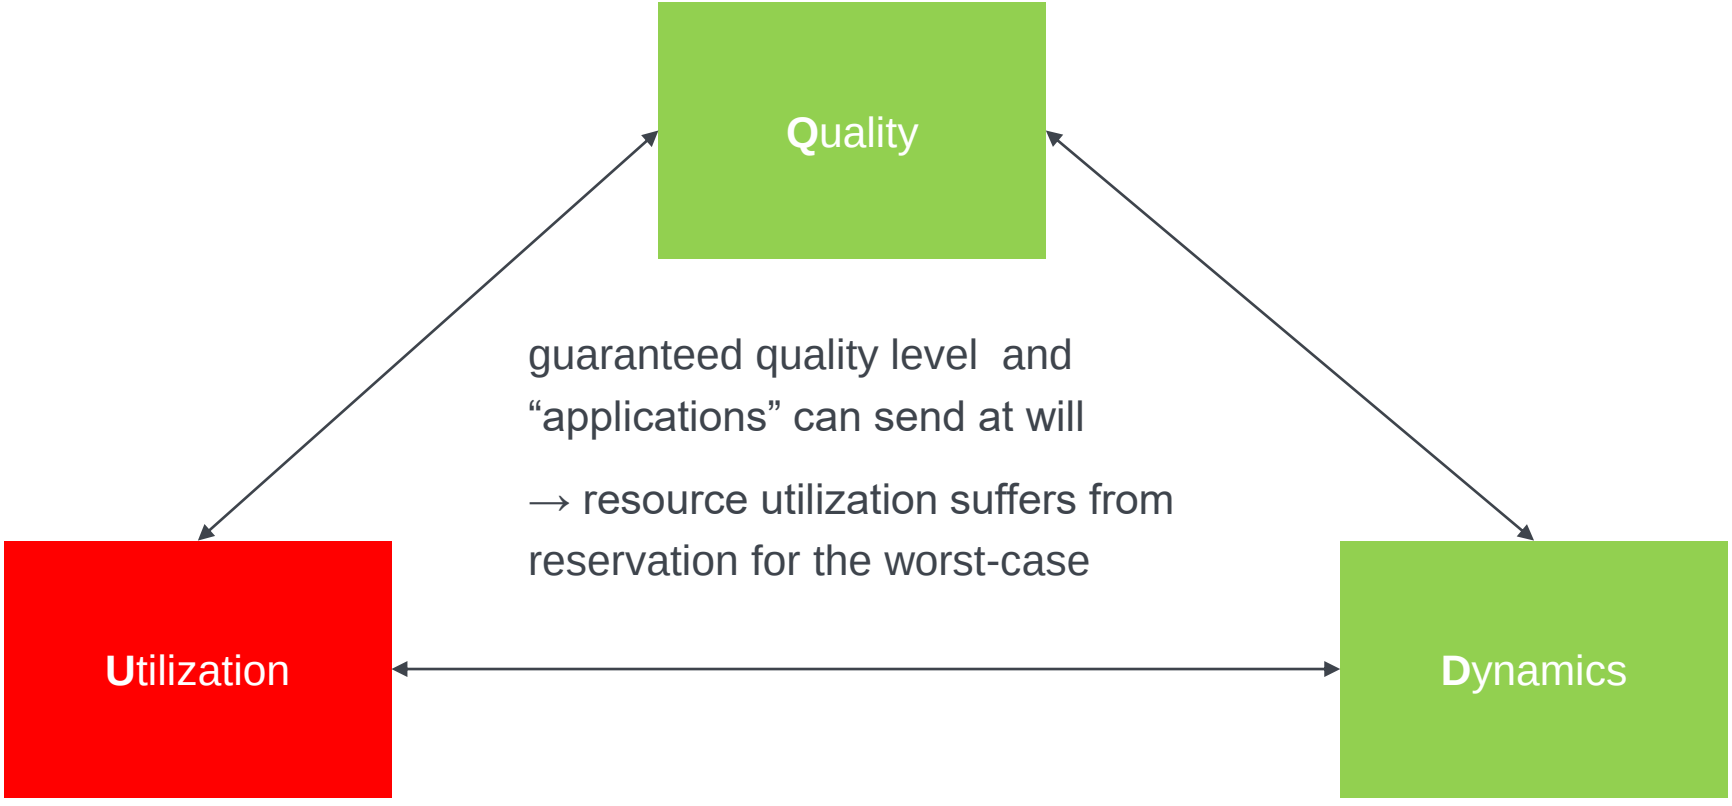

But! There’s a catch.

Conflicts for real-time communication in shared networks

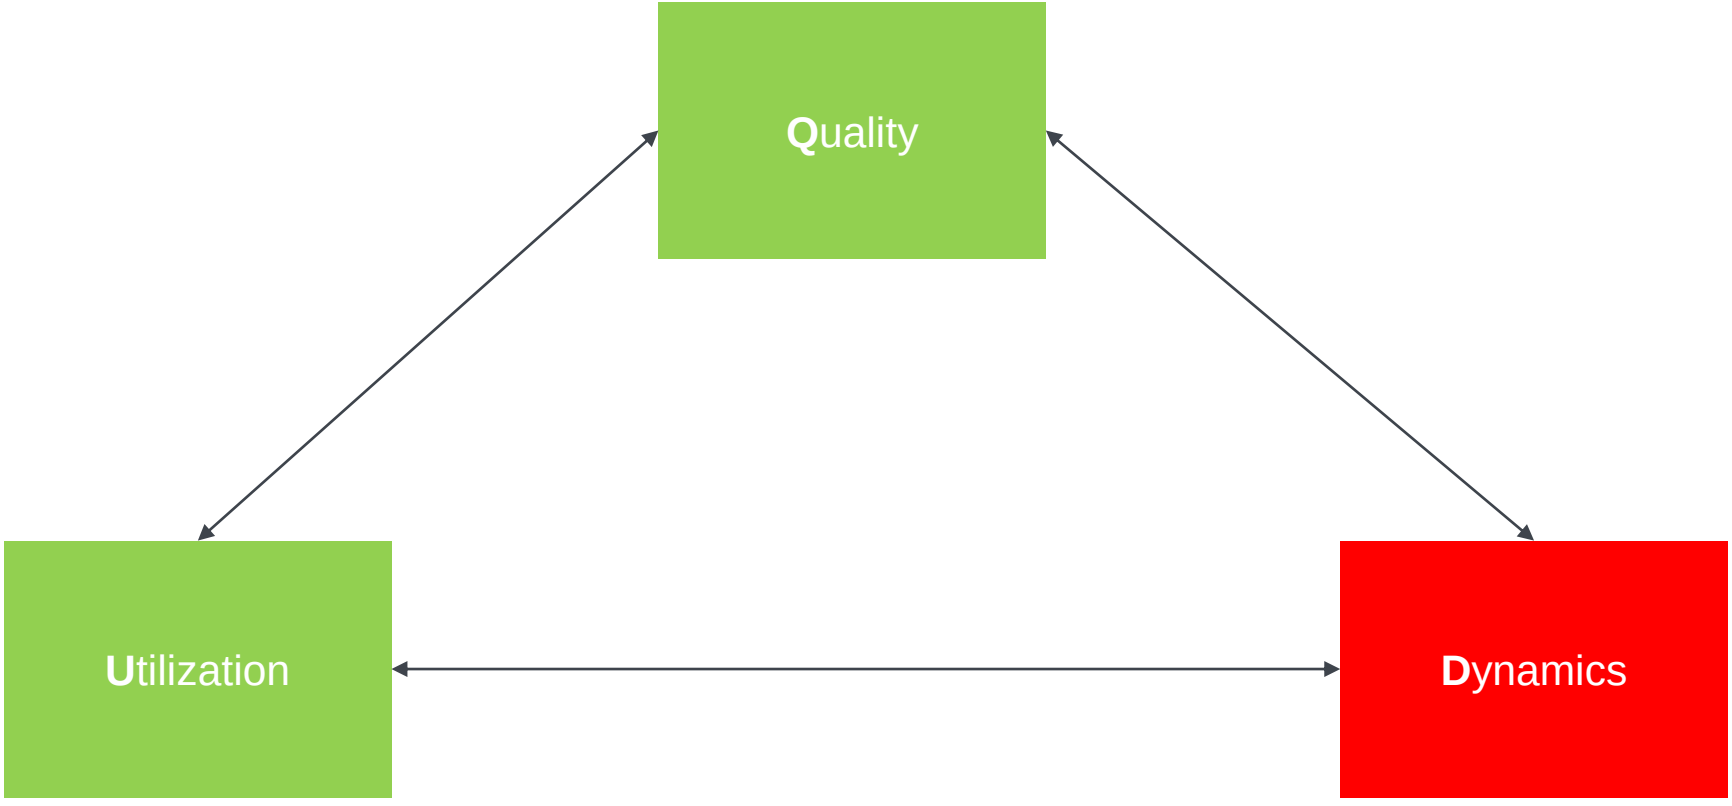

But! There’s a catch.

Conflicts for real-time communication in shared networks

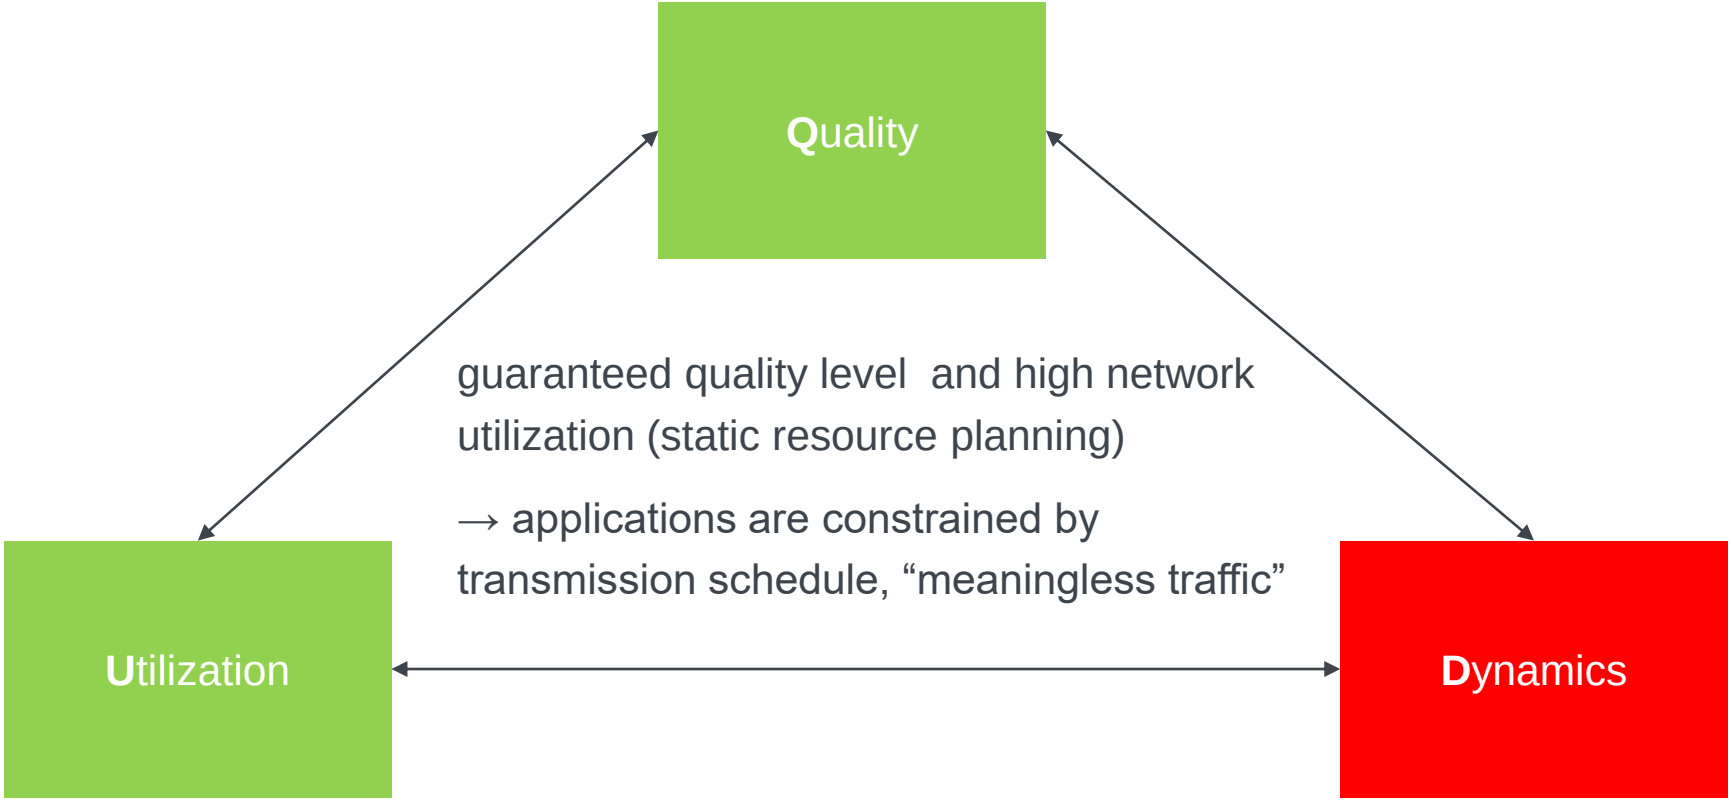

Contributions

Complemental flows

- guarantee (minimal) communication quality
- allow (some) dynamics to improve application performance
- account for “utility” of messages for the application via traffic metric

- How to route and schedule complemental flows?

Complemental flows

Quality

- Application:
  - periodic transmissions
  - certain level of application performance is guaranteed (e.g., stability of control-system)
- Network:
  - well-specified traffic (when, how much data)
  - delivery with bounded delay

deterministic traffic part

Complemental flows

Dynamics

opportunistic traffic part

- Application:
  - transmit, “when it makes sense”, e.g., unforeseen external disturbance
  - improve application performance beyond minimum
  - traffic metric
- Network:
  - dynamic traffic load
  - relaxed (or no) guarantees

Complemental flows

Applications?

- Control systems:
  - Linsenmayer, S., B. W. Carabelli, F. Dürr, J. Falk, F. Allgöwer, and K. Rothermel.. “Integration of Communication Networks and Control Systems Using a Slotted Transmission Classification Model.” CCNC '19

Complemental flows

Applications?

- Control systems:
  - Linsenmayer, S., B. W. Carabelli, F. Dürr, J. Falk, F. Allgöwer, and K. Rothermel.. “Integration of Communication Networks and Control Systems Using a Slotted Transmission Classification Model.” CCNC '19

deterministic and opportunistic messages are temporally and semantically interrelated!

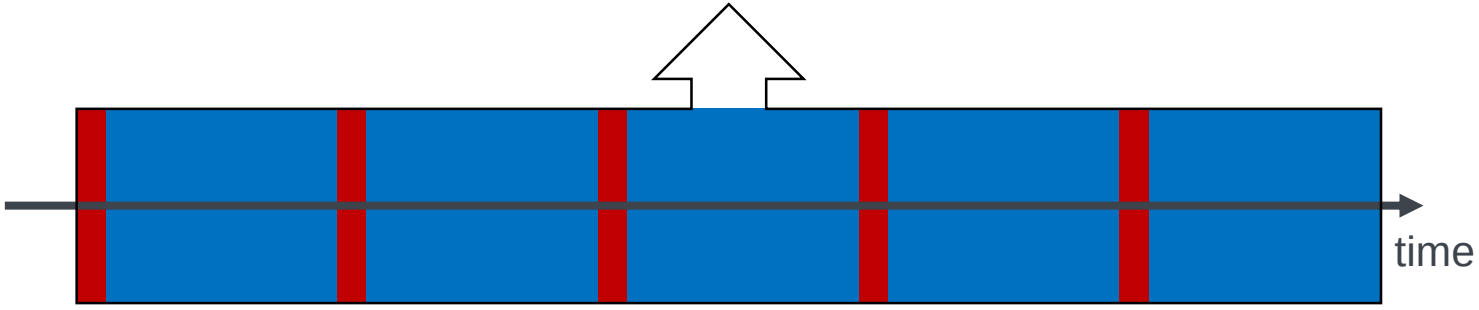

Complemental flows

Complemental flows?

opportunistic traffic part

best effort /
shaped traffic

deterministic traffic part

time-triggered traffic

Complemental flows?

opportunistic traffic part

best effort /
shaped traffic

deterministic traffic part

time-triggered traffic

Craciunas, et. al, “Scheduling Real-Time Communication in IEEE 802.1Qbv Time Sensitive Networks.”, RTNS ’16

Dürr, Frank, and Naresh Ganesh Nayak. “No-Wait Packet Scheduling for IEEE Time-Sensitive Networks (TSN).” RTNS ’16

Schweissguth, et. al., “ILP-Based Joint Routing and Scheduling for Time-Triggered Networks.”, RTNS ’17

Complemental flows?

Reisslein, M., K. W. Ross, and S. Rajagopal. “A Framework for Guaranteeing Statistical QoS.” IEEE/ACM Transactions on Networking 10, no. 1 (February 2002): 27–42.

Craciunas, et. al, “Scheduling Real-Time Communication in IEEE 802.1Qbv Time Sensitive Networks.”, RTNS ’16

Dürr, Frank, and Naresh Ganesh Nayak. “No-Wait Packet Scheduling for IEEE Time-Sensitive Networks (TSN).” RTNS ’16

Schweissguth, et. al., “ILP-Based Joint Routing and Scheduling for Time-Triggered Networks.”, RTNS ’17

Complemental flows?

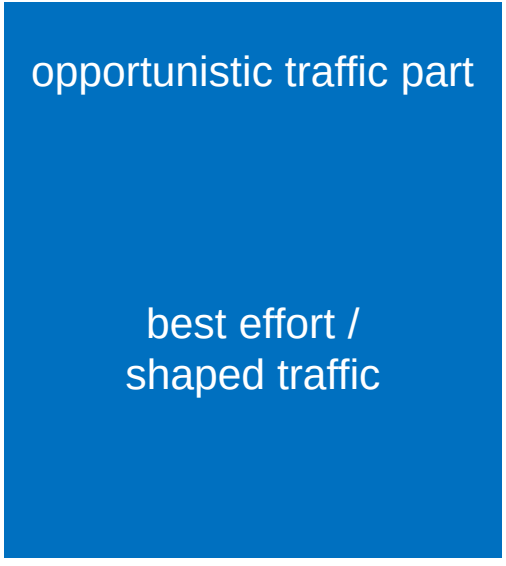

Reisslein, M., K. W. Ross, and S. Rajagopal. “A Framework for Guaranteeing Statistical QoS.” IEEE/ACM Transactions on Networking 10, no. 1 (February 2002): 27–42.

Jin, X., et al. “Real-Time Scheduling for Event-Triggered and Time-Triggered Flows in Industrial Wireless Sensor-Actuator Networks.” In IEEE INFOCOM 2019 - IEEE Conference on Computer Communications, 1684–92, 2019.

Steiner, W. “Synthesis of Static Communication Schedules for Mixed-Criticality Systems.” In 2011 14th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing Workshops, 11–18, 2011.

Pop, P., M. L. Raagaard, S. S. Craciunas, and W. Steiner. “Design Optimisation of Cyber-Physical Distributed Systems Using IEEE Time-Sensitive Networks.” IET Cyber-Physical Systems: Theory Applications 1, no. 1 (2016): 86–94.

Craciunas, et. al, “Scheduling Real-Time Communication in IEEE 802.1Qbv Time Sensitive Networks.”, RTNS ’16

Dürr, Frank, and Naresh Ganesh Nayak. “No-Wait Packet Scheduling for IEEE Time-Sensitive Networks (TSN).” RTNS ’16

Schweissguth, et. al., “ILP-Based Joint Routing and Scheduling for Time-Triggered Networks.”, RTNS ’17

Complemental flows?

opportunistic traffic part

best effort / shaped traffic

deterministic traffic part

time-triggered traffic

relation between time-triggered and non-time-triggered traffic?

Jin, X., et al. “Real-Time Scheduling for Event-Triggered and Time-Triggered Flows in Industrial Wireless Sensor-Actuator Networks.” In IEEE INFOCOM 2019 - IEEE Conference on Computer Communications, 1684–92, 2019.

Steiner, W. “Synthesis of Static Communication Schedules for Mixed-Criticality Systems.” 2011 14th IEEE International Symposium on Object/Component/Service-Oriented Real-Time Distributed Computing Workshops, 11–18, 2011.

Reisslein, M., K. W. Ross, and S. Rajagopal. “A Framework for Guaranteeing Statistical QoS.” IEEE/ACM Transactions on Networking 10, no. 1 (February 2002): 27–42.

Craciunas, et. al, “Scheduling Real-Time Communication in IEEE 802.1Qbv Time Sensitive Networks.”, RTNS ’16

Dürr, Frank, and Naresh Ganesh Nayak. “No-Wait Packet Scheduling for IEEE Time-Sensitive Networks (TSN).” RTNS ’16

Pop, P., M. L. Raagaard, S. S. Craciunas, and W. Steiner. “Design Optimisation of Cyber-Physical Distributed Systems Using IEEE Time-Sensitive Networks.” IET Cyber-Physical Systems: Theory Applications 1, no. 1 (2016): 86–94.

Schweissguth, et. al., “ILP-Based Joint Routing and Scheduling for Time-Triggered Networks.”, RTNS ’17

Traffic Engineering

Problem

Input:

- Network Topology (graph)
- Set of flows (including requirements and specifications for **deterministic** and **opportunistic** traffic part)

Output:

Optimal traffic configuration:

- route (connected sequence of edges from source to destination of flow)
- schedule (i.e., *phase* of deterministic transmissions)
- optimal with respect to expected opportunistic traffic

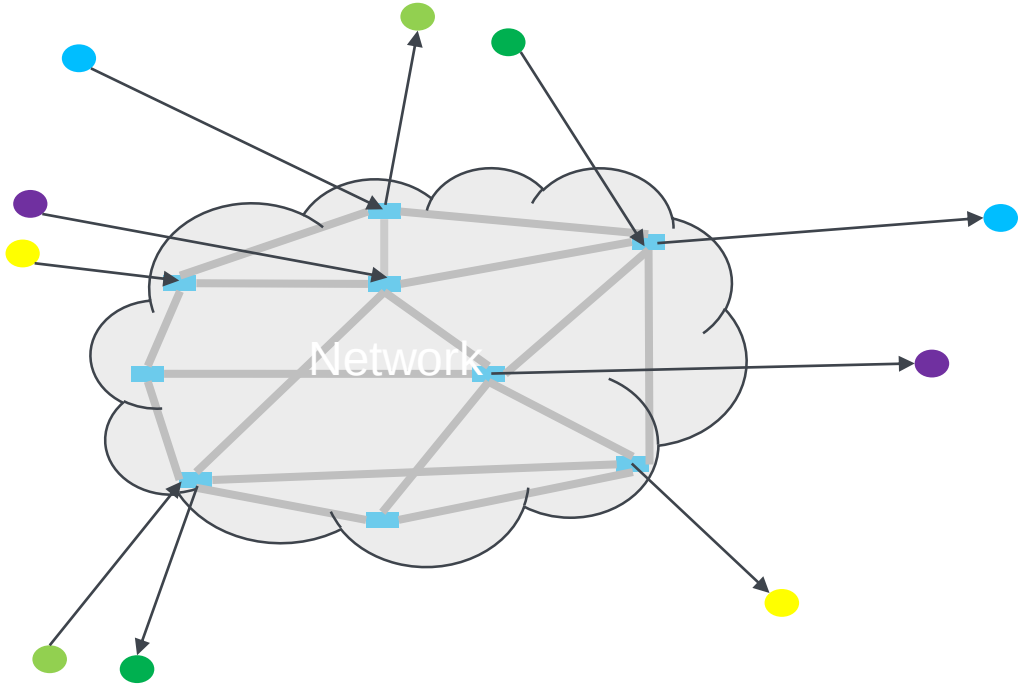

Traffic Engineering

Problem

Input:

- Network Topology (graph)
- Set of flows (including requirements and specifications for **deterministic** and **opportunistic** traffic part)

Output:

Optimal traffic configuration:

- route (connected sequence of edges from source to destination of flow)
- schedule (i.e., *phase* of deterministic transmissions)
- optimal with respect to expected opportunistic traffic

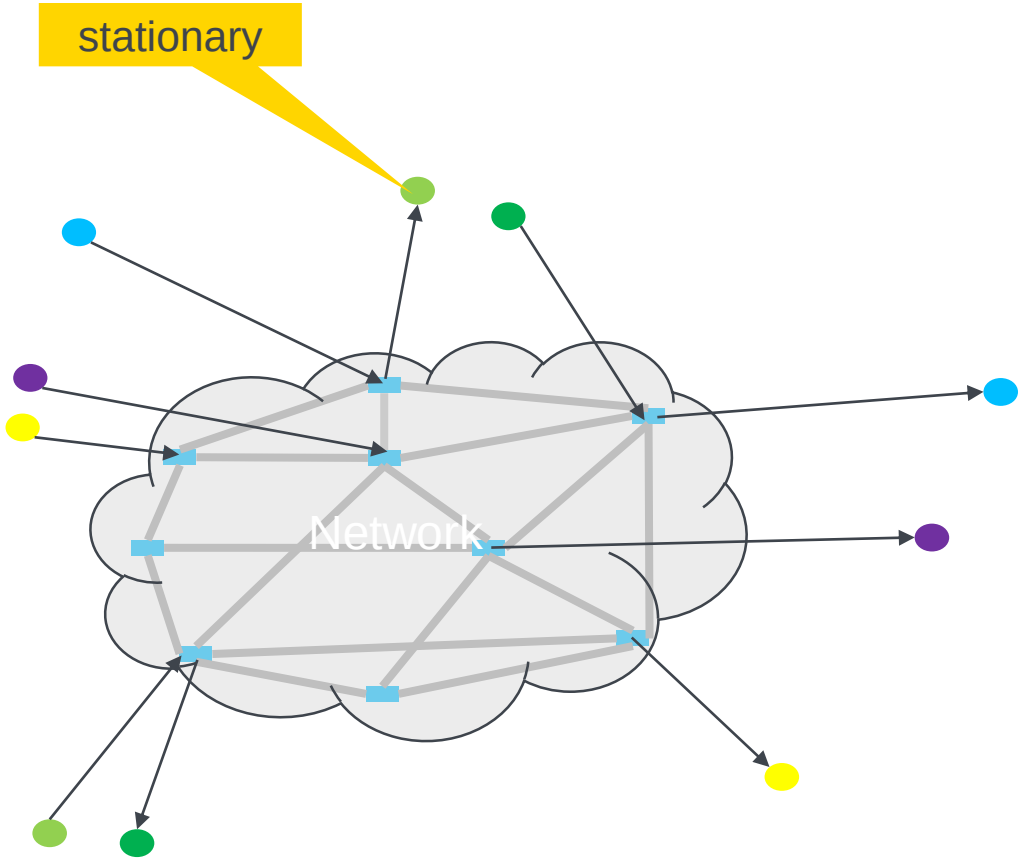

Traffic Engineering

Problem

Input:

- Network Topology (graph)
- Set of flows (including requirements and specifications for **deterministic** and **opportunistic** traffic part)

Output:

Optimal traffic configuration:

- route (connected sequence of edges from source to destination of flow)
- schedule (i.e., *phase* of deterministic transmissions)
- optimal with respect to expected opportunistic traffic

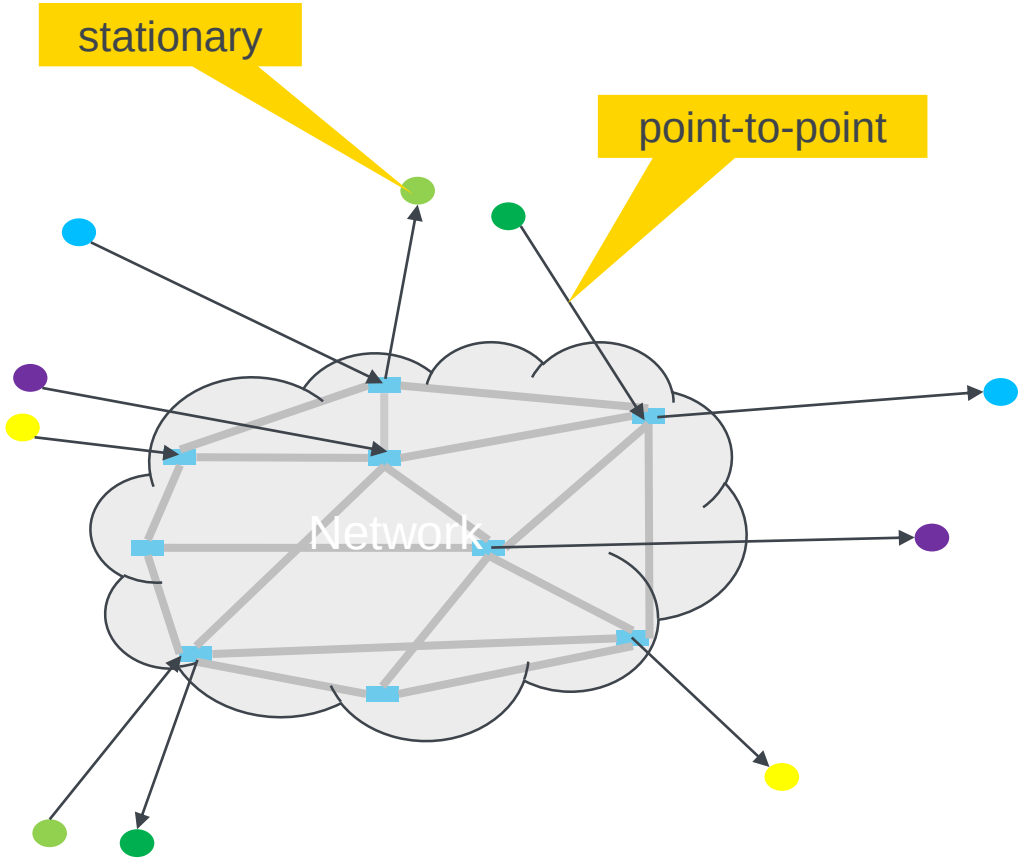

Traffic Engineering

Switch

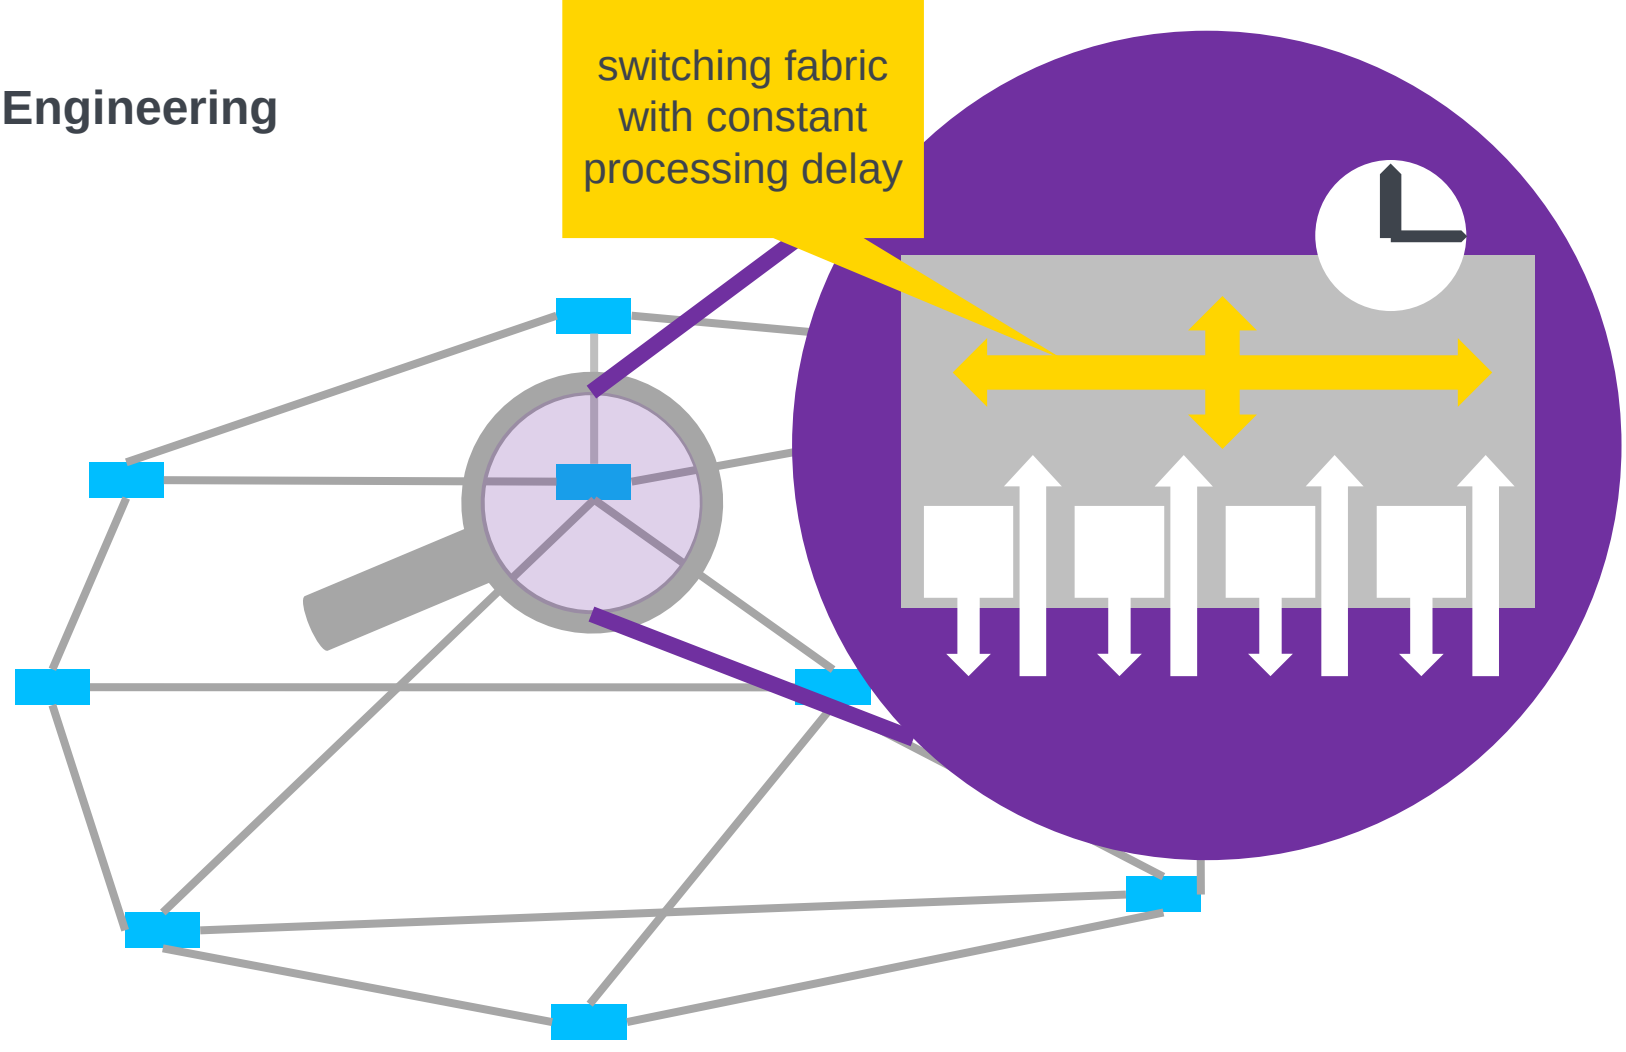

Traffic Engineering

Switch

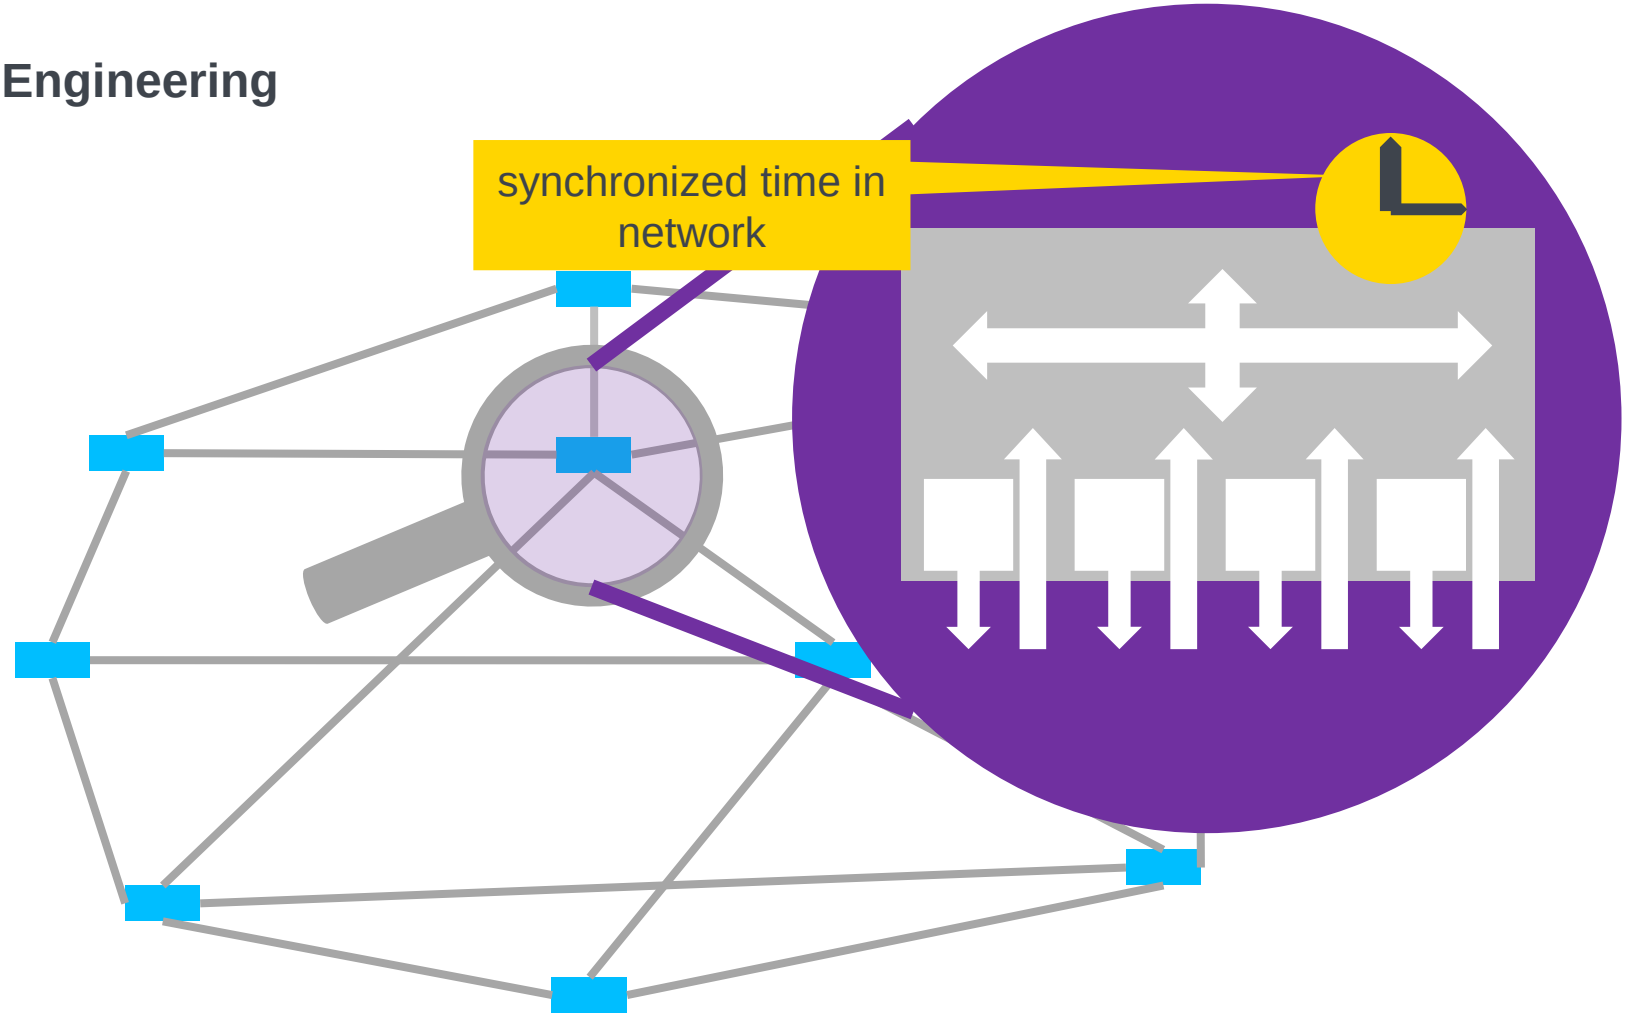

Traffic Engineering

Switch

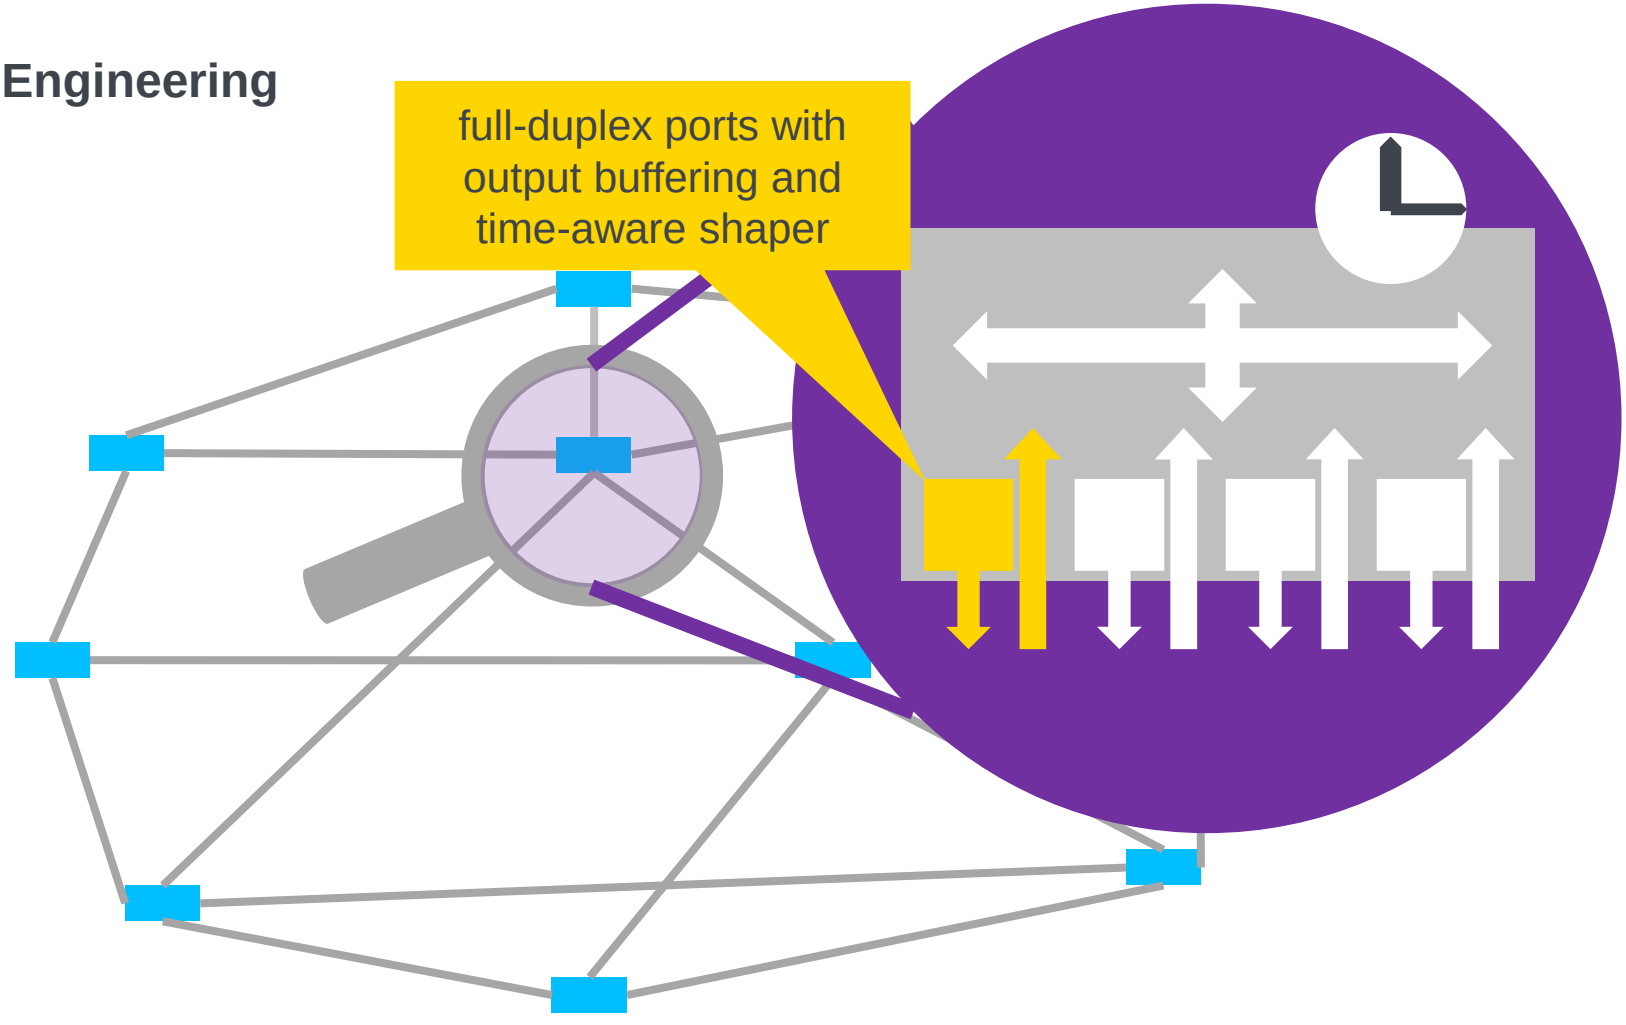

Traffic Engineering

Time-aware shaper

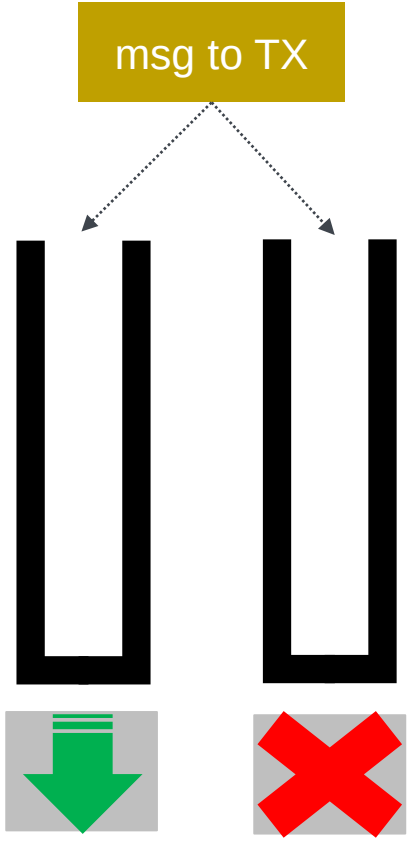

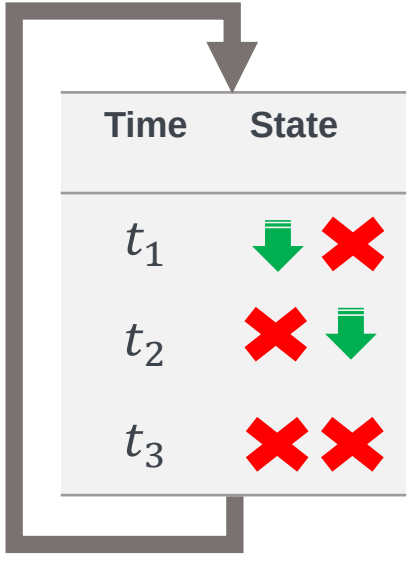

Traffic Engineering

Time-aware shaper

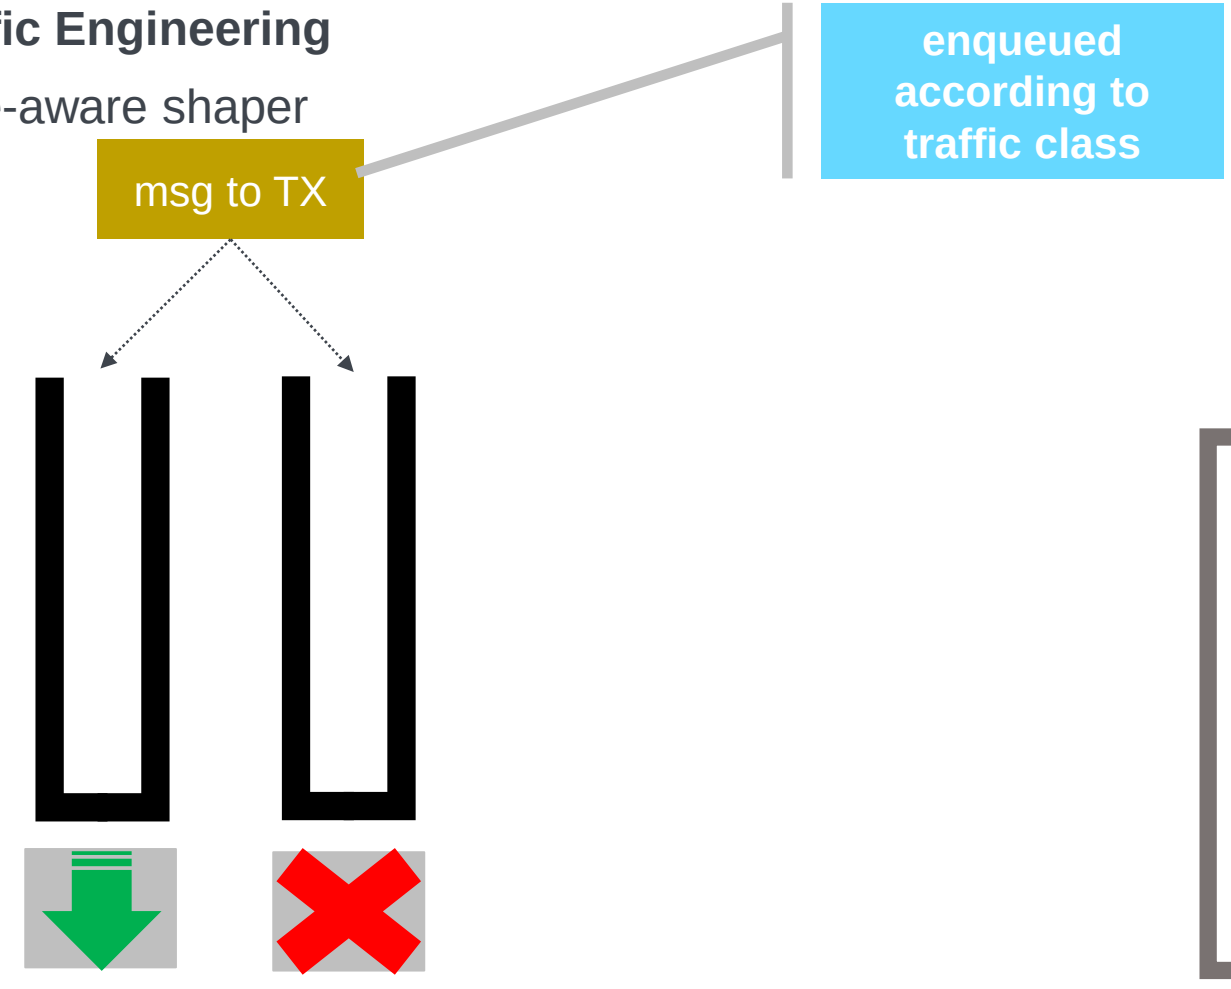

Traffic Engineering

Time-aware shaper

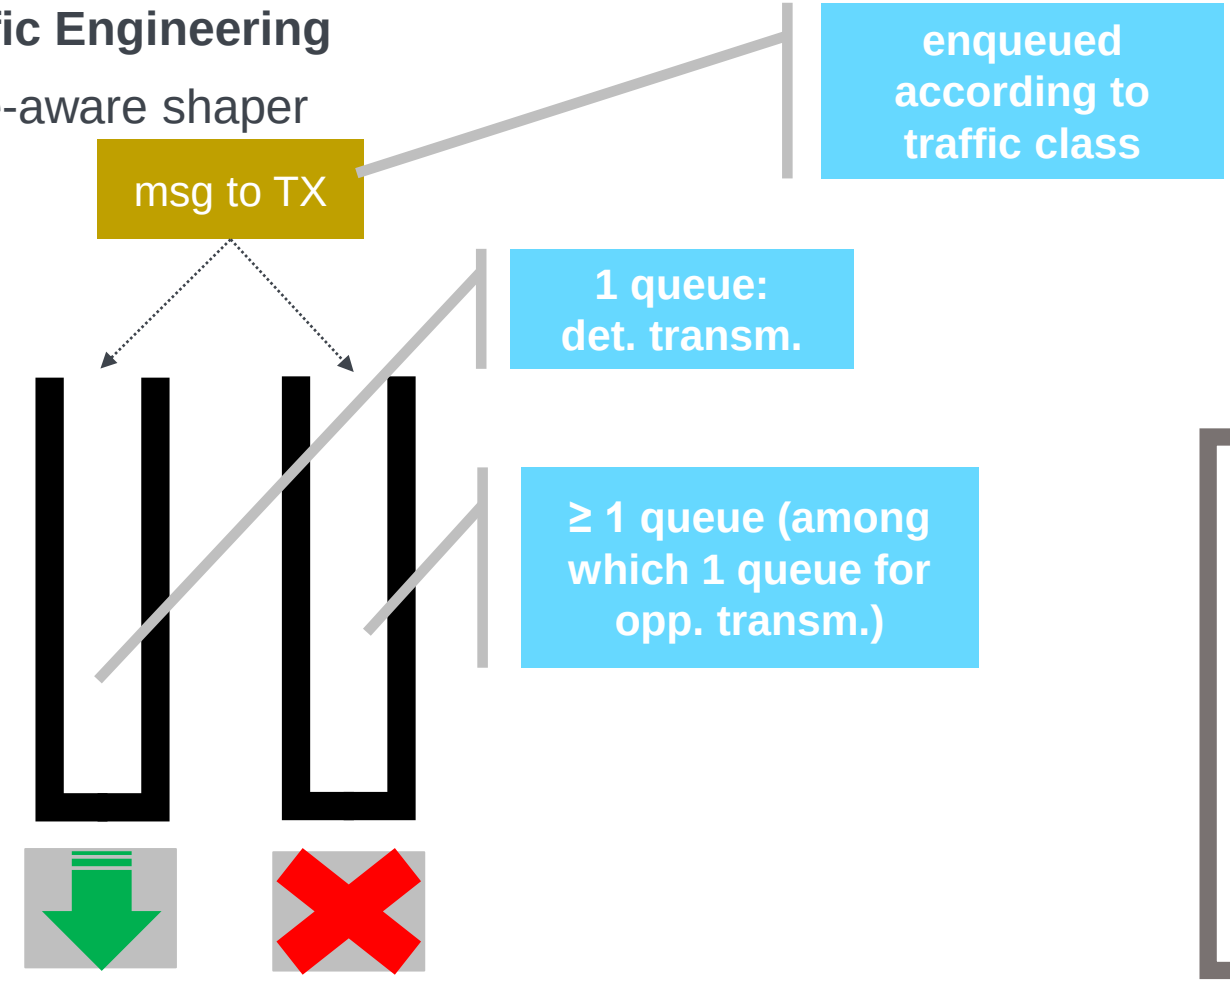

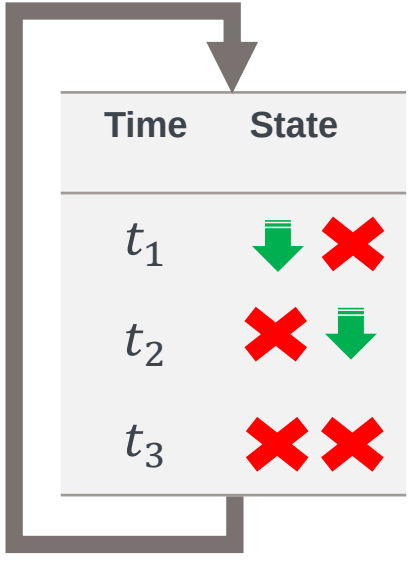

Traffic Engineering

Time-aware shaper

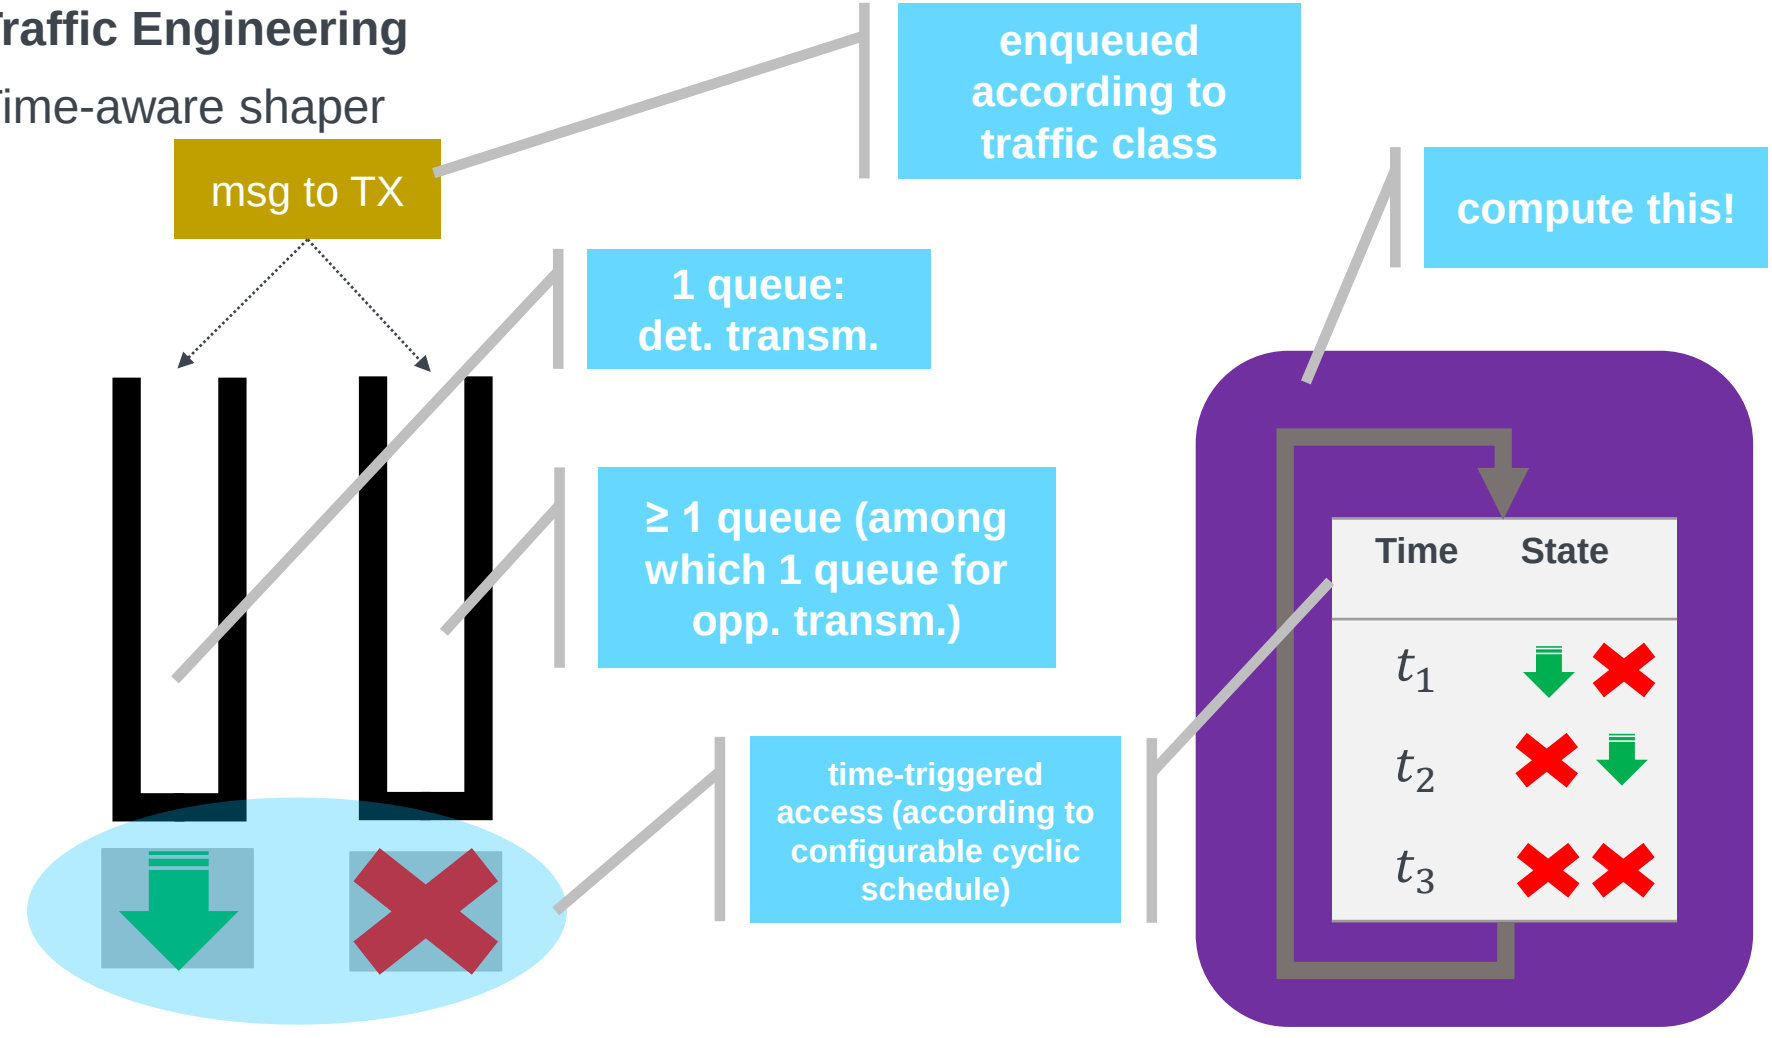

Traffic Engineering

Traffic metric for opportunistic traffic

- relates application performance to opportunistic messages
- offline traffic-engineering, i.e., traffic metric is required a-priori
  - estimation
  - measurement
  - analysis
- additive

- different temporal granularity

Traffic Engineering

Traffic metric for opportunistic traffic

time-independent

- e.g., average bandwidth of opportunistic messages
- “scalar” value

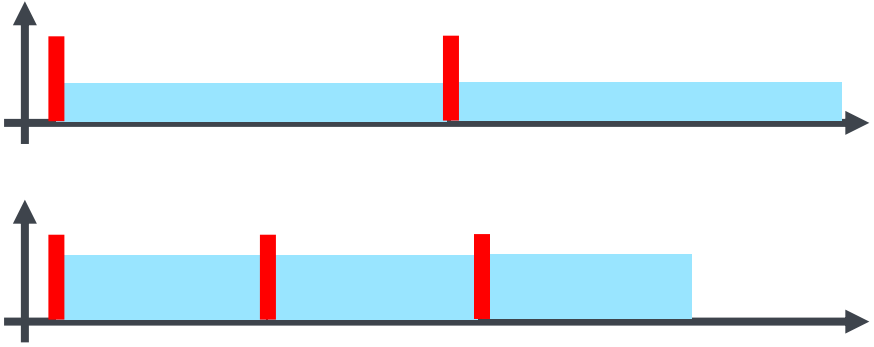

Traffic Engineering

Traffic metric for opportunistic traffic

time-independent

- e.g., average bandwidth of opportunistic messages
- “scalar” value

time-dependent

- e.g., average probability of transmission (expected traffic load)
- deterministic messages are ”renewal” points

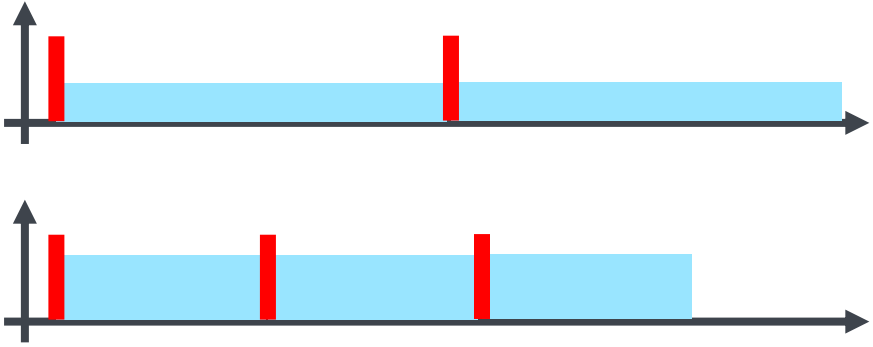

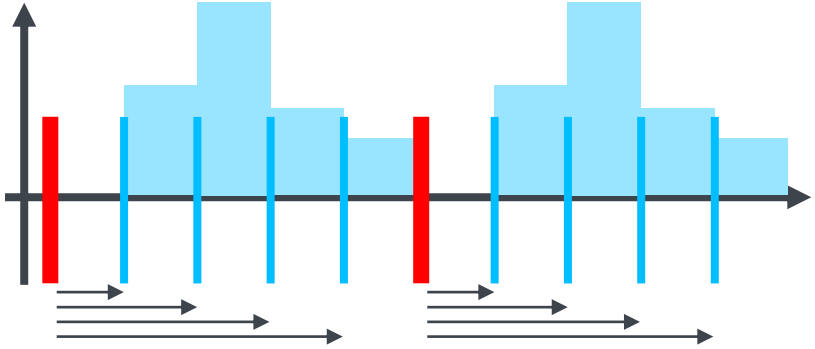

Traffic Engineering

Two approaches: mixed integer linear programming

Edge-granularity routing and time-independent traffic metric:

- MILP decides for each individual edge whether it is part of the route
- min-max to reduce globally highest value of the accumulated traffic metric (i.e., aiming for more even traffic distribution)

➢ freedom in routing

➢ deterministic transmissions impose constraints on **feasibility**

➢ opportunistic transmissions influence **optimality**

Traffic Engineering

Two approaches: mixed integer linear programming

Edge-granularity routing and time-independent traffic metric:

- MILP decides for each individual edge whether it is part of the route
- min-max to reduce globally highest value of the accumulated traffic metric (i.e., aiming for more even traffic distribution)

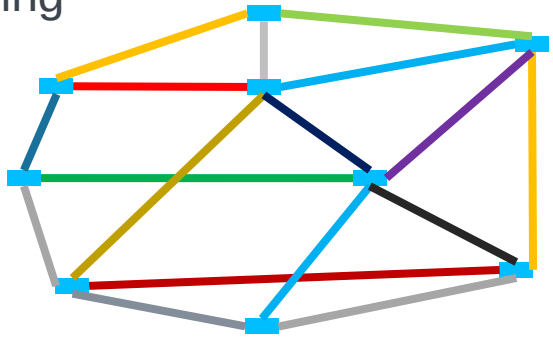

- ➢ freedom in routing
- ➢ deterministic transmissions impose constraints on **feasibility**
- ➢ opportunistic transmissions influence **optimality**

Traffic Engineering

Two approaches: mixed integer linear programming

Edge-granularity routing and time-independent traffic metric:

- MILP decides for each individual edge whether it is part of the route
- min-max to reduce globally highest value of the accumulated traffic metric (i.e., aiming for more even traffic distribution)

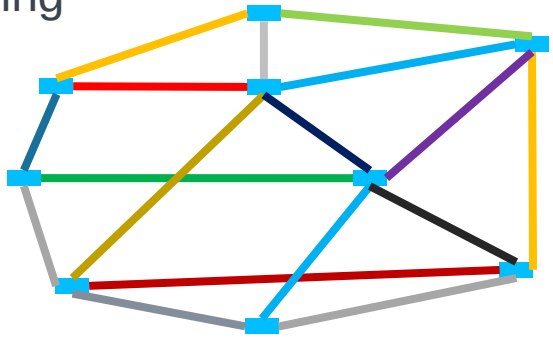

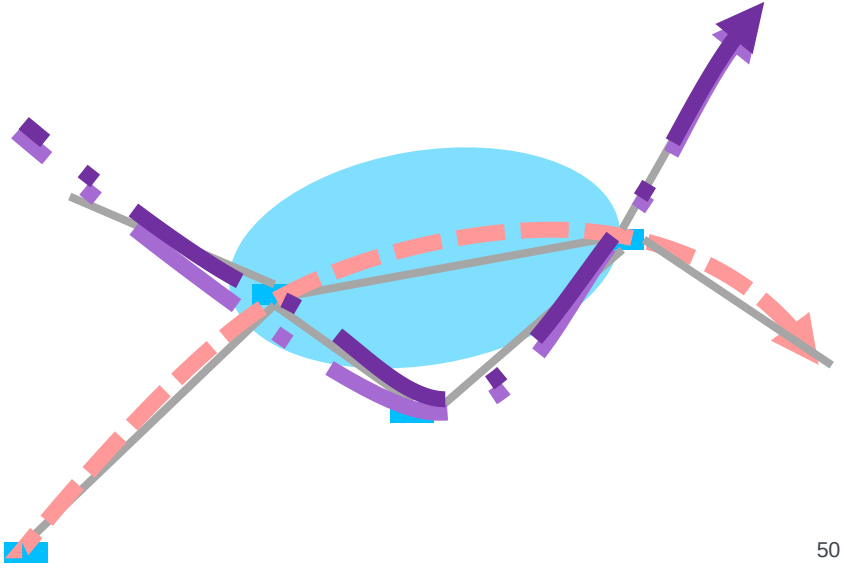

- freedom in routing
- deterministic transmissions impose constraints on **feasibility**
- opportunistic transmissions influence **optimality**

Traffic Engineering

Two approaches: mixed integer linear programming

Path-granularity routing and time-dependent traffic metric

- MILP chooses pre-computed path as route for each flow
- min-max optimization of the accumulated value of the traffic metric on any edge at any time

➢ more temporal information

➢ deterministic transmissions impose constraints on **feasibility**

➢ opportunistic transmissions influence **optimality**

Traffic Engineering

Two approaches: mixed integer linear programming

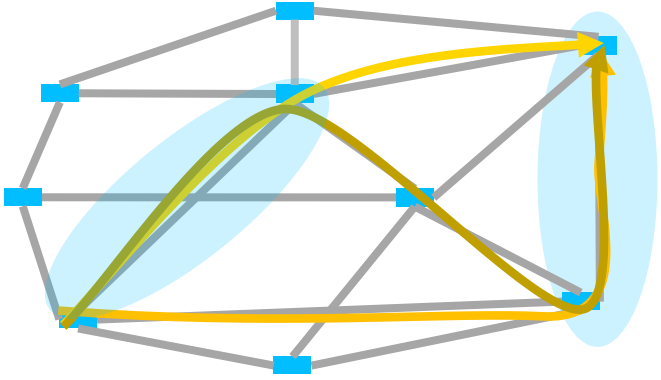

Path-granularity routing and time-dependent traffic metric

- MILP chooses pre-computed path as route for each flow
- min-max optimization of the accumulated value of the traffic metric on any edge at any time

➢ more temporal information

➢ deterministic transmissions impose constraints on **feasibility**

➢ opportunistic transmissions influence **optimality**

Traffic Engineering

Two approaches: mixed integer linear programming

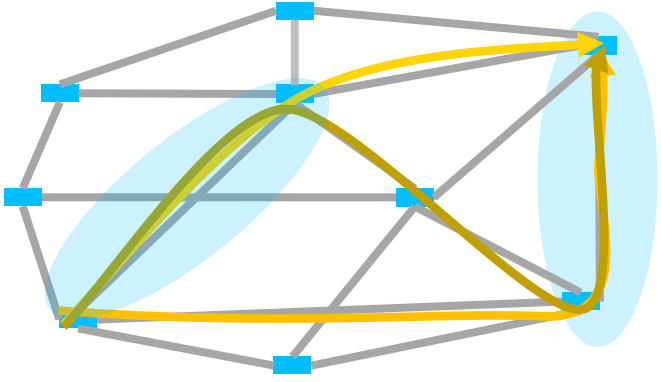

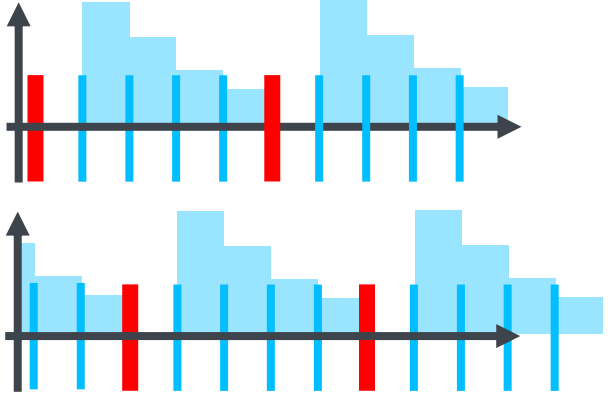

Path-granularity routing and time-dependent traffic metric

- MILP chooses pre-computed path as route for each flow
- min-max optimization of the accumulated value of the traffic metric on any edge at any time

➢ more temporal information

➢ deterministic transmissions impose constraints on **feasibility**

➢ opportunistic transmissions influence **optimality**

Traffic Engineering

Two approaches: mixed integer linear programming

Edge-granularity routing and time-independent traffic metric:

- MILP decides for each individual edge whether it is part of the route
- min-max to reduce globally highest value of the accumulated traffic metric (i.e., aiming for more even traffic distribution)

➢ possibly more freedom in routing

➢ deterministic transmissions impose constraints on **feasibility**

➢ opportunistic transmissions influence **optimality**

Path-granularity routing and time-dependent traffic metric

- MILP chooses pre-computed path as route for each flow
- min-max optimization of the accumulated value of the traffic metric on any edge at any time

➢ more temporal information

➢ deterministic transmissions impose constraints on **feasibility**

➢ opportunistic transmissions influence **optimality**

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Min-max objective

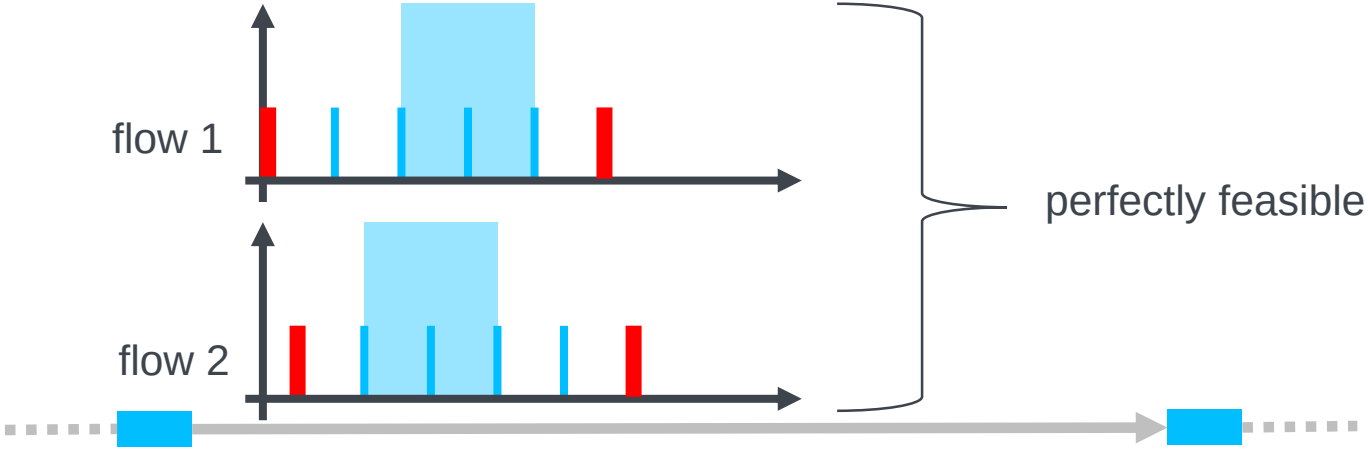

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Min-max objective

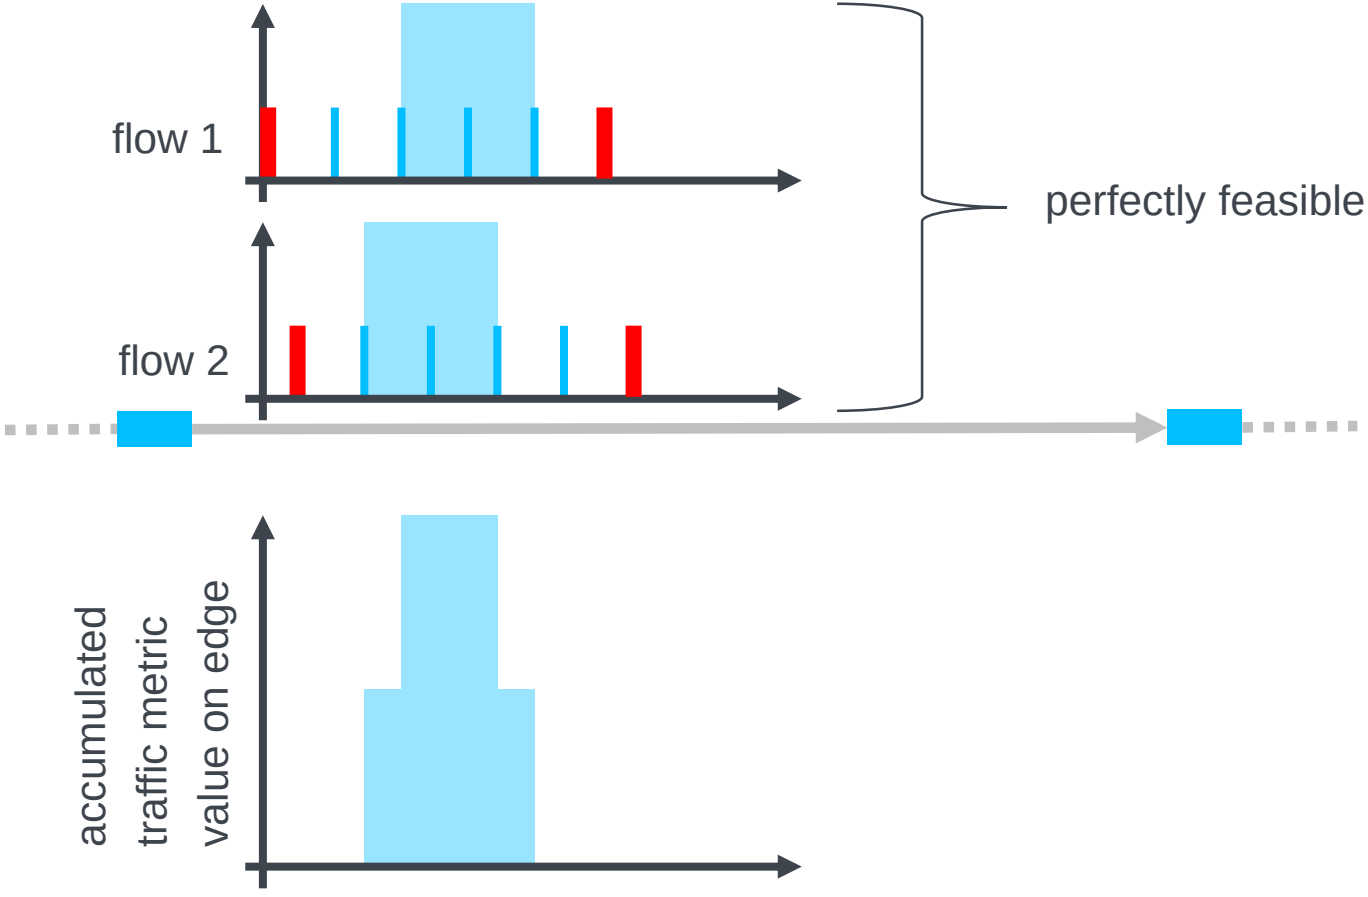

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Min-max objective

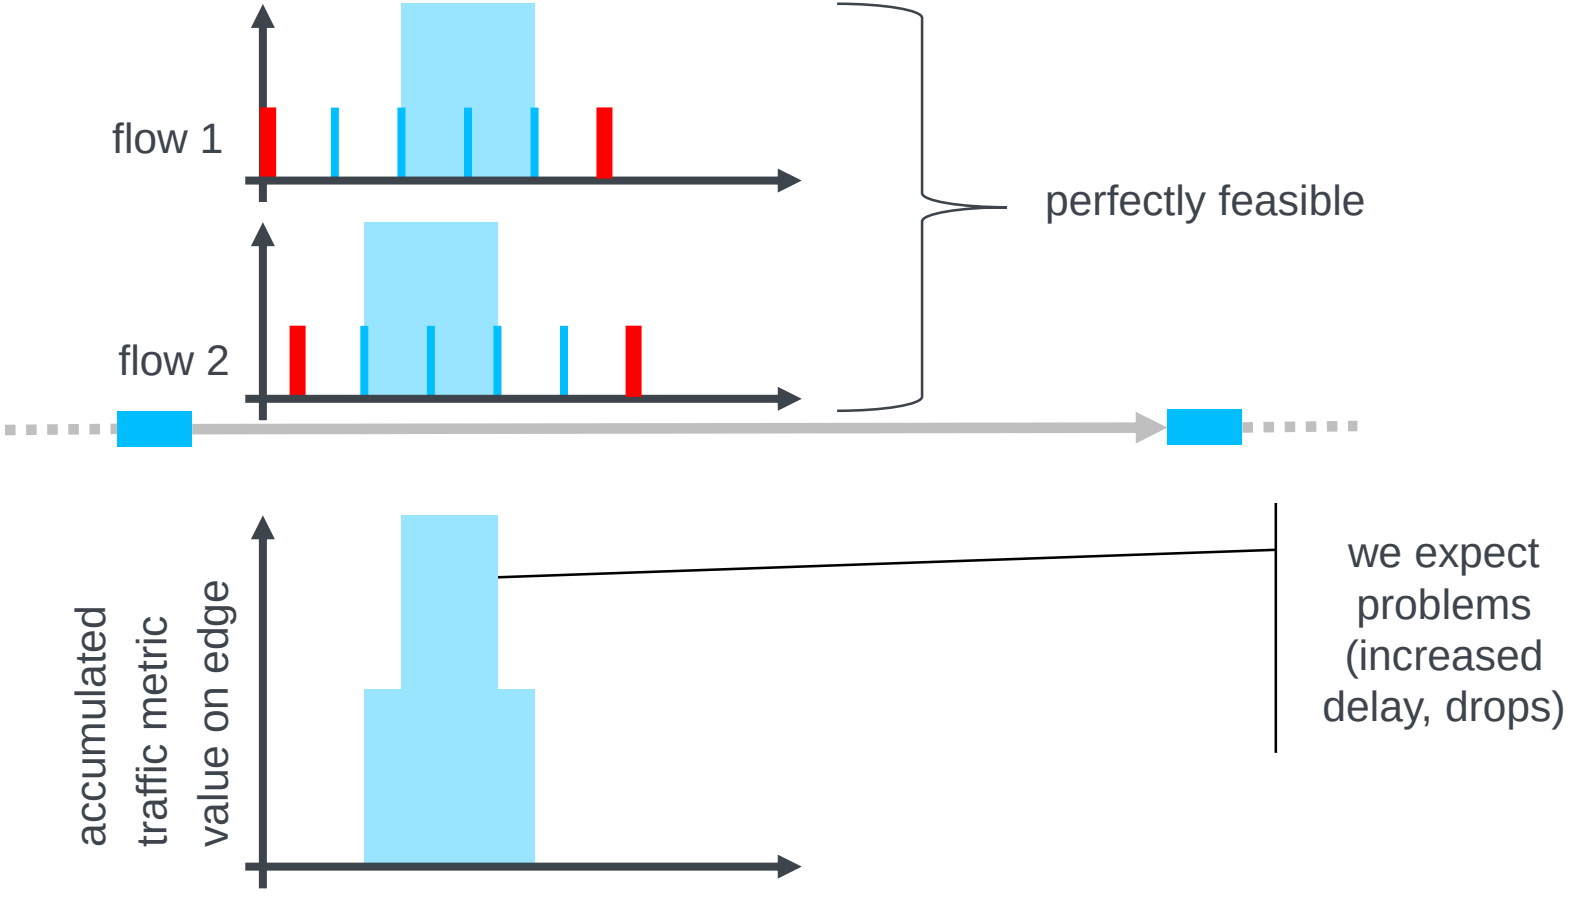

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Min-max objective

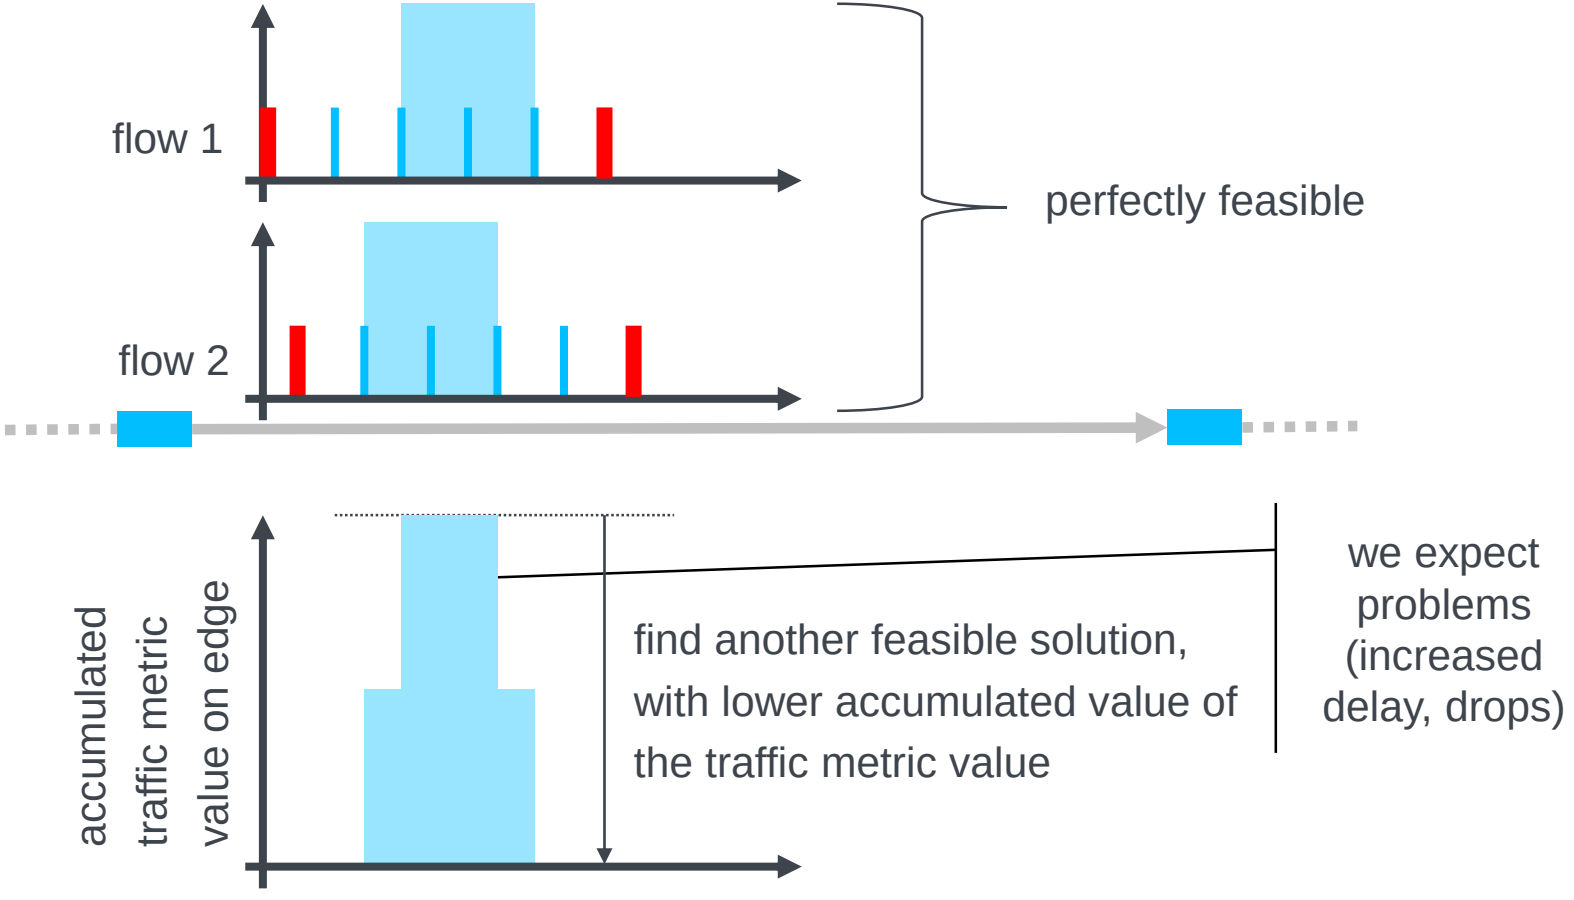

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Modeling space.

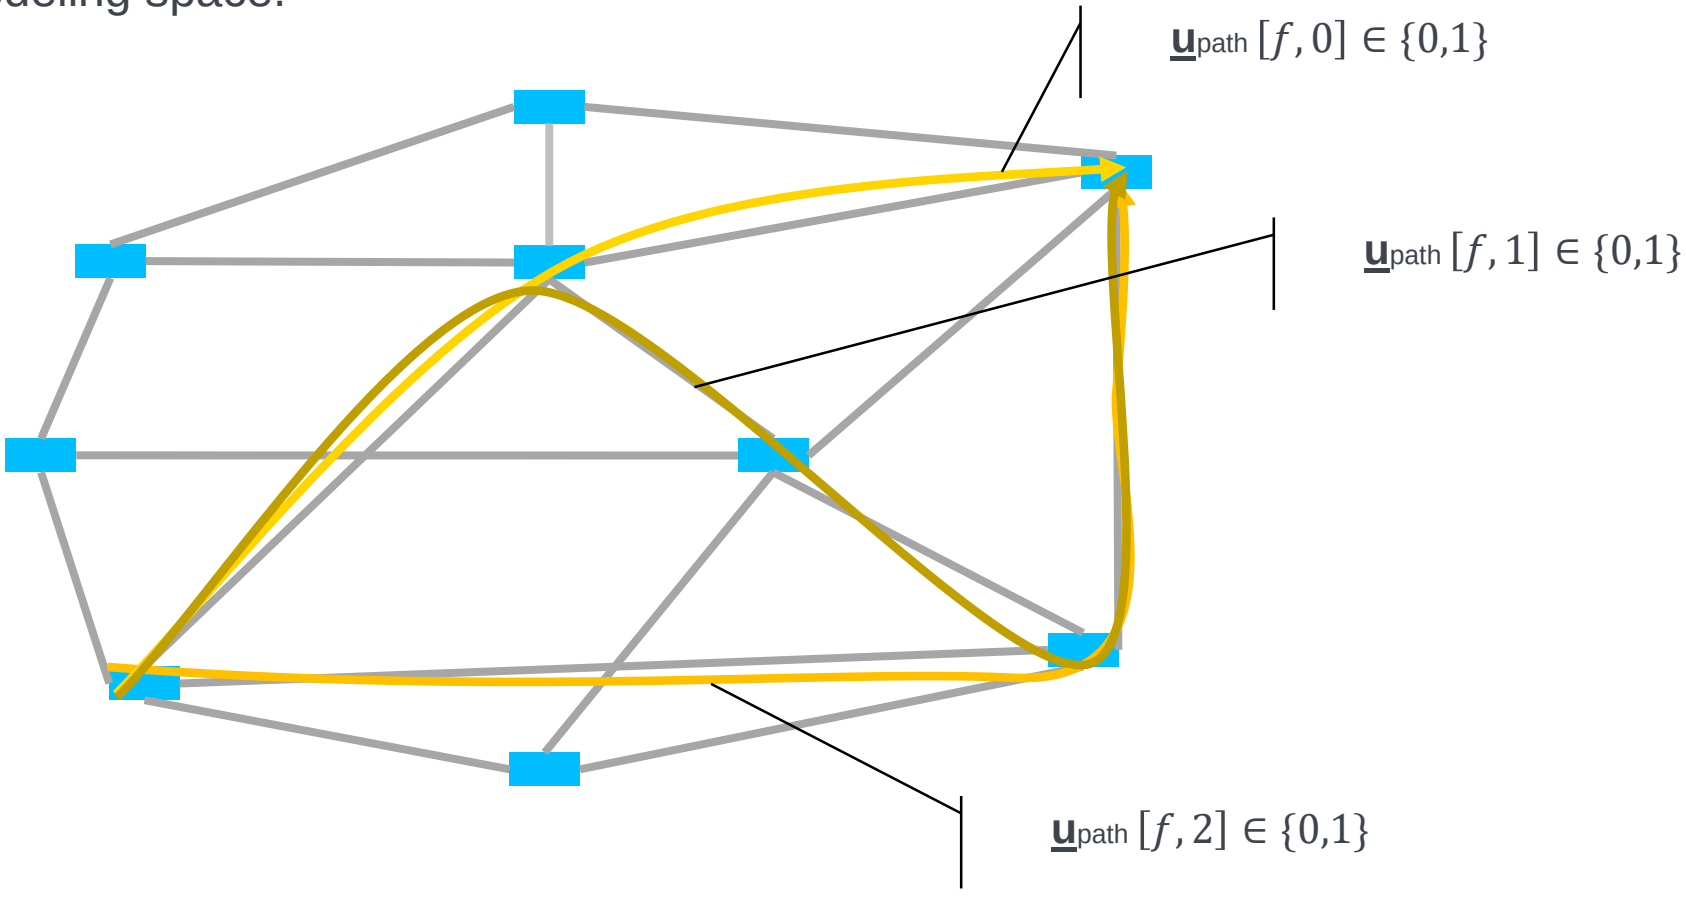

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Modeling time.

scheduling variable (per edge, per flow),

derivation of schedules from phase

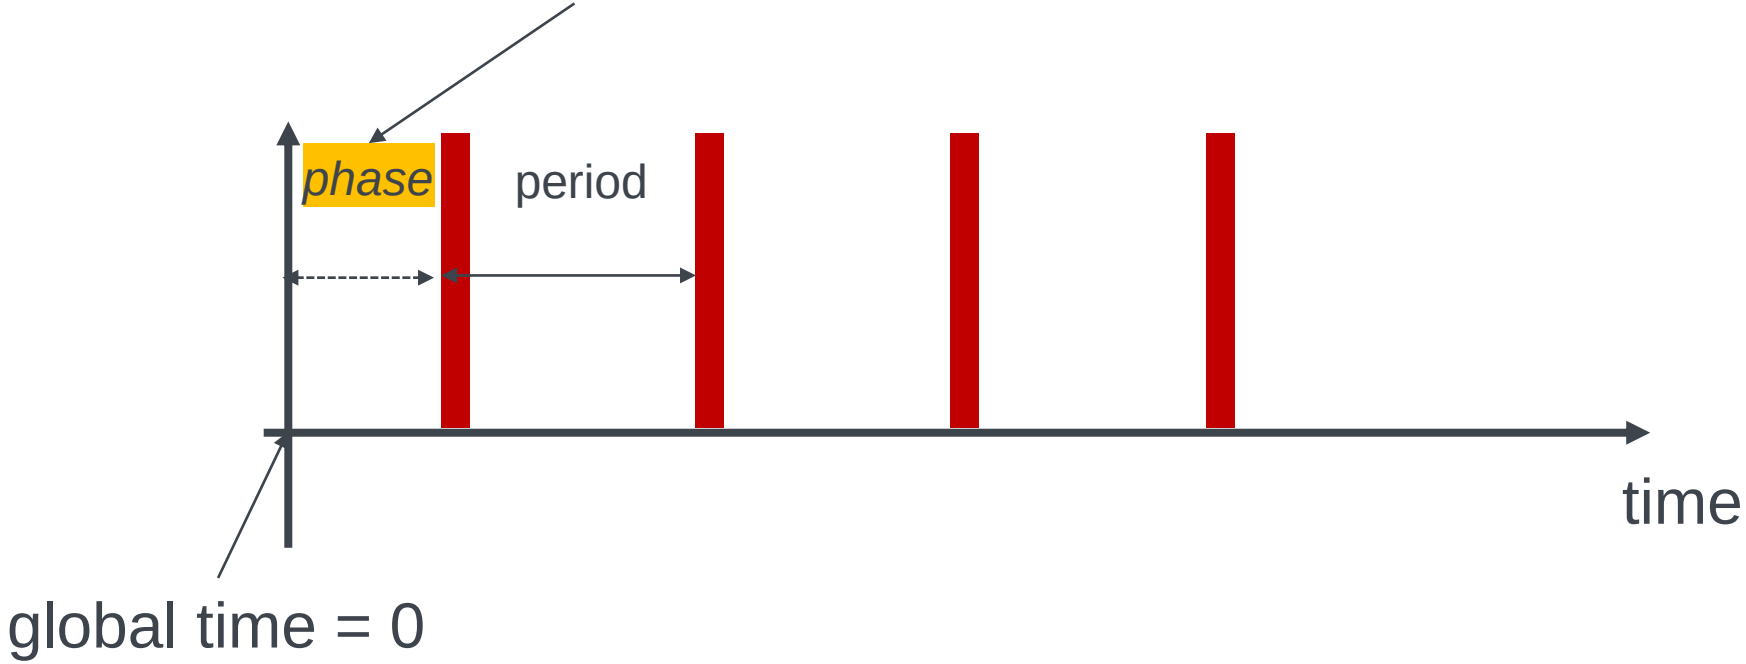

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Discrete time.

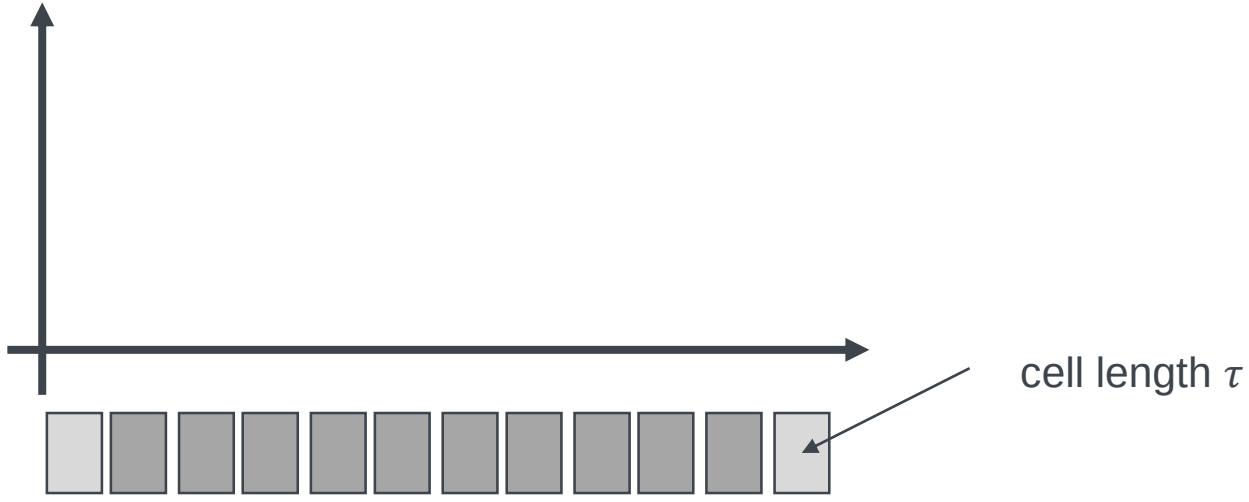

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Discrete time.

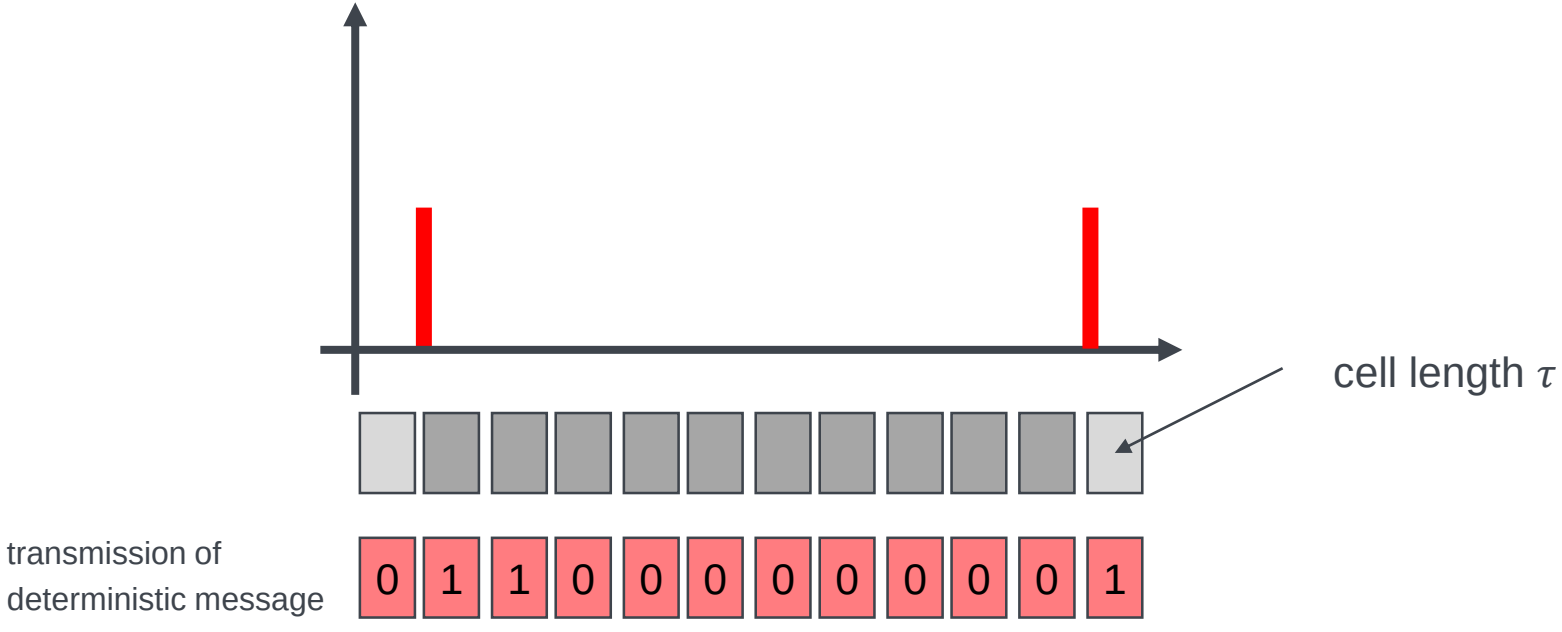

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Discrete time.

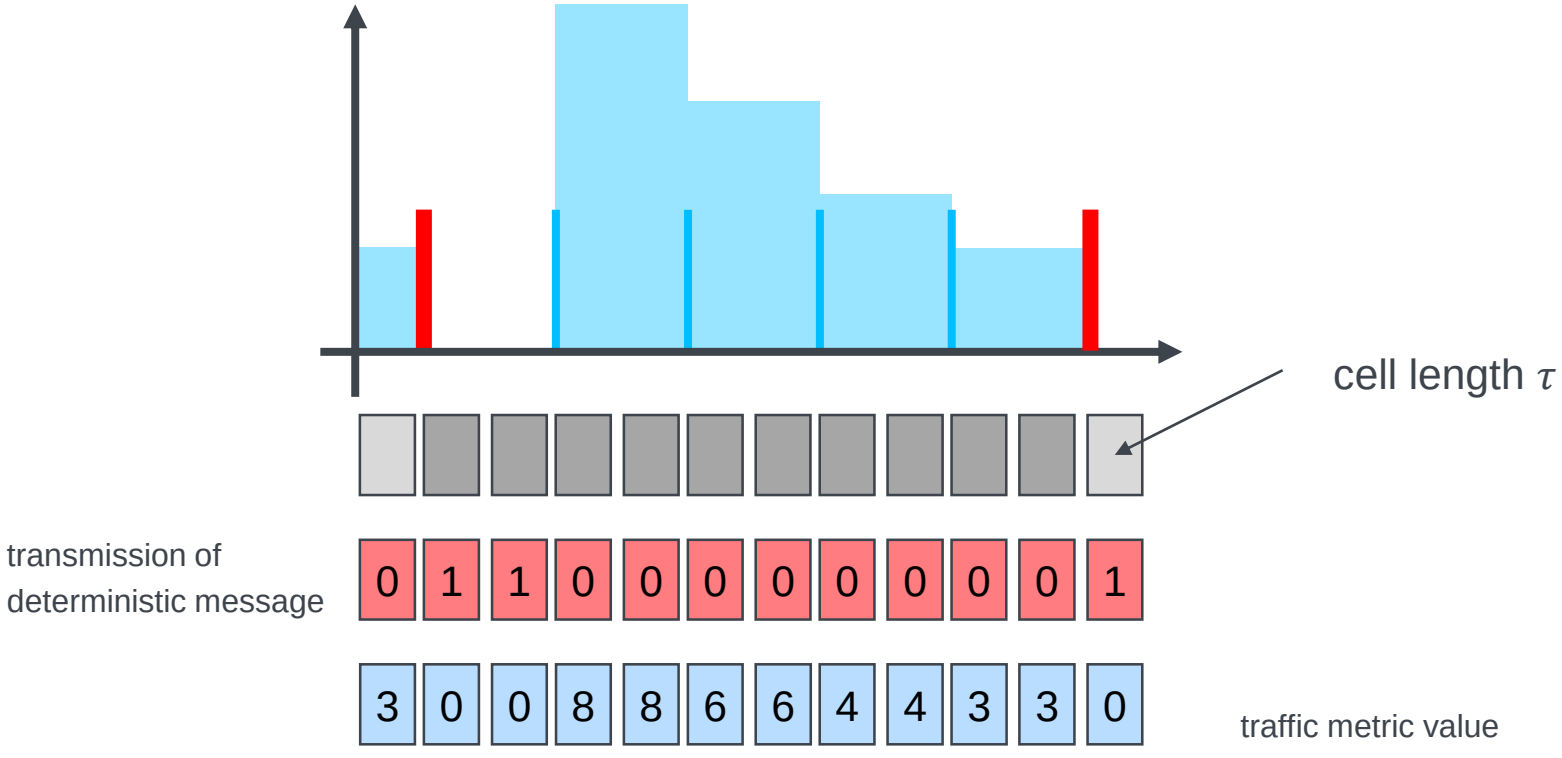

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Phase = 0

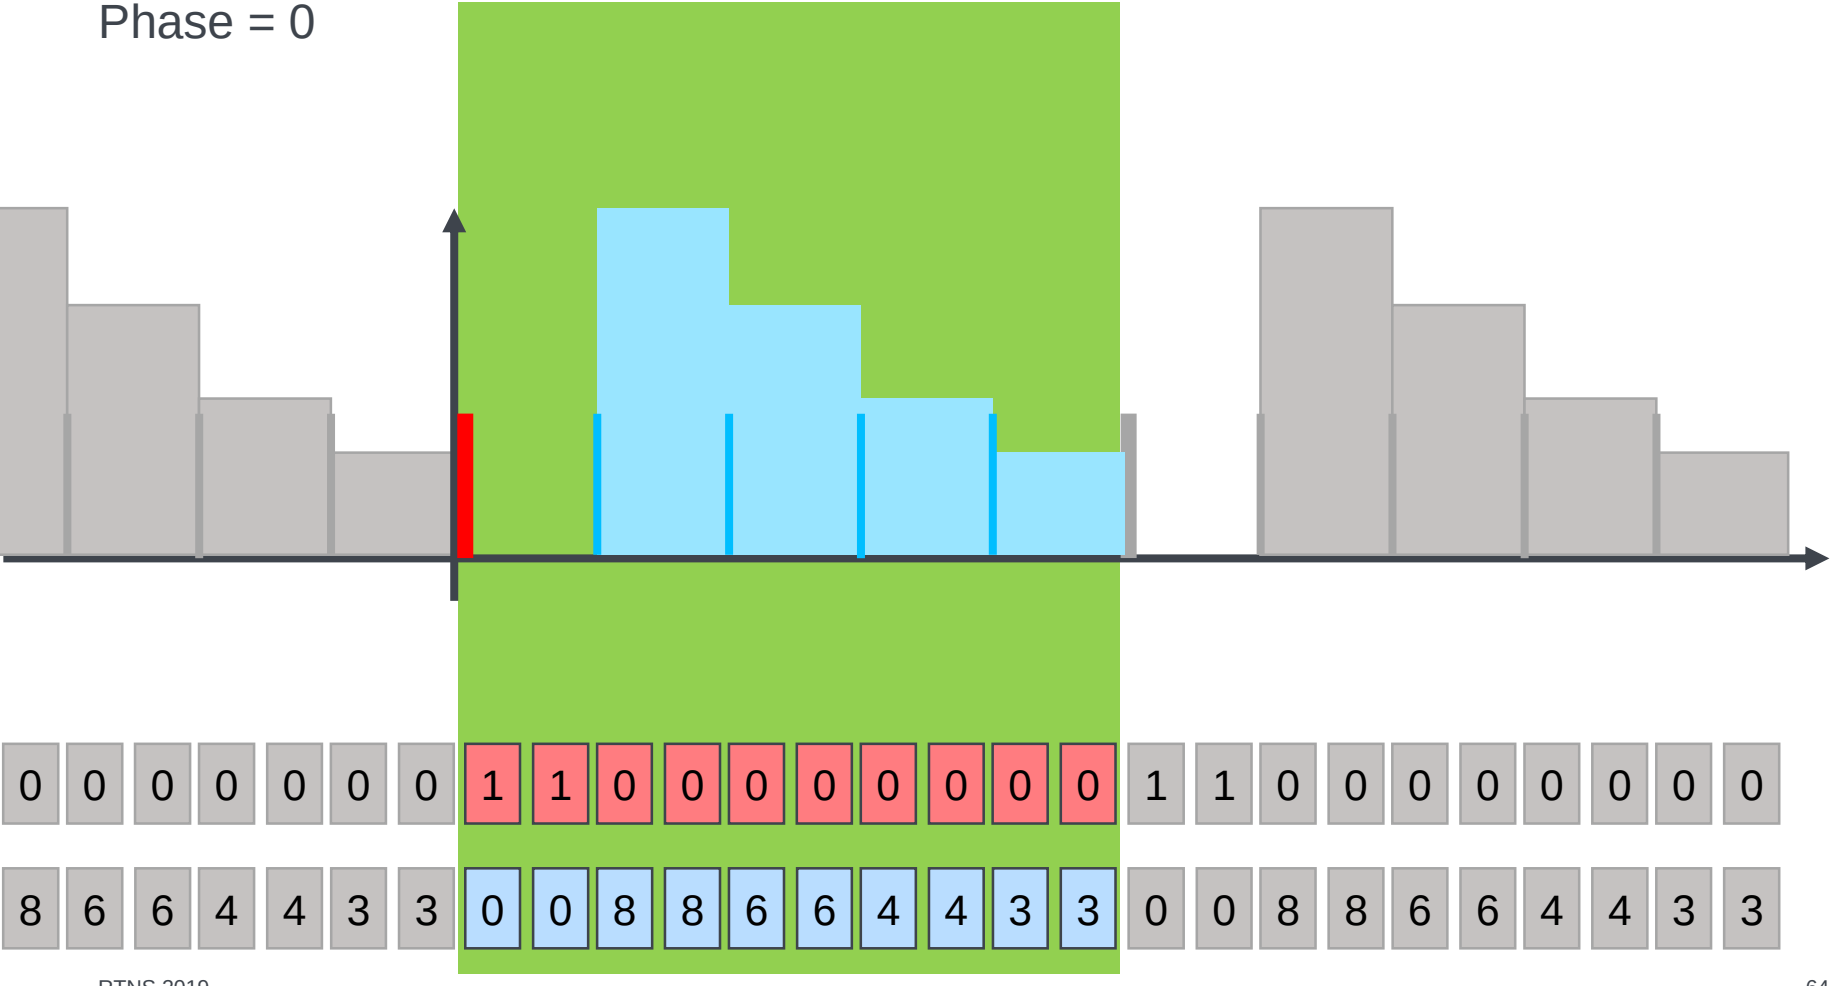

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Phase = $1 \cdot \tau$

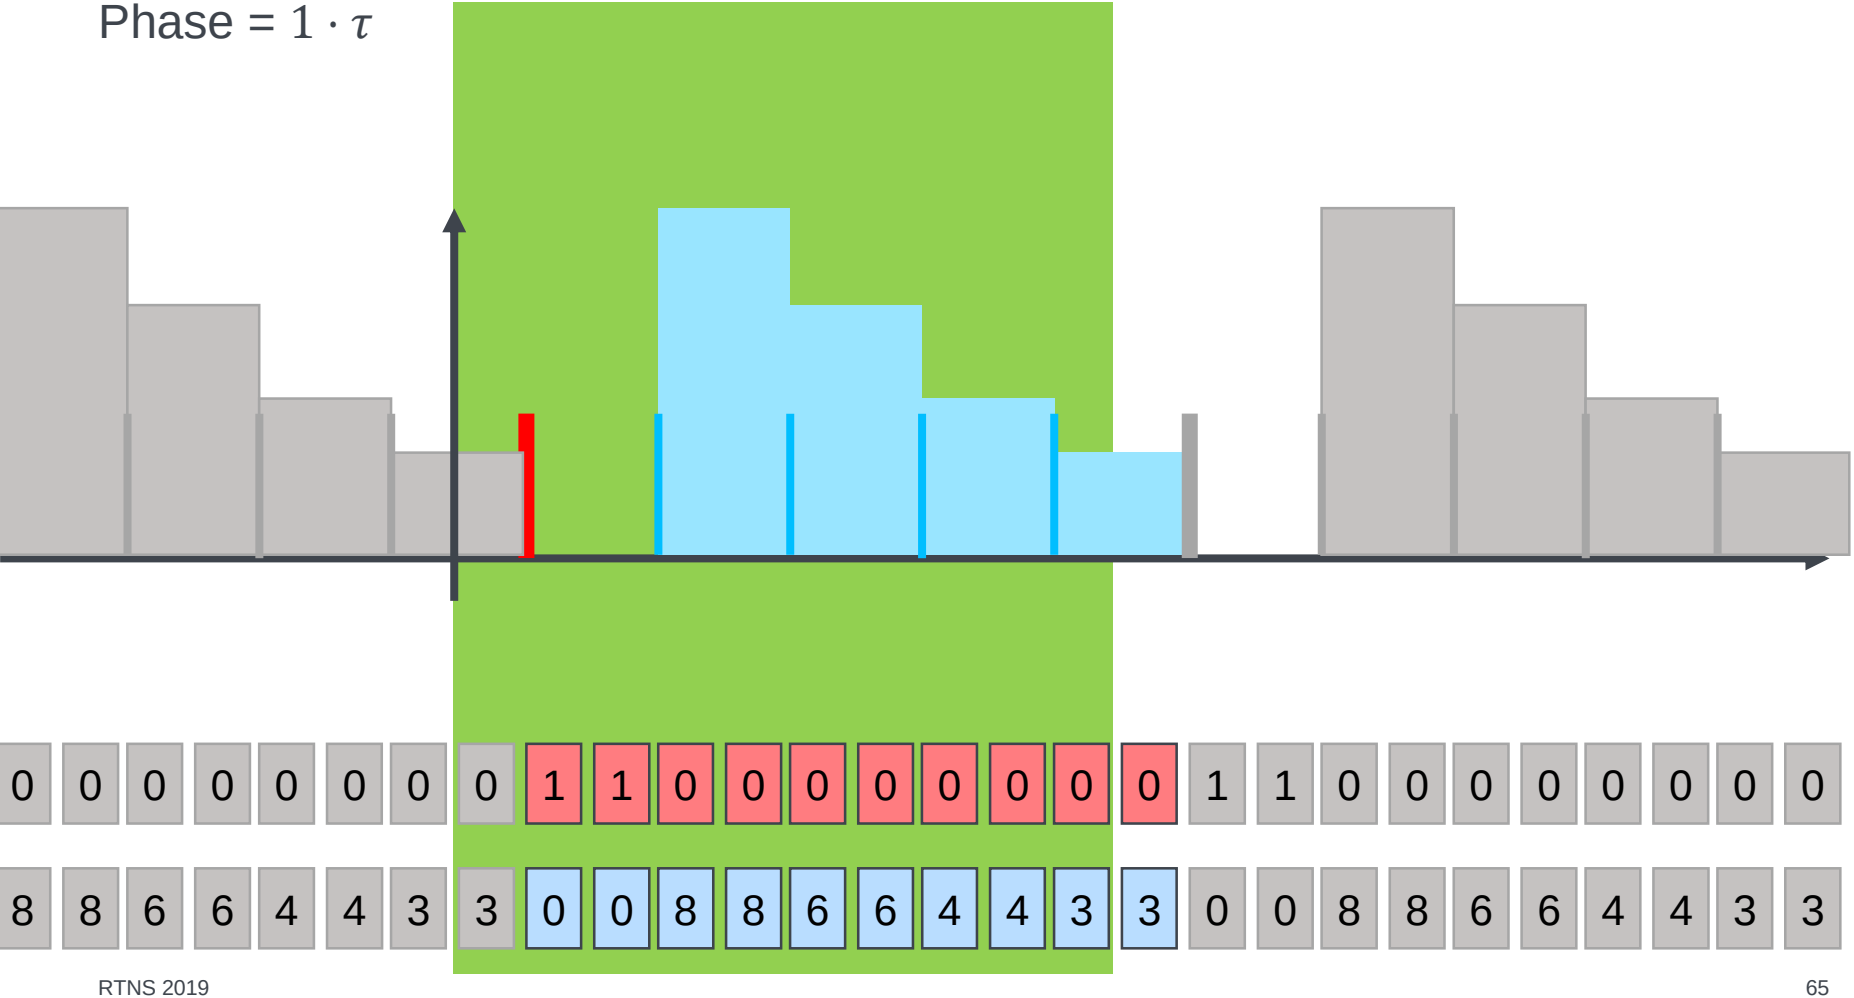

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Phase = $2 \cdot \tau$

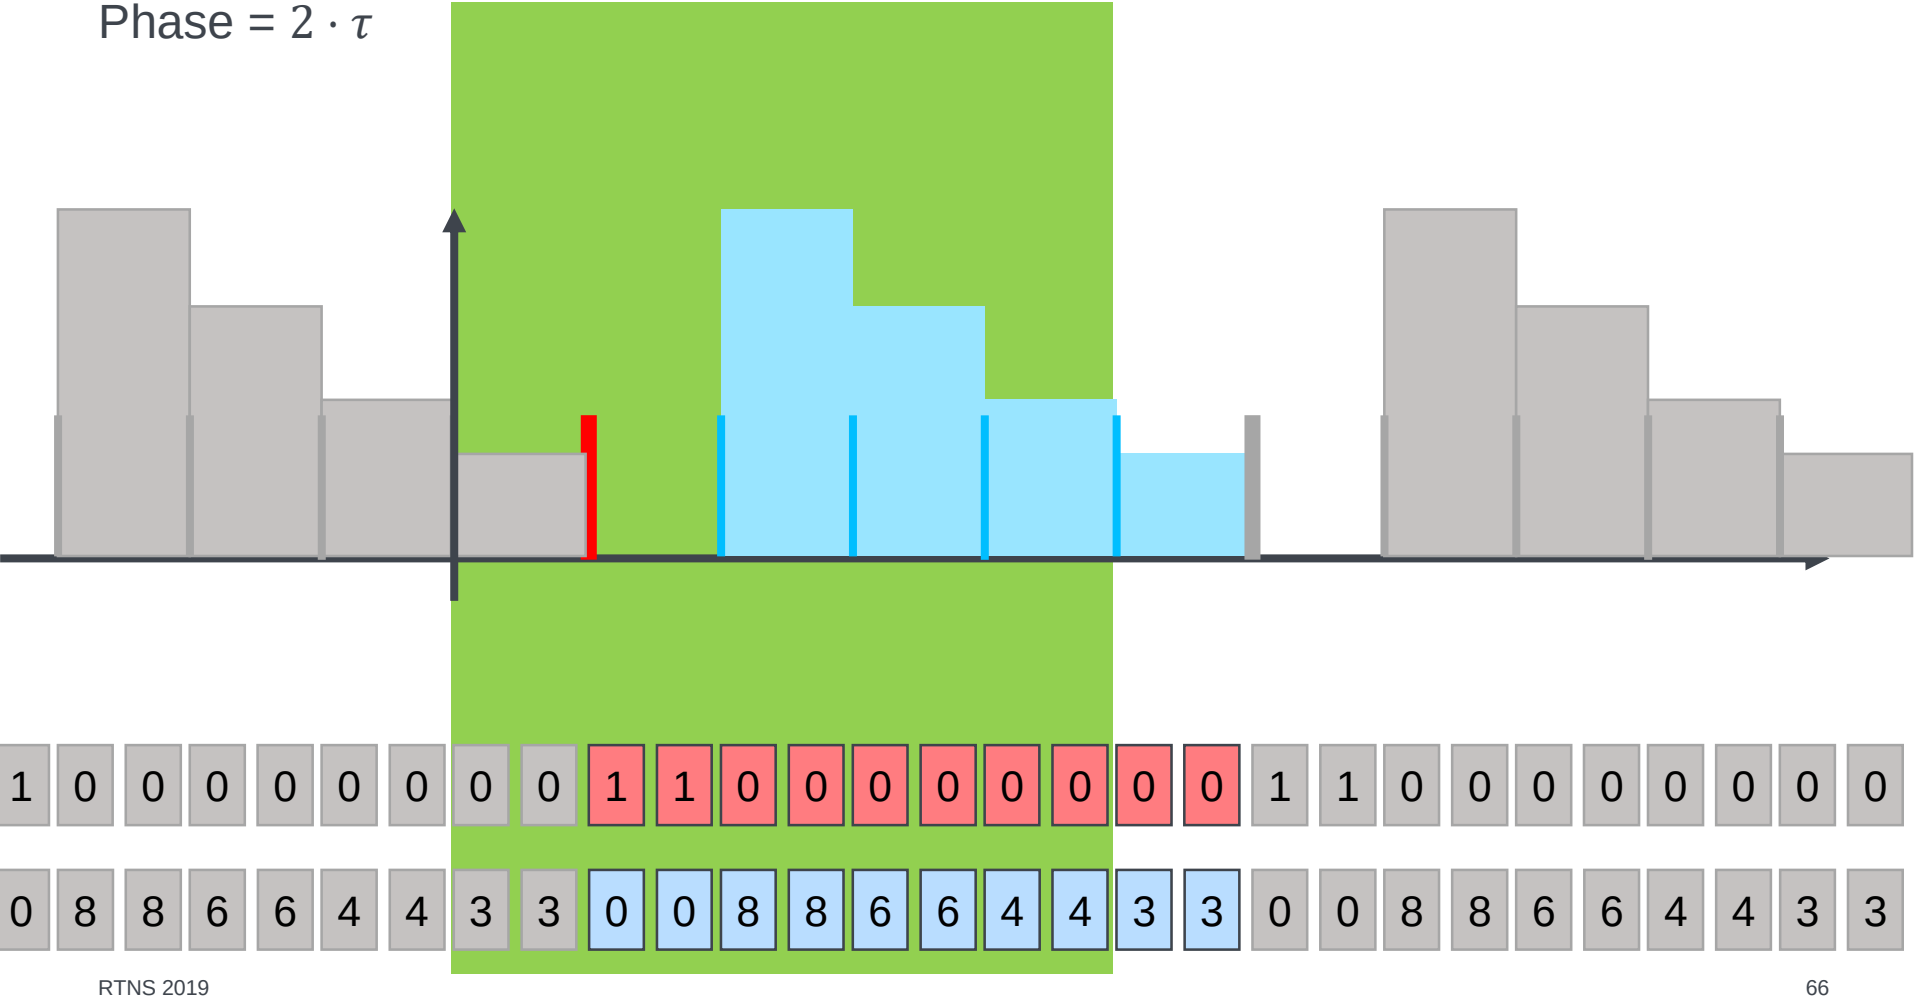

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Phase = $2 \cdot \tau$

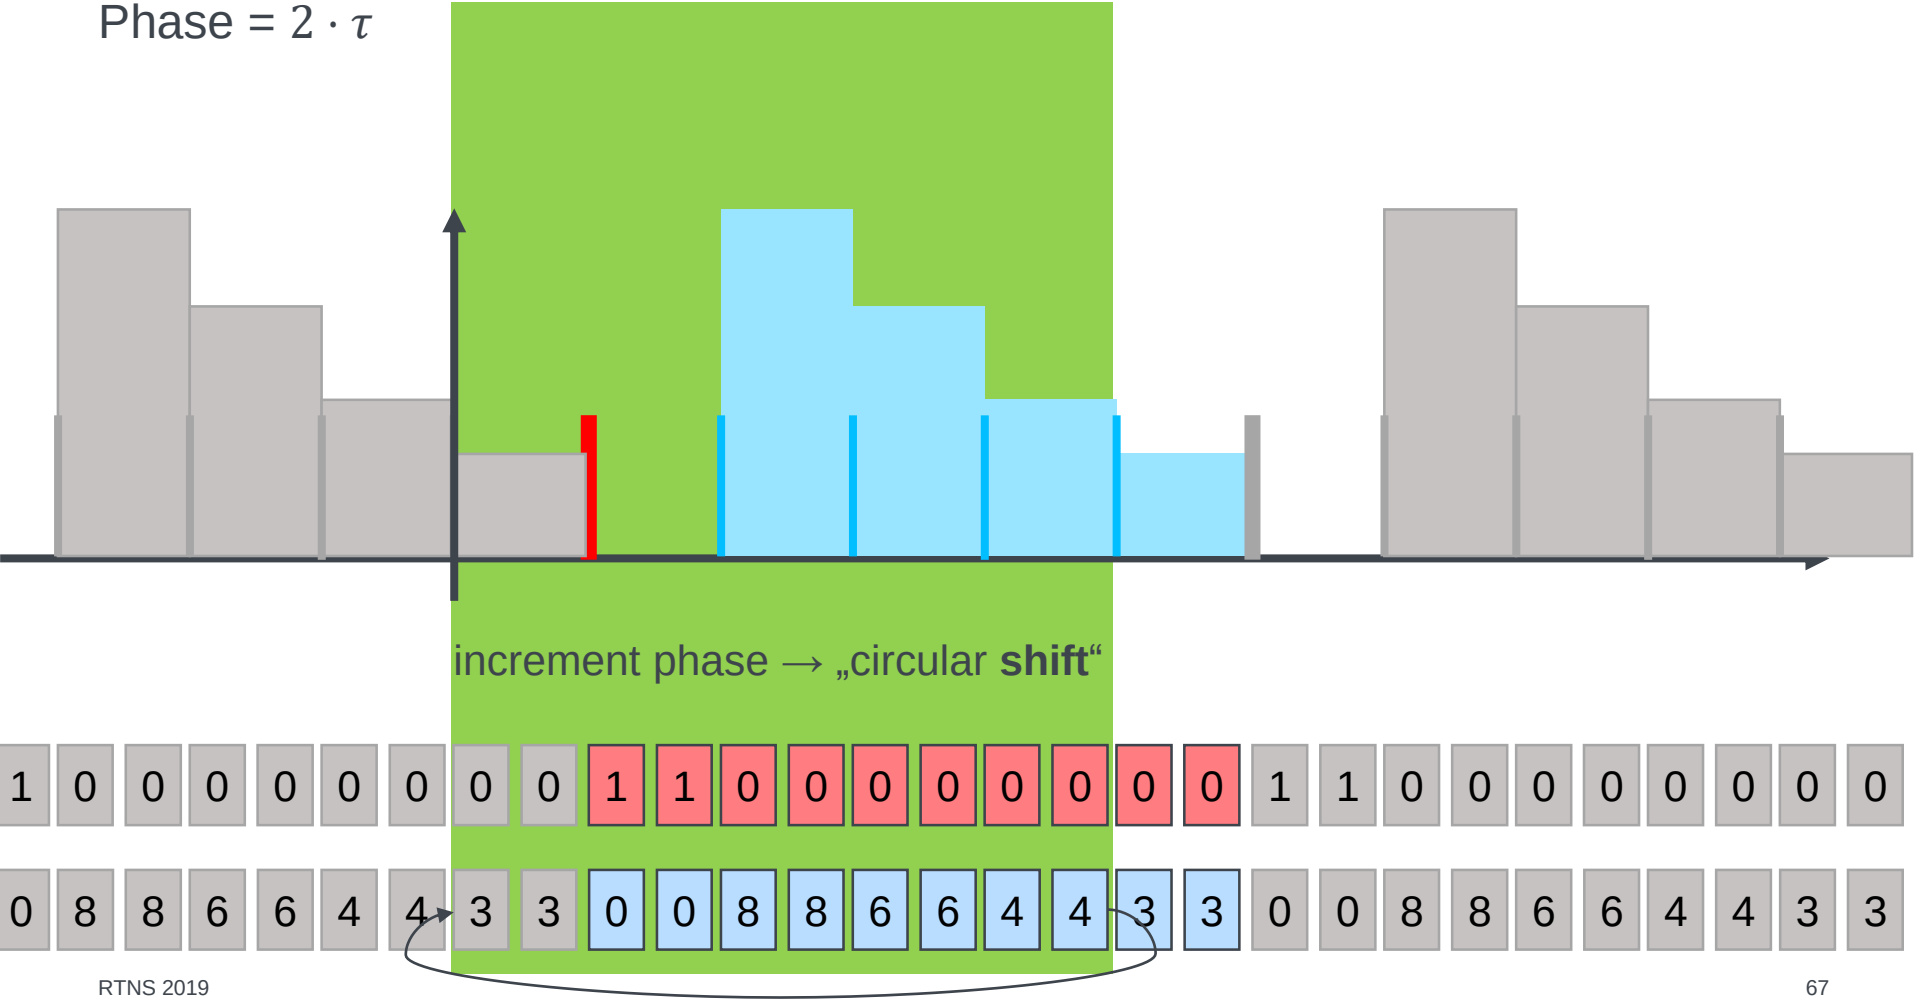

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Combining the scheduling variable and the cell arrays to model temporal properties.

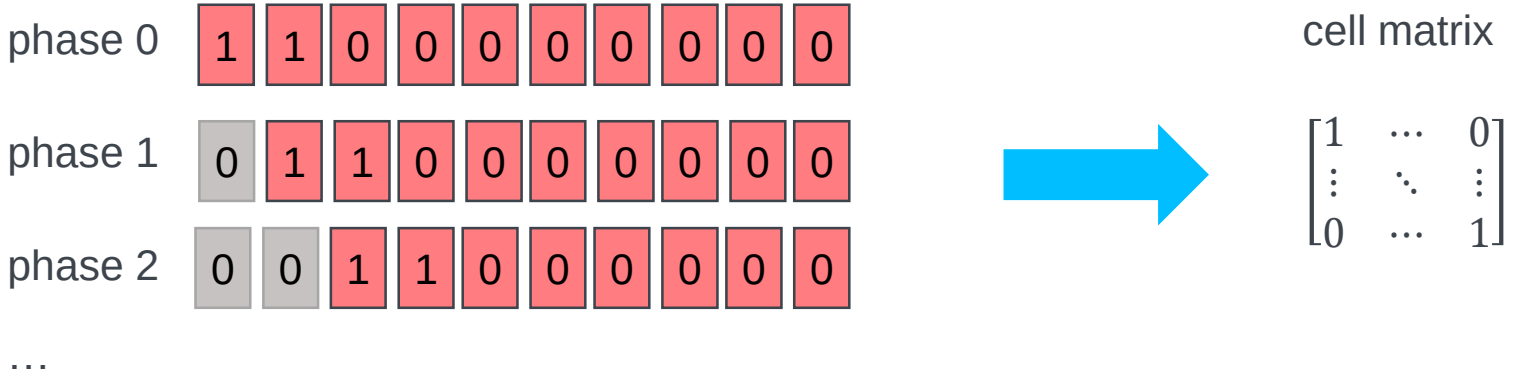

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Combining the scheduling variable and the cell arrays to model temporal properties.

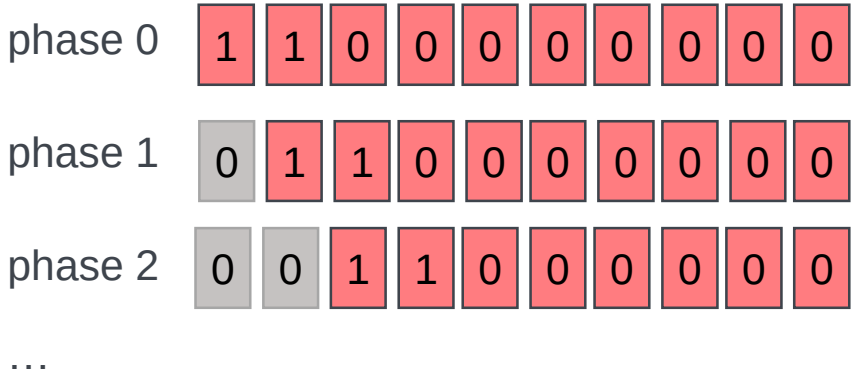

cell matrix

$$\begin{bmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{bmatrix}$$

1 in i’th position

$$[0 \cdots 0 \quad 1 \quad 0 \cdots 0] \cdot \begin{bmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{bmatrix} = \text{i'th row}$$

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Combining the scheduling variable and the cell arrays to model temporal properties.

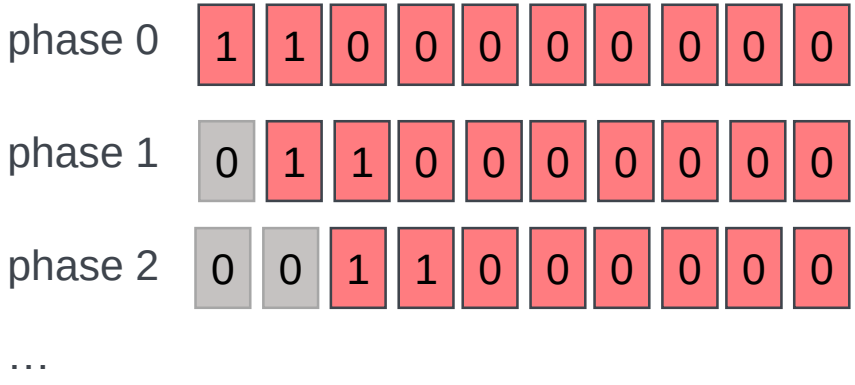

cell matrix

$$\begin{bmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{bmatrix}$$

MILP scheduling variables

1 in i’th position

$$[0 \cdots 0 \quad 1 \quad 0 \cdots 0] \cdot \begin{bmatrix} 1 & \cdots & 0 \\ \vdots & \ddots & \vdots \\ 0 & \cdots & 1 \end{bmatrix} = \text{i'th row}$$

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Intraflow scheduling and zero-queuing

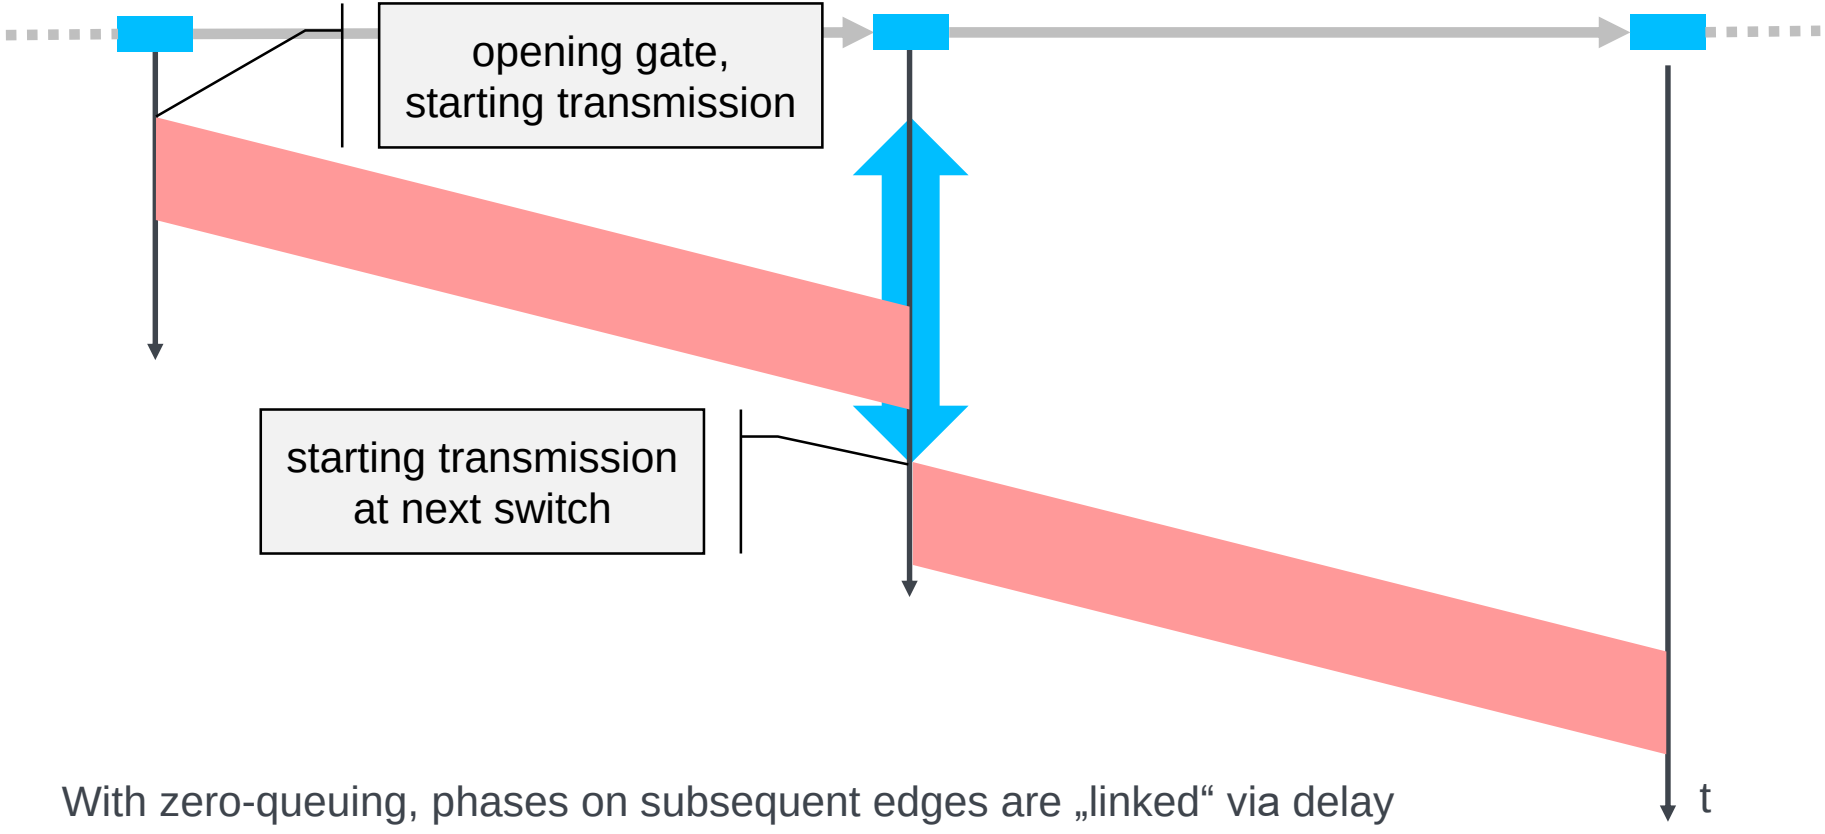

With zero-queuing, phases on subsequent edges are „linked“ via delay

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Intraflow scheduling and zero-queuing

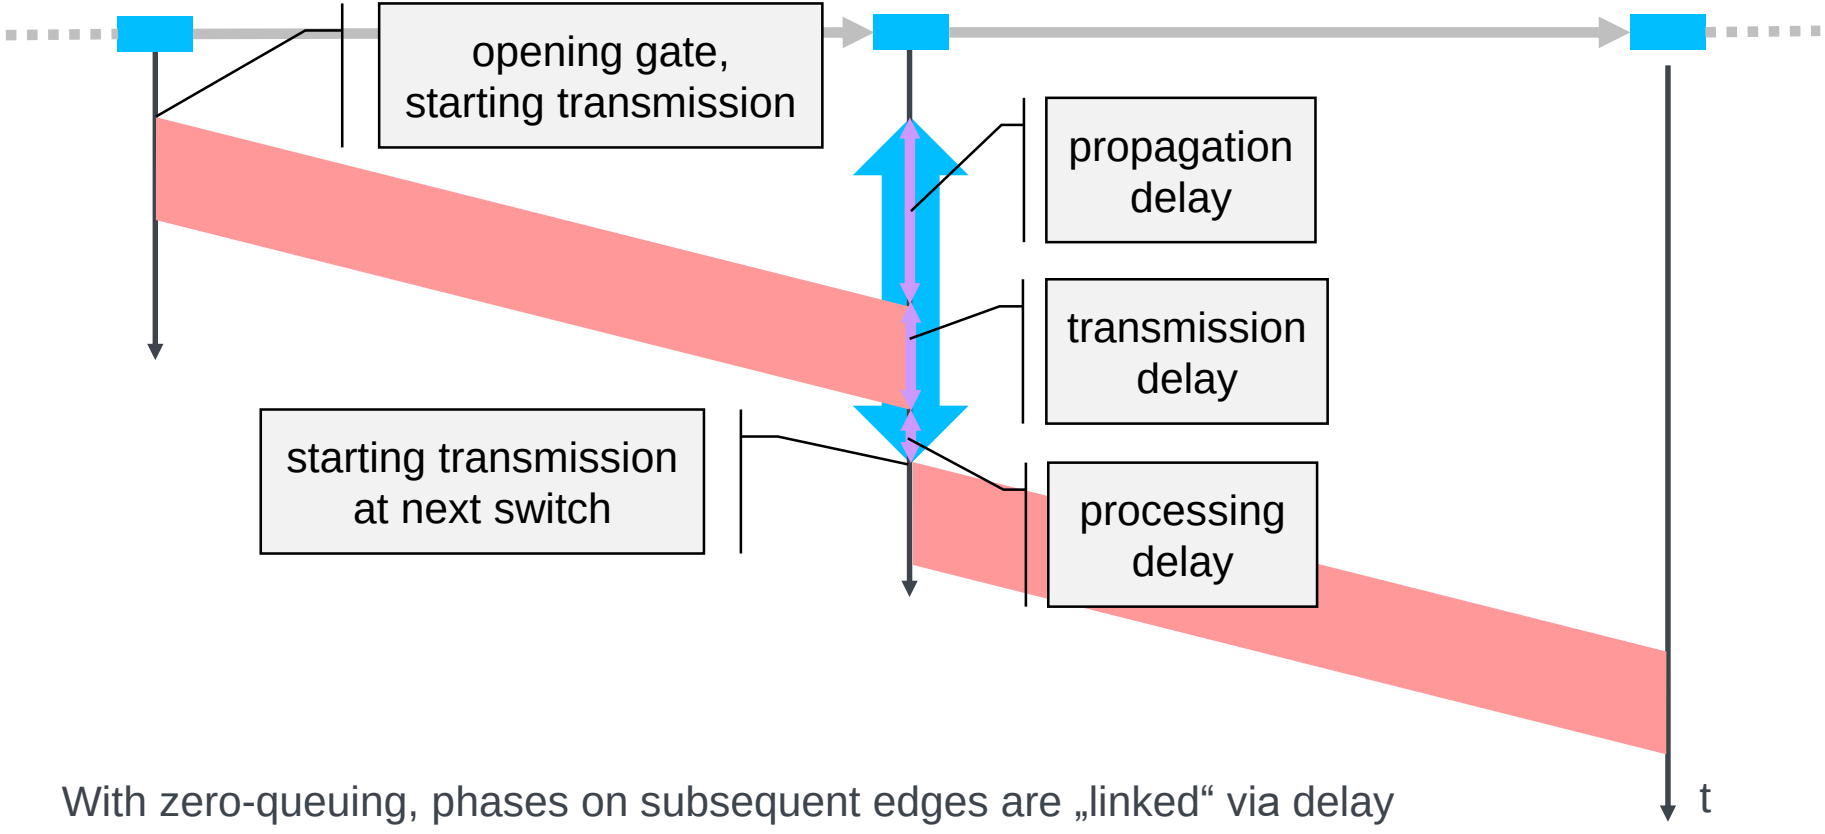

With zero-queuing, phases on subsequent edges are „linked“ via delay

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Exemplary constraints: Interflow scheduling and temporal exclusion

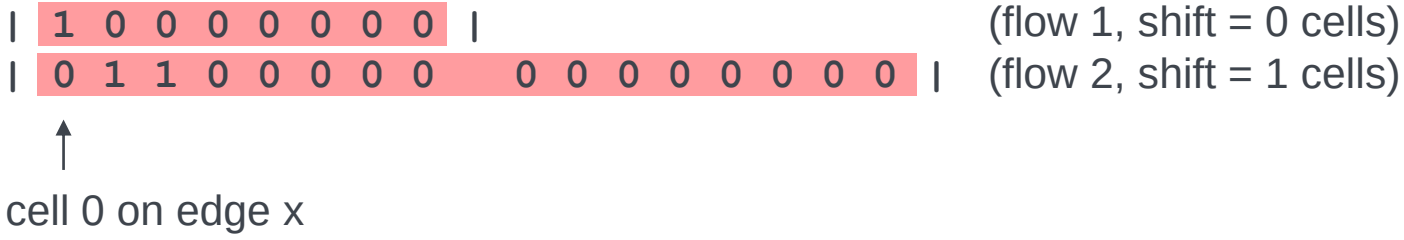

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Exemplary constraints: Interflow scheduling and temporal exclusion

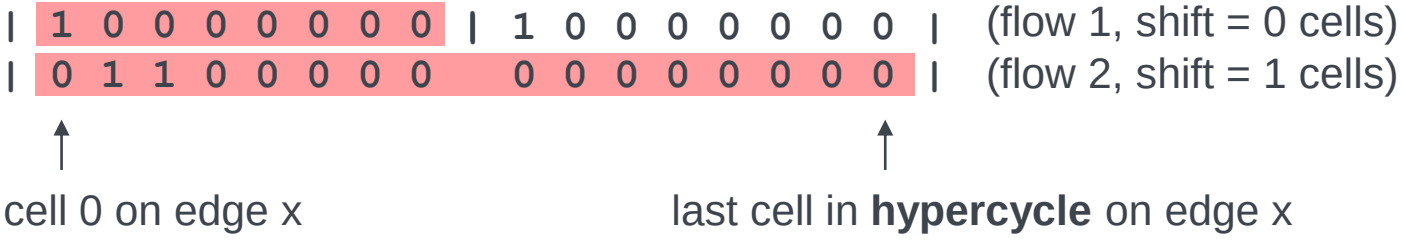

MILP: Path-Granularity Routing and Time-Dependent Traffic Metric

Exemplary constraints: Interflow scheduling and temporal exclusion

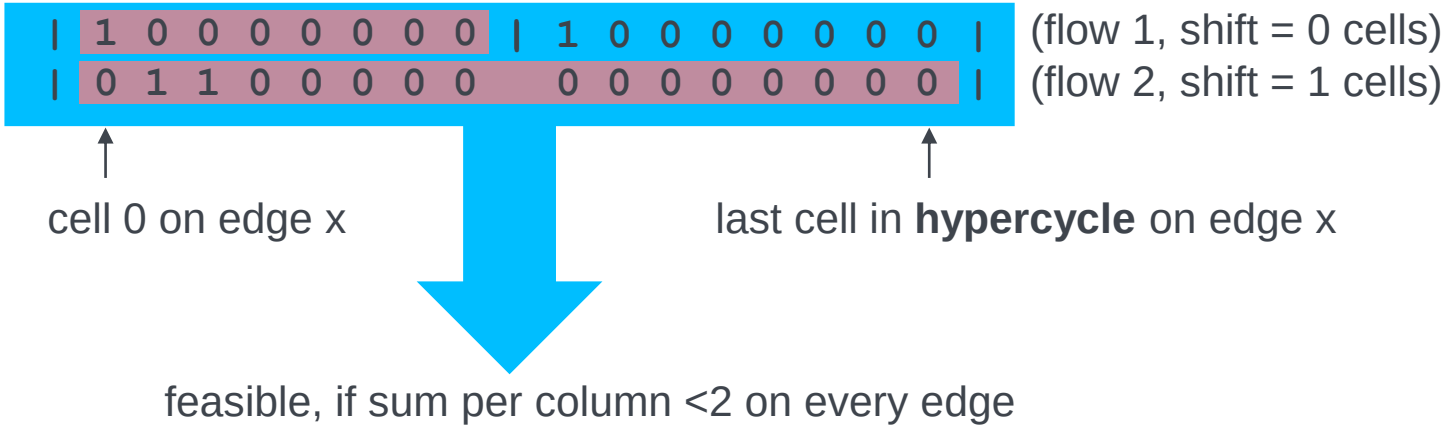

Evaluation

4-way comparison

- edge-gran., min-max
- edge-gran., no obj.
- path-gran., min-max
- path-gran., no obj.

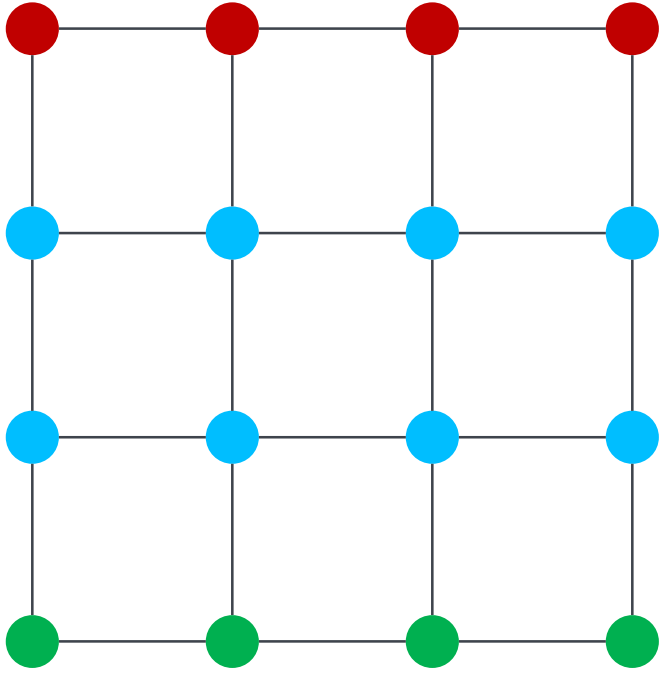

- time limit: 30 min
- container, pyomo + gurobi 8.1.0
- 4x Intel Xeon E7-4850, 2.1 GHz, 1 TB RAM, Linux 4.19.4

Evaluation

Varying number of flows

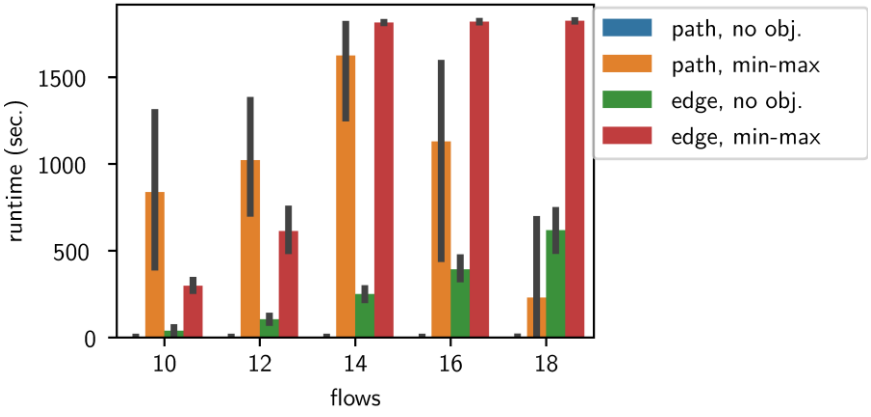

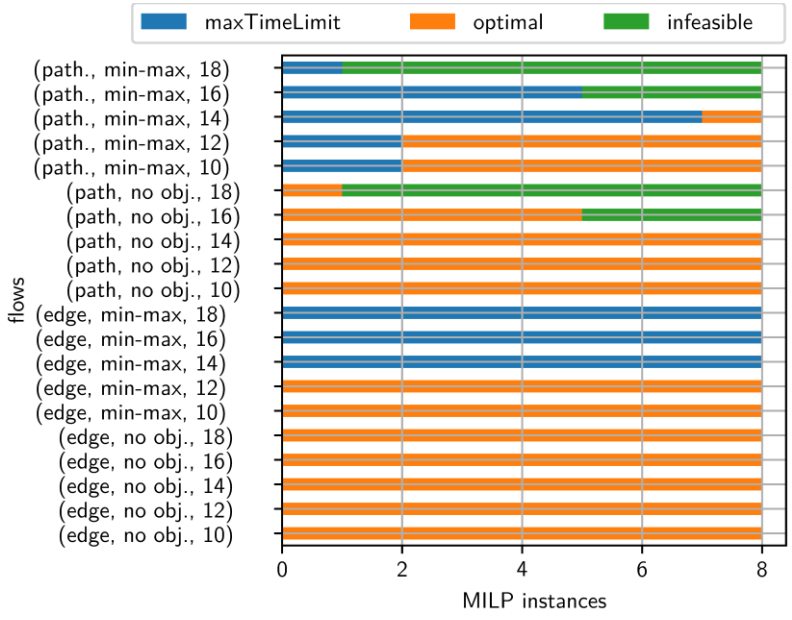

4 precomputed paths for path no obj./min-max

- edge, no obj./min-max yields larger MILP (order of ~10^5 constraints/variables) compared to path, no obj./min-max (order of ~10^4 constraints/variables)
- path, min-max reaches runtime limit earlier than edge, min-max (cf. 2 instances at 10 flows)
- infeasibility discovered quickly

Evaluation

Varying number of paths

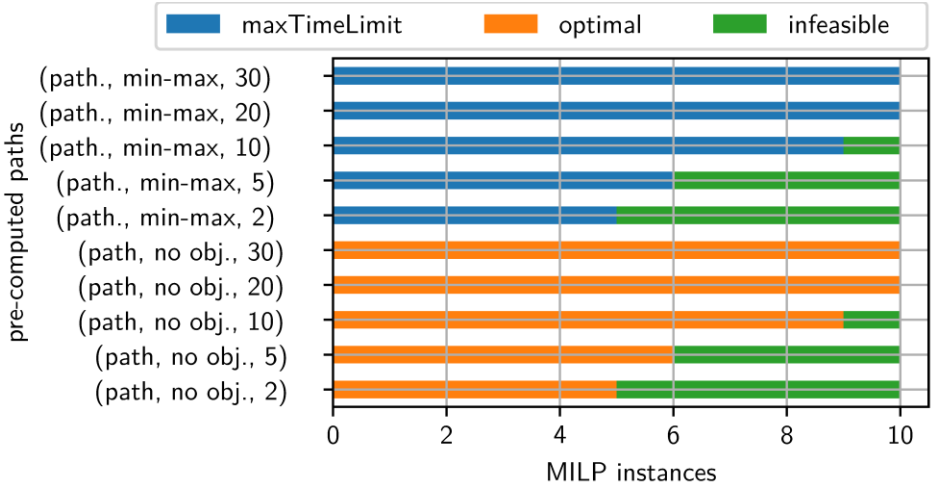

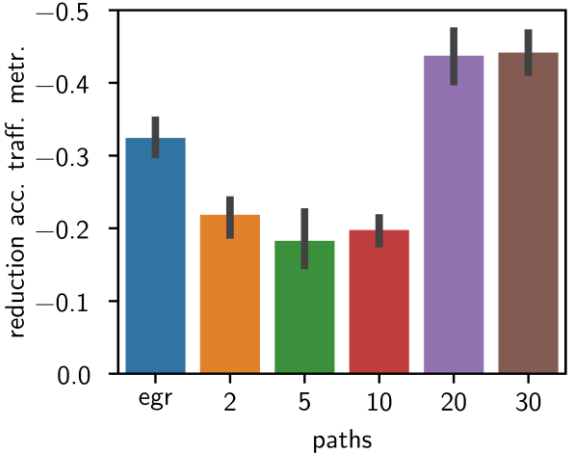

16 flows

- more paths: increases feasibility and possible improvement, (edge, no obj./min-max always return feasible solution)
- MILP with objective yields improvement over “feasible” schedule

Final remarks

- Complemental flows to improve application performance
- Routing + scheduling of complemental flows

- Different MILP formulations
  - trade-off routing vs. scheduling fidelity

- Open Question
  - Consideration of queueing effects for opportunistic messages?

La fin.

Traffic Engineering

Application

Traffic Engineering

Application

- stationary
- directed, point-to-point data stream
- **traffic metric** for opp. traffic

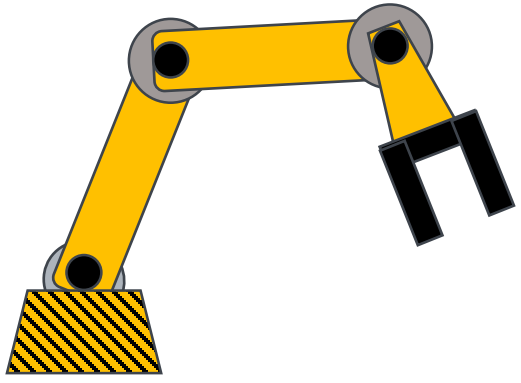

Traffic Engineering

Application

- stationary
- directed, point-to-point data stream
- **traffic metric** for opp. traffic

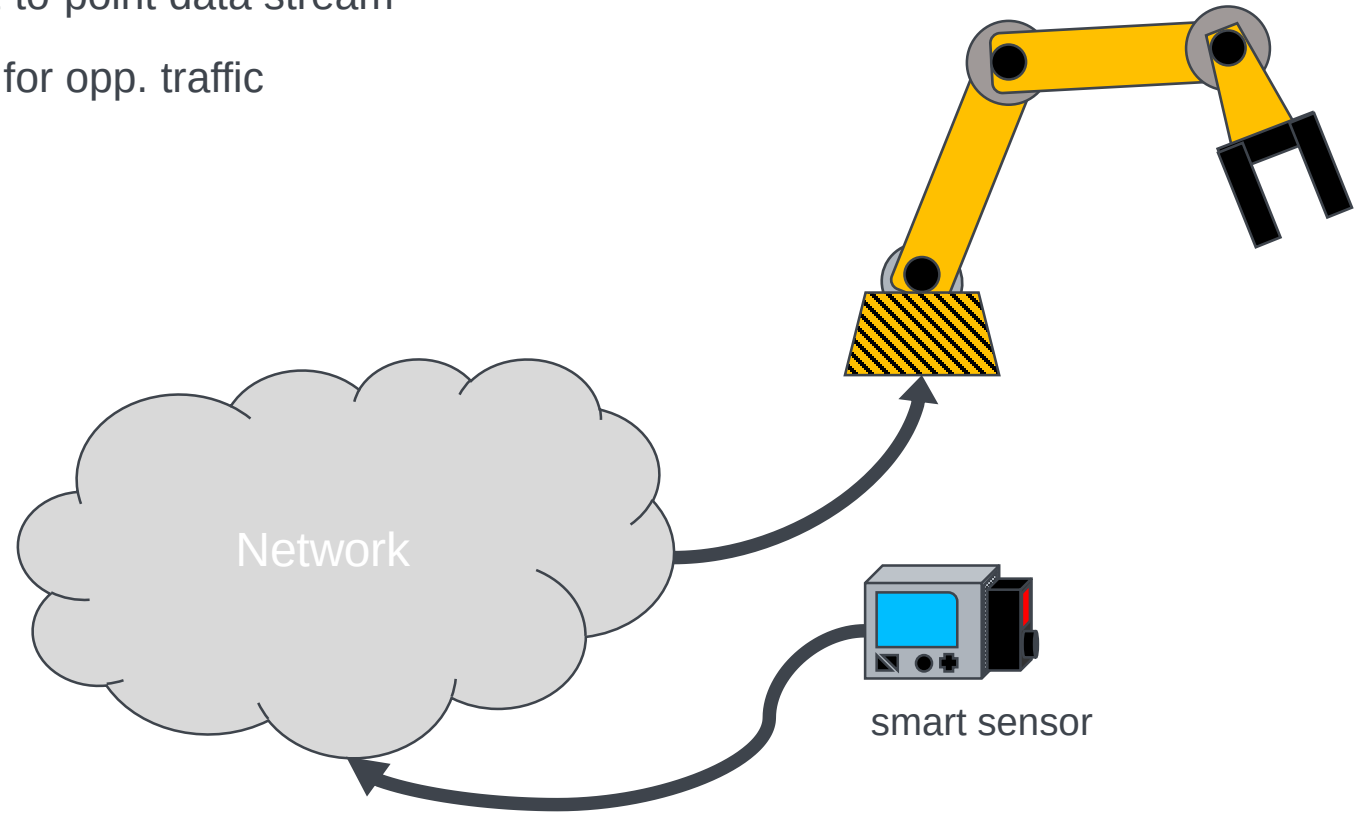

Traffic Engineering

Application

- stationary
- directed, point-to-point data stream
- **traffic metric** for opp. traffic

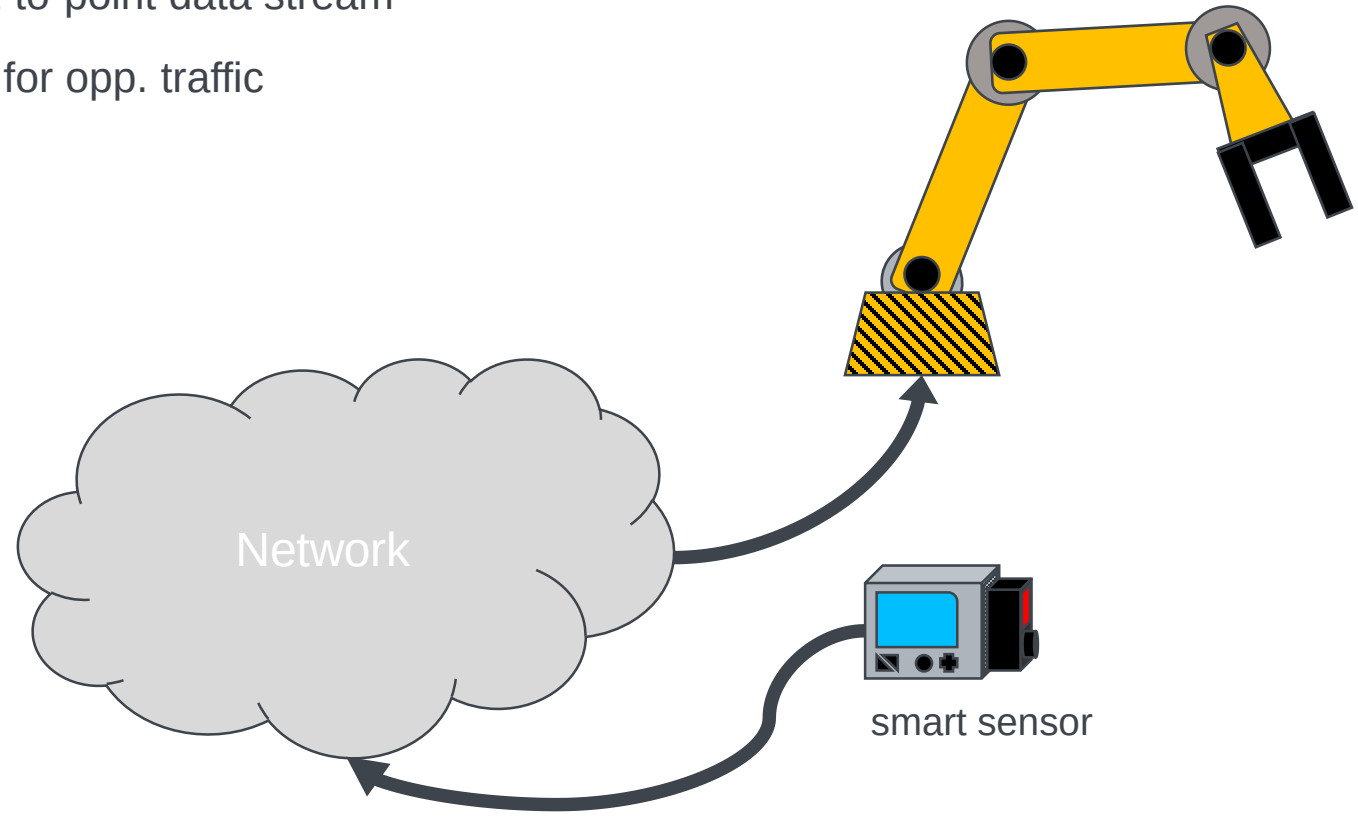

Traffic Engineering

Application

- stationary
- directed, point-to-point data stream
- **traffic metric** for opp. traffic