# End-to-end Real-time Guarantees in Wireless Cyber-physical Systems


Romain Jacob
Marco Zimmerling
Pengcheng Huang
Jan Beutel
Lothar Thiele



RTSS 16 - IoT and Networking Session
December 1, 2016



ETH zürich
TECHNISCHE UNIVERSITÄT DRESDEN

# Predictability is key!

# A Cyber-physical System goes beyond a real-time wireless network

Predictability

Real-time guarantees > *Deadlines are met*

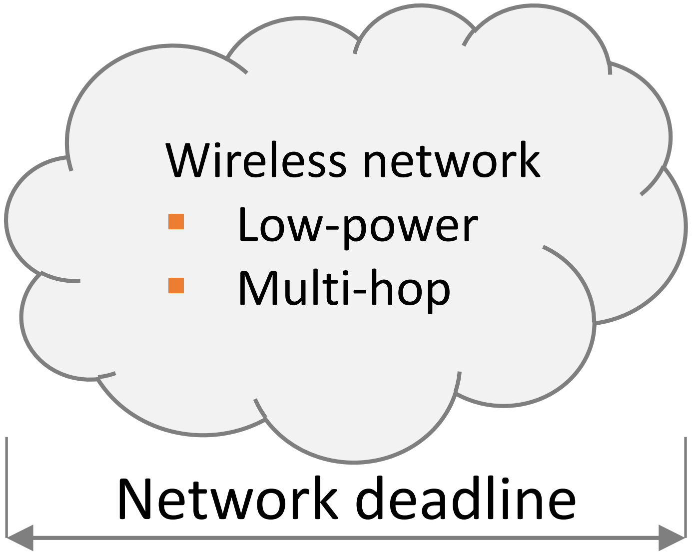

# A Cyber-physical System goes beyond a real-time wireless network

Predictability

Real-time guarantees > *End-to-end deadlines are met*

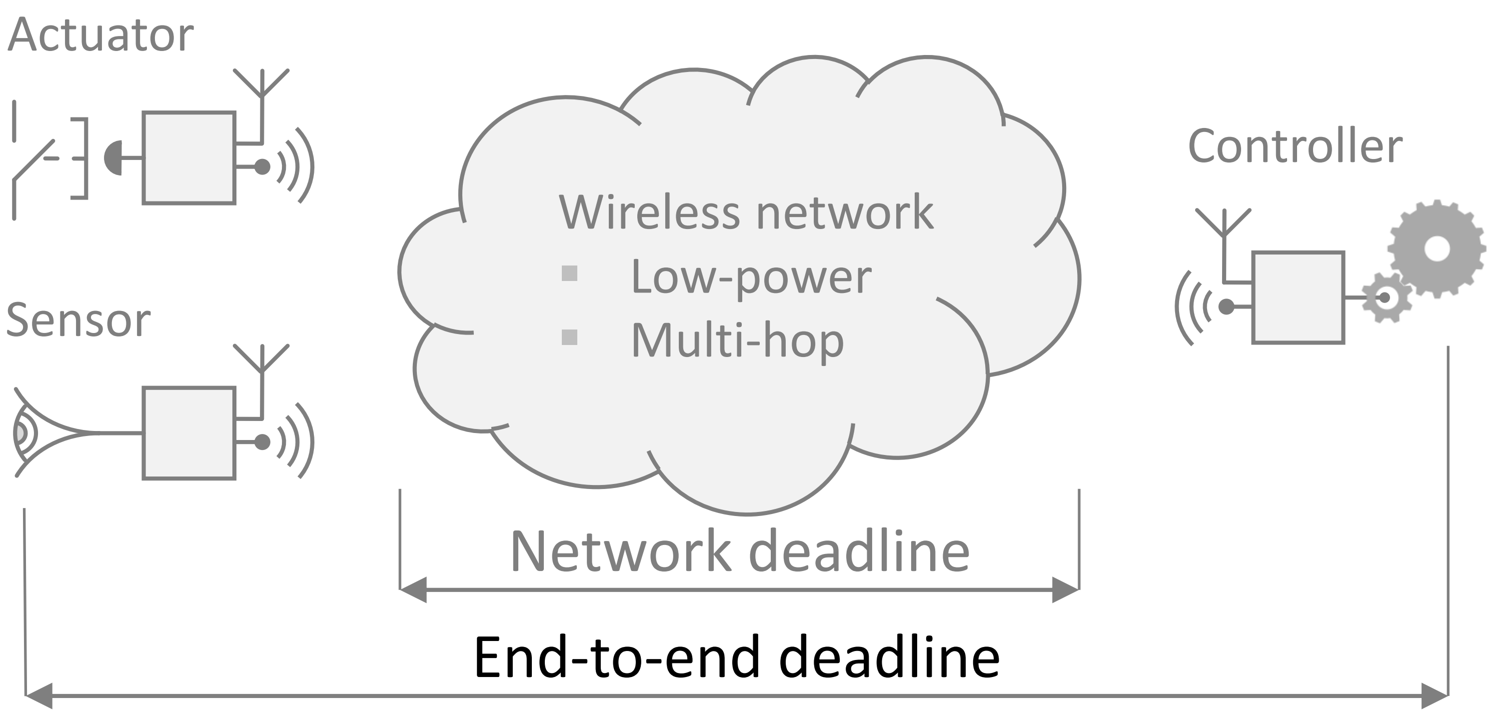

# A Cyber-physical System goes beyond a real-time wireless network

Predictability

Real-time guarantees > *End-to-end deadlines are met*

Buffer management > *No buffer overflow*

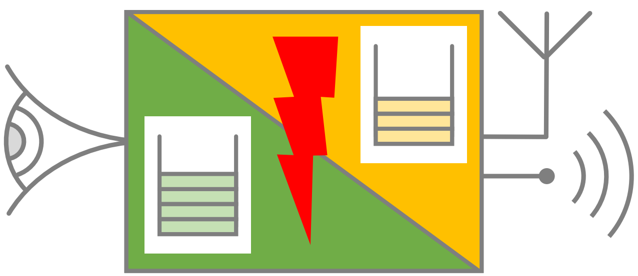

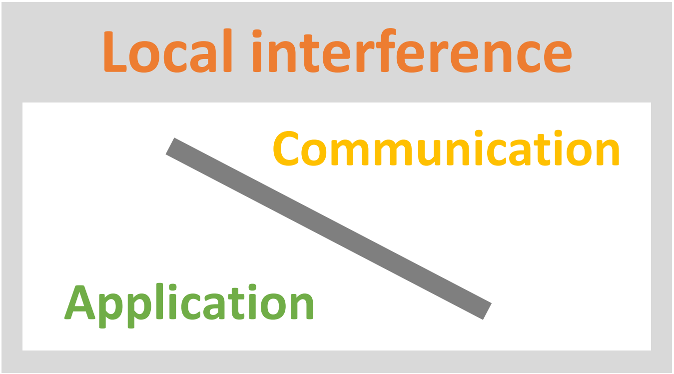

# Design goals

| | Network | Device | System |
|---|---|---|---|
| *Predicatability*<br><br>Real-time guarantees<br><br>Buffer management | | | |
| *Adaptability* | | | |
| Efficiency | | | |
| Composability | | | |

First piece

# A *predictable* and *adaptive* wireless network

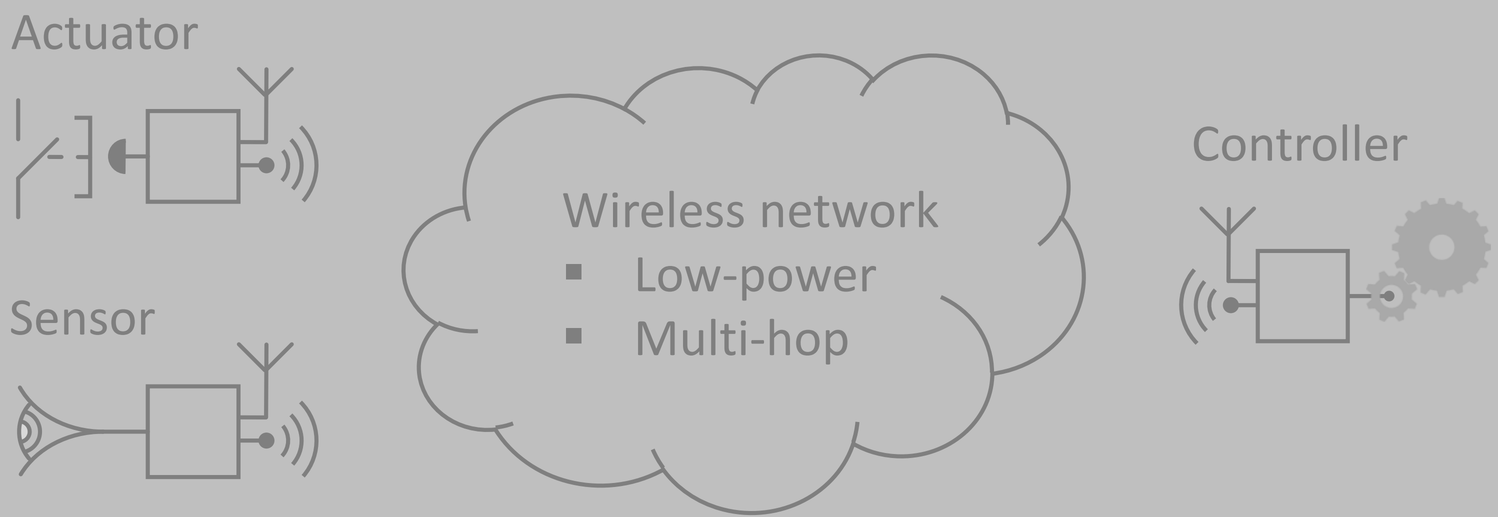

First piece

# A *predictable* and *adaptive* wireless network

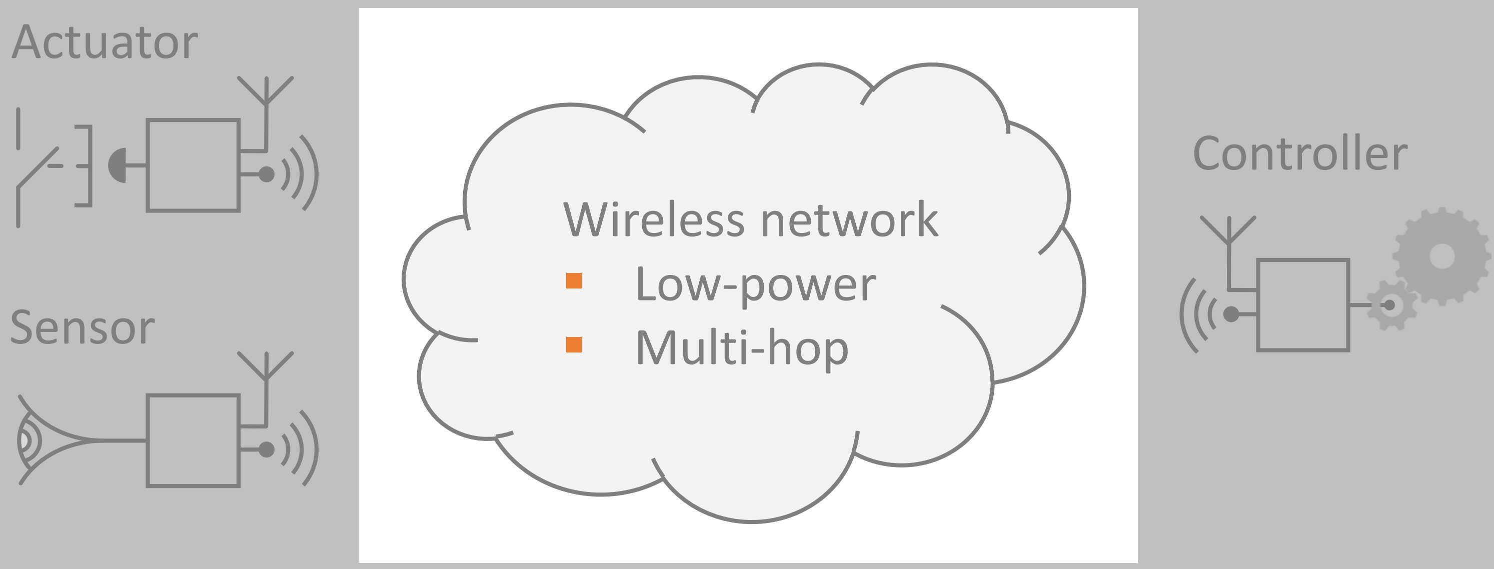

# State-of-the-art wireless protocols are *predictable* or *adaptive*

| | |
| --- | --- |
| Splash, RAP | Efficient<br>Adaptive<br>> *not Predictable* |
| WirelessHART | Predictable<br>Efficient<br>> *not Adaptive* |

# Blink A real-time, reliable and adaptive wireless protocol

[1]

| | |
|---|---|
| Adaptive | Based on Glossy<br>> *Flooding primitive* |
| Reliable | Average 99.97% reception rate<br>> *Multiple testbeds*<br>> *Tested up to 94 nodes* |
| Real-time | Online scheduling<br>> *EDF-based Lazy Scheduling* |

[1] Zimmerling M. et al., *Adaptive Real-time Communication for Wireless Cyber-physical Systems*
To appear in ACM Transactions on Cyber-Physical Systems, **2016**

# Blink [1] A real-time, reliable and adaptive wireless protocol

Abstraction — Low-power Wireless Bus (LWB)
> *MAC protocol*

Communication in rounds — TDMA-based
> *Sleep between rounds*
> *Wireless yields sync!*

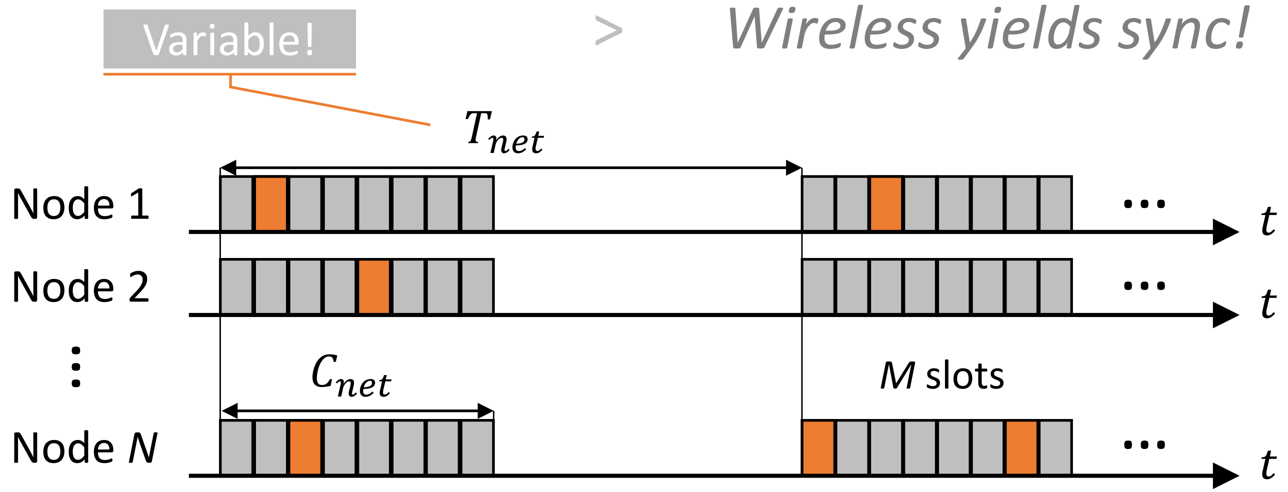


[1] Zimmerling M. et al., *Adaptive Real-time Communication for Wireless Cyber-physical Systems*
To appear in ACM Transactions on Cyber-Physical Systems, **2016**

# Blink [1] A real-time, reliable and adaptive wireless protocol

Abstraction — Low-power Wireless Bus (LWB)
> *MAC protocol*

Communication in rounds — TDMA-based
> *Sleep between rounds*
> *Wireless yields sync!*

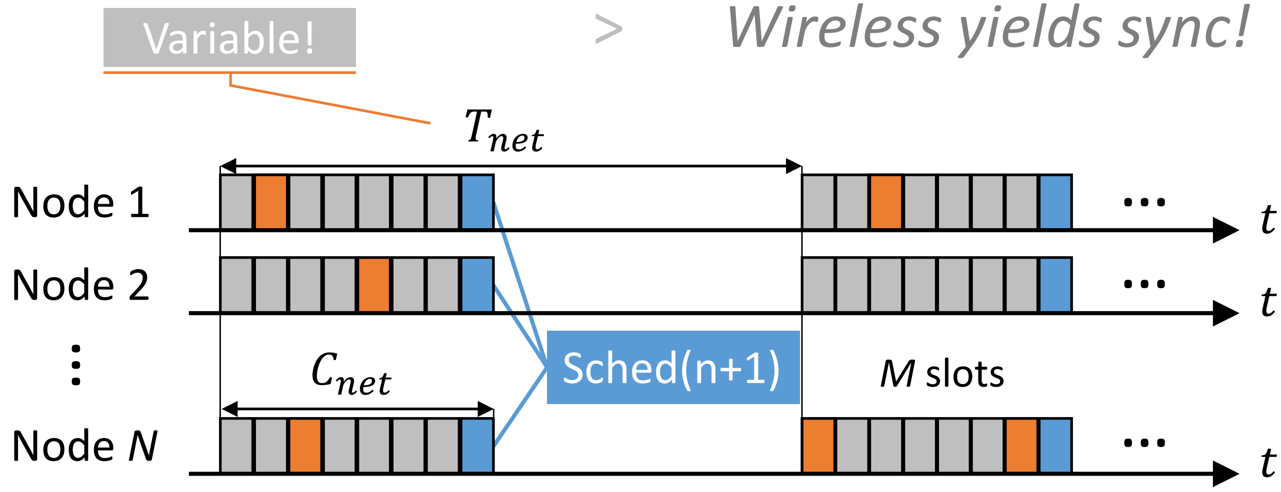


[1] Zimmerling M. et al., *Adaptive Real-time Communication for Wireless Cyber-physical Systems*
To appear in ACM Transactions on Cyber-Physical Systems, **2016**

# Design goals

| | Blink | | |
|---|---|---|---|
| | Network | Device | System |
| Real-time guarantees | ✓ | | |
| Buffer management | na | | |
| Adaptability | ✓ | | |
| Efficiency | ✓ | | |
| Composability | na | | |

Second piece

# A *dual-processor architecture* to mitigate local interference

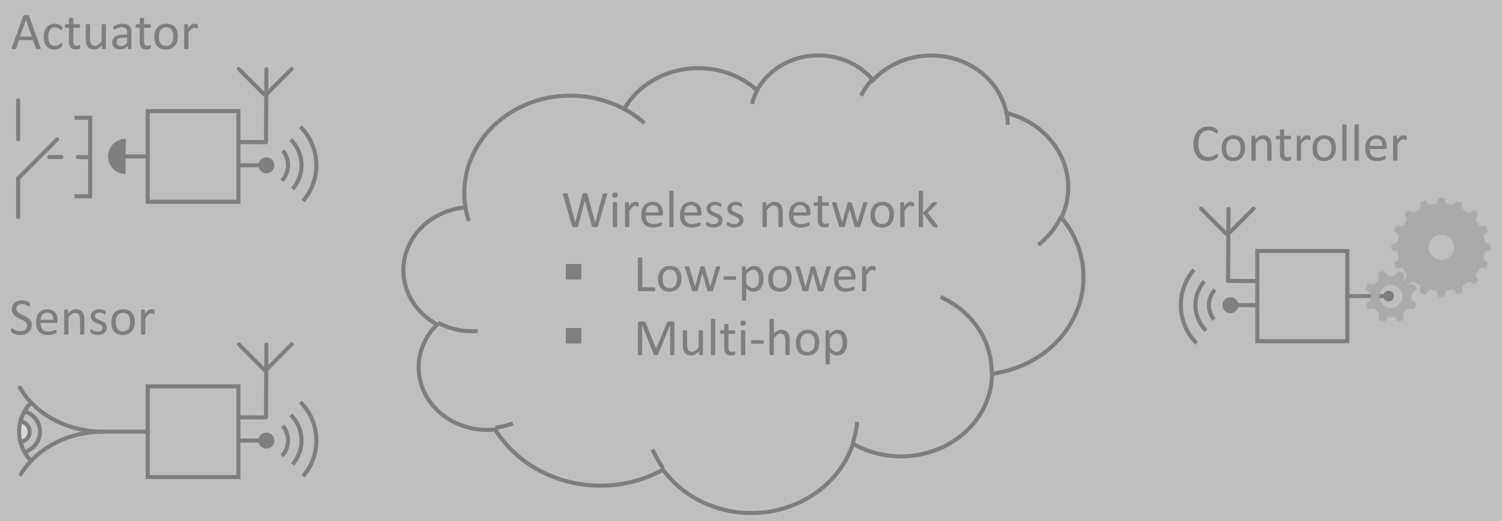

# Second piece — A *dual-processor architecture* to mitigate local interference

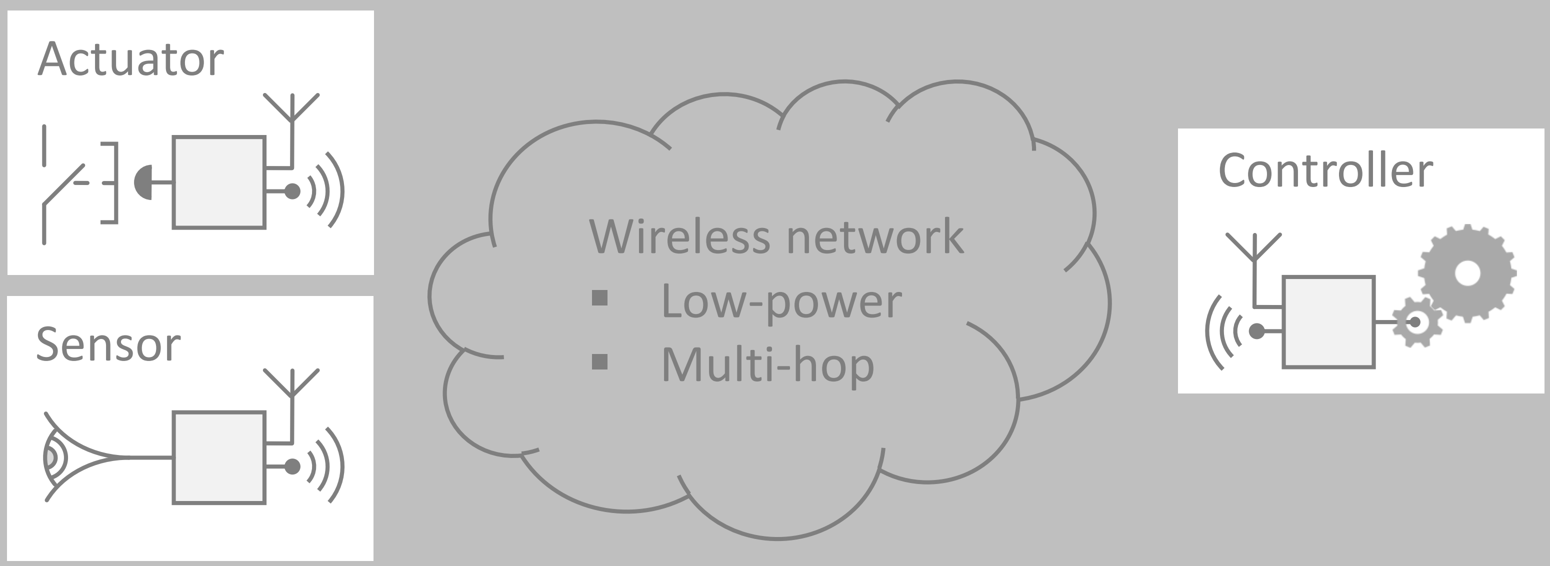

# Bolt

[2]

A *dual-processor architecture* to mitigate local interference

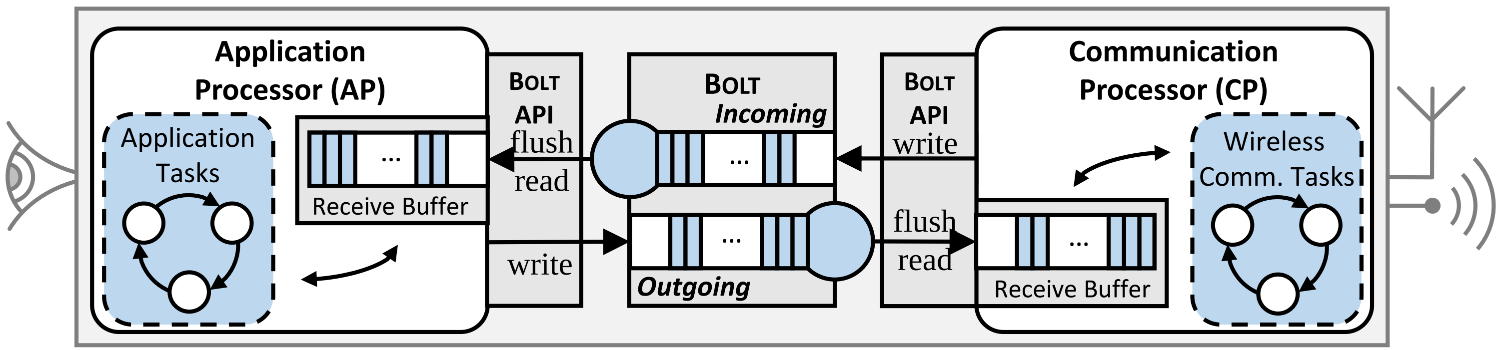


| | |
|---|---|
| Real-time behavior | *Formally verified* *Implemented, tested, deployed* |
| Efficient | *µW > sleep* *mW > active* |
| Composable | *Hardware/Software free composition* |


[2] Sutton F. et al., *Bolt: A Stateful Processor Interconnect*, SenSys'15, **2015**

# Design goals

| | Blink | Bolt | |
|---|---|---|---|
| | Network | Device | System |
| Real-time guarantees | ✓ | ✓ | |
| Buffer management | na | | |
| Adaptability | ✓ | ✓ | |
| Efficiency | ✓ | ✓ | |
| Composability | na | ✓ | |

# Design goals → *How can we fill the blanks?*

| | Blink | Bolt | |
|---|---|---|---|
| | Network | Device | System |
| Real-time guarantees | ✓ | ✓ | ? |
| Buffer management | na | ? | ? |
| Adaptability | ✓ | ✓ | ? |
| Efficiency | ✓ | ✓ | ? |
| Composability | na | ✓ | ? |

Third piece

# Distributed Real-time Protocol (DRP)

| | Blink | Bolt | DRP |
|---|---|---|---|
| | Network | Device | System |
| Real-time guarantees | ✓ | ✓ | ? |
| Buffer management | na | ? | ? |
| Adaptability | ✓ | ✓ | ? |
| Efficiency | ✓ | ✓ | ? |
| Composability | na | ✓ | ? |

# DRP is based on three main concepts

Communication is constrained within *registered flows* only

*Global requirements are splited* across distributed components

Interaction between components is based on *contracts*

# DRP is based on three main concepts

Communication is constrained within *registered flows* only

*Global requirements are splited* across distributed components

Interaction between components is based on *contracts*

# Communication is constrained within *registered flows* only

Flow $i$ $F_i = ( n_i^s, n_i^d, T_i , J_i , \mathbf{D}_i )$

Source
Destination
Min. release interval
Jitter
End-to-end deadline

Release model *Sporatic* with jitter

# DRP is based on three main concepts

Communication is constrained within *registered flows* only

*Global requirements are splited* across distributed components

Interaction between components is based on *contracts*

# *Global requirements are splited*
# across distributed components

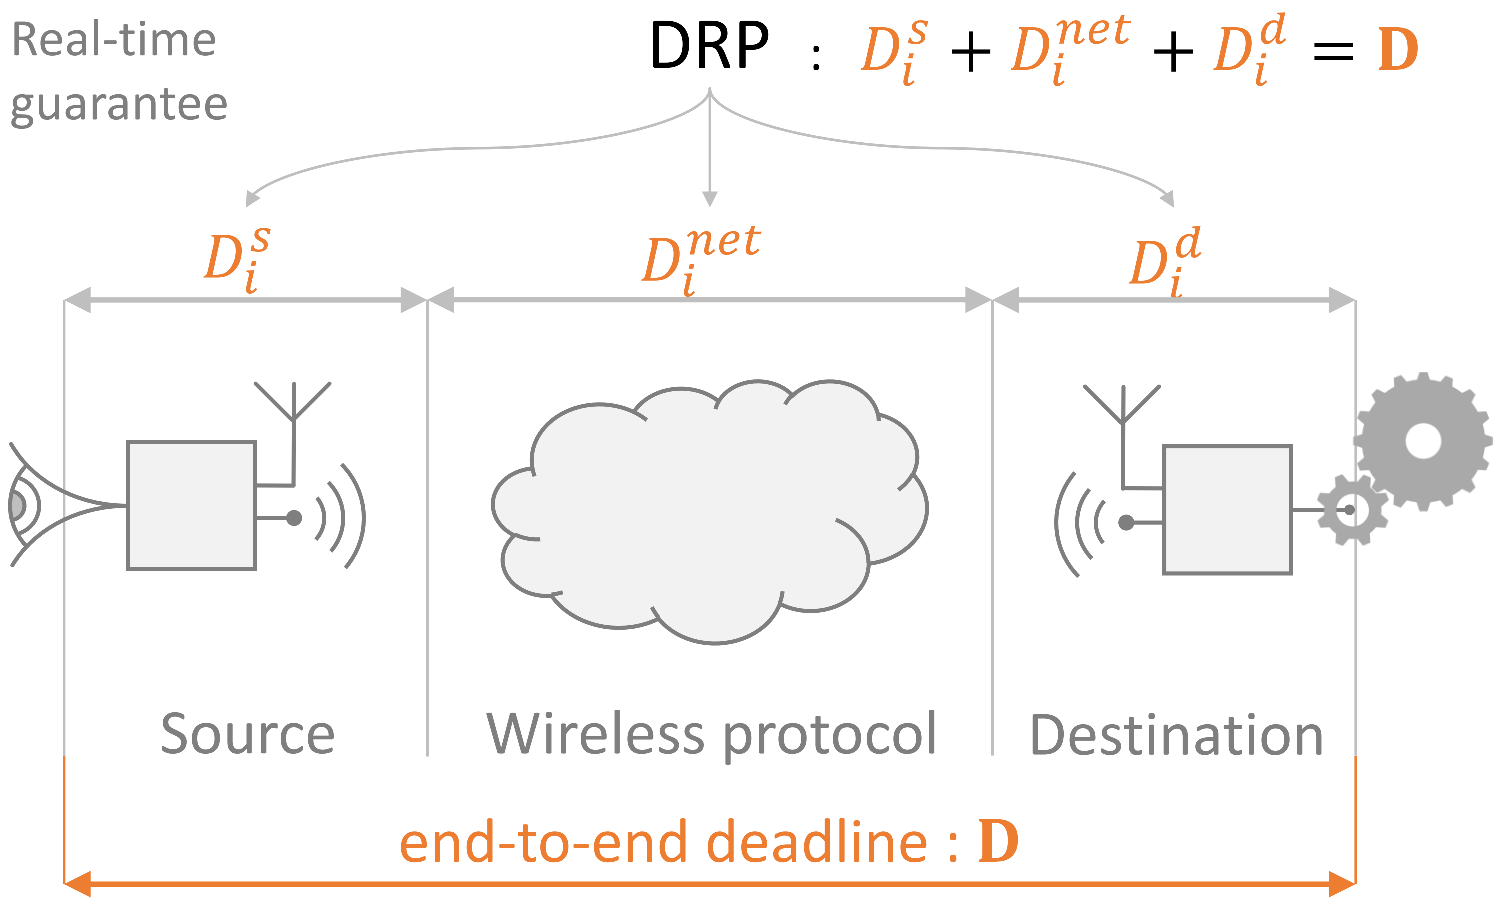

# DRP is based on three main concepts

Communication is constrained within *registered flows* only

*Global requirements are splited* across distributed components

Interaction between components is based on *contracts*

# Interaction between components is based on *contracts*


> Node ↔ Network

> Max resource demand
> Min service delivered

> Predictability    Performance

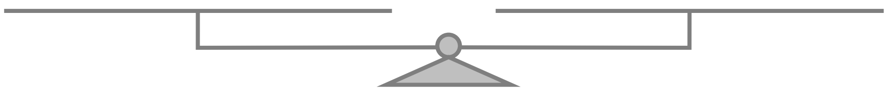

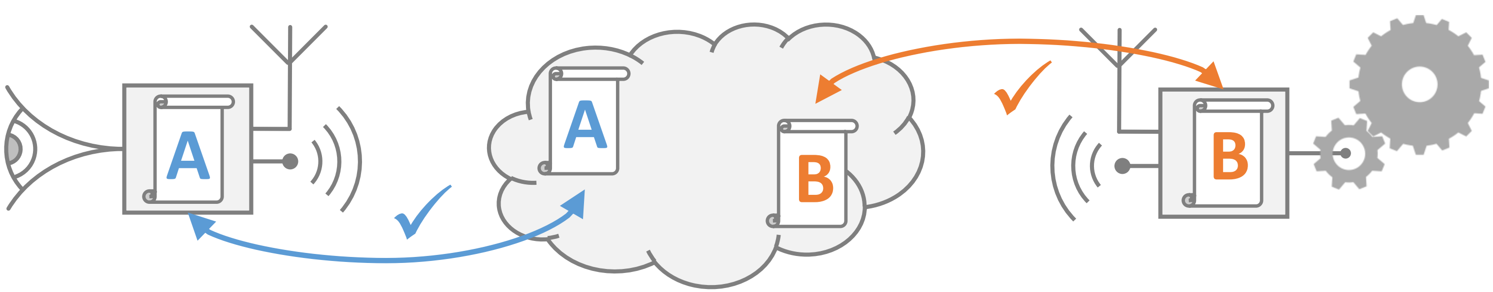

Example

# Scheduling of Communication Processors (CP)

| | |
|---|---|
| Real-time guarantees | *Participate in rounds*<br>> $T_{net}$ Blink schedule |
| | *Flush before rounds*<br>*Write after rounds*<br>> $T_{net}$ Blink schedule |
| Buffer management | *Flush Bolt regularly*<br>> $T_f^S$ Comm. Processor schedule |

Example

# Scheduling of Communication Processors (CP)

*How can we guarantee that all tasks are schedulable?*

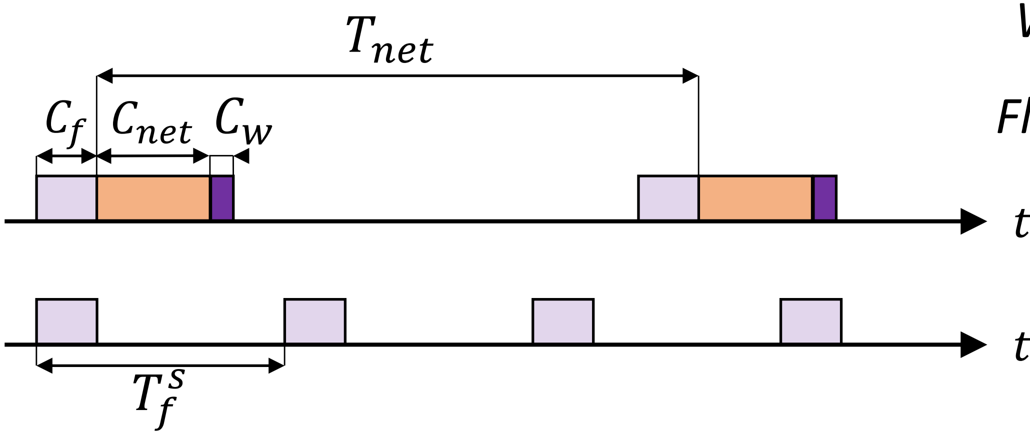


*Participate in rounds*

*Flush before rounds*

*Write after rounds*

*Flush Bolt regularly*

Example

# Scheduling of Communication Processors (CP)

*How can we guarantee that all tasks are schedulable?*

*Participate in rounds*

*Flush before rounds*

*Write after rounds*

*Flush Bolt regularly*

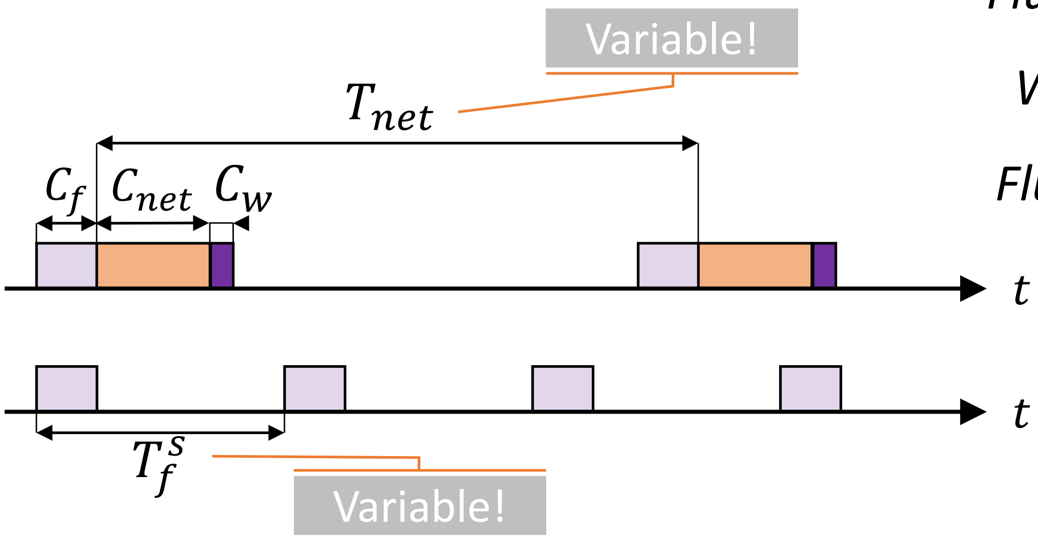

Example

# Scheduling of Communication Processors (CP)

*How to guarantee that all tasks are schedulable?*

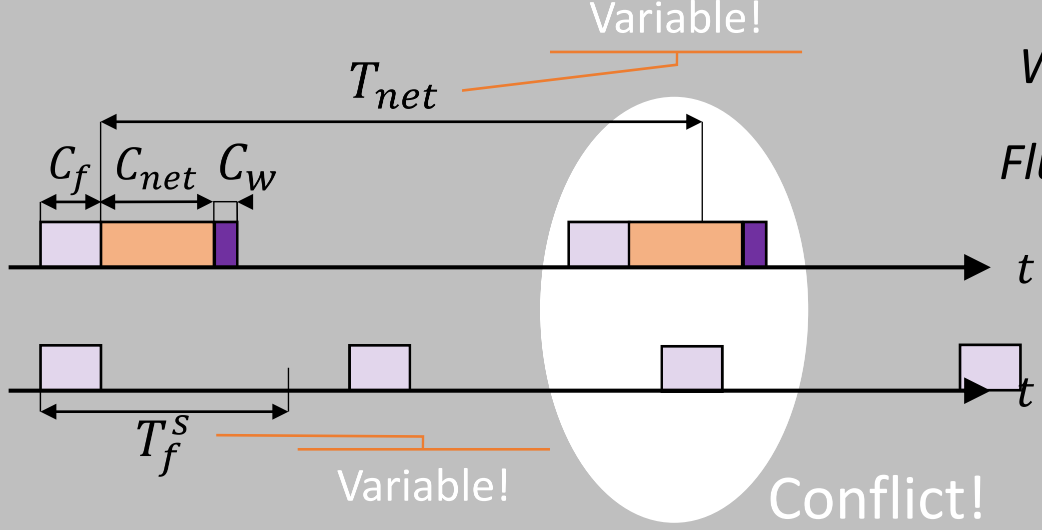


*Participate in rounds*

*Flush before rounds*

*Write after rounds*

*Flush Bolt regularly*

Example

# Scheduling of Communication Processors (CP)

*How to guarantee that all tasks are schedulable?*

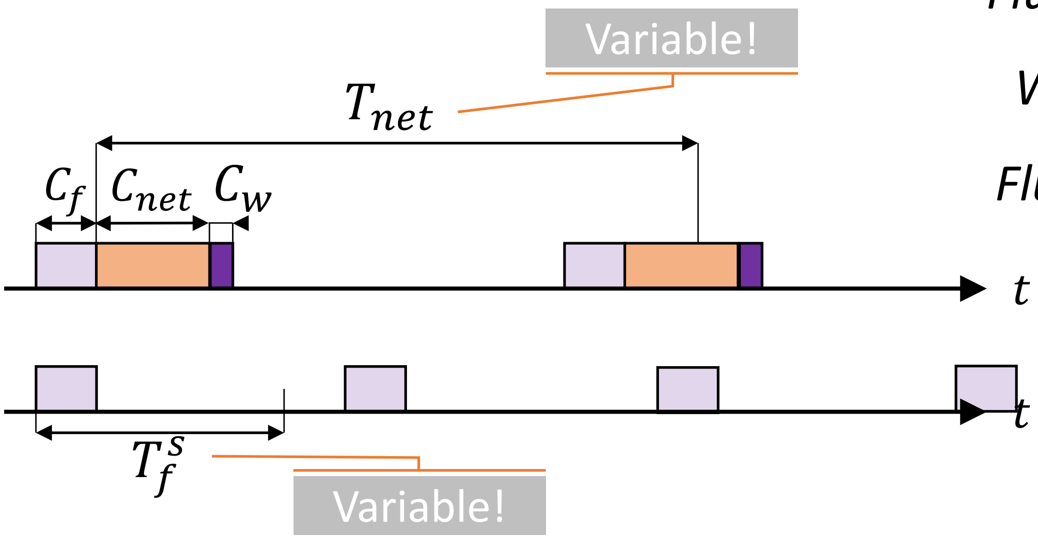


*Participate in rounds*

*Flush before rounds*

*Write after rounds*

*Flush Bolt regularly*

*Solution* TDMA approach

$$C_{CP} \triangleq C_f + C_{net} + C_w$$

Example

# Scheduling of Communication Processors (CP)

*How to guarantee that all tasks are schedulable?*

*Participate in rounds*

*Flush before rounds*

*Write after rounds*

*Flush Bolt regularly*

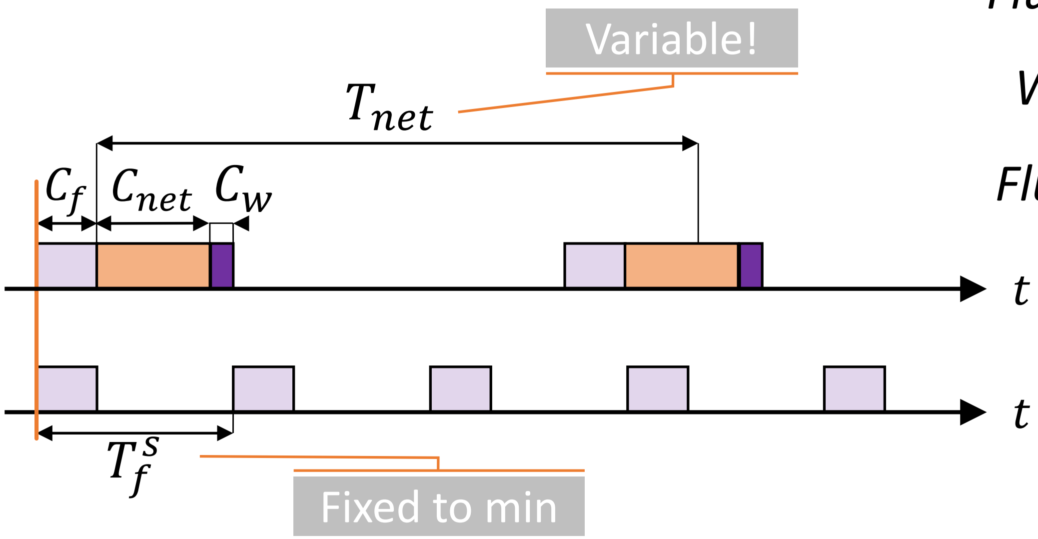


*Solution* TDMA approach

$$C_{CP} \triangleq C_f + C_{net} + C_w$$

$$> \quad T_f^S \triangleq C_{CP}$$

Example

# Scheduling of Communication Processors (CP)

*How to guarantee that all tasks are schedulable?*

*Participate in rounds*

*Flush before rounds*

*Write after rounds*

*Flush Bolt regularly*

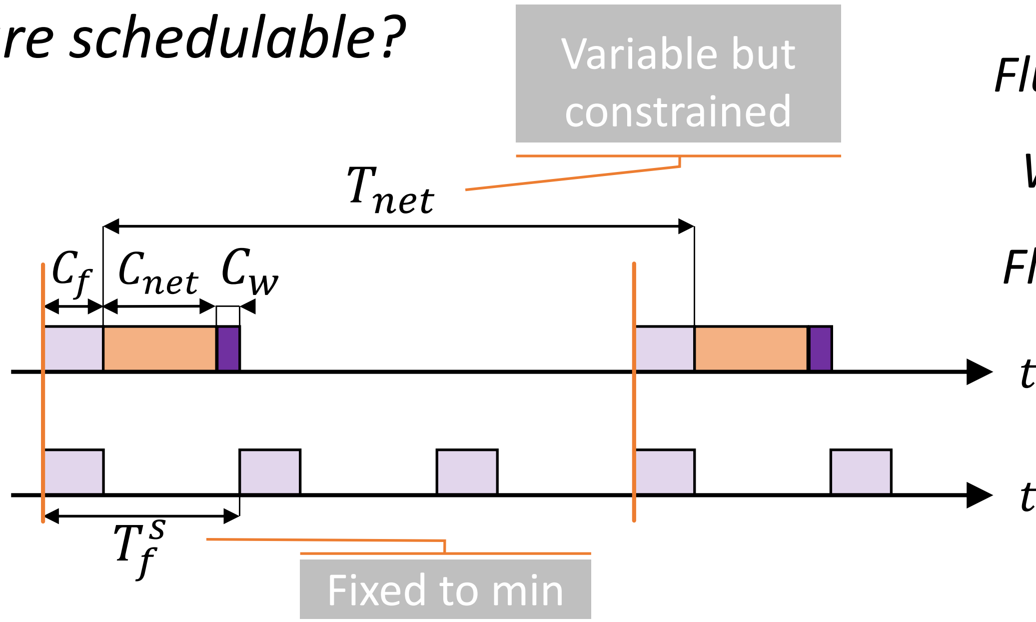


*Solution* TDMA approach

$$C_{CP} \triangleq C_f + C_{net} + C_w$$

> $T_f^S \triangleq C_{CP}$

> $T_{net} \triangleq k \cdot C_{CP}$

Predictability of Network
\+ Predictability of Devices
\+ DRP contracts
= *System predictability*

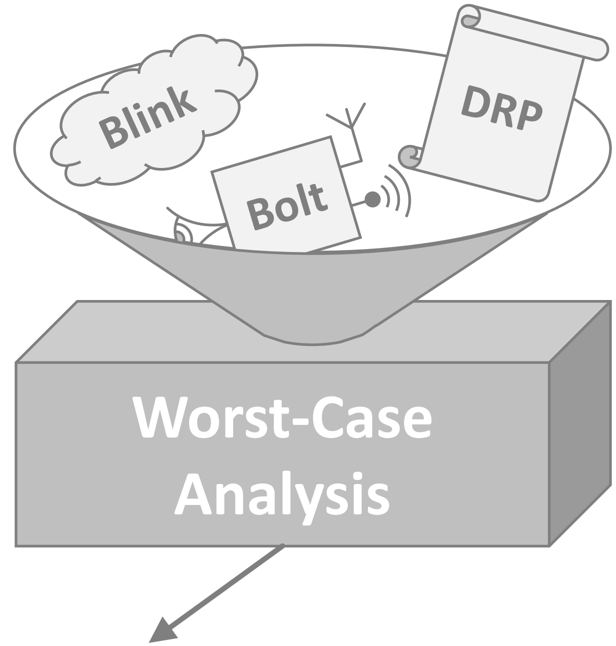


*Ultimately*

End-to-end latency | $= f($ Local parameters $)$
Buffer size |

*Can we set local parameters such that…* $\leq$ End-to-end deadline
Memory space

| | |
|---|---|
| Flush interval | $T_f^s$ |
| Flush interval | $T_f^d$ |
| Network deadline | $D_i^{net}$ |

If **YES** System predictability *by design*

# Design goals

| | Blink | Bolt | DRP |
|---|---|---|---|
| | Network | Device | System |
| Real-time guarantees | ✓ | ✓ | ? |
| Buffer management | na | ? | ? |
| Adaptability | ✓ | ✓ | ? |
| Efficiency | ✓ | ✓ | ? |
| Composability | na | ✓ | ? |

# Design goals

| | Blink | Bolt | DRP |
|---|---|---|---|
| *Predicatability* | Network | Device | System |
| Real-time guarantees | ✓ | ✓ | ✓ |
| Buffer management | na | ✓ | ✓ |
| Adaptability | ✓ | ✓ | ? |
| Efficiency | ✓ | ✓ | ? |
| Composability | na | ✓ → | ✓ |

# From there, *adaptability* is one (close) step away

*Can we set local parameters such that…*

**ADMISSION TESTS**

| | | |
|---|---|---|
| End-to-end latency | $\leq$ | End-to-end deadline |
| Buffer size | $\leq$ | Memory space |

# From there, *adaptability* is one (close) step away

Example

## Communication Processor

**ADMISSION TESTS**

$$S_{Bolt} \geq \sum_{\substack{F_i \in \mathcal{F}_{new},\\ n=n_i^s}} \left\lceil \frac{T_f^s + C_w + C_r + J_i}{T_i} \right\rceil \quad \text{and}$$

$$S_{CP} \geq \sum_{\substack{F_i \in \mathcal{F}_{new},\\ n=n_i^s}} 1 + \left\lceil \frac{D_i + \overline{J_i} + C_f}{T_i} \right\rceil + \sum_{\substack{F_i \in \mathcal{F}_{new},\\ n=n_i^d}} 1$$

Depends only on *local parameters!*

# Adaptability is achieved via a *distributed registration scheme*

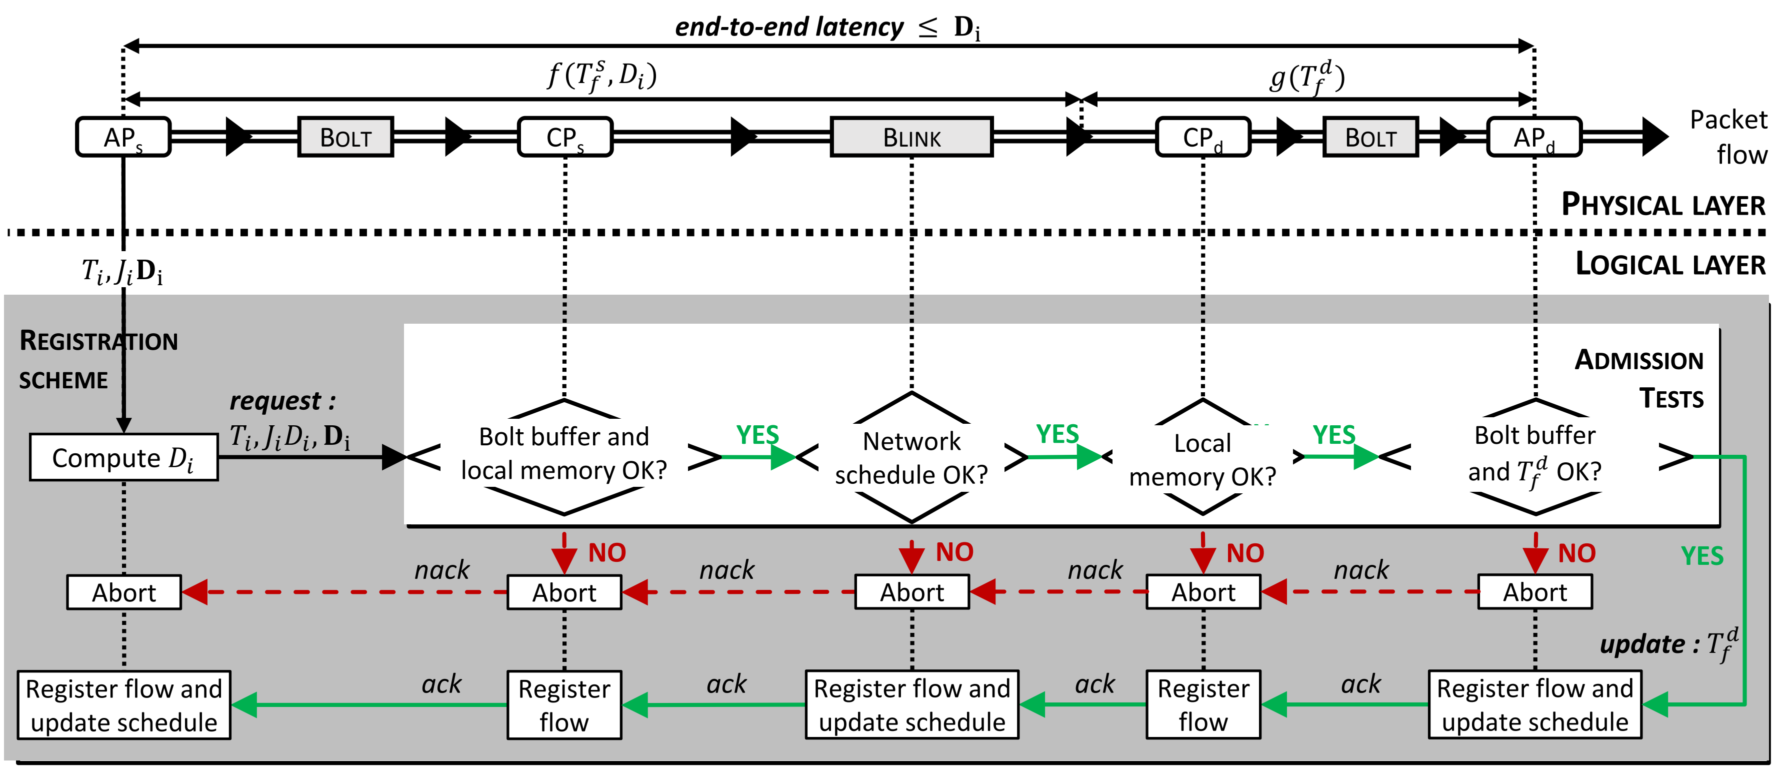

# Design goals

| | Blink | Bolt | DRP |
|---|---|---|---|
| | Network | Device | System |
| *Predicatability* | | | |
| Real-time guarantees | ✓ | ✓ | ✓ |
| Buffer management | na | ✓ | ✓ |
| *Adaptability* | ✓ | ✓ | ✓ |
| Efficiency | ✓ | ✓ | ? |
| Composability | na | ✓ | ✓ |

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T_{f,min}^{d}$$

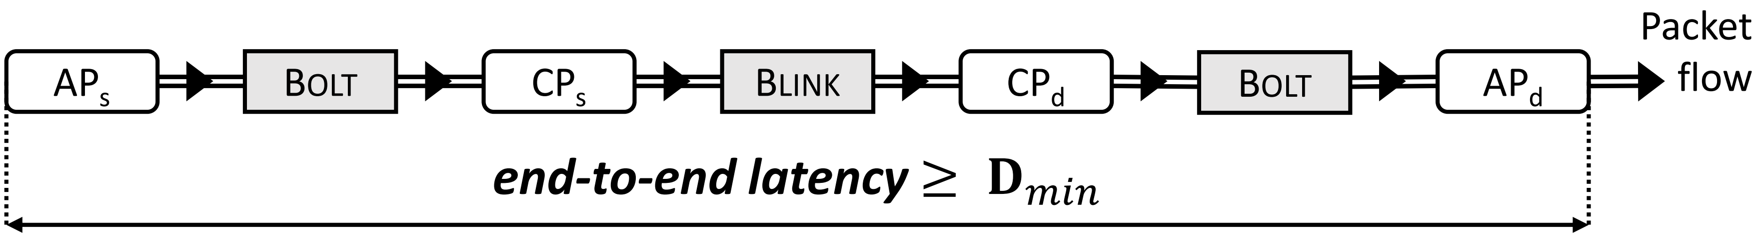

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T_{f,min}^{d}$$

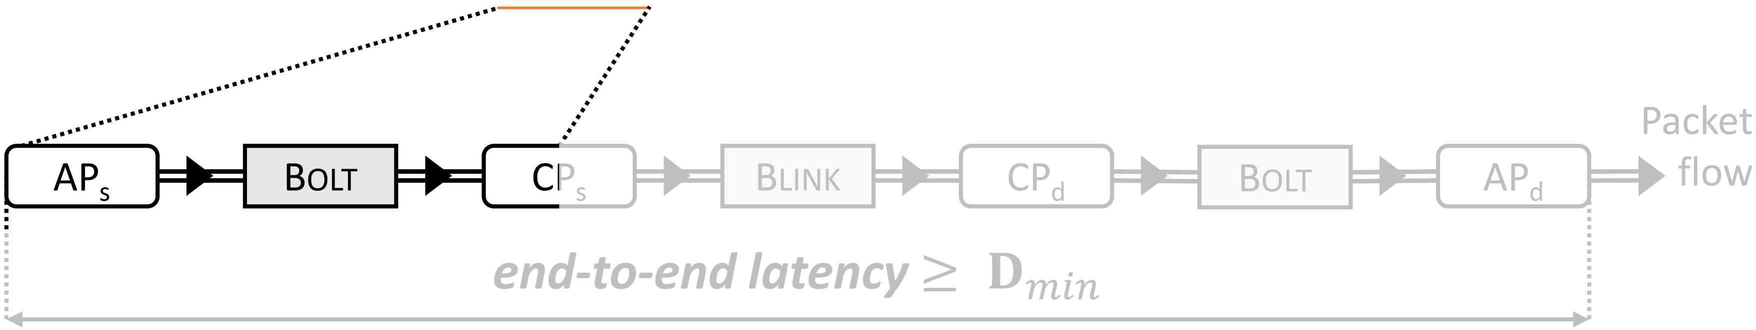


Source delay

- *Message release by the source*
- *Available for communication*

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T_{f,min}^d$$

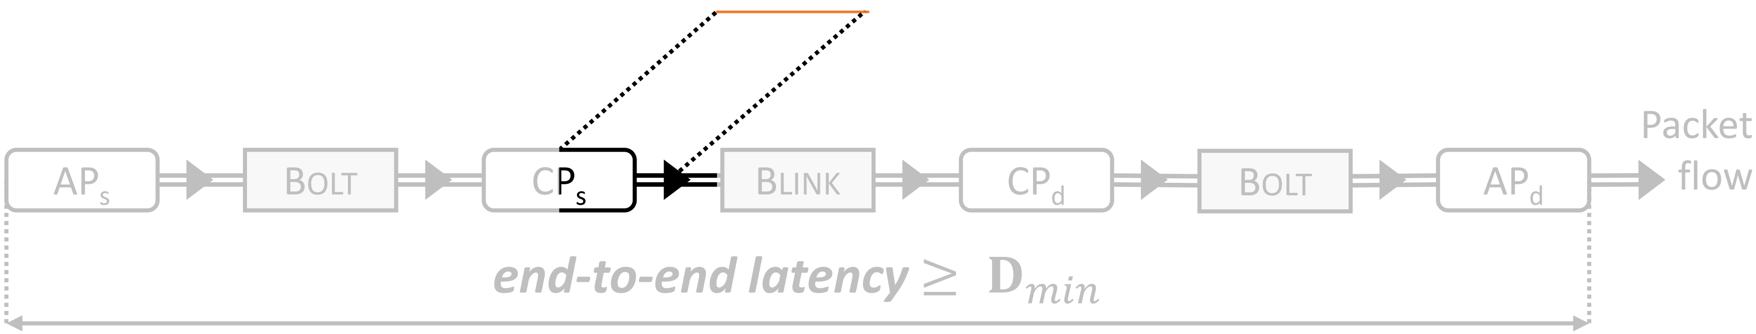


Network delay: Waiting time

> *Message available*

> *Message processed by the network*

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T^{d}_{f,min}$$

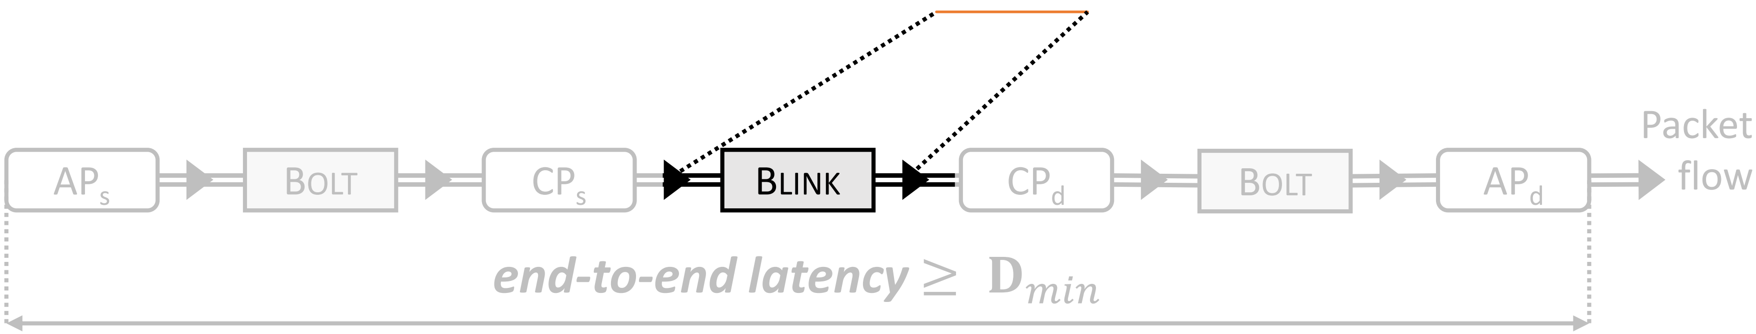


Network delay: Transmission

- *Message processed by the network*
- *Message transmitted to the destination node*

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T_{f,min}^{d}$$

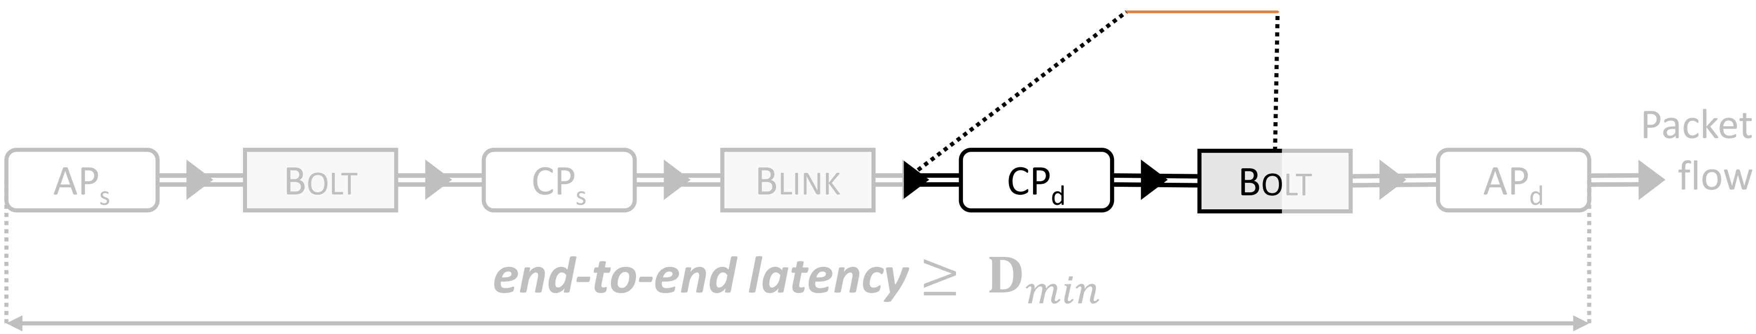


Destination delay: Bolt

- *Message transmitted to the destination node*
- *Message available in the Bolt queue*

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T_{f,min}^{d}$$

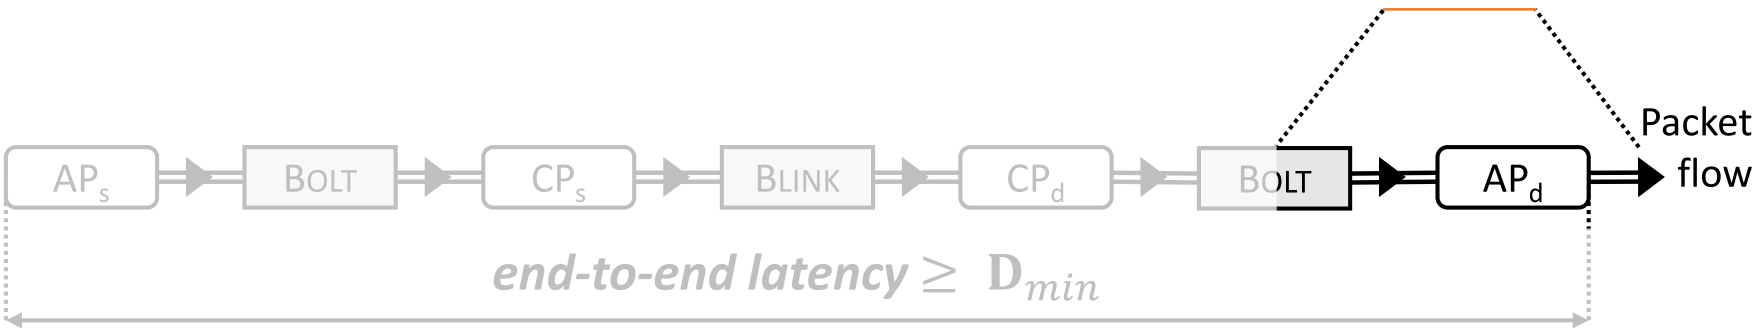


Destination delay: Application

- *Message available in the Bolt queue*
- *Message retrieved by the destination application*

# The detailed system analysis allows for *performance optimization*

Minimal admissible end-to-end deadline

$$\mathbf{D}_{min} = \delta_f^{const} + T_{min} + T_{min} + \delta_g^{const} + T_{f,min}^{d}$$

| | | | |
|---|---|---|---|
| Given | Packet size | 32 | Bytes |
| | $C_{net}$ | $1$ | $s$ |
| | $T_{f,min}^{d}$ | $0.1$ | $s$ |
| > | $\mathbf{D}_{min}$ | $3.46$ | $s$ |
| | Max data rate | $29.7$ | $Bps$ per flow |

# Design goals

| | Blink | Bolt | DRP |
|---|---|---|---|
| *Predicatability* | Network | Device | System |
| Real-time guarantees | ✓ | ✓ | ✓ |
| Buffer management | na | ✓ | ✓ |
| *Adaptability* | ✓ | ✓ | ✓ |
| Efficiency | ✓ | ✓ | ✓ |
| Composability | na | ✓ | ✓ |

Using *flooding primitives*

enables the design of both

*adaptive* AND *predictable*

Wireless Cyber-Physical Systems

# End-to-end Real-time Guarantees in Wireless Cyber-physical Systems

Romain Jacob
Marco Zimmerling
Pengcheng Huang
Jan Beutel
Lothar Thiele

RTSS 16 - IoT and Networking Session
December 1, 2016

ETH zürich
TECHNISCHE UNIVERSITÄT DRESDEN

# Simulation correlates closely with the analysis

Typical simulation trace result