LARN
Latency- and Resilience-Aware Networking

Latency- and Resilience-Aware Networking

SPP 1914: "Cyber-Physical Networking"

http://larn.systems

Andreas Schmidt, Pablo Gil Pereira, Thorsten Herfet
Telecommunications Lab
Saarland Informatics Campus - Saarbrücken

Stefan Reif, Timo Hönig, Wolfgang Schröder-Preikschat
Department of Computer Science 4 (Distributed Systems and Operating Systems)
Friedrich-Alexander-Universität Erlangen-Nürnberg

November 30th, 2018

Recapitulation: Goals

- Tailor PRRT to control applications
- Improve PRRT's performance
- OS for predictable communication
- Tailor network to control applications

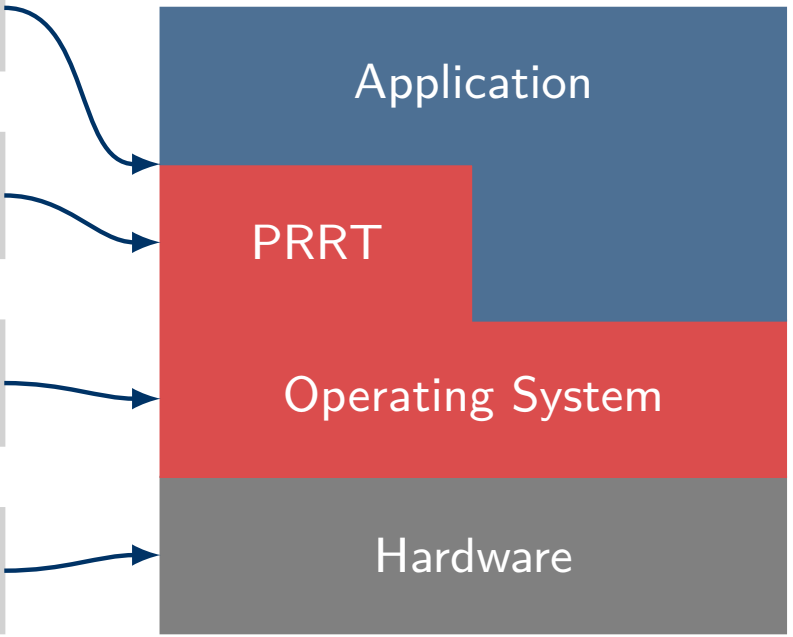

LARN

Software, Hardware & Algorithms

PRRT Predictable Reliable Real-time Transport protocol

- Improved Packet Loss Measurement
- Coding Configuration Optimization
- Cross-Layer Pacing + Rate Control
- Released stable version 0.3.1

LARN

Software, Hardware & Algorithms

PRRT — Predictable Reliable Real-time Transport protocol

- Improved Packet Loss Measurement
- Coding Configuration Optimization
- Cross-Layer Pacing + Rate Control
- Released stable version 0.3.1

X-LAP — Cross-Layer Timing Analysis

ΔELTA — System Modification Analysis

LARN

Software, Hardware & Algorithms

PRRT	Predictable Reliable Real-time Transport protocol
	- Improved Packet Loss Measurement
	- Coding Configuration Optimization
	- Cross-Layer Pacing + Rate Control
	- Released stable version 0.3.1
X-LAP	Cross-Layer Timing Analysis
ΔELTA	System Modification Analysis
RNA	Reliable Networking Atom
	- Drone
	- Quality of Control Metrics
	- Preparing release 0.9.0

LARN

LARN

Changelog

- v0.3.1 released in November
  - Congestion control inspired by BBR 1.x
  - Error control with reliable timing for ARQ
  - Improved network estimation capabilities
- v0.2.0 released in May
  - API enhancements.
  - Significant extension of the Wiki, README and other pieces of documentation.

Changelog

- v0.3.1 released in November
  - Congestion control inspired by BBR 1.x
  - Error control with reliable timing for ARQ
  - Improved network estimation capabilities
- v0.2.0 released in May
  - API enhancements.
  - Significant extension of the Wiki, README and other pieces of documentation.

Components under Development

- Rate Control & Pacing:
  Reduced latencies and jitter.
- Improved loss estimation
- Performance optimizations

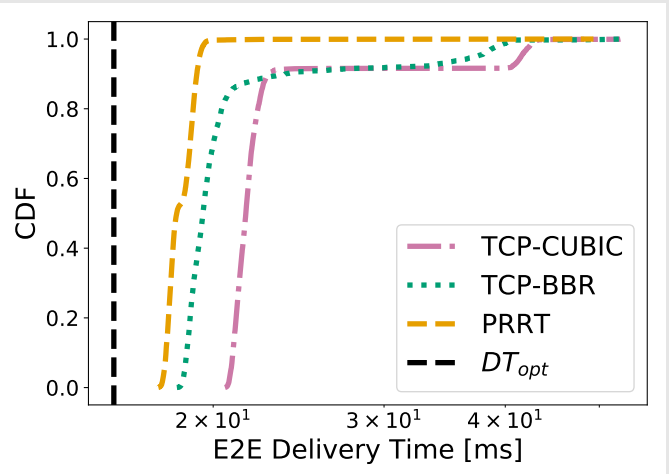

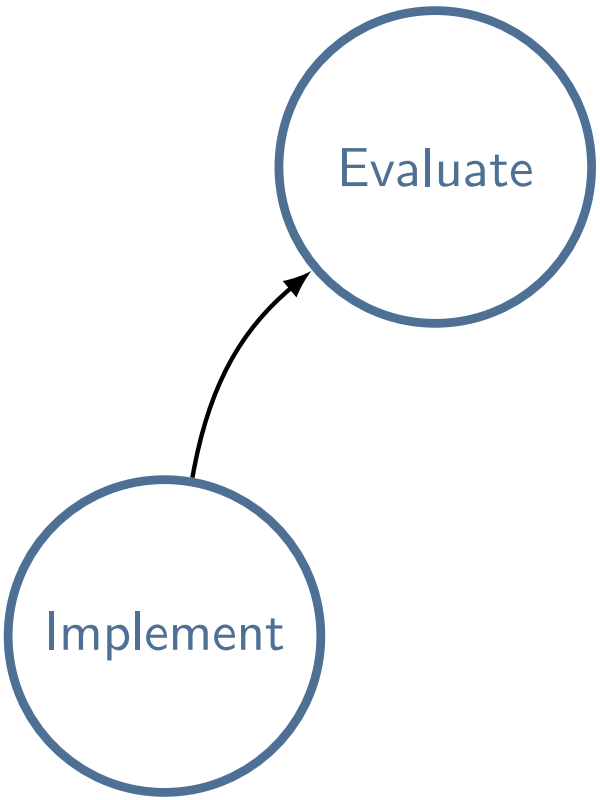

Approach:

- Measurement-based protocol evaluation

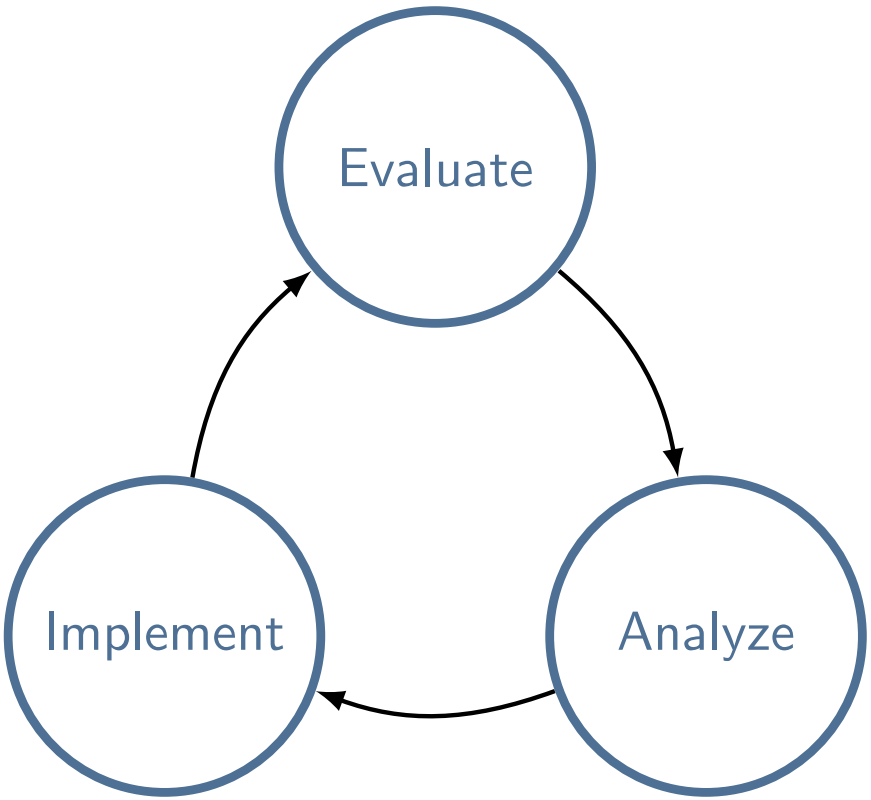

Approach:

- Measurement-based protocol evaluation
- Derive meaningful information
- Support protocol design

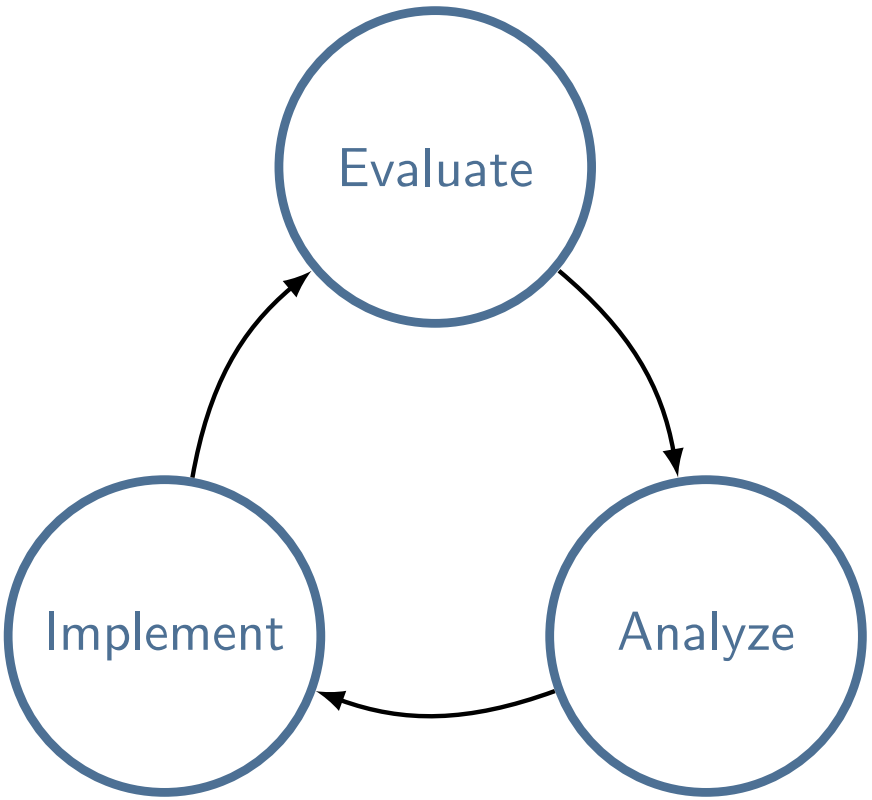

Approach:

- Measurement-based protocol evaluation
- Derive meaningful information
- Support protocol design

Challenges:

- ☐ Identify *critical code parts*
  - → Latency, Jitter, Energy
- ☐ Analyze code modifications
- ☐ Automate feedback

LARN

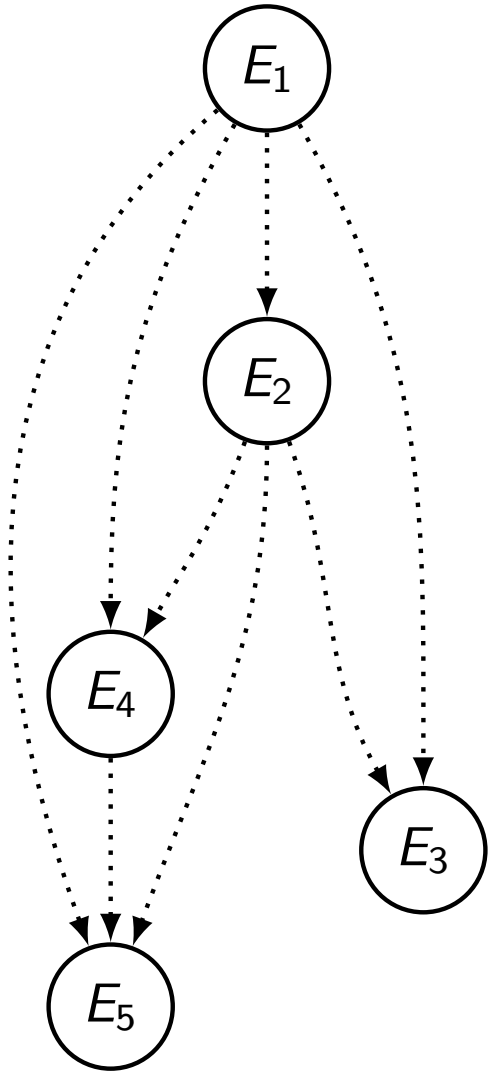

Goal:

- Reconstruct protocol information from timestamps

Control flow reconstruction:

1. Compute *happens-before* relation of events

LARN

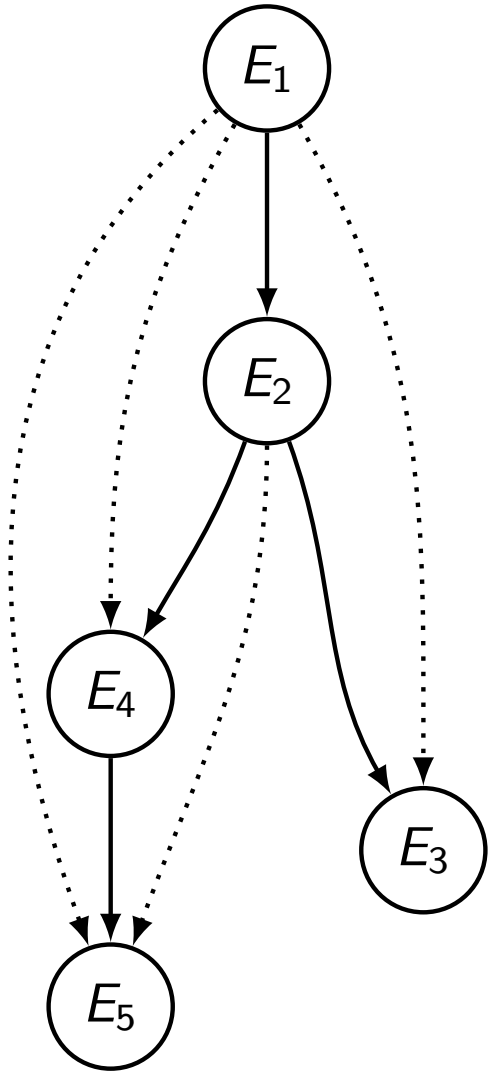

Goal:

- Reconstruct protocol information from timestamps

Control flow reconstruction:

1. Compute *happens-before* relation of events
2. Compute *happens-directly-before* relation of events

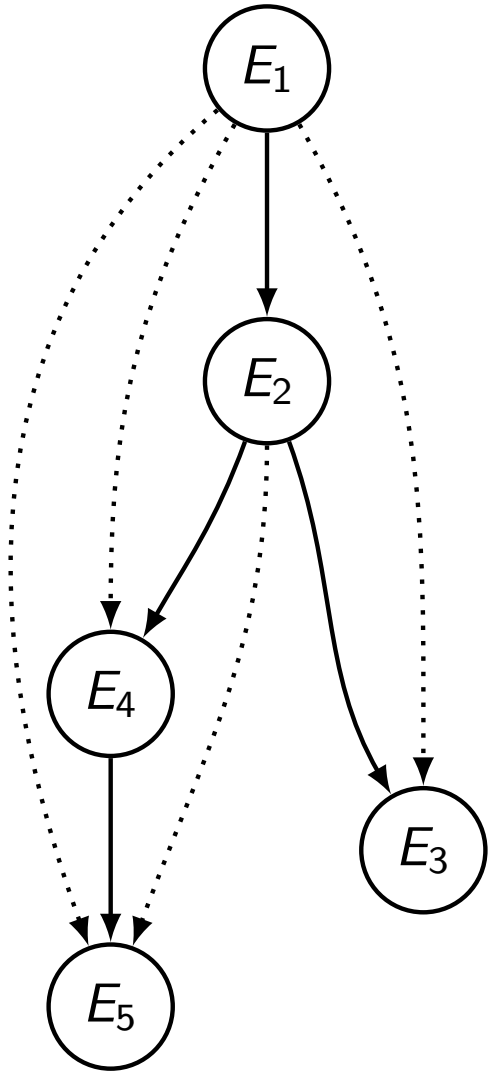

Goal:

- Reconstruct protocol information from timestamps

Control flow reconstruction:

1. Compute *happens-before* relation of events
2. Compute *happens-directly-before* relation of events

Edges represent:

- Code segments $\langle E_i, E_j \rangle$
- Local communication
- False-positives

LARN

Goals:

- Filter out false-positive segments
- Identify code segments worth optimizing

Detect *relevant* code segments:

1. Compute the correlation between the segment latency and the E2E latency
2. Ignore segments with negative correlation

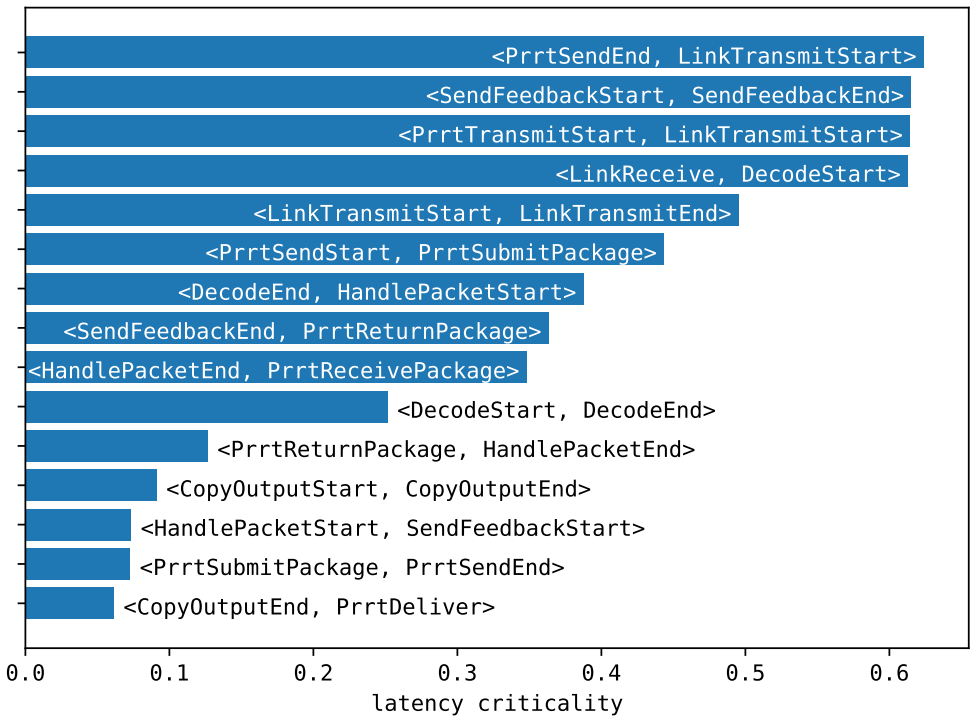

Goals:

- Filter out false-positive segments
- Identify code segments worth optimizing

Detect *relevant* code segments:

1. Compute the correlation between the segment latency and the E2E latency
2. Ignore segments with negative correlation

Goals:

- Evaluate the influence of code changes
- Interferences are complex

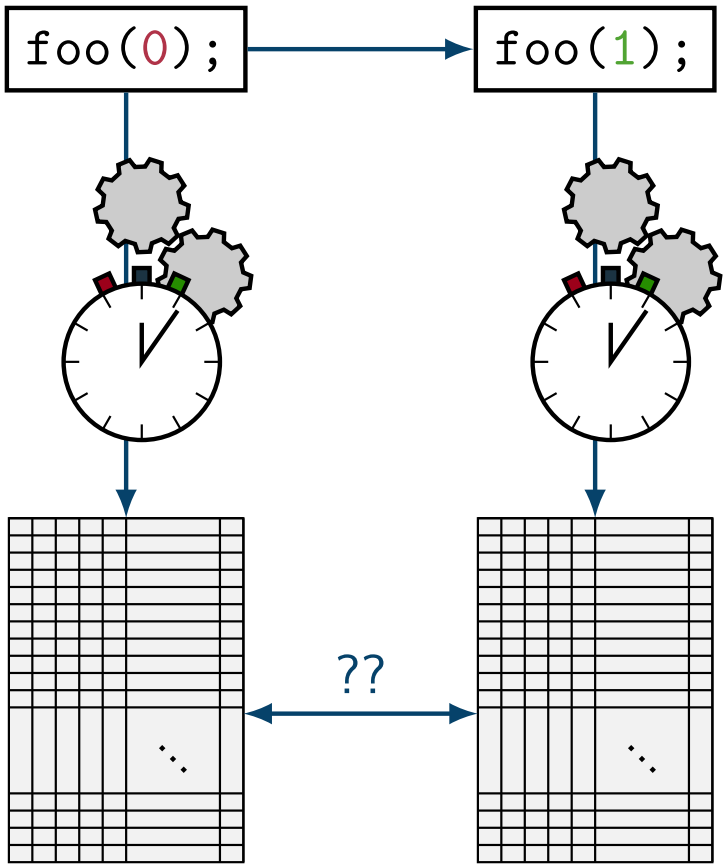

Goals:

- Evaluate the influence of code changes
- Interferences are complex

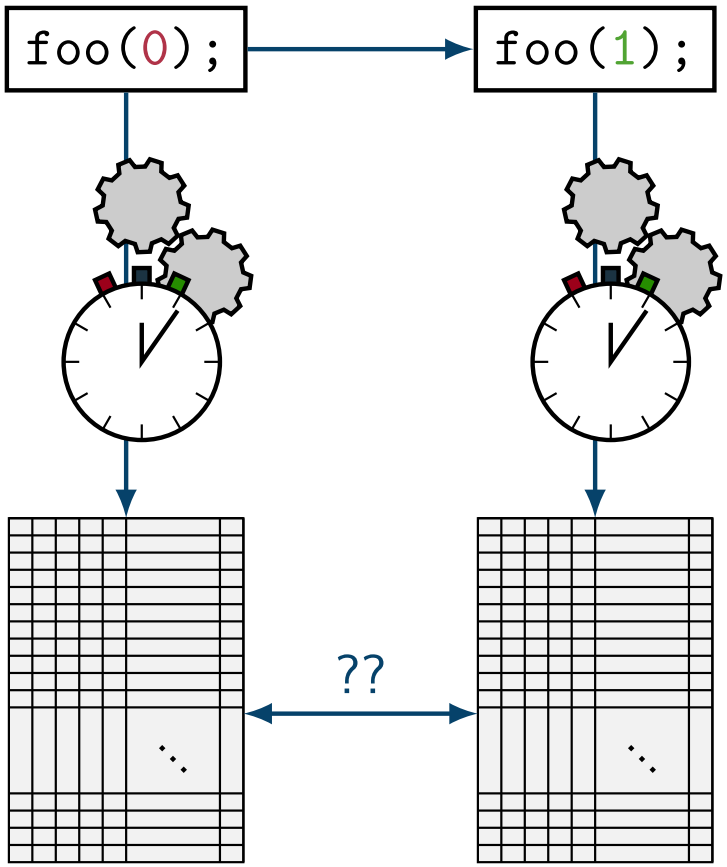

Goals:

- Evaluate the influence of code changes
- Interferences are complex

Modification impact analysis:

1. Evaluate the original and the modified versions of the protocol
2. Apply the *k-sample Anderson-Darling-Test* on the traces
3. For each code segment, test decides: “similar” or “different”

LARN

```
<PrrtSendPacketStart,PrrtSendPacketEnd> has not changed
<CopyOutputStart,CopyOutputEnd>         has not changed
<SendFeedbackStart,SendFeedbackEnd>     has not changed
...

<PrrtSubmitPackage,PrrtSendEnd> has changed:
  2.85±0.91 µs → 4.94±0.58 µs

<LinkReceive,DecodeStart> has changed:
  2.36±3.80 µs → 3.56±4.88 µs
...
```

LARN

Lack of Data

- Hardware complexity complicates interrupt latency predictability
- Unknown overhead for software interrupt handling mechanisms

Lack of Understanding for Delay Sources

- Unknown reasons for delays
- Reasons for delays vary among interrupt handling mechanisms

Lack of Tool Support

- Tool for precise and flexible overhead measurements
- Support of different hardware architectures

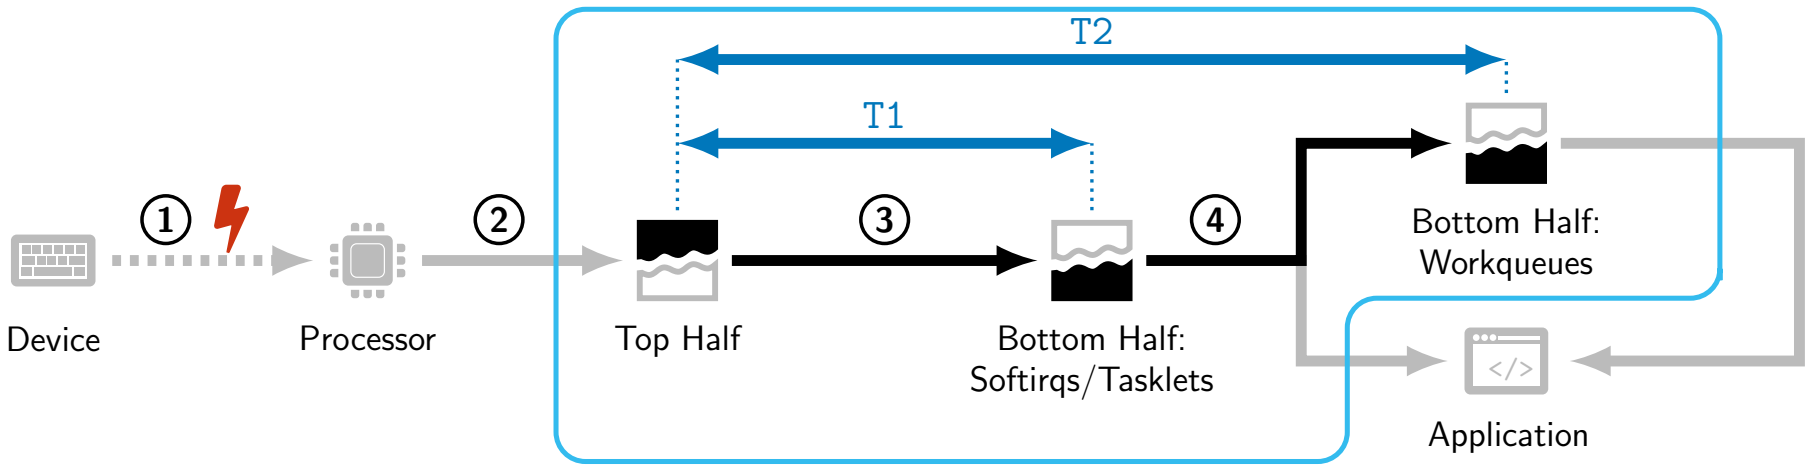

Linux Interrupt Handling

- Top half
- Bottom half
  - → Softirq, Tasklet, Workqueue
  - → Variation in **flexibility**

INTSPECT Analysis

- Latency and Jitter analysis
- Software overhead
  - → Latency `T1` and `T2`
  - → Variation in **latency** & **jitter**

LARN

Kernel Module

- Support all bottom half mechanisms
- Minimal run-time interference
- Portable

User-Space Component

- Analysis & Control framework
- Plugin interface
  - → load generators, hardware states

Hardware Support

- x86
- ARM → Atmel SAMA5D3 Xplained

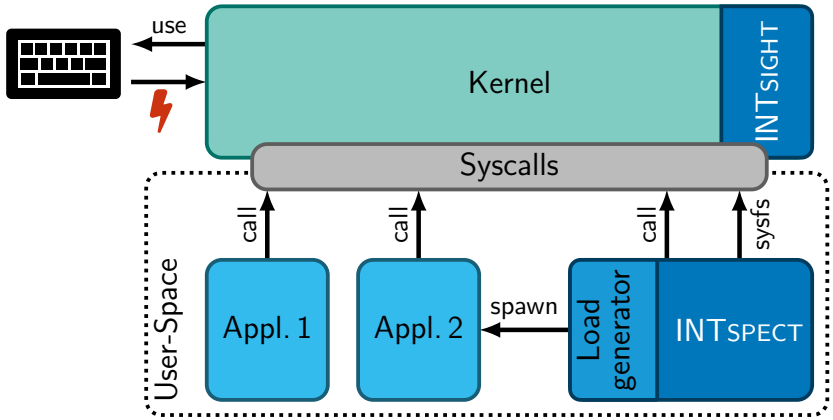

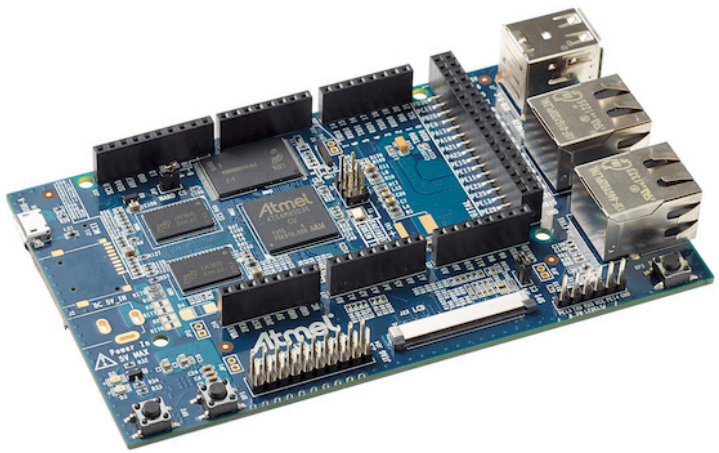

Comparison of Bottom Handler Mechanisms

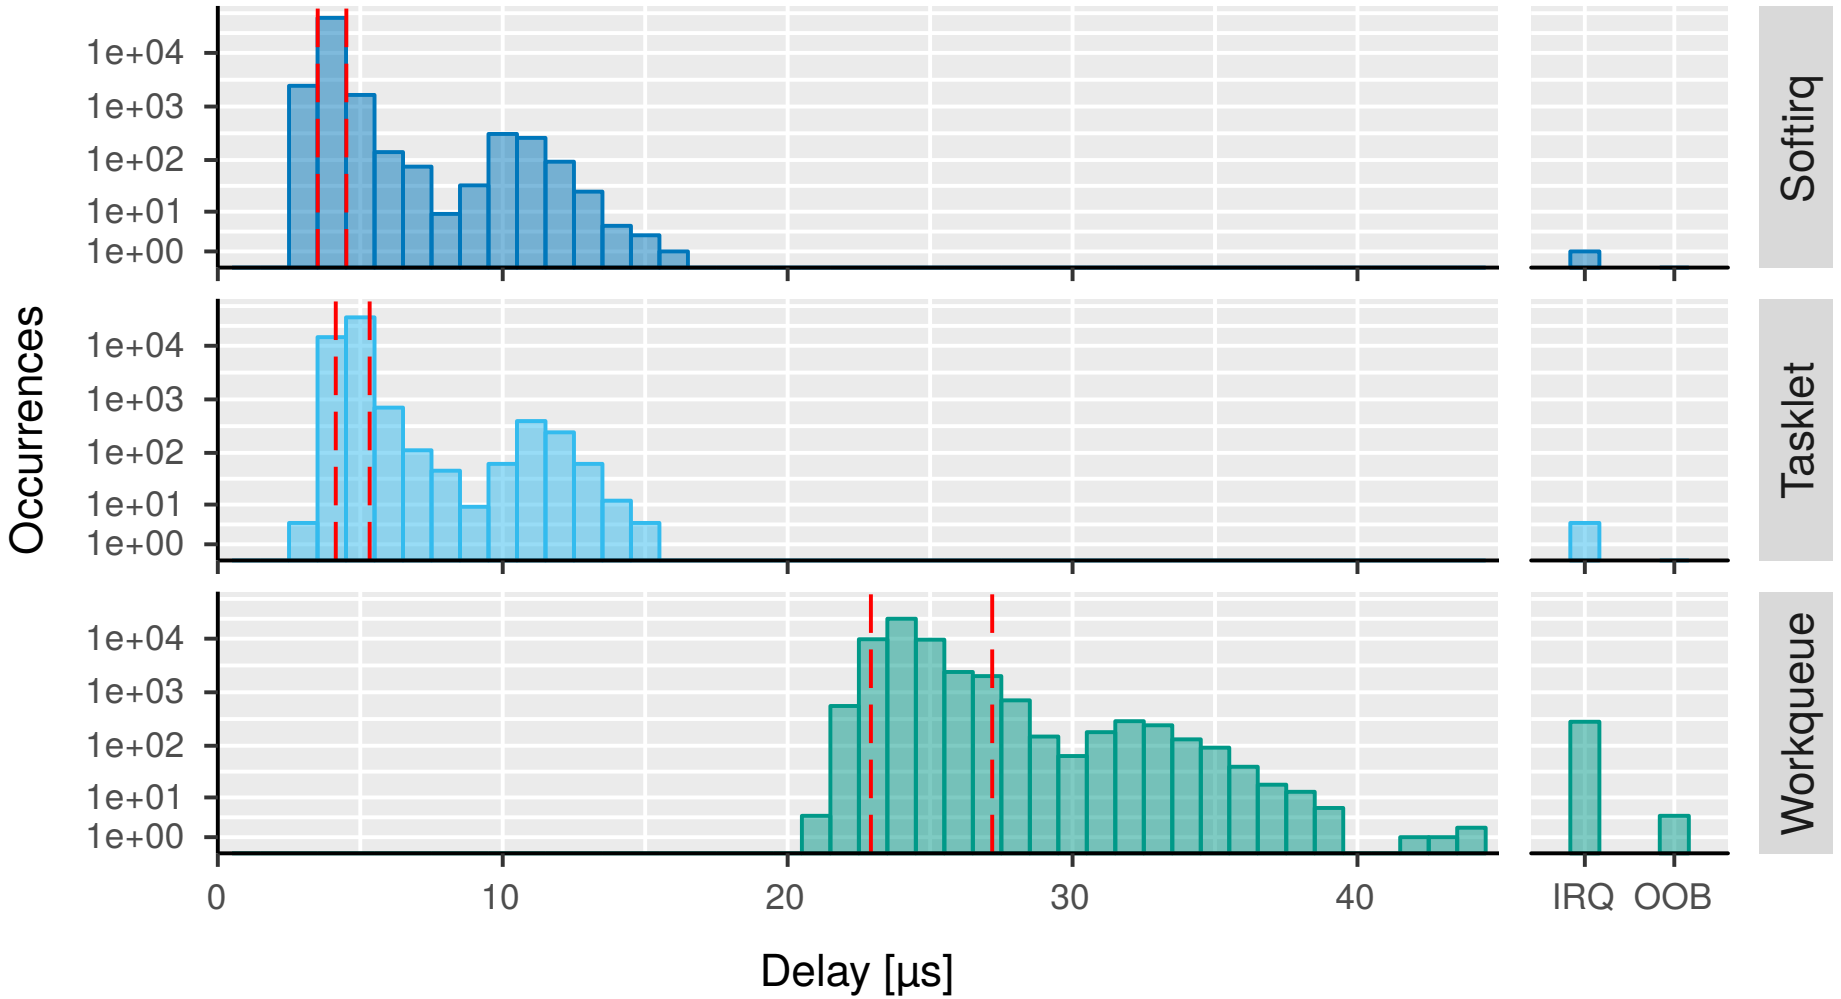

Sources of Interference

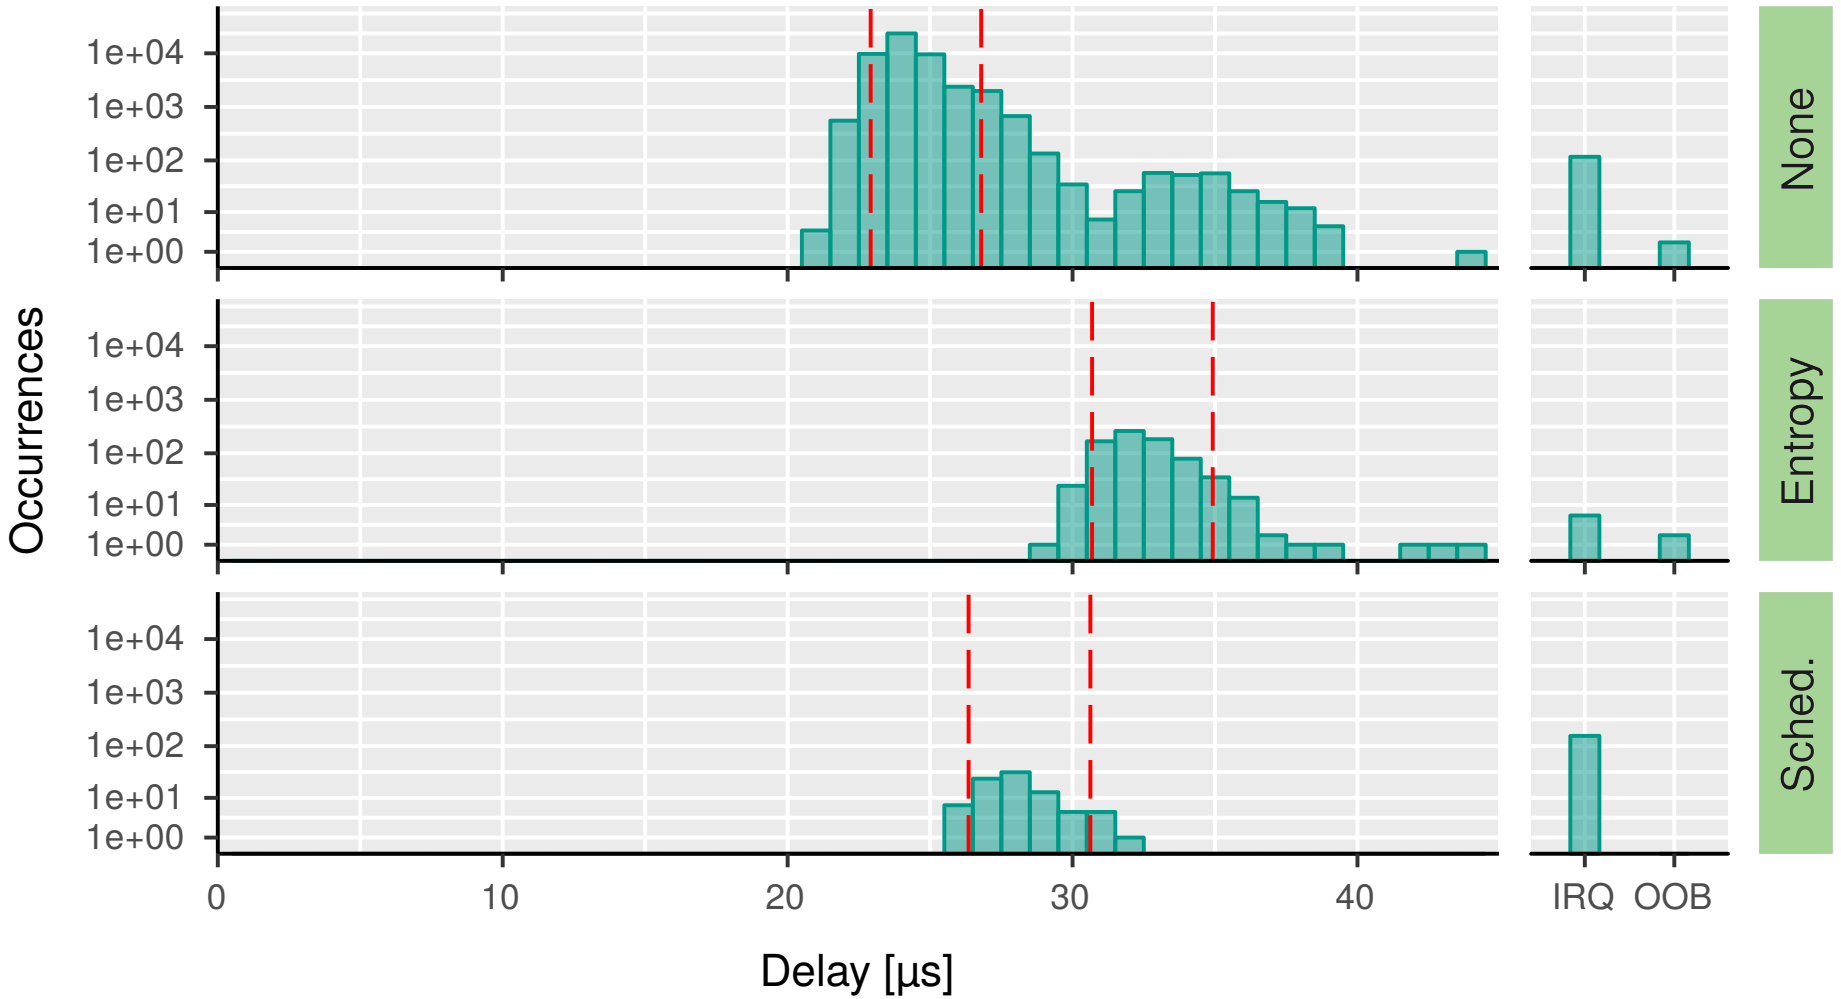

LARN

Accepted Publications

- Reif, Stefan; Schmidt, Andreas; Hönig, Timo; Herfet, Thorsten; Schröder-Preikschat, Wolfgang: **"Δelta: Differential Energy-Efficiency, Latency, and Timing Analysis for Real-Time Networks"**, International Workshop on Real-Time Networks (RTN), 2018
- Herzog, Benedict; Gerhorst, Luis; Heinloth, Bernhard; Reif, Stefan; Hönig, Timo; Schröder-Preikschat, Wolfgang: **INTSPECT: Interrupt Latencies in the Linux Kernel**, Brazilian Symposium on Computing Systems Engineering (SBESC), 2018.
- Reif, Stefan; Gerhorst, Luis; Bender, Kilian; Hönig, Timo: **"Towards Low-Jitter and Energy-Efficient Data Processing in Cyber-Physical Information Systems"**, Hawaii International Conference on System Sciences (HICSS), 2019

LARN

Accepted Publications

- Reif, Stefan; Schmidt, Andreas; Hönig, Timo; Herfet, Thorsten; Schröder-Preikschat, Wolfgang: **"Δelta: Differential Energy-Efficiency, Latency, and Timing Analysis for Real-Time Networks"**, International Workshop on Real-Time Networks (RTN), 2018
- Herzog, Benedict; Gerhorst, Luis; Heinloth, Bernhard; Reif, Stefan; Hönig, Timo; Schröder-Preikschat, Wolfgang: **INTSPECT: Interrupt Latencies in the Linux Kernel**, Brazilian Symposium on Computing Systems Engineering (SBESC), 2018.
- Reif, Stefan; Gerhorst, Luis; Bender, Kilian; Hönig, Timo: **"Towards Low-Jitter and Energy-Efficient Data Processing in Cyber-Physical Information Systems"**, Hawaii International Conference on System Sciences (HICSS), 2019

Publications Under Review

- Schmidt, Andreas; Reif, Stefan; Gil Pereira, Pablo; Hönig, Timo; Herfet, Thorsten; Schröder-Preikschat, Wolfgang: **"Blinded"**, Blinded, 2019.

LARN

Accepted Publications

- Reif, Stefan; Schmidt, Andreas; Hönig, Timo; Herfet, Thorsten; Schröder-Preikschat, Wolfgang: **"Δelta: Differential Energy-Efficiency, Latency, and Timing Analysis for Real-Time Networks"**, International Workshop on Real-Time Networks (RTN), 2018
- Herzog, Benedict; Gerhorst, Luis; Heinloth, Bernhard; Reif, Stefan; Hönig, Timo; Schröder-Preikschat, Wolfgang: **INTSPECT: Interrupt Latencies in the Linux Kernel**, Brazilian Symposium on Computing Systems Engineering (SBESC), 2018.
- Reif, Stefan; Gerhorst, Luis; Bender, Kilian; Hönig, Timo: **"Towards Low-Jitter and Energy-Efficient Data Processing in Cyber-Physical Information Systems"**, Hawaii International Conference on System Sciences (HICSS), 2019

Publications Under Review

- Schmidt, Andreas; Reif, Stefan; Gil Pereira, Pablo; Hönig, Timo; Herfet, Thorsten; Schröder-Preikschat, Wolfgang: **"Blinded"**, Blinded, 2019.

Publications in the Writing

- Joint publication based on BarCamp IV results.

LARN

Network

- ✓ PRRT – Predictable protocol
- ✓ Network pacing

Operating System

- ✓ X-Lap, ΔELTA, INTspect
- ✓ Wait-free synchronisation

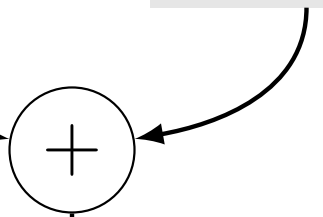

Reliable Network Atom (RNA)
(Embedded Edge-Server, Car, Drone, ...)

LARN

Network

- ✓ PRRT – Predictable protocol
- ✓ Network pacing

Operating System

- ✓ X-Lap, Δelta, INTspect
- ✓ Wait-free synchronisation

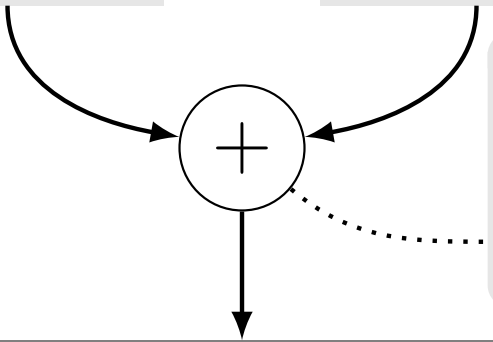

- Integration
- Optimisation
- Cross-layer pacing

Reliable Network Atom (RNA)

(Embedded Edge-Server, Car, Drone, ...)

LARN

Network

- ✓ PRRT – Predictable protocol
- ✓ Network pacing

Operating System

- ✓ X-Lap, ΔELTA, INTspect
- ✓ Wait-free synchronisation

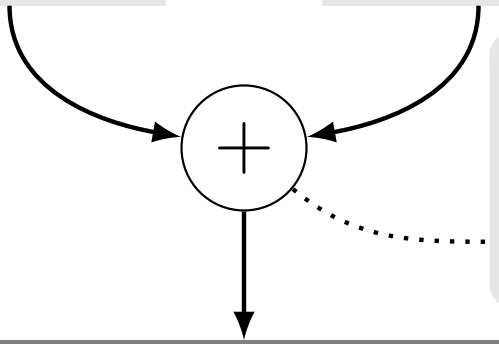

- Integration
- Optimisation
- Cross-layer pacing

Reliable Network Atom (RNA)

(Embedded Edge-Server, Car, Drone, ...)

Thank you for your attention. Questions?