Jessica Müller

PhD Candidate @ PZAI

Hochschule Darmstadt

Bosch Rexroth AG

h_da
HOCHSCHULE DARMSTADT
UNIVERSITY OF APPLIED SCIENCES

Promotionszentrum
Angewandte Informatik
HAW Hessen

Dissertation title: "*Possibilities for Integrating a Virtual Controller into a Production Environment with High Requirements for Real-Time and Functional Safety*"

Establishing Trust in Multi-Virtual Machine-Environments: VM Status Legitimization for Industry 4.0
DOI: 10.1109/DCOSS-IoT65416.2025.00067

Virtualization in Industrial Production - A Survey Focusing on Virtual and Virtualized Industrial Controls
DOI: 10.1109/ETFA61755.2024.10710668

Comparison of Implementation Options for Control Systems in the Energy Sector: Safety Requirements, Efficiency, and Effectiveness at Centralized and Decentralized Levels [German]
DOI: 10.1365/s40702-024-01091-2

Extended kernel density estimation for anomaly detection in streaming data
DOI: 10.1016/j.procir.2022.09.065

Load Profile Cycle Recognition for Industrial DC Microgrids with Energy Storage Systems
DOI: 10.1109/isie45063.2020.9152432

rexroth
A Bosch Company

An Interference-Aware vCPU Scheduling Framework for Paravirtualized Real-Time Industrial Control Systems

Jessica Müller
Bosch Rexroth AG
Lohr am Main, Germany
0009-0007-6478-7307

Steven Dietrich
Bosch Rexroth AG
Lohr am Main, Germany
Steven.Dietrich@boschrexroth.de

Michael Massoth
Hochschule Darmstadt
Darmstadt, Germany
michael.massoth@h-da.de

Multi-VM-System

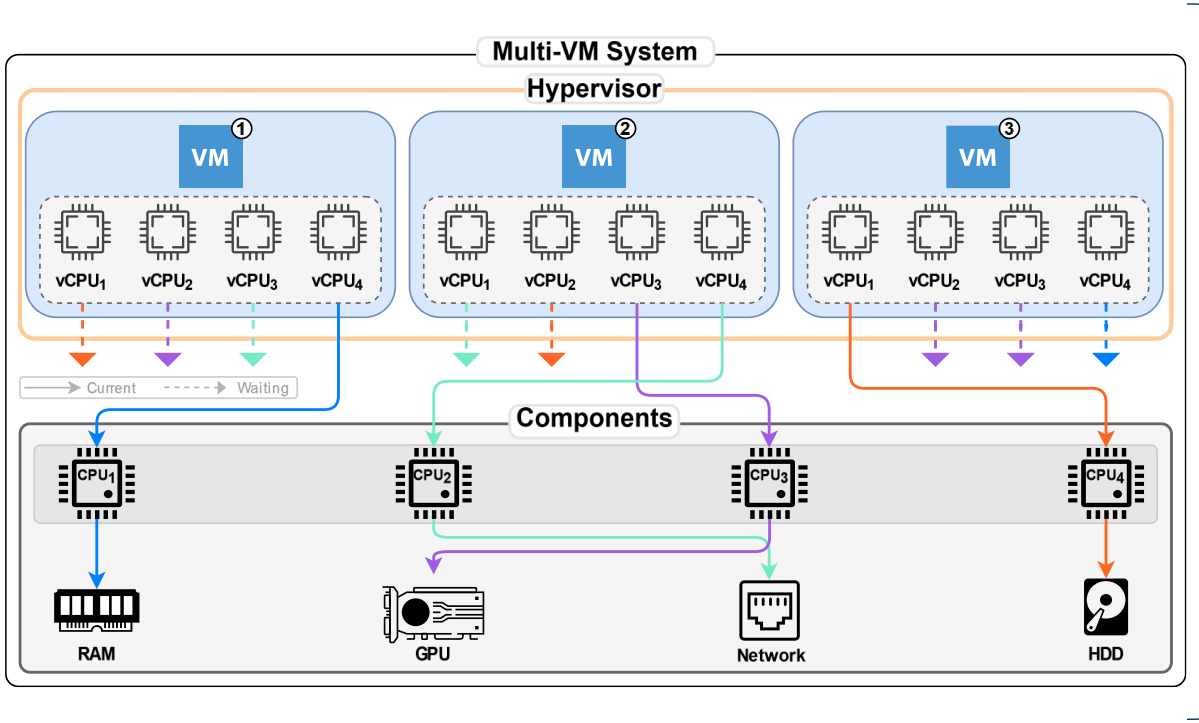

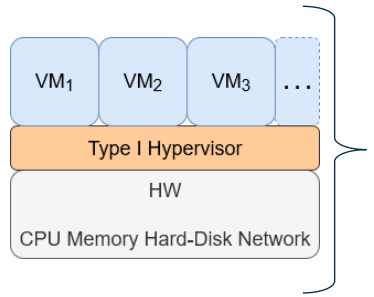

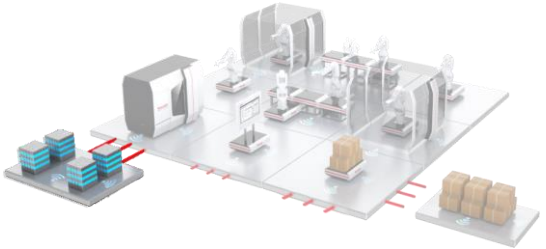

Picture: https://www.boschrexroth.com/en/hu/industries/industry-4-0-solutions/techtalk/

rexroth
A Bosch Company

Shortcomings and Interference Issues

Realtime + isolation ≠ interference-free I/O

Core binding protects compute, not the I/O path

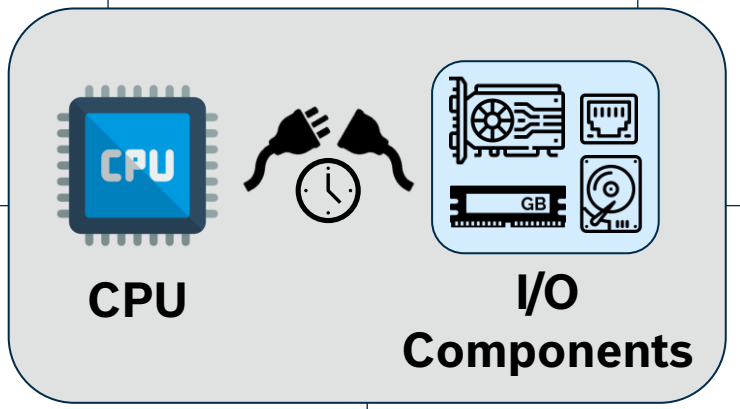

Load spikes
→ context / interrupt thrash

No joint CPU & I/O scheduling
→ lost timing

Shared/interrupt-driven I/O
→ hidden latency

jitter & missed hard real-time cycles

rexroth
A Bosch Company

Interference-Aware Scheduling via Paravirtualized VM Profiling

1 Classification

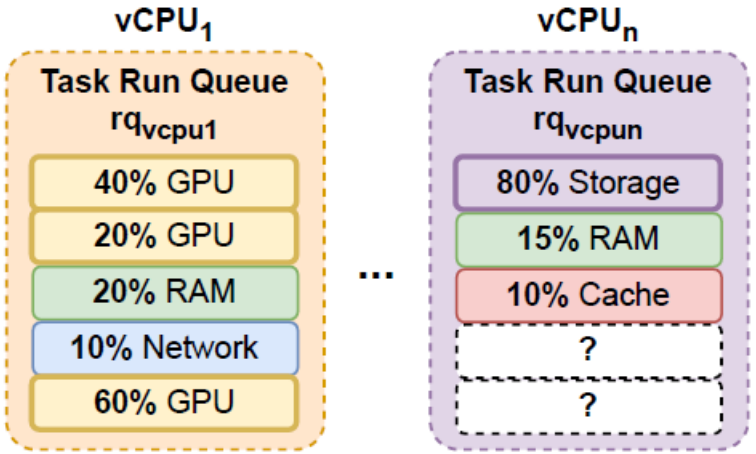

X% = I/O Utilization, Color = I/O Category

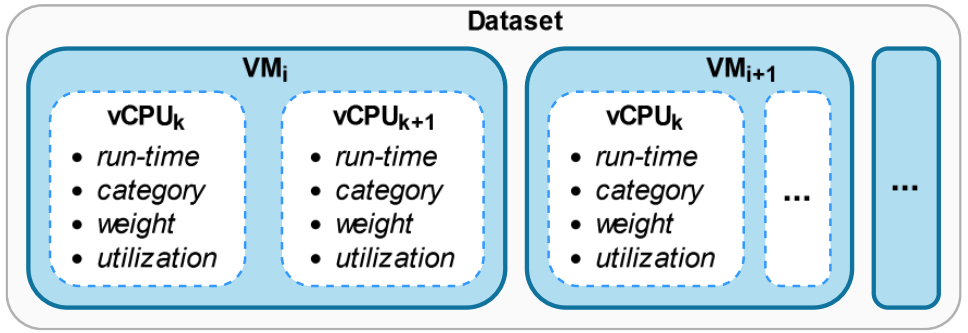

2 Scheduling

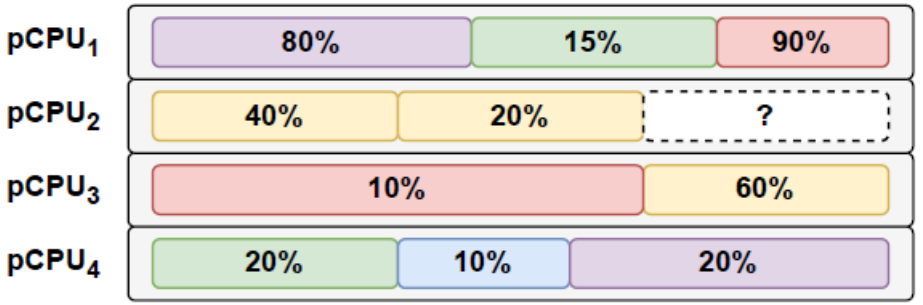

X% = I/O Utilization, Color = I/O Category

Integration into Virtualized Environments

Integration Layer	Kernel Integration	sched_ext Framework
Operation Mode	Stand Alone	Refinement
Transparency / Telemetry	compact run-queue summaries via hypercall / paravirtual channel	
Overhead	only on vCPU ready events	

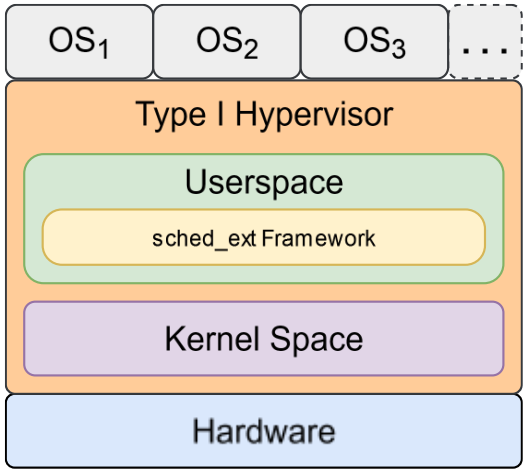

rexroth
A Bosch Company

Scenario Overview and Scheduling Behavior

Only Greedy

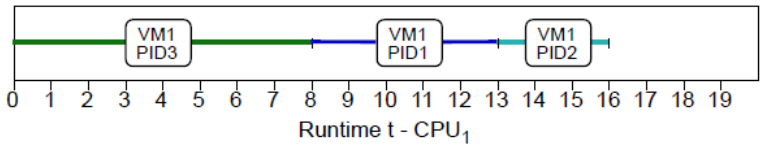

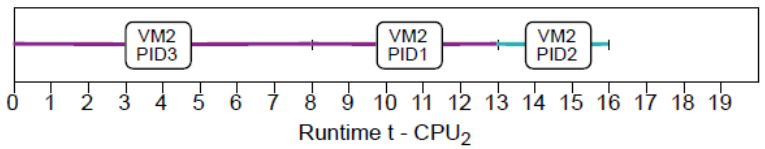

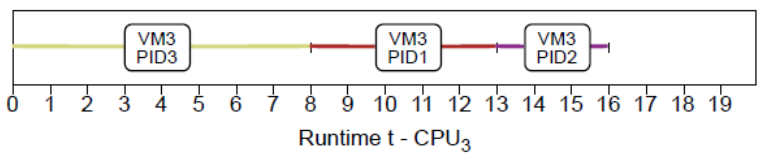

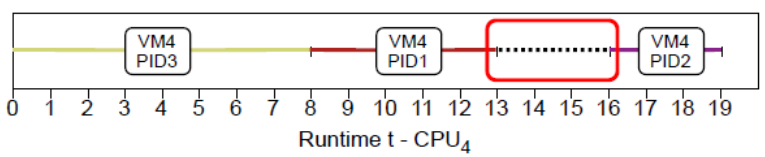

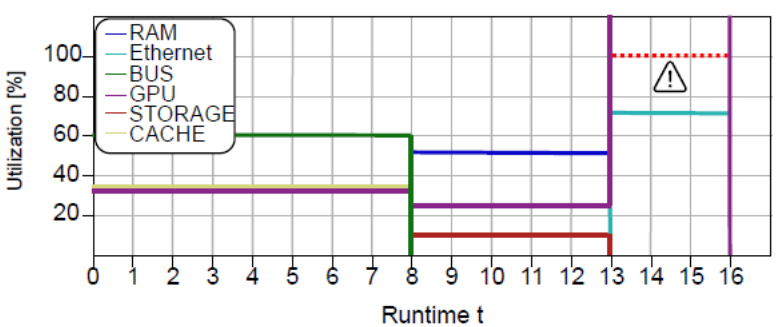

GPU utilization over **100%** not possible

Idle Time during the waiting period

Plan I/O utilization during scheduling

Interference-awareness

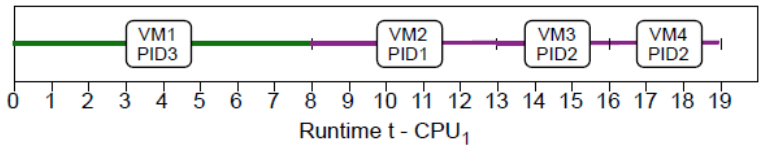

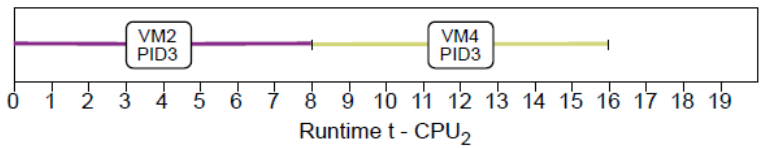

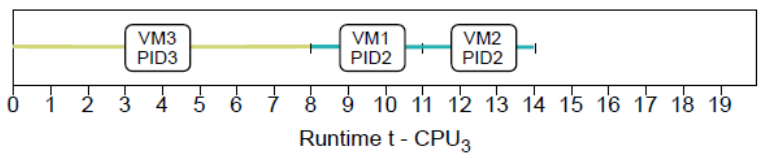

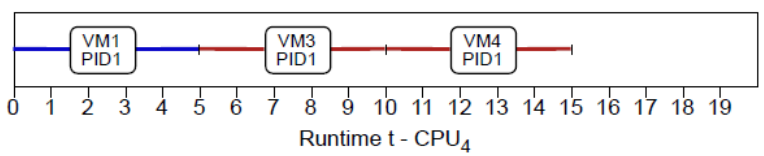

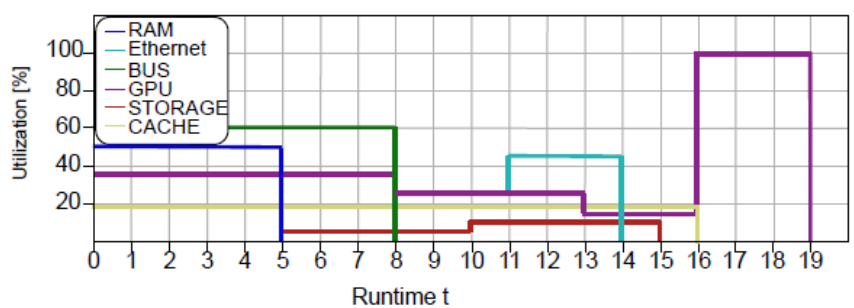

rexroth
A Bosch Company

Quantitative Metrics and Performance Comparison

Flatter I/O utilization over time

- Avg ETHERNET: 12.35% → 10.5% (≈−15%)
- Avg CACHE: 16.0% → 13.6% (≈−15%)

Lower peak loads

- Max ETHERNET: 70% → 45% (−36%)
- Max CACHE: 34% → 17% (−50%)
- Max STORAGE: ↓ by ~33%

Optimization

➔ **No overloads after optimization**

- GPU peaks up to 113% → **<100%** at all times

➔ **CPU load distribution**

- Per-core loads more differentiated (by design)

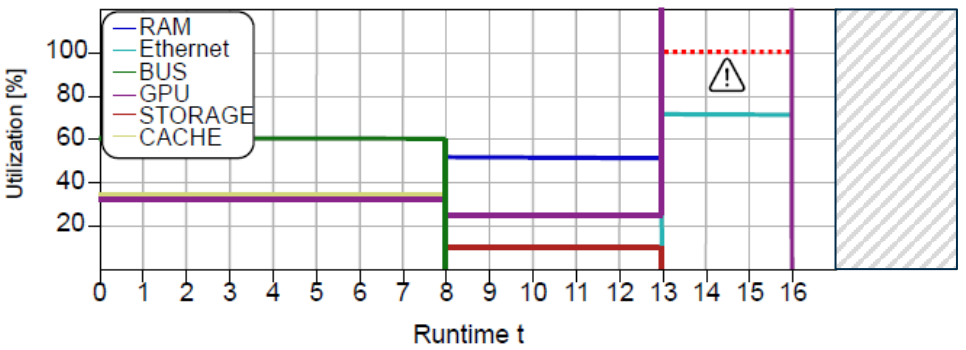

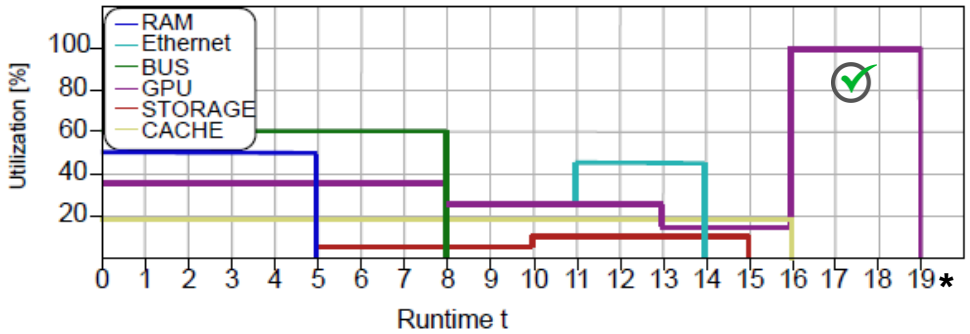

* Slightly staggered core finish times in synthetic test; irrelevant with continuous workloads

rexroth
A Bosch Company

Implications for Industrial Deployments

Predictable cycles
Safety-ready I/O

Deterministic timing
Smooth utilization

Lean sizing
Simpler verification

Robust under load
Continuous workloads

High availability
No overload events

Better thermals
Longer hardware life

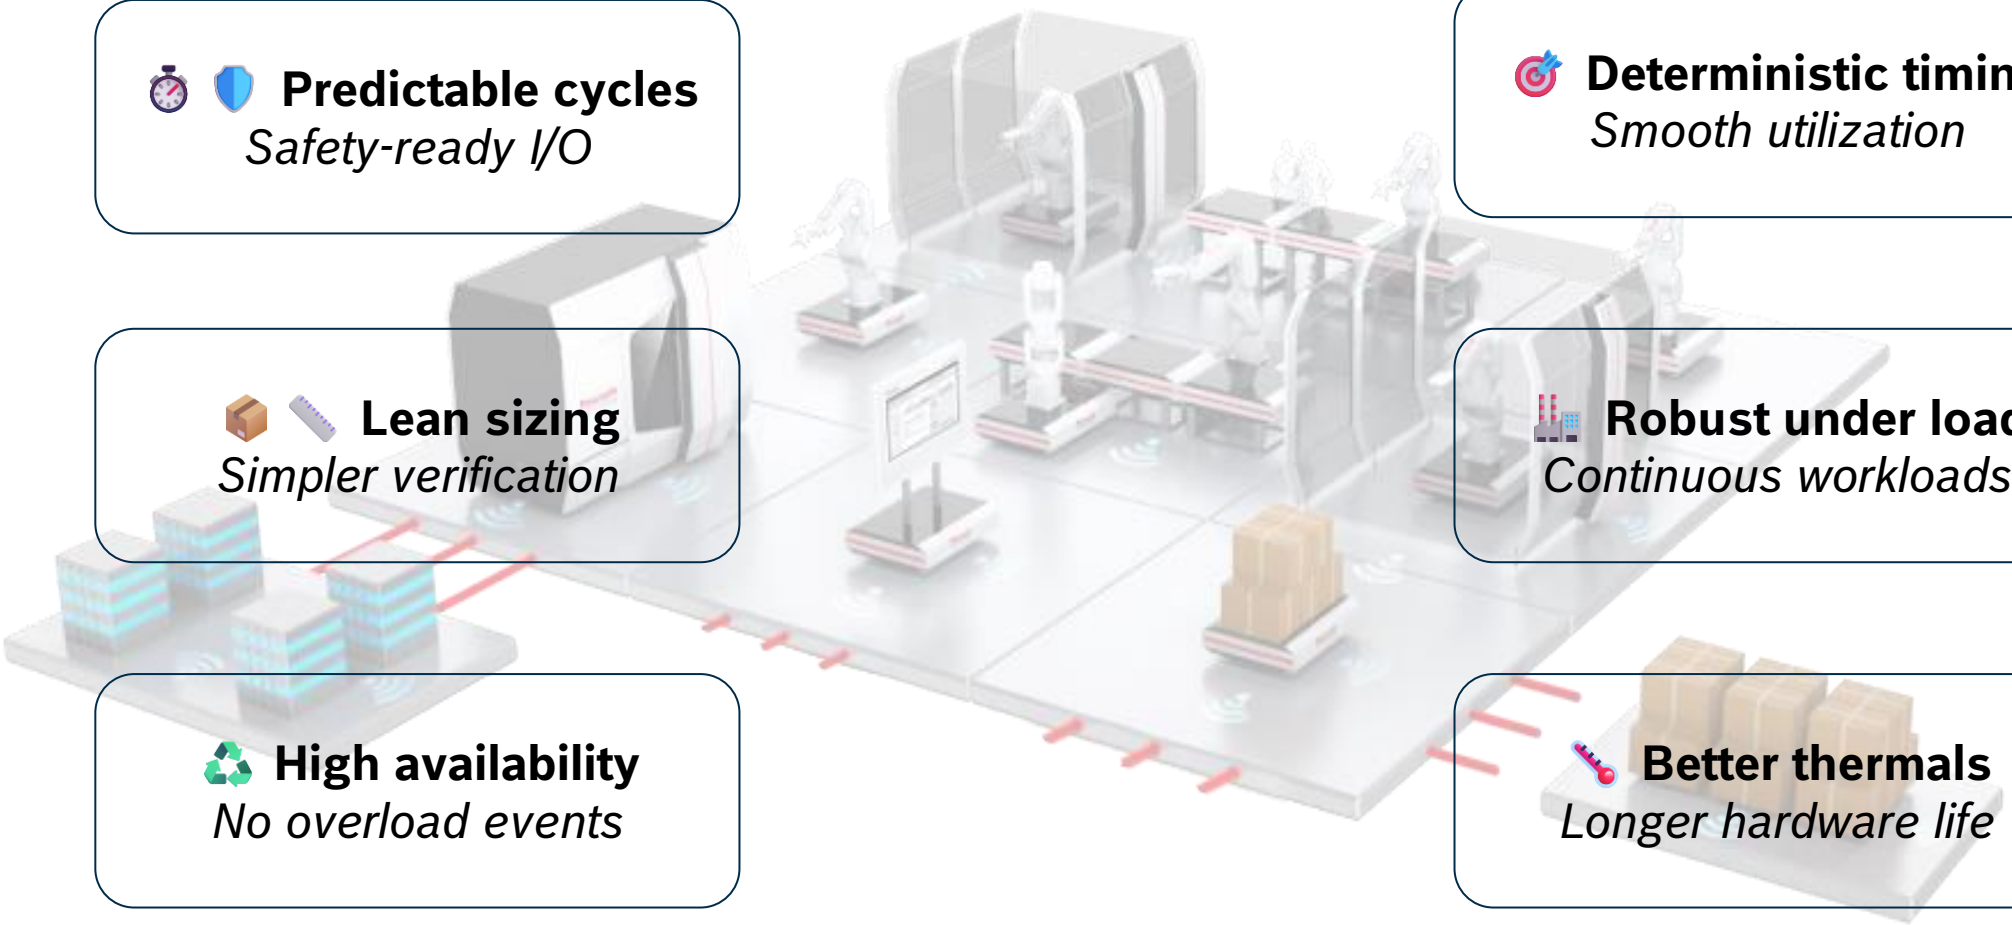

rexroth
A Bosch Company

Takeaways and Future Research

Takeaways

- **Interference-aware vCPU scheduling**: classify by dominant I/O and avoid co-running same classes to reduce contention.
- **Steadier resource use**: utilization is more even, peak stress is damped, and components stay below saturation.
- **More predictable timing**: reduced jitter and fewer cycle-time violations for real-time control loops.
- **Balanced CPU load**: distribution across cores improves; small per-core finish differences are non-critical in continuous workloads.
- **Lightweight integration**: works standalone or as a refinement to existing RT schedulers with minimal guest telemetry.

Future Research

- Kernel-level implementation & field validation
- Adaptive, criticality-aware policies
- Joint CPU/I-O coordination
- Operational impacts on energy, thermals, etc.
- Tooling & resilience

rexroth
A Bosch Company

Thank you!
Any questions?

Jessica Müller
jessica.mueller4@boschrexroth.de
0009-0007-6478-7307

rexroth
A Bosch Company