SILEXICA

SLX Analytics & Testing Platform

Benjamin Goldschmidt
goldschmidt@silexica.com
Stuttgart, 12.12.2019

Silexica Facts

Est. 2014 after a decade of research

Team of world leading software and hardware experts

60 people worldwide, engineering HQ in Germany

3 offices and worldwide local support engineers

Motivation

- Autoware ROS-based AD stack (Moriyama example, v1.8)
- Problem: Positioning of the car is sometimes unstable
- Many active modules in the example + large codebase:
  - Where to start?!
  - How to make behavior reproducible!?

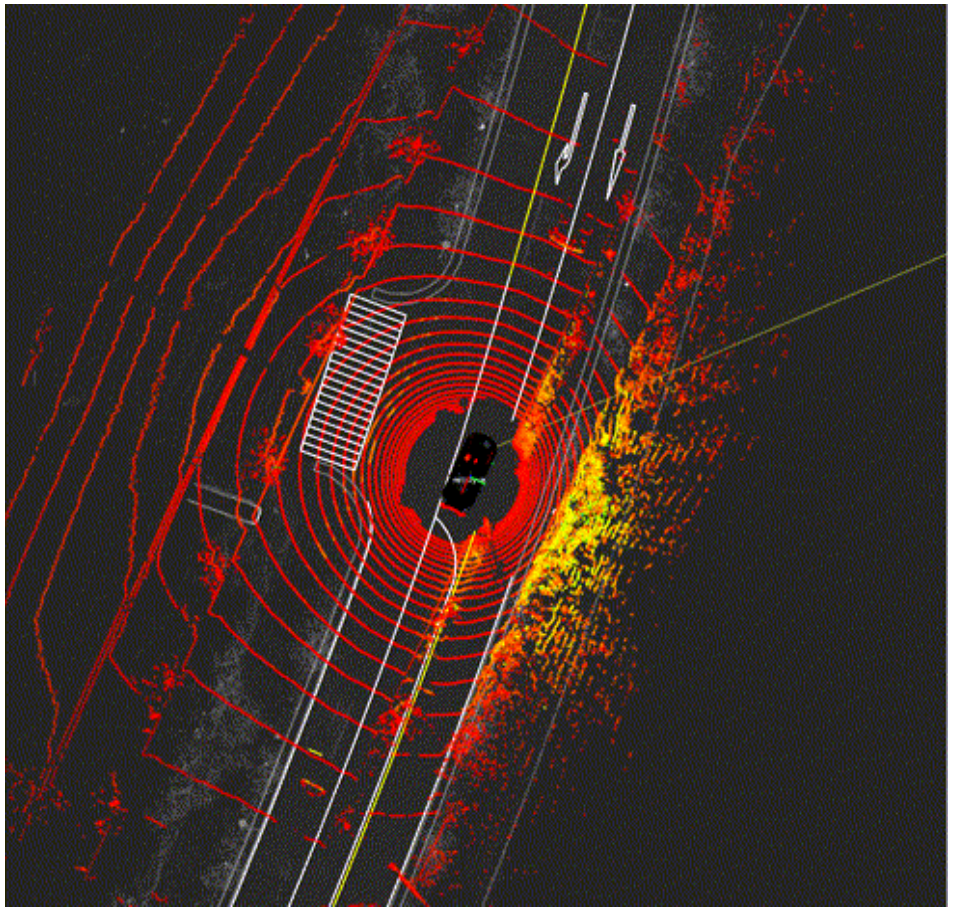

Motivation

Two Runs of the Same Scenario

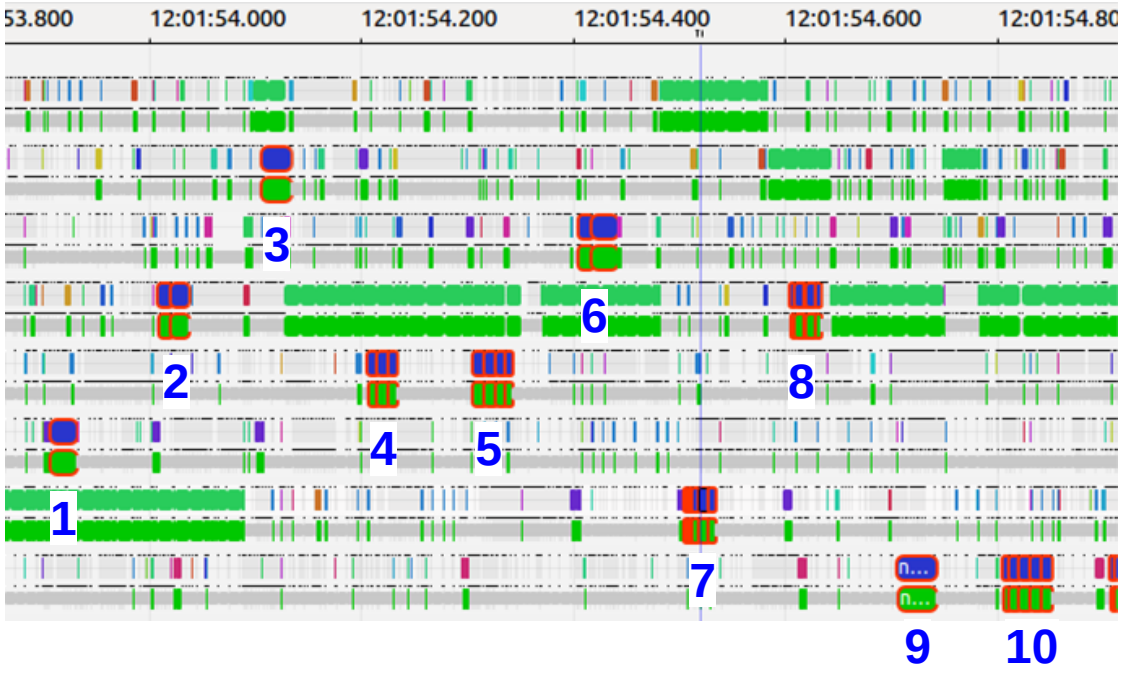

Successful run:

- Positioning module (marked red) runs in regular intervals with acceptable runtime variation

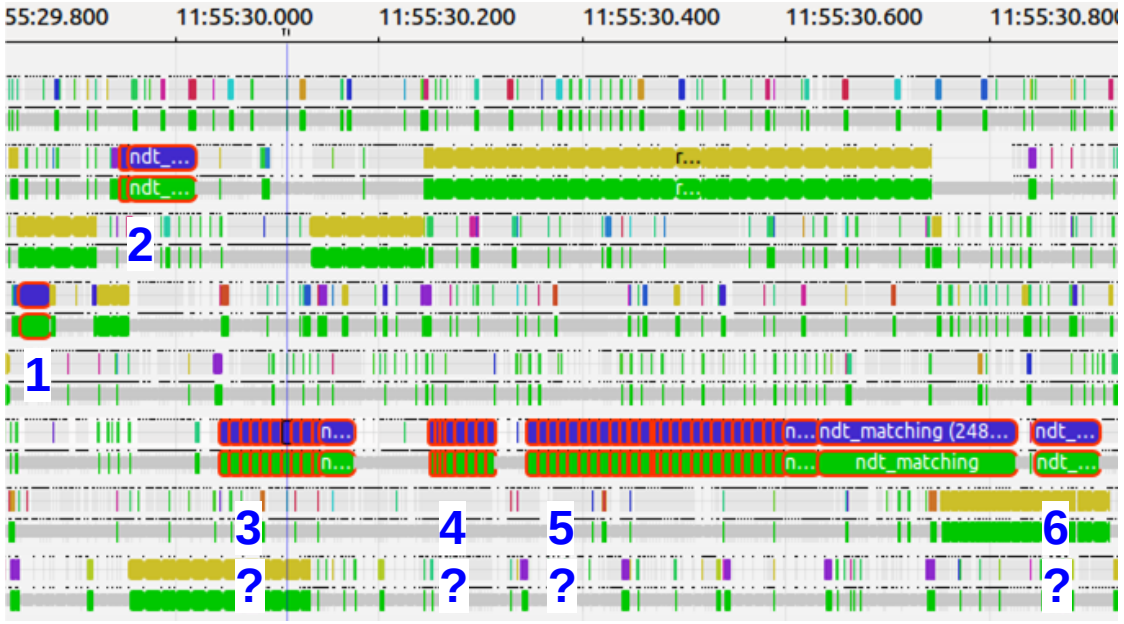

Unsuccessful run:

- Positioning module (marked red) does not run in regular time intervals and has a highly variable runtime

Problem Statement

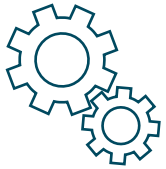

Missing system overview → Engineers are fixing, not developing

No statistical profiling of system → High risk of (creeping) system degradation

No automated testing of system metrics → Last-minute fire-fighting / delay of release

No scalability of testing solution → Limited test coverage to achieve high release rate

SLX Analytics & Testing Platform

- Automated outlier detection:
  - Multi-level/run tracing + analysis
  - CI integration
- Interactive root cause analysis
- Highly scalable:
  - Cloud, on-premise, desktop
  - Open API

Multi-Level Analysis

- System-level:
  - Acquire kernel/system events → global timeline of the system
- Middleware-level (e.g. ROS{1, 2}, Adaptive AUTOSAR, etc.):
  - Acquire middleware-specific data (e.g. from the communication stack) to provide semantic context
- Module-level:
  - Acquire module-specific data (static/dynamic) for in-depth analysis

Connect

Multi-Level Analysis

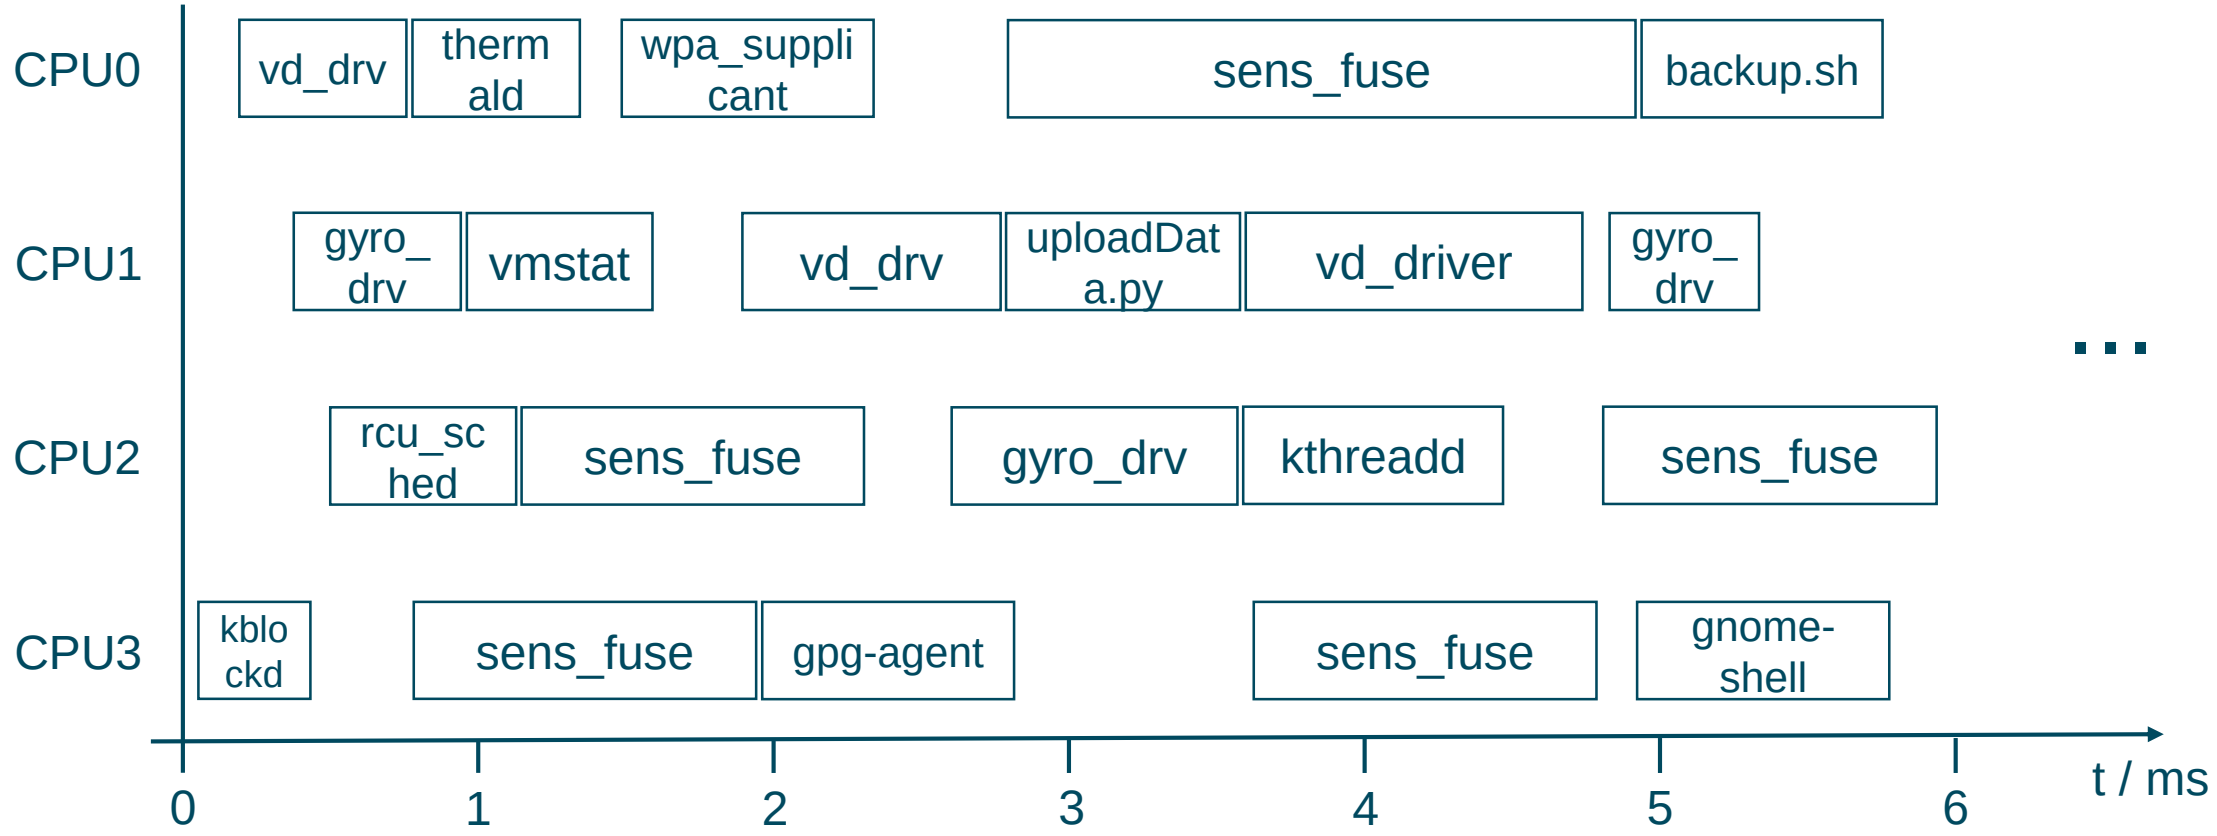

System-level data: Global CPU timeline (proccesses)

Multi-Level Analysis

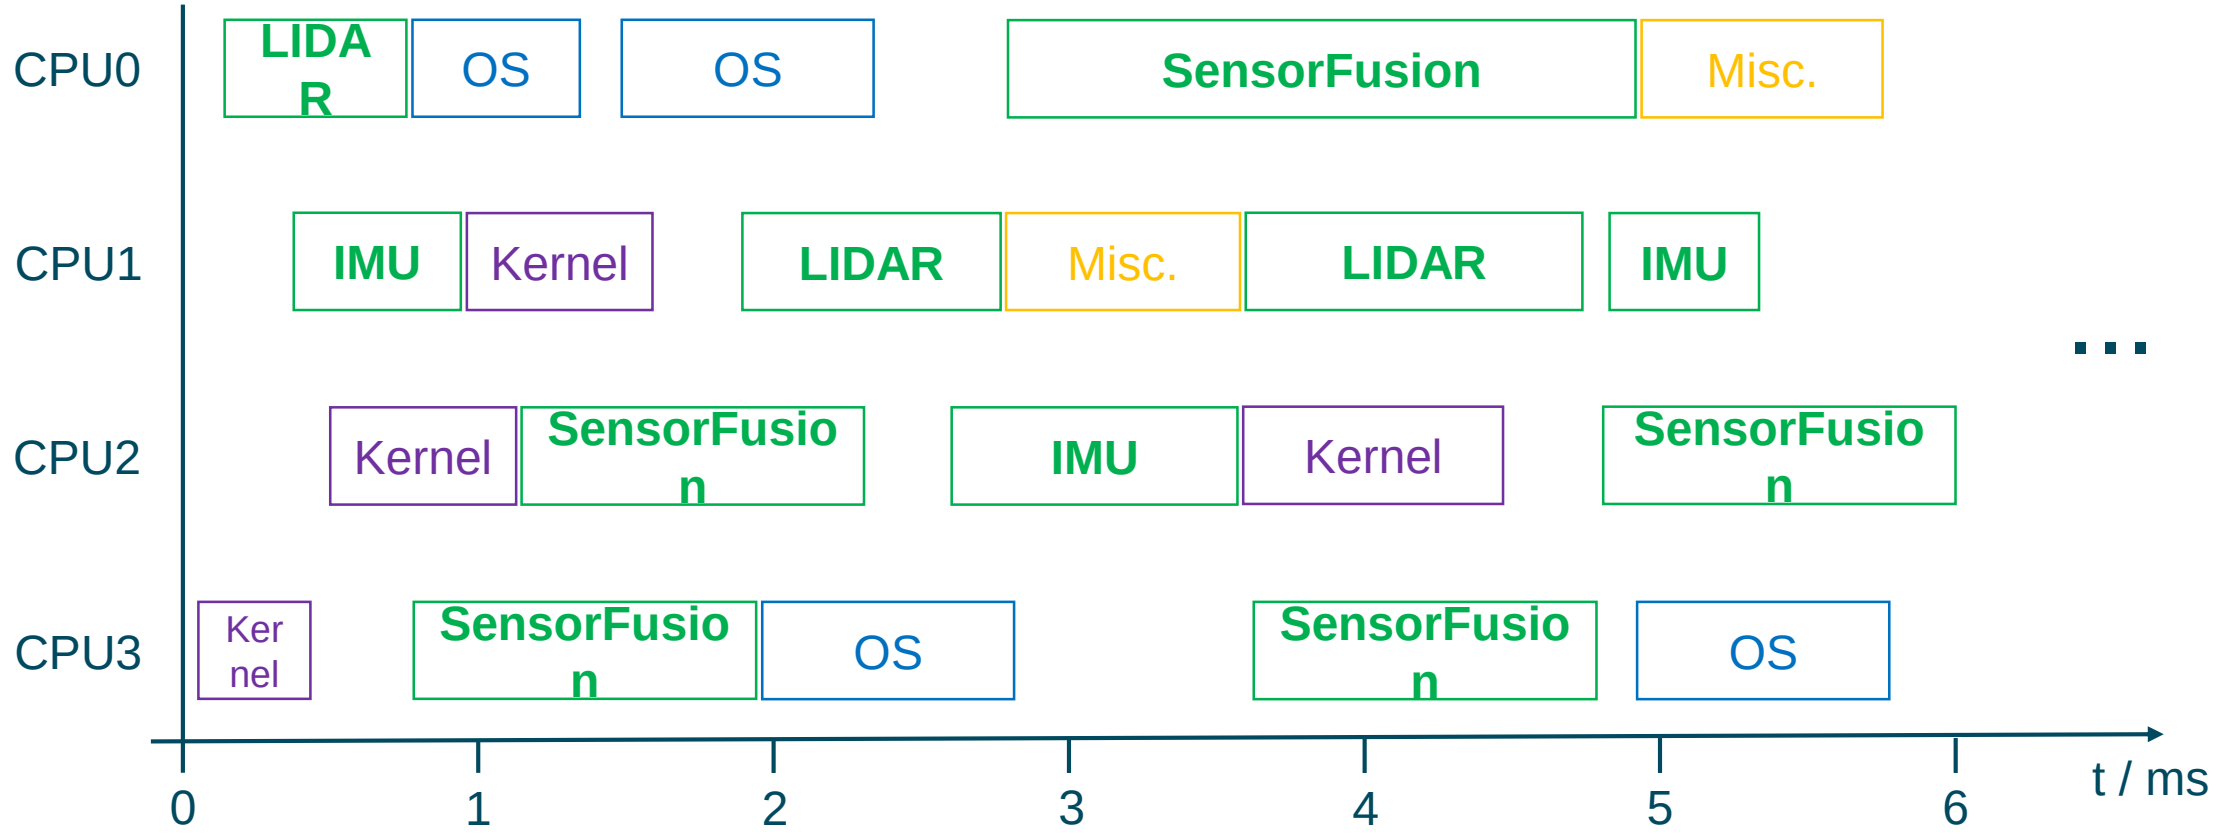

Use middleware-level data to highlight relevant processes

Multi-Level Analysis

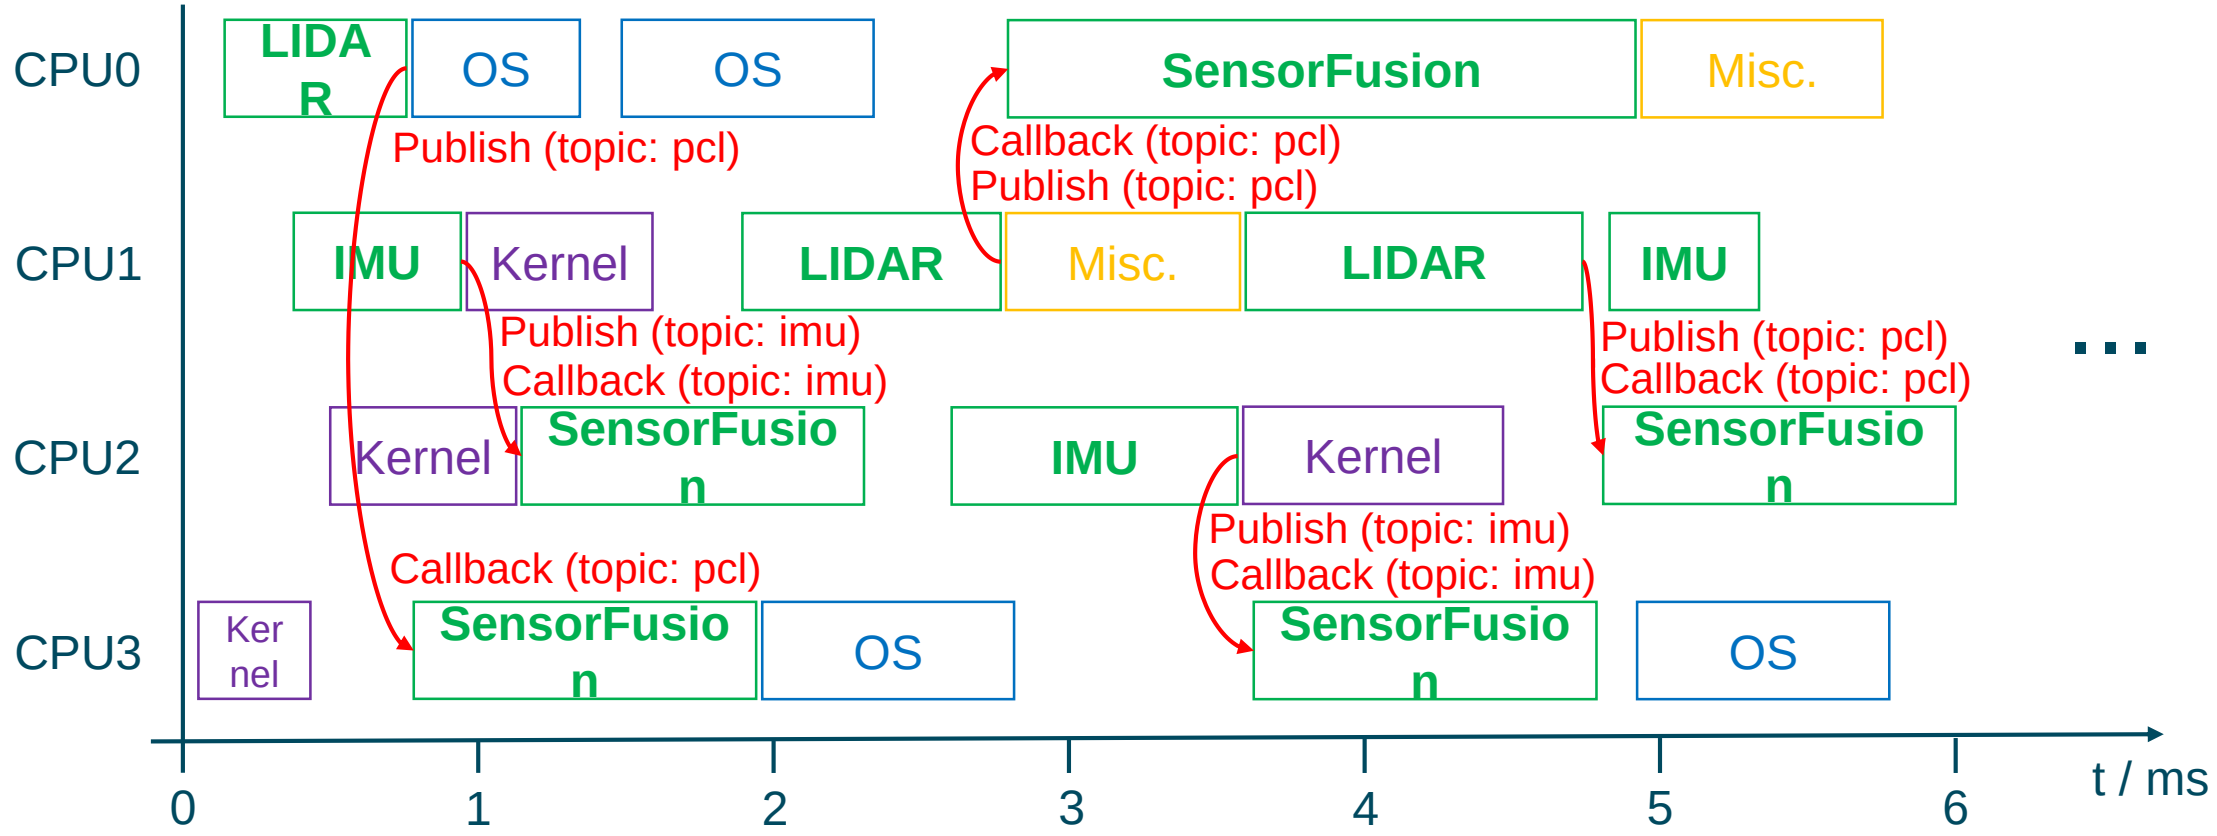

Add middleware communication events to global timeline

Multi-Run Analysis

- Run test scenarios multiple times to detect outliers
- Configurable constraints for automated tests
- Full results history:
  - Compare code revisions
  - Compare same/different inputs

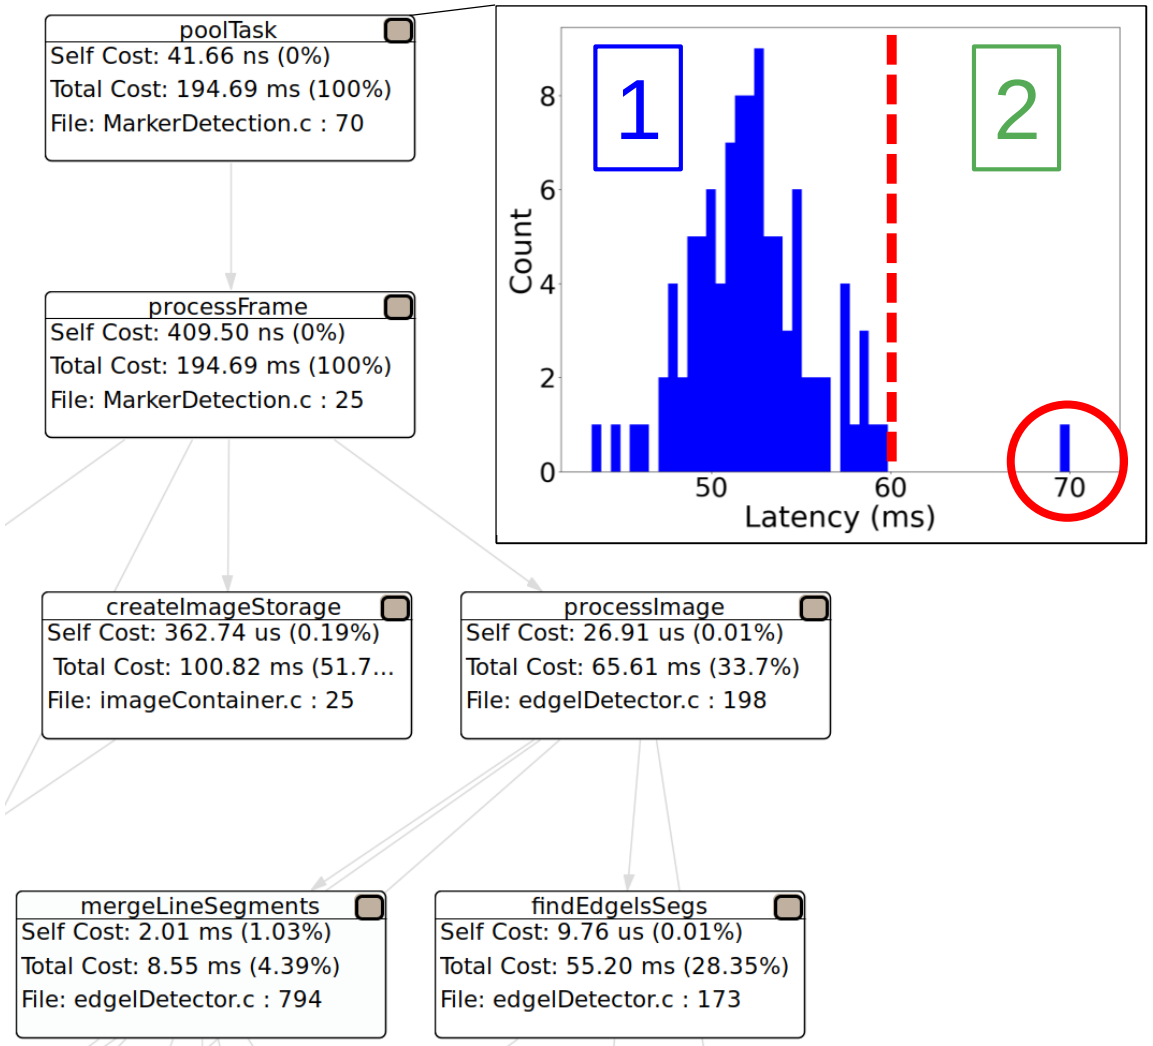

Module-Level Analysis

- Detect memory allocations after startup phase
- Detect system calls with difficult to predict runtime behavior
- Detect TCP, UDP, and middleware communication

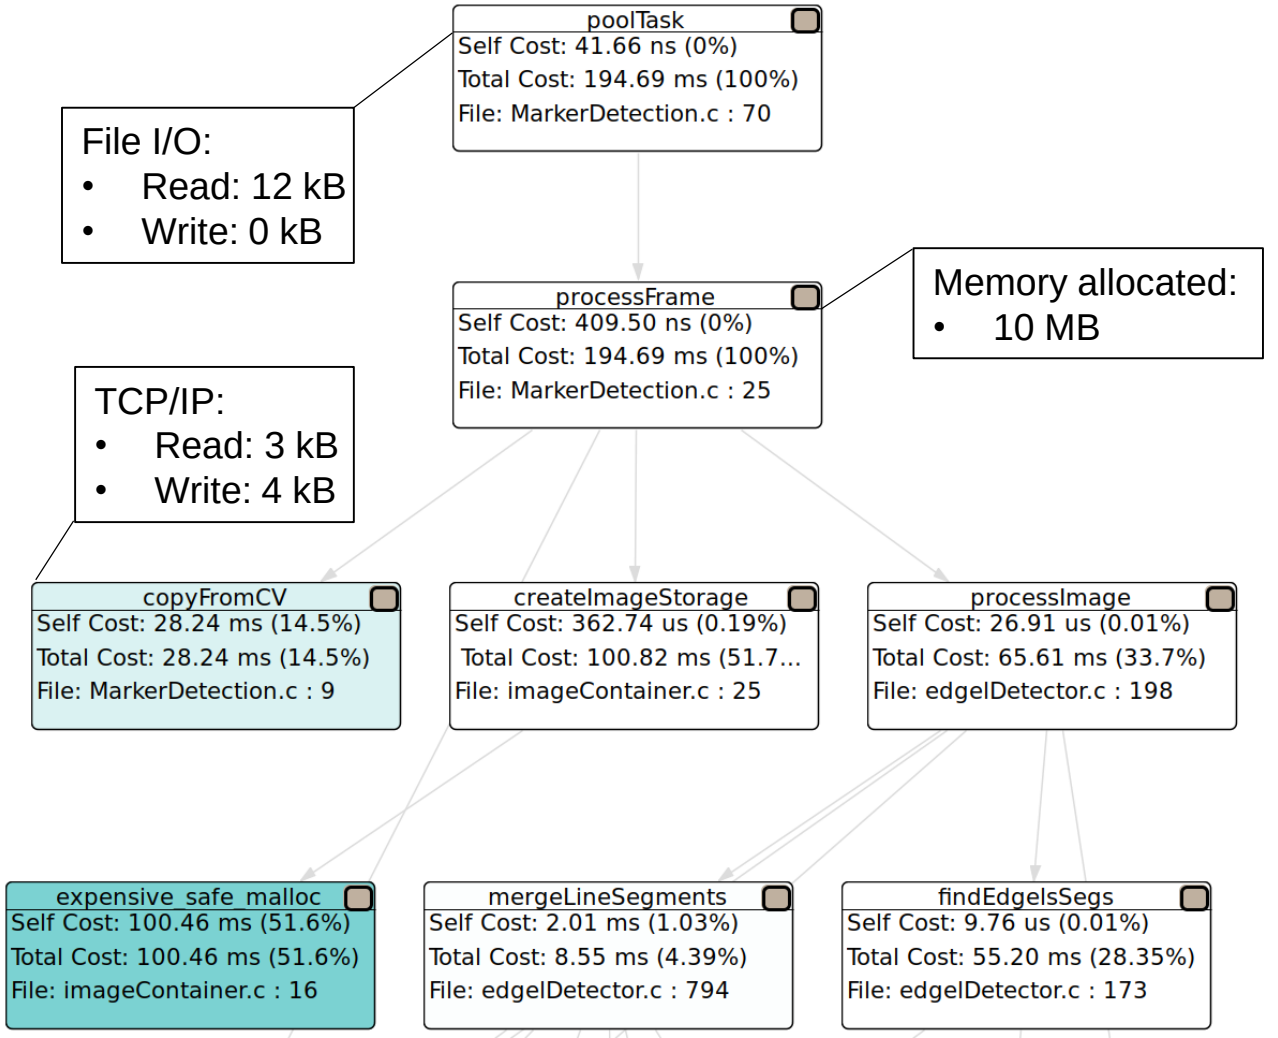

Stress Tests

- Highly configurable stress tests: CPU, memory, cache, etc.
- Loadable/storable stress profiles (stress over time)
- Chaos engineering: Randomly stop processes during execution

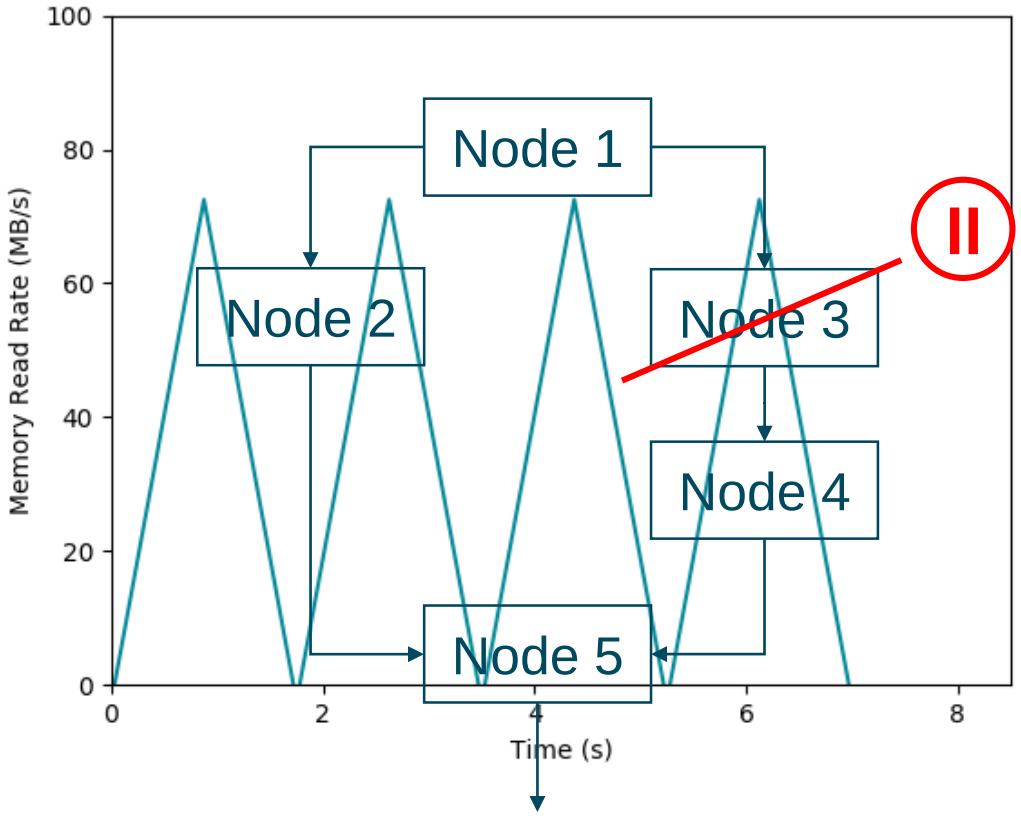

Interactive Root Cause Analysis

Latency Test History

Constraint violated!
Investigate (click!)

Constraint:
850 ms

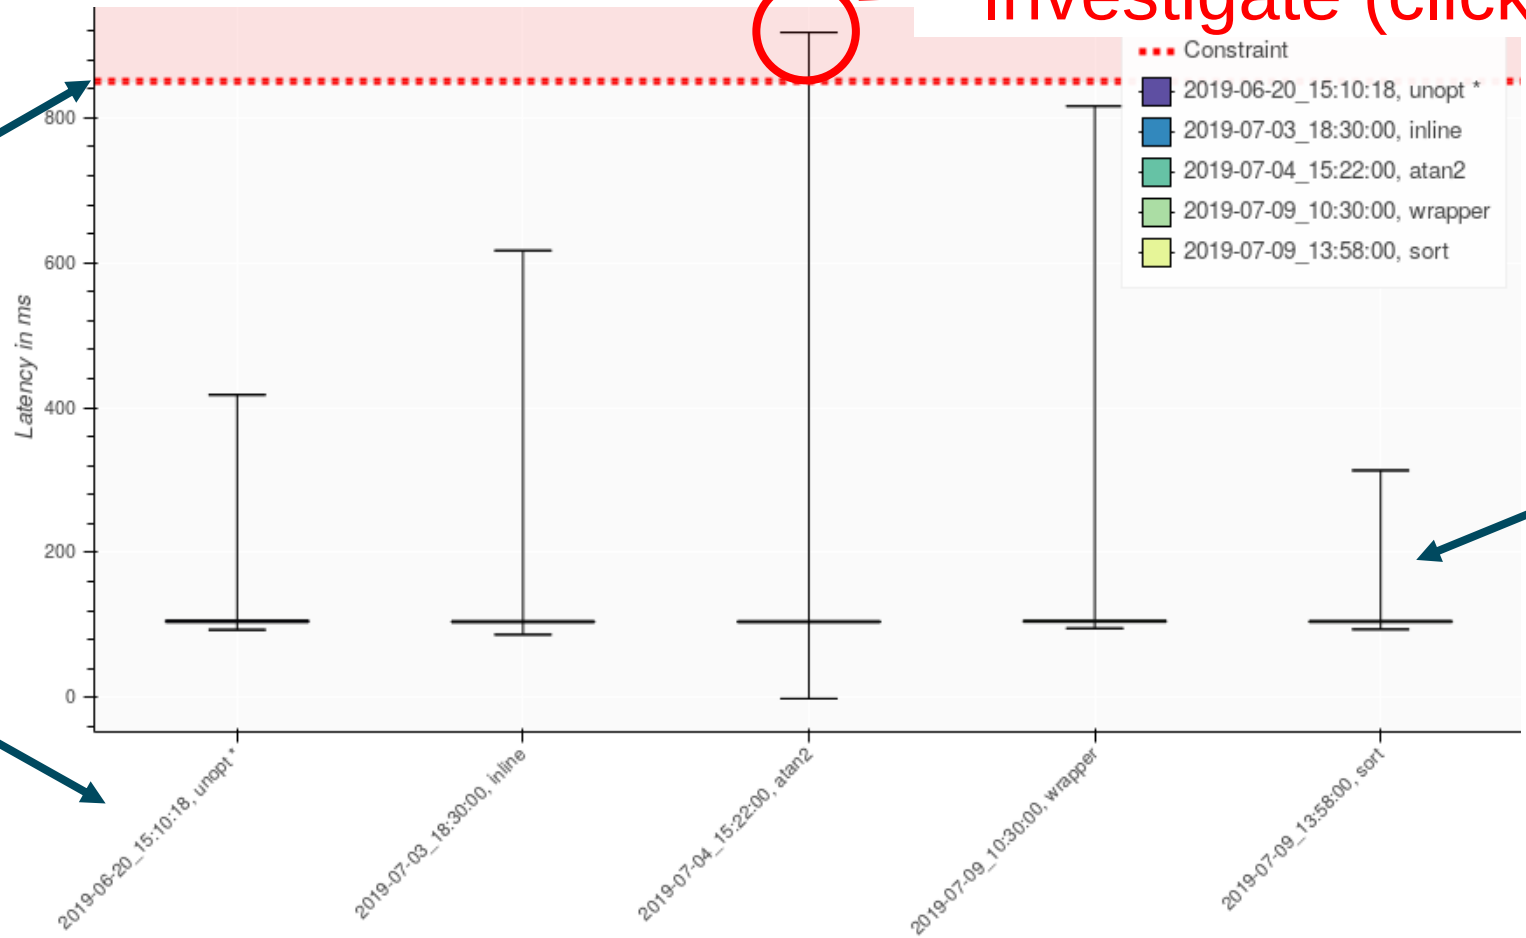

Measured
latency
distributions

Commits

Level 1

Interactive Root Cause Analysis

Latency Histogram

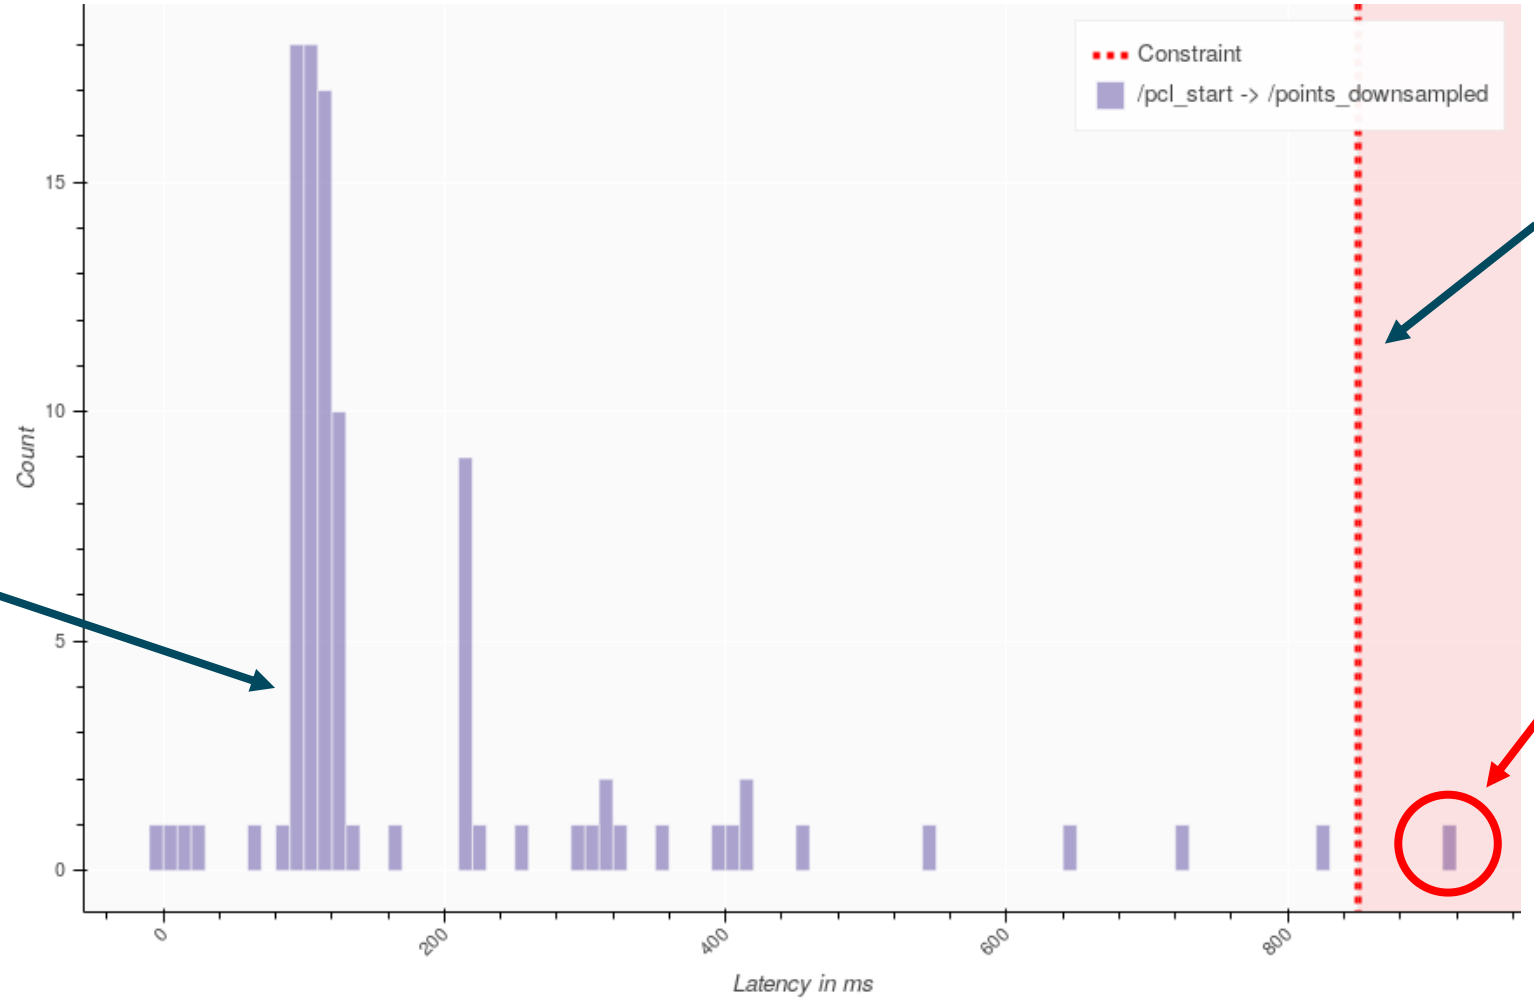

Constraint:
850 ms

Measured
latencies

Investigate
further (click!)

Level 2

Interactive Root Cause Analysis

Latency Over Time

Investigate further (click!)

Constraint:
850 ms

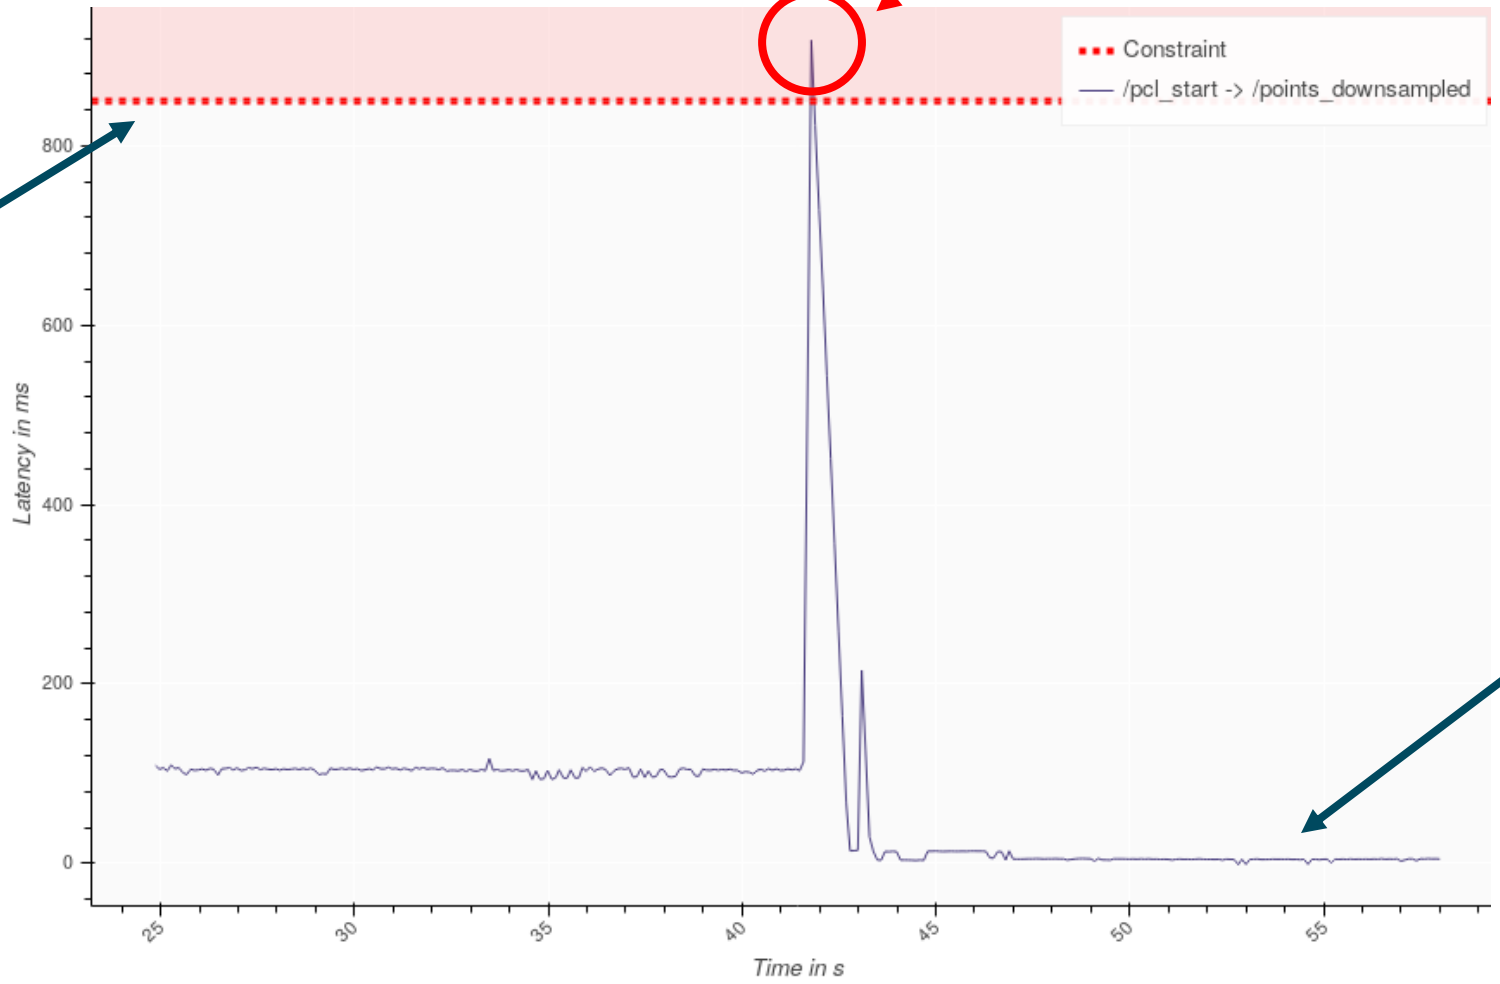

Measured
latency
over time

Level 3

Interactive Root Cause Analysis

CPU Gantt Chart

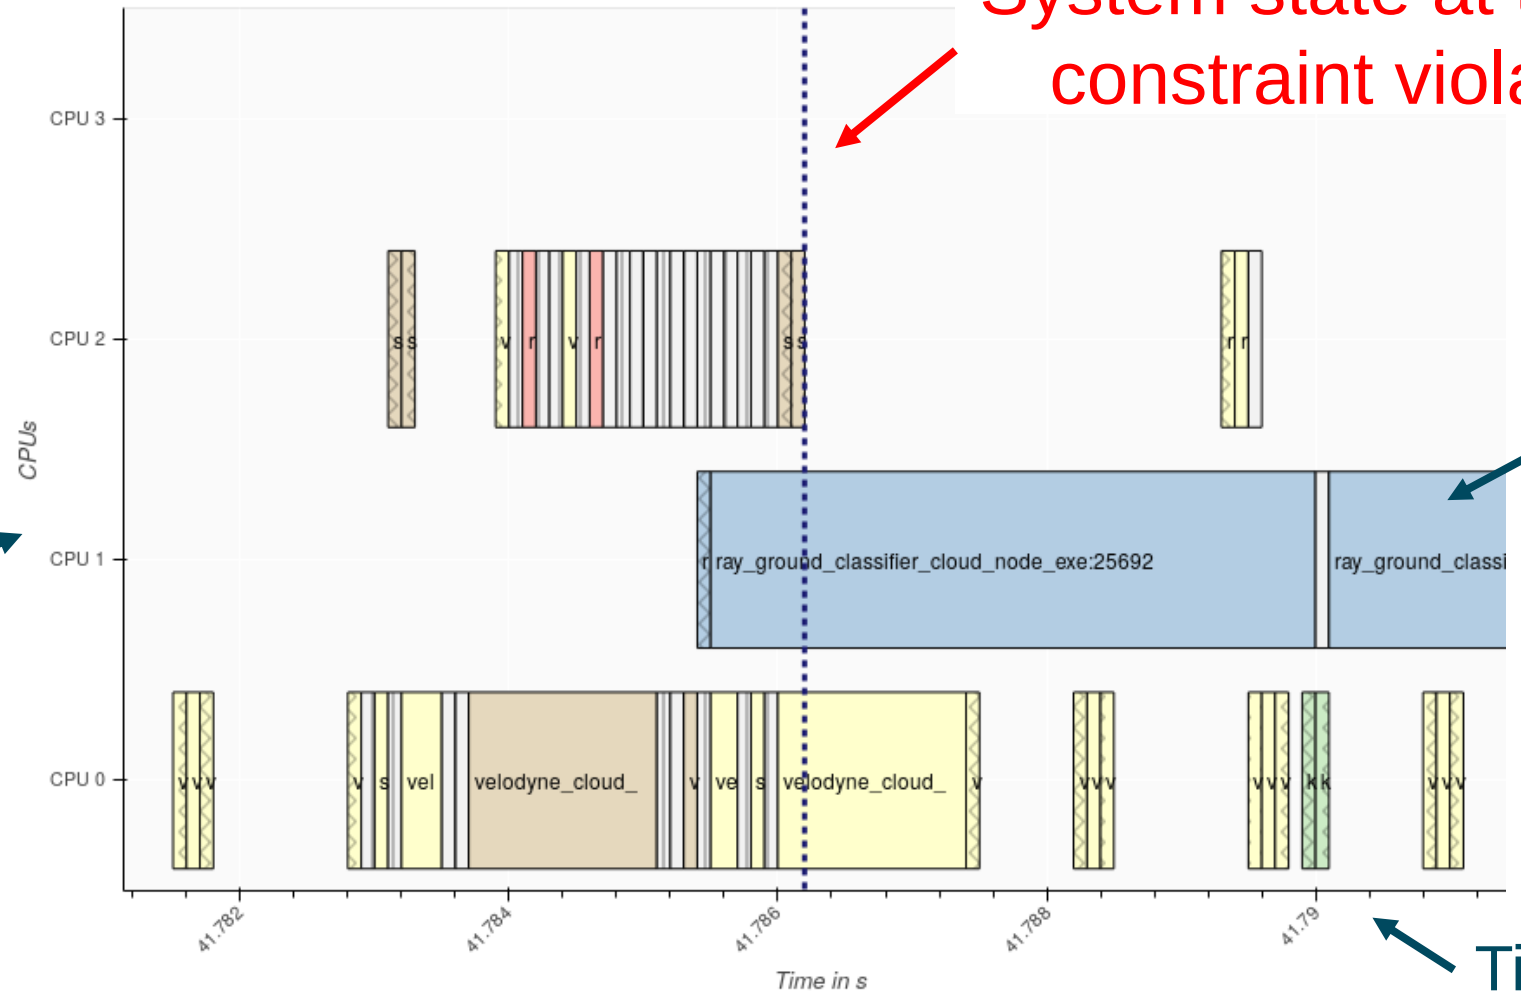

Level 4

Open API / Jupyter Notebooks

Select Results

```
import slx.api
engine = slx.api.db_engine()
slx.api.init_plots()

config_aarch64 = {
    'arch': 'aarch64',
    'OS': 'Ubuntu 18.04 bionic',
    'scenario': 'perception',
    'tag': None
}
exp_ids_aarch64, exp_configs_aarch64, exp_infos_aarch64 = slx.cmd.list_experiments(engine, config_aarch64)

# Create a plot for the given metric
plot_cfg = slx.api.make_plot_config_experiments(slx.api.Metric.CPU_LOAD_BY_CPU, hist_avg=False, line_all=True)
slx.api.plot_experiments(engine, plot_cfg, exp_ids_aarch64, exp_configs_aarch64)
```

Select Metric

Visualize

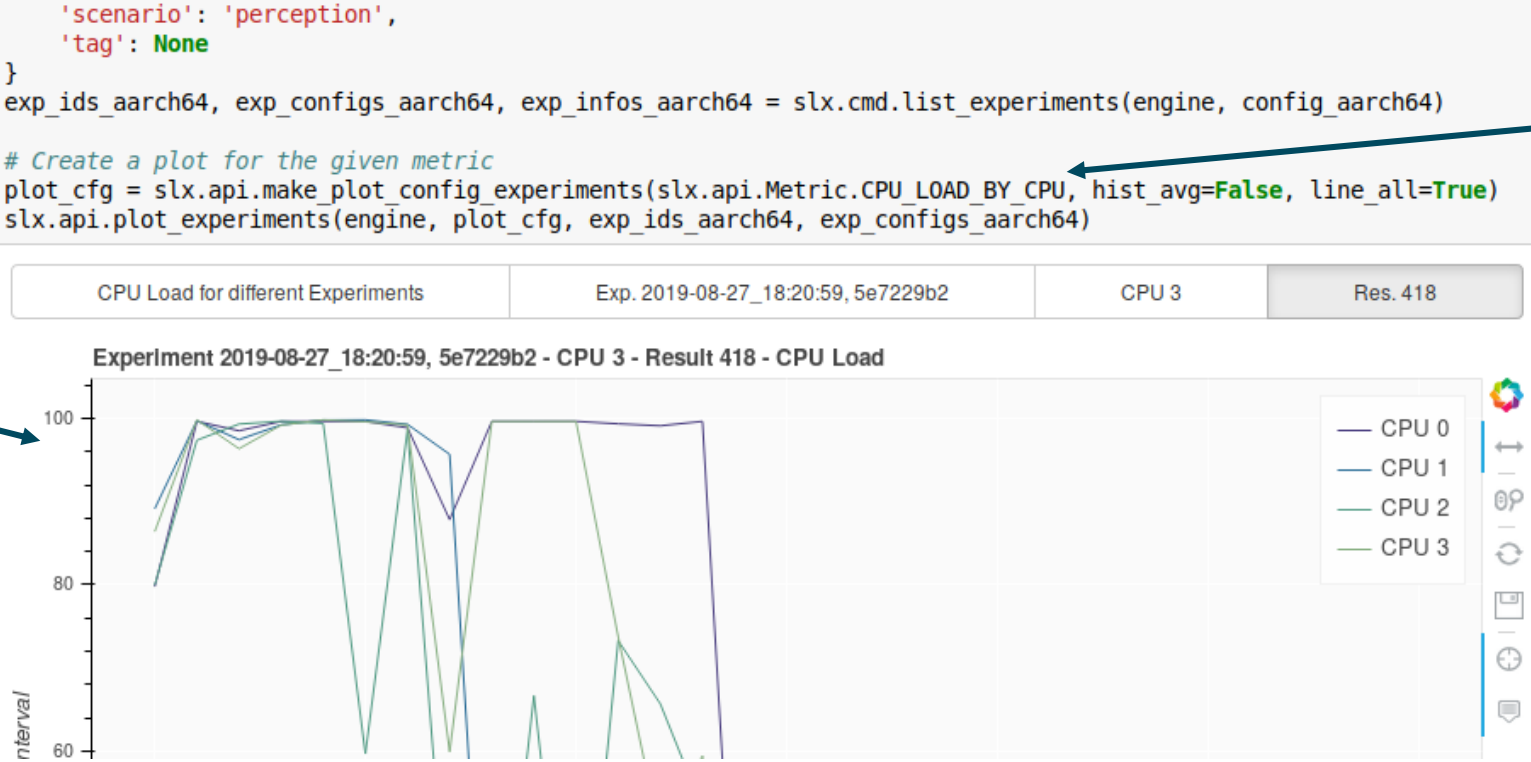

Thanks

1st product release in January 2020

silexica