High Frequency Performance Monitoring via Architectural Event Measurement
--- on ARM Processors

CHEN LIU

DEPARTMENT OF ELECTRICAL AND COMPUTER ENGINEERING

CLARKSON UNIVERSITY

POTSDAM, NEW YORK, USA

CLIU@CLARKSON.EDU

May 25, 2023 @ AHUG 2023

Motivation

- Significant demand to improve various metrics for modern computing systems

  - performance, energy efficiency, etc.

- Performance monitoring counters (PMCs) built into modern processors

  - Can provide fine-grain performance details
  - ARM Neoverse processor: 1 fixed + 6 programmable performance counters

Applications of performance counter profiling

- Performance Optimization
- Malware Detection
- OS Scheduling
- Dynamic power estimation

Motivation

- Many tools have been developed to provide a high-level API to access the low-level performance counters.

Limitations

Kernel - Lineage of Event Behavior (K-LEB)

A performance-counter-based profiling tool that utilizes a kernel space collection system to produce ***precise, non-intrusive, low overhead, high periodicity*** performance counter data.

K-LEB System Model

Controller process

- Control the kernel module from user space.

Kernel module

- Access PMU to collect performance counter data.

PMU

- Special hardware unit used to monitor the hardware events.

Monitored process

- Process being monitored.

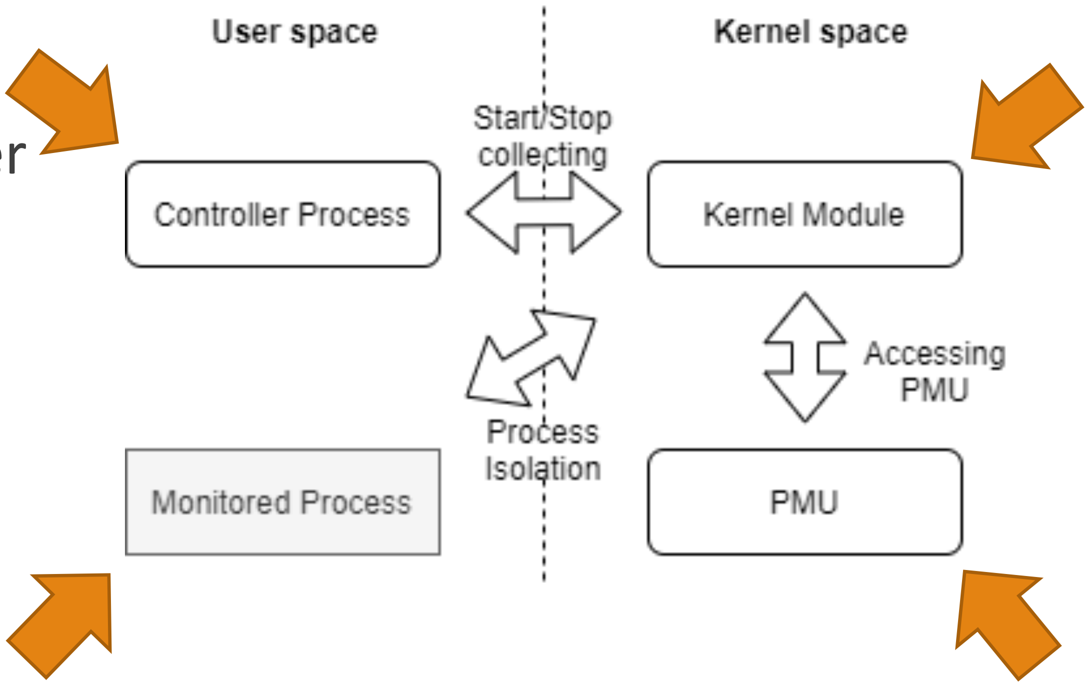

Process Flows

5 phases

1) Module initialization
2) Start monitoring
3) Stop monitoring
4) Module de-initialization
5) Logging

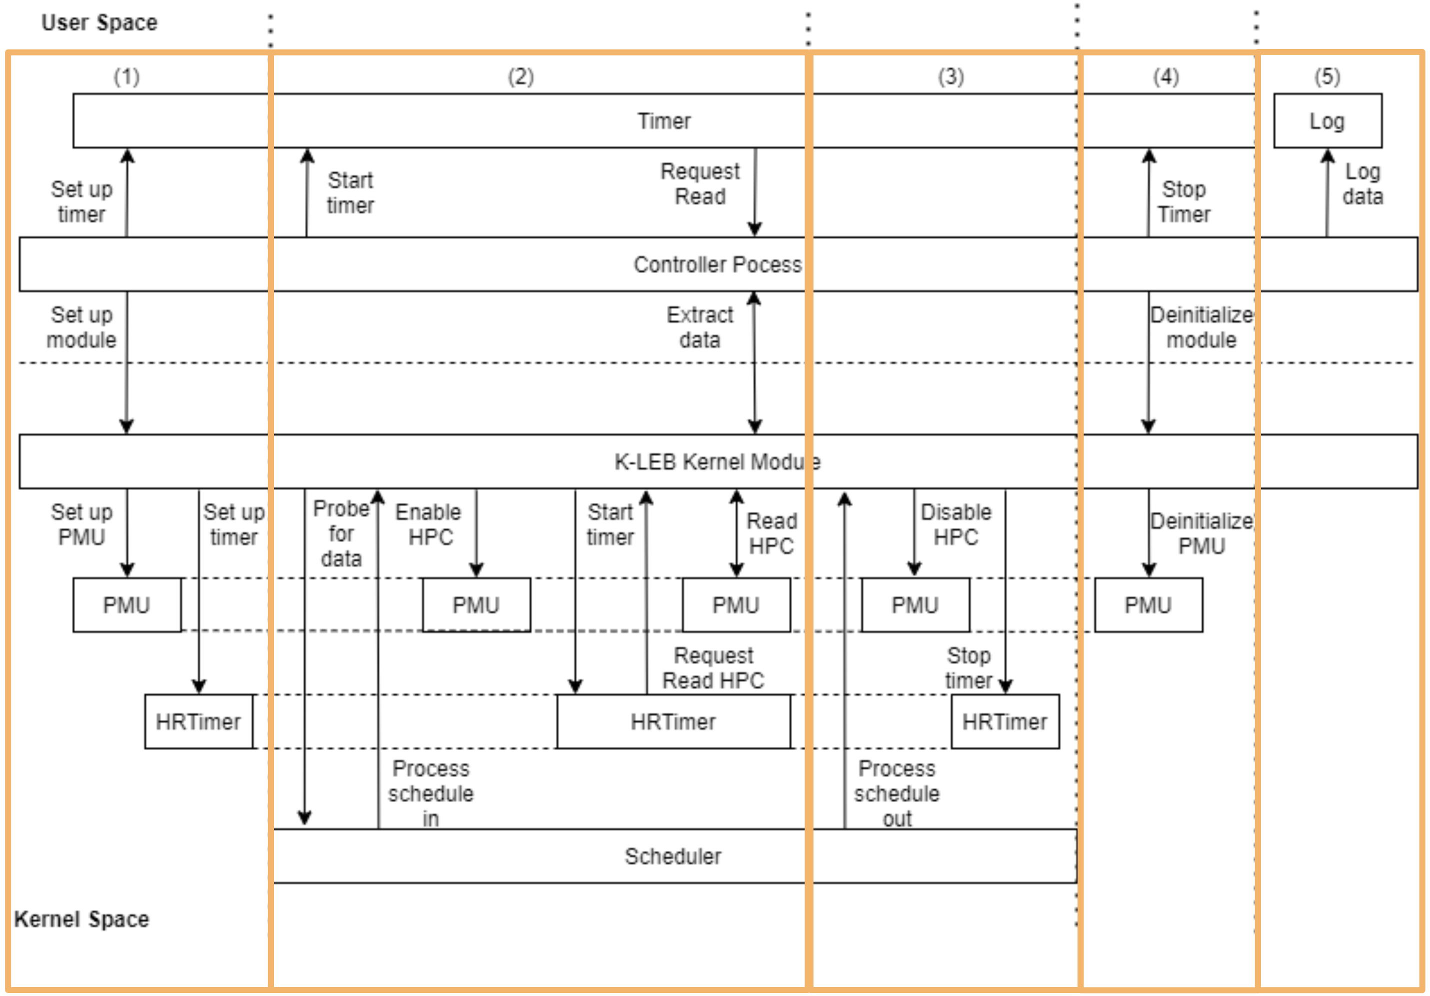

Process Interaction

Interaction between K-LEB and the monitored process.

- 5 phases
  1) Module initialization
  2) Start monitoring
  3) Stop monitoring
  4) Module de-initialization
  5) Logging

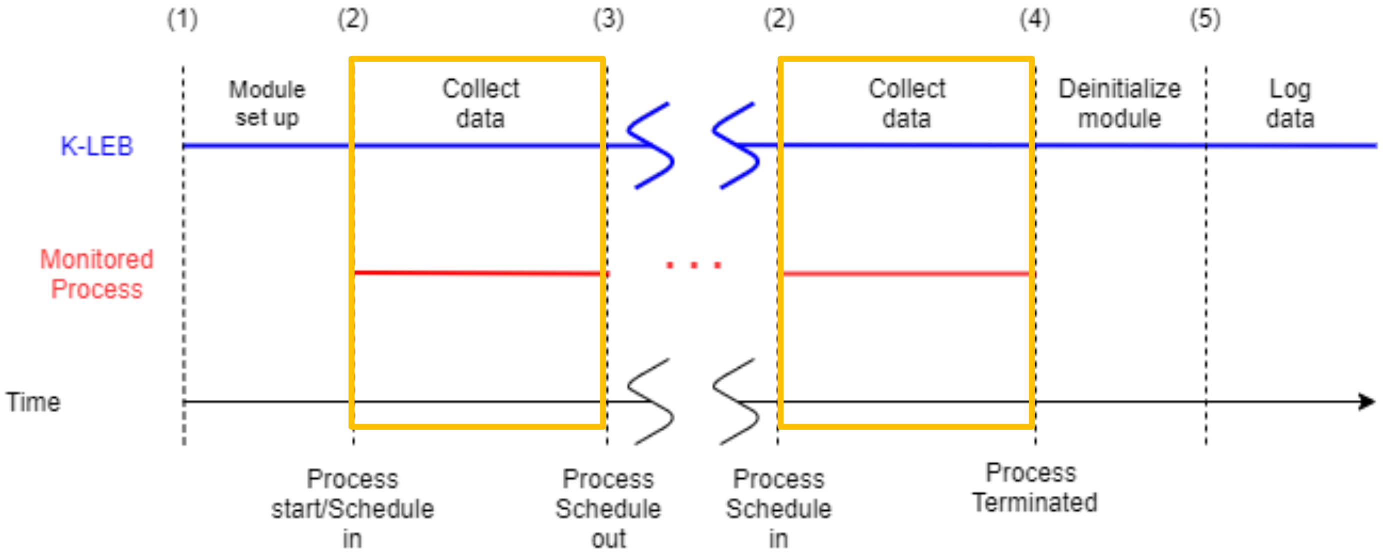

Experiment setup

AWS Graviton Processors feature 64-bit Arm Neoverse running Ubuntu 20.04.2 LTS with Linux kernel version 5.4.0-1041-aws

Raspberry Pi 4 Model B Broadcom BCM2711 @ 1.5GHz processor running Debian 10 with Linux kernel version 5.4.72-v8+

Case study 1: Linpack

- Capture phase behavior.
- K-LEB has a very small FLOPS loss of 0.04% in comparison with 0.98% from Perf.
- No source code require.

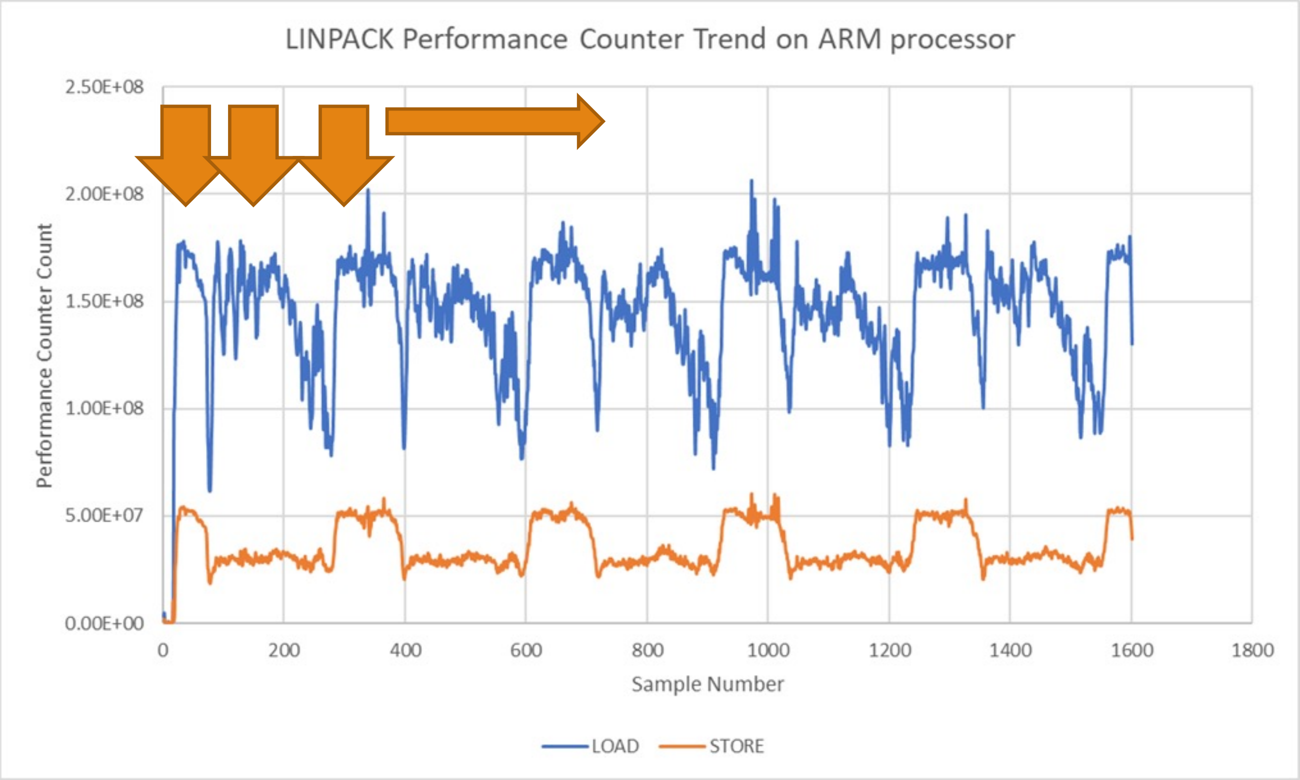

Case study 2: Docker

- ❑Workload characterization.
- ❑Computation/Memory intensive.
- ❑Non-intrusive to a running program.

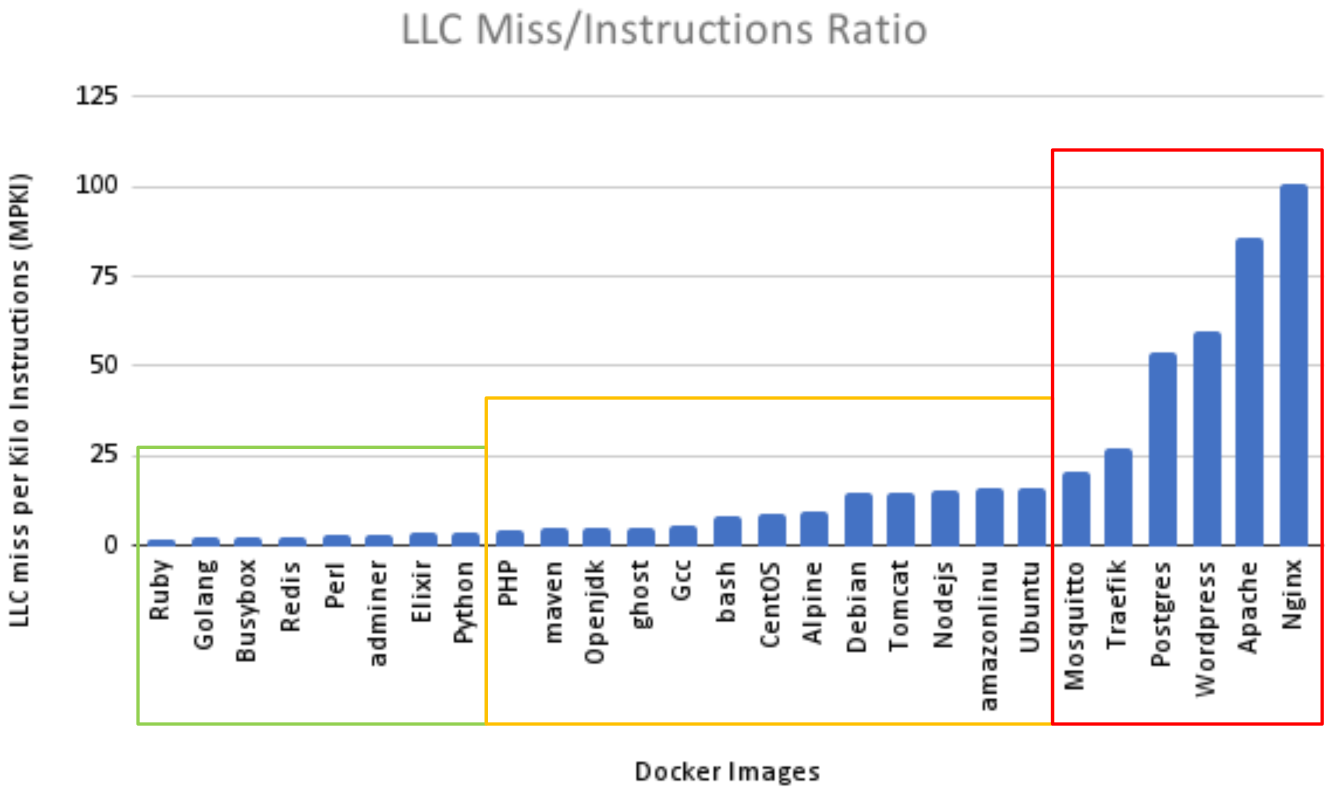

Case study 3: Spectre Attack

- Anomaly detection.
- High frequency timer.
- Monitor program with short execution time.

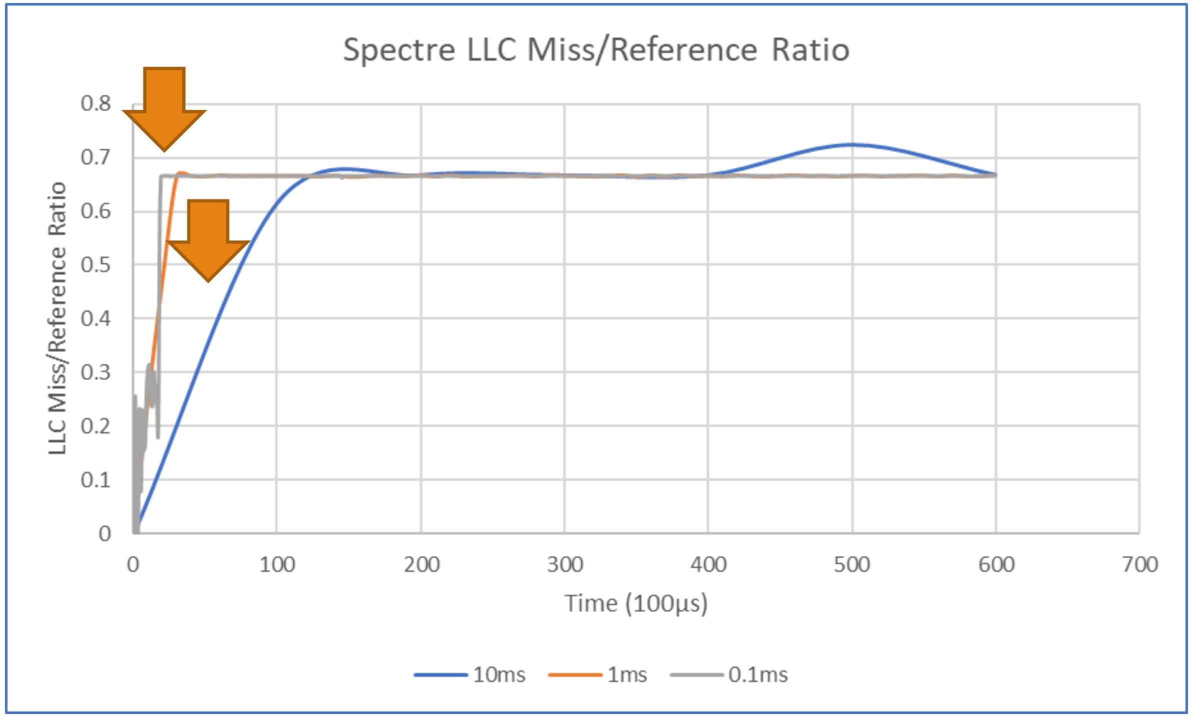

Performance overhead comparison

- Test on matrix multiplication program.
- Percentage performance overhead for each profiling tool.

Sample Rate	K-LEB	Perf stat	Perf record
10 ms	0.16	0.41	0.31
1 ms	0.43	N/A	1.52
0.1 ms	2.82	N/A	10.08

Performance overhead comparison

- Test on Linpack benchmark.
- Percentage GFlops loss for each profiling tools.

Sample Rate	K-LEB	Perf stat	Perf record
10 ms	0.04	0.98	0.47
1 ms	0.31	N/A	2.97
0.1 ms	2.18	N/A	15.378

Conclusions

We introduce K-LEB, a kernel module-based approach for performance counter collection with following features:

- Non-intrusive to the program being monitored.
- Provide high frequency sampling rate up to 100µs, which is 100 times finer granularity than current available tools.
- Very low overhead.
- Can benefit program analysis, malware detection and scheduling techniques.

Acknowledgement

Chutitep Woralert, Ph.D. candidate, Clarkson University

James Bruska

Dr. Lok Yan

Dr. Gidlo Torres

C. Woralert, J. Bruska, C. Liu and L. Yan, "High Frequency Performance Monitoring via Architectural Event Measurement," *2020 IEEE International Symposium on Workload Characterization (IISWC)*, Beijing, China, 2020, pp. 114-122, doi: 10.1109/IISWC50251.2020.00020.