

Protecting Real-Time GPU Kernels on Integrated CPU-GPU SoC Platforms (Artifact)*

Waqar Ali

University of Kansas
Lawrence, USA
wali@ku.edu

Heechul Yun

University of Kansas
Lawrence, USA
heechul.yun@ku.edu

Abstract

This artifact is based on **BWLOCK++**, a software framework to protect the performance of GPU kernels from co-scheduled memory intensive CPU applications in platforms containing integrated GPUs. The

artifact is designed to support the claims of the companion paper and contains instructions on how to build and execute **BWLOCK++** on a target hardware platform.

2012 ACM Subject Classification Software and its engineering → Real-time schedulability, Computer systems organization → Heterogeneous (hybrid) systems, Computer systems organization → Processors and memory architectures

Keywords and phrases GPU, memory bandwidth, resource contention, CPU throttling, fair scheduler

Digital Object Identifier 10.4230/DARTS.4.2.3

Related Article Waqar Ali and Heechul Yun, “Protecting Real-Time GPU Kernels on Integrated CPU-GPU SoC Platforms”, in Proceedings of the 30th Euromicro Conference on Real-Time Systems (ECRTS 2018), LIPIcs, Vol. 106, pp. 19:1–19:22, 2018.

<http://dx.doi.org/10.4230/LIPIcs.ECRTS.2018.19>

Related Conference 30th Euromicro Conference on Real-Time Systems (ECRTS 2018), July 3–6, 2018, Barcelona, Spain

1 Scope

The artifact is designed to support the following claims made in the companion paper:

- In integrated CPU-GPU SoC platforms, the performance of GPU kernels suffers due to co-scheduled memory intensive CPU applications
- Throttling memory intensive CPU applications negatively impacts CFS scheduling and Throttle-Fair Scheduler (TFS) of **BWLOCK++** resolves this problem
- **BWLOCK++** is effective in minimizing the interference experienced by GPU kernels due to memory intensive CPU applications
- TFS improves the overall system utilization in an overloaded system containing a mix of memory intensive and compute intensive CPU applications

* This work was partly supported by NSF CNS 1718880.



3:2 Protecting Real-Time GPU Kernels on Integrated CPU-GPU SoC Platforms (Artifact)

2 Content

The artifact package includes:

- A description file containing step-by-step instructions on how to build and run BWLOCK++
- A place-holder directory for benchmarks used in evaluation in the companion paper and a shell script for fetching the benchmarks from the web
- A place-holder directory for Linux kernel and a shell script for fetching the kernel from the web and patching it with the changes required for BWLOCK++
- A folder containing the source code of BWLOCK++ kernel module
- A folder containing shell scripts for configuring the system for BWLOCK++ and running all the experiments from the companion paper

3 Getting the artifact

The artifact endorsed by the Artifact Evaluation Committee is available free of charge on the Dagstuhl Research Online Publication Server (DROPS). In addition, the artifact is also available at: <https://github.com/wali-ku/BWLOCK-GPU/tree/ECRTS18-AE>.

4 Tested platforms

The artifact was evaluated on NVIDIA Jetson TX-2 board.

5 License

MIT (<https://opensource.org/licenses/MIT>)

6 MD5 sum of the artifact

b125fd21860f3dc8dd7b9798d16f1b3d

7 Size of the artifact

68 KiB