

PMAM 2023 Workshop Program

(26th Feb, 2023)

08:00 – 08:10 (EDT): Opening Remarks

Quan Chen, Zhiyi Huang, Min Si

08:10 – 09:00: Keynote

Keynote Speaker: Tao B. Schardl

Title: OpenCilk: Architecting a Task-Parallel Software Infrastructure for Modularity, Extensibility, and Performance

Abstract:

OpenCilk is a new open-source software infrastructure to support Cilk task-parallel programming, especially for researchers and teachers. OpenCilk provides a full-featured implementation of the Cilk and an integrated infrastructure that includes three main components: a compiler, based on the award-winning Tapir/LLVM compiler; a streamlined and efficient work-stealing runtime scheduler; and a framework for productivity-tool development, based on compiler instrumentation for fork-join parallel computations. OpenCilk’s architecture allows for substantial code reuse and easy exploration of design choices in task-parallel language abstraction, compilation strategy, runtime mechanism, and productivity-tool development.

OpenCilk is modular — modifying one component for the most part does not necessitate modifications to other components — and easy to extend — its construction naturally encourages code reuse. For example, we extended OpenCilk to support multiple additional parallel runtime systems, including Cilk Plus, OpenMP, and oneTBB. Supporting each new parallel runtime required fewer than 2000 new lines of code. At the same time, OpenCilk produces high-performing executables that achieve high work-efficiency and parallel scalability.

Biography:

Tao B. Schardl is a Research Scientist in the MIT Computer Science and Artificial Intelligence Laboratory (CSAIL) and Chief Architect of OpenCilk. His research combines algorithms and systems to develop technologies that support principled, scientific approaches to writing fast code. He has previously worked on parallel programming models, theories of performance, compilers, runtime systems, diagnostic tools, parallel algorithms, and the future of computer performance. He received the US Department of the Air Force Artificial Intelligence Accelerator Scientific Excellence Award in 2022 for his work on OpenCilk. His work on the Tapir/LLVM compiler earned the best paper award at the ACM SIGPLAN Symposium

on Principles and Practice of Parallel Programming in 2017. His work on computer performance in the post-Moore's Law era was published in Science and has been spotlighted in two Turing-award lectures. Dr. Schardl earned his S.B., M.Eng, and Ph.D. in Computer Science and Engineering from MIT.

09:00 – 10:00: Session 1 Chaired by Sid Raskar

- Jonathan D. Wapman, Serban D. Porumbescu, John D. Owens. “Harmonic CUDA: Asynchronous Programming on GPUs”.
- Baodi Shan, Mauricio Araya-Polo, Abid M. Malik, and Barbara Chapman. “MPI-based Remote OpenMP Offloading: A More Efficient and Easy-to-use Implementation”.
- Kewei Yan, Yaying Shi, and Yonghong Yan. “Exploring OpenMP GPU Offloading for Implementing Convolutional Neural Networks”.

10:00- 10:20: Morning Break

10:20 – 11:40: Session 2 Chaired by Baodi Shan

- Jay Hwan Lee, Yeonsoo Kim, Yonghyun Ryu, Wasuwee Sodsong, Hyunjun Jeon, Jinsik Park, Bernd Burgstaller, and Bernhard Scholz. “Julia Cloud Matrix Machine: Dynamic Matrix Language Acceleration on Multicore Clusters in the Cloud”.
- Yoshiki Kawanishi, Patrick Finnerty, Tomio Kamada, and Chikara Ohta. “Distributed Cell Set: A Library for Space-Dependent Communication/Computation Overlap on Manycore Cluster”.
- Siddhisanket Raskar, Thomas Appencourt, Kalyan Kumaran, and Guang Gao. “Towards Maximum Throughput of Dataflow Software Pipeline under Resource Constraints”.
- Aristeidis Mastoras, and Albert-Jan N. Yzelman. “Studying the expressiveness and performance of parallelization abstractions for linear pipelines”.

11:40: Closing Remarks

Quan Chen, Zhiyi Huang, Min Si