

革新的コンピューティング技術の開拓
2020年度採択研究代表者

2022年度
年次報告書

WAHIB MOHAMED

理化学研究所 計算科学研究センター
チームリーダー

「ムーアの法則を超えた並列プログラミング」

§ 1. 研究成果の概要

In this fiscal year we end-to-end capability of our autonomous execution model. We proposed a new method for autonomously communicating between accelerators, which are widely used in HPC scientific simulations. Our method removes the CPU from the critical path of communication in GPU-accelerated systems. By running scientific application where communication is driven by the GPUs, in a persistent way, we can reduce the communication latency to minimum (the main bottleneck in strong scaling of GPU-accelerated systems) and achieve speedups in software that are higher than any state-of-the-art production communication libraries. Additionally, the achieved results make us well-positioned to enter the final fiscal year with the clear goal of applying our autonomous execution environment on real-world application on full supercomputer scale.

【代表的な原著論文情報】

[1] Ismayil Ismayilsoy, Javid Baydamirli, Doğan Sağbili, Mohamed Wahib, Didem Unat: Multi-GPU Communication Schemes for Iterative Solvers: When CPUs are Not in Charge, ACM ICS'23

[2] Lingqi Zhang, Mohamed Wahib, Peng Chen, Jintao Meng, Xiao Wang, Toshio Endo, Satoshi Matsuoka:
Revisiting Temporal Blocking Stencil Optimizations, ACM ICS'23

[3] Lingqi Zhang, Mohamed Wahib, Peng Chen, Jintao Meng, Xiao Wang, Toshio Endo, Satoshi Matsuoka: PERKS: a Locality-Optimized Execution Model for Iterative Memory-bound GPU Applications, ACM ICS'23

[4] Matsuoka S, Domke J, Wahib M, Drozd A, Hoefler T. Myths and legends in high-performance computing. The International Journal of High Performance Computing Applications. 2023;0(0). doi:10.1177/10943420231166608

[5] Xiao Wang, Aristeidis Tsaris, Debangshu Mukherjee, Mohamed Wahib, Peng Chen, Mark Oxley, Olga Ovchinnikova, Jacob D. Hinkle:
Image Gradient Decomposition for Parallel and Memory-Efficient Ptychographic Reconstruction. SC 2022: 1-13

[6] Satoshi Matsuoka, Jens Domke, Mohamed Wahib, Aleksandr Drozd, Andrew A. Chien, Raymond Bair, Jeffrey S. Vetter, John Shalf: Preparing for the Future - Rethinking Proxy Applications. Comput. Sci. Eng. 24(2): 85-90 (2022)

[7] Anshu Dubey, Klaus Weide, Jared O'Neal, Akash Dhruv, Sean M. Couch, James Austin Harris, Tom Klosterman, Rajeev Jain, Johann Rudi, Bronson Messer, Michael Pajkos, Jared Carlson, Ran Chu, Mohamed Wahib, Saurabh Chawdhary, Paul M. Ricker, Dongwook Lee, Katie Antypas, Katherine M. Riley, Christopher S. Daley, Murali K. Ganapathy, Francis X. Timmes, Dean M. Townsley, Marcos Vanella, John Bachan, Paul M. Rich, Shravan Kumar, Eirik Endeve, W. Raphael Hix, Anthony Mezzacappa, Thomas Papatheodore: Flash-X: A multiphysics simulation software instrument. *SoftwareX* 19: 101168 (2022)

[8] Truong Thao Nguyen, François Trahay, Jens Domke, Aleksandr Drozd, Emil Vatai, Jianwei Liao, Mohamed Wahib, Balazs Gerofi: Why Globally Re-shuffle? Revisiting Data Shuffling in Large Scale Deep Learning. *IPDPS 2022*: 1085-1096