PMAM 2021 Workshop Program

(27th February, 2021)

09:00 – 09:10 (EST): Opening Remarks

Quan Chen, Zhiyi Huang, Min Si

09:10 – 10:00: Keynote

Keynote Speaker: Prof. Hailong Yang, Beihang University, China

Title: High Performance Across-stack Optimization

Abstract: High performance computing has become the driving force to push the frontier of scientific discoveries. However, to achieve desiring merits of high performance, high scalability and high portability for scientific applications, various optimizations need to be applied across stack of the high performance computing system. In this talk, I will discuss our recent work on performance optimization techniques from the aspects of numerical algorithm design, runtime system, architecture adaption and performance analysis tool. We demonstrate that, due to the increasing complexity of both hardware and software, across-stack optimization becomes inevitable to achieve the three good merits of high performance computing. Looking into the future, we believe domain specific programming model and profile-guided optimization would be essential to enter the era of exascale computing.

BIO: Hailong Yang is an associate professor at Beihang university, Beijing, China. His research interest includes high performance computing, numerical algorithm, performance analysis tool and energy efficiency. Dr. Yang published several papers in top-tier conferences and journals such as SC, ISCA, ASPLOS, ICS, ICPP, TPDS and TACO. Dr. Yang received IEEE TCSC Awards for Excellence (Early Career Researcher). Dr. Yang currently serves as architecture area co-chair for CLUSTER21. Dr. Yang also leads the student supercomputing team in Beihang university, which has won several international competition awards.

10:00- 10:15: Morning Break

10:15 – 11:30: Session 1

Session Chair: Quan Chen, Shanghai Jiao Tong University

- Sandra Catalán, Francisco D. Igual, Rafael Rodríguez-Sánchez, José R. Herrero, and Enrique S. Quintana-Ortí. "A New Generation of Task-Parallel Algorithms for Matrix Inversion in Many-Threaded CPUs".

- Weicheng Li, Rui Wang, and Depei Qian. “CompactNet: Platform-Aware Automatic Optimization for Convolutional Neural Networks”.
- Eric Raut, Jonathon Anderson, Mauricio Araya-Polo, and Jie Meng. “Porting and Evaluation of a Distributed Task-driven Stencil-based Application”.

11:30: Closing Remarks

Quan Chen, Zhiyi Huang, Min Si