Program Overview

The program consists of 1 full day of workshops (Mon, July 28) and a 3.5 days of conference activities, consisting of 3 keynotes, 42 full papers and 13 Brief Announcements.
Each full paper is allocated a 20 minute talk time, 2 of which is reserved for questions.

Each brief announcement is allocated a 14 minute talk time, 1 of which is reserved for questions.

Mon, July 28: Workshops

8am - 10:30am: **Tutorial:** Concurrent Data Structures in RDMA

10:30am - 11am: **Morning break**

11am-12:30pm: **Tutorial:** Parallelizing Sequential Iterative Algorithms

12:30pm-2pm: **Lunch** break (on your own)

2pm - 3:30pm: **Tutorial:** Parallel Clustering

3:30pm - 4pm: **Afternoon break**

4pm - 7pm: **Workshop:** Highlights of Parallel Computing (HOPC)

7pm - 10pm: **Reception** + HOPC poster session

Tue, July 29: Conference Day 1

Opening Remarks: 8:30 - 8:45

Session 1: 8:45am - 10:30am

8:45 - 9:05 *Parallel Minimum-Cost Flow in Near-Linear Work and Square Root Depth for Dense Instances*
J. van den Brand (GT), H. Gholizadeh (KIT), Y. Jiang (MPI-INF), T. de Vos (TUG)

9:06 - 9:26
Applying Fast Fourier Transforms to Accelerate Spatially and Temporally Inhomogeneous Stencil Computations
R. Bentley (SBU), R. Chowdhury (SBU), A. Gregory (SBU), M. Santomauro (SBU), A. Gregory (SBU)

9:27 - 9:47
Towards Zero Spawn Overhead: Work Stealing Without Deques
A. Handleman (WashU), K. Singer (MIT), T.B. Schardl (MIT), I.A. Lee (WashU)

9:48 - 10:08
A Deterministic Work-Depth Tradeoff for Strongly Connected Components
A. Karczmarz (MIM UW), B. Lewandowski (MIM UW)

10:09 - 10:29 *ABD-HFL: Byzantine-resistant Decentralized Hierarchical Federated Learning*
T. An (Inria), S. Fdida (SU), M. Potop-Butucaru (SU, CNRS, LIP6), S. Tixeuil (SU), T. An (Inria, IMT-Atlantique)

Break: 10:30am - 11am

Keynote 1: 11am - 12:15pm

SPAA Parallel Computing Award: Uzi Vishkin (UMD)

Lunch 12:15pm - 2pm

Session 2: 2pm - 3:30pm

2:05 - 2:25
Approximate Agreement Algorithms for Byzantine Collaborative Learning
M. Cambus (Aalto), D. Melnyk (TUBerlin), T. Milentijevifá (TUBerlin), S. Schmid (TUBerlin)

2:26 - 2:46
Dataflow-Specific Algorithms for Resource-Constrained Scheduling and Memory Design
A. Bhattacharjee (Yale), Q.C. Liu (Yale), R. Manohar (Yale), R.P. Pothukuchi (Yale), M. Ugur (Yale), Q.C. Liu (UCB)

2:47 - 3:07
TLF: Transactional Lock Fusion
G. Blelloch (CMU), Z. Kent (CMU), Y. Wei (UBC)

3:08 - 3:28
Approximation Hardness of Resource Scheduling
R. Das (UH), H. Sun (UH)

Break 3:30pm - 4pm

Session 3: 4pm - 5:45

4:00 - 4:20 *Minimizing Processor Count for Fault Tolerant Toom-Cook Algorithms*
R. Nissim (HUJI), O. Schwartz (HUJI), Y. Spiizer (TAU)

4:21 - 4:41 *Parallel Point-to-Point Shortest Paths and Batch Queries*
X. Dong (UCR), A. Li (UCR), Y. Gu (UCR), Y. Sun (UCR)

4:42 - 5:02
Eliminating Bank Conflicts in GPU Mergesort

K. Berney (UH), N. Sitchinava (UH)

5:03 - 5:23
Managing High-Bandwidth Memory is a Parallel Scheduling Problem K. Agrawal (WashU), M.A. Bender (SBU), B. Moseley (CMU), K. Pruhs (Pitt), C. Stein (Columbia)

5:24 - 5:44
Optimal quantum sampling on distributed databases
L. Chen (QICI HKU; SKL-NST), J. Liu (SKL-NST), P. Yao (SKL-NST; HFNL), L. Chen (SKL-NST)

Business Meeting: 6-9pm

Wed, July 30: Conference Day 2

Session 4: 8:45am - 10:30am

8:45 - 9:05
ReFINE: A Reactive and Fine-Grained Scheduling Framework For Concurrency on General Purpose GPUs
G. Gilman (WPI), R.J. Walls (WPI)

9:06 - 9:26
A New Algorithm for Online Scheduling of Rigid Task Graphs with Near-Optimal Competitive Ratio
L. Perotin (VU), H. Sun (KU), P. Raghavan (VU)

9:27 - 9:47
Parallel Batch-Dynamic Coreness Decomposition with Worst-Case Guarantees
M. Ghaffari (MIT), J. Koo (MIT)

9:48 - 10:08
Tight Bounds for Online Balanced Partitioning in the Generalized Learning Model
H. Räcke (TUM), S. Schmid (TU Berlin), R. Zabrodin (TUM)

10:09 - 10:29
Decoupled Fallback: A Portable Single-Pass GPU Scan
T. Smith (Google), R. Levien (Google), J.D. Owens (UCD)

Break 10:30am - 11am

Keynote 2: 11am - 12:15pm

Michael Garland, NVIDIA

Lunch 12:15pm - 2pm

Session 5: 2pm - 3:30pm

2:05 - 2:25
POPQC: Parallel Optimization for Quantum Circuits
P. Liu (CMU), J. Arora (CMU), M. Xu (CMU), U.A. Acar (CMU)

2:26 - 2:46
Listening Efficient Contention Resolution for Semi-batch Arrivals without Collision Detection Y. Xie (NJU), C. Zheng (NJU)

2:47 - 3:07
Parallel Batch-Dynamic Algorithms for Spanners, and Extensions
M. Ghaffari (MIT), J. Koo (MIT)

3:08 - 3:28
Reconciling Hardware Transactional Memory and Persistent Programming with Buffered Durability
M. Du (UR), Z. Su (University or Rochester), M. Scott (UR)

Break 3:30pm - 4pm

Session 6: 4pm - 5:45

4:00 - 4:20
CLEANN: Lock-Free Augmented Trees for Low-Dimensional k-Nearest Neighbor Search
M.D. Manohar (CMU & Microsoft), Y. Wei (UBC), G. Blelloch (CMU), M. Manohar (CMU)

4:21 - 4:41
Faster MPC Algorithms for Approximate Allocation in Uniformly Sparse Graphs
J. Łącki (Google), S. Mitrović (UCD), S. Ramachandran (UCD), W. Sheu (UCD)

4:42 - 5:02
Optimal Batch-Dynamic kd-trees for Processing-in-Memory with Applications
Y. Zhao (CMU), H. Kang (THU), Y. Gu (UCR), G.E. Blelloch (CMU), L. Dhulipala (UMD), C. McGuffey (Reed), P.B. Gibbons (CMU)

5:03 - 5:23
Dispersion is (Almost) Optimal under (A)synchrony
A. Kshemkalyani (UIC), M. Kumar (IITM), A.R. Molla (ISI), D. Pattanayak (uOttawa), G. Sharma (KSU)

5:24 - 5:44
A Study of Synchronization Methods for Concurrent Size
H. Kas-Sharir (Technion), G. Sela (EPFL), E. Petrank (Technion)

Banquet Dinner 7pm - 10pm

Thu, July 31: Conference Day 3

Session 7: 8:45am - 10:30am

8:45 - 9:05
Fully-Dynamic Parallel Algorithms for Single-Linkage Clustering
Q. De Man (UMD), L. Dhulipala (UMD), K.N. Gowda (UMD)

9:06 - 9:26
Persistent HyTM via Fast Path Fine-Grained Locking
G. Coccimiglio (UWaterloo), T. Brown (UWaterloo), S. Ravi (USC)

9:27 - 9:47
Parallel Batch-Dynamic Maximal Matching with Constant Work per Update
G. E. Blelloch (CMU), A. C. Brady (CMU)

9:48 - 10:08
A framework for boosting matching approximation: parallel, distributed, and dynamic
S. Mitrovifá (UCD), W. Sheu (UCD)

10:09 - 10:29 *Non-preemptive Throughput Maximization under Time-varying Capacity* A. Murhekar (UIUC), M. Purohit (Google), Z. Svitkina (Google), E. Vee (Google), J. Wang (Google)

Break 10:30am - 11am

Session 8: 11am - 12:15pm

11:00 - 11:14
Brief Announcement: Optimality Conditions for Parallel Communication-Avoiding Matrix Multiplication with Overlapped Communication
M. Isaev (Nvidia), S. Eswar (Argonne), R. Vuduc (GT)

11:15 - 11:29 *Brief Announcement: Delta-Snap: Snapshotting the Differential*
P. Jayanti (Dartmouth), S. Jayanti (Dartmouth)

11:30 - 11:44
Brief Announcement: Distributed Butterfly Analysis using Mobile Agents
P.K. Chand (ISI), A. Das (BITS Pilani), A.R. Molla (ISI)

11:45 - 11:59
Brief Announcement: Algorithms for Distance Sensitivity Oracles on the PRAM
V. Manoharan (UTAUS), V. Ramachandran (UTAUS)

12:00 - 12:14
Brief Announcement: Improved Understanding of Landlord Under Suffix Analysis
C. McGuffey (Reed), B. Schuijff (Reed), A. Snelling (Reed)

Lunch 12:15pm - 2pm

Session 9: 2pm - 3:30pm

2:05 - 2:25
Parallel Batch Queries on Dynamic Trees: Algorithms and Experiments
H. Ikram (CMU), A. Brady (CMU), D. Anderson (CMU), G. Blelloch (CMU)

2:26 - 2:46
Formal Machine-Verification of MemSnap: An Efficient, Far-Future Linearizable Snapshot Algorithm
S. Jayanti (Dartmouth), U.Y. Yavuz (BU), S. Jayanti (Dartmouth)

2:47 - 3:07
Don't Melt Your Cache: Low-Associativity with Heat-Sink
M. Bender (SBU), A. Conway (CTech), D. DeLayo (SBU), M. Farach-Colton (NYU), J. Han (UNC-CH), L. He (Rutgers), R. Johnson (VMWare), S. Kannan (Rutgers), W. Kuszmaul (CMU), D. Porter (UNC-CH), E. West (SBU)

3:08 - 3:28 *Bounded Memory in Distributed Networks*
R.B. Basat (UCL), K. Censor-Hillel (Technion), Y. Chang (NUS), W. Han (UCL), D. Leitersdorf (Technion), G. Schwartzman (JAIST)

Break 3:30pm - 4pm

Session 10: 4pm - 5:45

4:00 - 4:20 *Scaling Parallel Algorithms to Massive Datasets using Multi-SSD Machines*
H. Li (UMD), J. Khan (UMD), L. Dhulipala (UMD)

4:21 - 4:41
RACE: Operator Choreography for Inference Acceleration in Personalized Recommender System
S. Li (SKL-NST), X. Yang (SKL-NST), J. Liu (SKL-NST), Q. Qi (SKL-NST), J. Liao (SKL-NST), J. Huang (Meituan), J. Wang (SKL-NST)

4:42 - 5:02
JOVS: Joint Optimization of Vectorization and Scheduling for DNN on AI DSPs
Y. Han (BUAA), H. Jiang (BUAA), R. Zhang (BUAA), R. She (BUAA)

5:03 - 5:23
The Impact of Partial Computations on the Red-Blue Pebble Game
P. Papp (ZRC), A. Sobczyk (ZRC), A.N. Yzelman (ZRC)

Friday, August 1: Conference Day 4

Session 11: 8:45am - 10:30am: Brief Announcements I

8:45 - 8:59
Brief Announcement: A Novel Integrated Parallel Accelerator for an Irregular Killer App
Y. Song (UMD), M. Franklin (UMD), U. Vishkin (UMD)

9:00 - 9:14 *Brief Announcement: Fine-Grained Computation in 3-Space: Matrix Multiplication, Paths, and Mazes*
Q. Stout (U-M)

9:15 - 9:29
Brief Announcement: SCOT: Fix non-blocking data structures, not memory reclamation
M. A. H. Arovi (PSU), R. Nikolaev (PSU)

9:30 - 9:44
Brief Announcement: Fully-Distributed Construction of Byzantine-Resilient Dynamic Peer-to-Peer Networks
A. Gupta (UH), G. Pandurangan (UH)

9:45 - 9:59 *Brief Announcement: Load Balancing with Duration Predictions*
Y. Azar (TAU), N. Buchbinder (TAU), T. Epshtein (TAU)

10:00 - 10:14
Brief Announcement: Stochastic Parallel Scheduling with Bandit Feedback
G. Benadè (BU), R. Das (UH), T. Lavastida (UTD)

10:15 - 10:29
Brief Announcement: Minimizing Communication for Parallel Symmetric Tensor Times Same Vector Computation
H. Al Daas (RAL), G. Ballard (WFU), L. Grigori (EPFL; PSI), S. Kumar (Inria), K. Rouse (Inmar), M. Verite (EPFL)

Break 10:30am - 11am

Session 12: 11am - 11:30pm: Brief Announcements II

11:00 - 11:14 *Brief Announcement: LCTree: A Fast Hardware BVH Constructor for Real-Time Ray Tracing*
R. Yan (NUDT), S. Yin (NUDT), H. Guo (NUDT), Y. Wang (NUDT), G. Chen (SYSU), N. Xiao (NUDT), L. Huang (NUDT)

11:15 - 11:29
Brief Announcement: Accelerating Distributed Search System with a In-network Computation
P. Zhang (SXU)