

SPAA 2022 Tentative Schedule

Monday July 11th

09:00 - 10:30	Tutorial: ParlayLib: Fast and Scalable Parallelism for Everyone
10:30 - 12:00	Tutorial: Recent Progress in Practical State-Machine Replication Systems: Geo-Replication, Randomization, RDMA in parallel with
09:00 - 12:00	Tutorial: Fast Fourier and Gauss Transforms as Fundamental Components of Fast Stencil Computations
12:00 - 14:00	Lunch Break
14:00 - 17:00	Workshop on Large-Scale Parallel Graph Processing
18:30 - 20:30	Welcome Reception

Tuesday July 12th

08:45 - 08:50	Opening Remarks
Session 1: Distributed Computing on Graphs, Tim Kaler	
08:50 - 09:13	Deterministic Distributed Sparse and Ultra-Sparse Spanners and Connectivity Certificates
09:13 - 09:36	Fully Polynomial-Time Distributed Computation in Low-Treewidth Graphs
09:36 - 09:59	Adaptive Massively Parallel Algorithms for Cut Problems
09:59 - 10:22	Massively Parallel Algorithms for b-Matching
10:22 - 10:45	Average Awake Complexity of MIS and Matching
10:45 - 10:52	Brief Announcement: Distributed Lightweight Spanner Construction for Unit Ball Graphs in Doubling Metrics
10:52 - 11:20	AM Coffee Break
Keynote Talk I, Chair: I-Ting Angelina Lee	
11:20 - 12:20	Keynote Talk: Large Scale Parallel Sparse Matrix Streaming Graph/Network Analysis by Jeremy Kepner, MIT Lincoln Laboratory Supercomputing Center
12:20 - 13:50	Lunch Break
Session 2: Distributed Algorithms, Chair: Guy Blelloch	
13:50 - 14:13	Preparing for Disaster: Leveraging Precomputation to Efficiently Repair Graph Structures Upon Failures
14:13 - 14:36	The Energy Complexity of Las Vegas Leader Election
14:36 - 14:59	A Fully-Distributed Peer-to-Peer Protocol for Byzantine-Resilient Distributed Hash Tables
14:59 - 15:06	Brief Announcement: The (Limited) Power of Multiple Identities: Asynchronous Byzantine Reliable Broadcast with Improved Resilience through Collusion
15:06 - 15:13	Brief Announcement: Composable dynamic secure emulation
15:13 - 15:40	AM Coffee Break
Session 3: Networks and Communications, Chair: Phillip B. Gibbons	
15:40 - 16:03	Robust and Optimal Contention Resolution without Collision Detection
16:03 - 16:26	Contention Resolution for Coded Radio Networks
16:26 - 16:49	Achieving Sublinear Complexity under Constant T in T-interval Dynamic Networks
16:49 - 16:56	Brief Announcement: Fast(er) Construction of Round-optimal n-Block Broadcast Schedules
Session 4: Cache and Memory, Chair: Yan Gu	
17:00 - 17:23	Automatic HBM Management: Models and Algorithms
17:23 - 17:46	Competitive Algorithms for Block-Aware Caching
17:46 - 17:53	Brief Announcement: Spatial Locality and Granularity Change in Caching
18:00 - 19:00	Business Meeting

Wednesday July 13th

Session 5: Best Paper Session, Chair: Tao B. Schardl

09:00 - 09:23	Parallel Shortest Paths with Negative Edge Weights
09:23 - 09:46	Parallel Batch-Dynamic Algorithms for k-Core Decomposition and Related Graph Problems
09:46 - 10:09	Online Parallel Paging with Optimal Makespan
10:09 - 10:32	PREP-UC: A Practical Replicated Persistent Universal Construction
10:32 - 11:00	AM Coffee Break

Keynote Talk II, Chair: I-Ting Angelina Lee

11:00 - 12:00 Keynote Talk: Algorithm Improvement: How Fast Has It Been and How Much Farther Can It Go?
by Neil Thompson, Massachusetts Institute of Technology

12:00 - 13:30 Lunch Break

Session 6: Parallel Algorithms and Data Structures, Chair: Laxman Dhulipala

13:30 - 13:53 Parallel Batch-Dynamic Minimum Spanning Forest and the Efficiency of Dynamic Agglomerative Graph Clustering

13:53 - 14:16 Scalable Fine-Grained Parallel Cycle Enumeration Algorithms

14:16 - 14:39 Parallel Cover Trees and their Applications

14:39 - 15:02 Many Sequential Iterative Algorithms Can Be Parallel and (Nearly) Work-efficient

15:02 - 15:09 Brief Announcement: A Parallel (Δ, T) -Stepping Algorithm for the Constrained Shortest Path Problem

15:09 - 15:16 Brief Announcement: Faster Stencil Computations using Gaussian Approximations

15:16 - 15:45 PM Coffee Break

Session 7: Concurrency and Synchronization, Chair: Julian Shun

15:45 - 16:08 NUMA-Aware Recoverable Mutex Lock

16:08 - 16:31 wCQ: A Fast Wait-Free Queue with Bounded Memory Usage

16:31 - 16:54 HybriDS: Cache-Conscious Concurrent Data Structures for Near-Memory Processing Architectures

16:54 - 17:17 Performance Analysis and Modelling of Concurrent Multi-access Data Structures

Session 8: Scheduling, Chair: Kunal Agrawal

17:20 - 17:43 The k-Server with Preferences Problem

17:43 - 18:06 Permutation Predictions for Non-Clairvoyant Scheduling

18:06 - 18:29 Balancing Flow Time and Energy Consumption

18:29 - 18:36 Brief Announcement: Nested Active-Time Scheduling

18:36 - 18:43 Brief Announcement: Towards a More Robust Algorithm for Flow Time Scheduling with Predictions

Thursday July 14**Session 9: More Scheduling, Chair: Rathish Das**

09:00 - 09:23 Balanced Allocations in Batches: Simplified and Generalized

09:23 - 09:46 Approximate Dynamic Balanced Graph Partitioning

09:46 - 10:09 Bamboo Trimming Revisited: Simple Algorithms Can Do Well Too

10:09 - 10:16 Brief Announcement: Tight Bounds for Repeated Balls-into-Bins

10:16 - 10:45 AM Coffee Break

Session 10: Matrix-Based Algorithms and I/O Bounds, Chair: Rezaul A. Chowdhury

10:45 - 11:08 I/O-Optimal Algorithms for Symmetric Linear Algebra Kernels

11:08 - 11:31 Sparse matrix multiplication in the low-bandwidth model

11:31 - 11:38 Brief Announcement: Tight Memory-Independent Parallel Matrix Multiplication Communication Lower Bounds

Brief Announcement: On the I/O complexity for Hybrid Integer Multiplication Algorithms in the Sequential and Parallel Distributed Memory Model