Contributed Presentations

Title: Optimal Speedup on a Low-Degree Multi-Core Parallel Architecture (LoPRAM)

Reza Dorrigiv, Alejandro López-Ortiz, and Alejandro Salinger, University of Waterloo

Abstract: Over the last five years, major microprocessor manufacturers have released plans for a rapidly increasing number of cores per microprossesor, with upwards of 64 cores by 2015. In this setting, a sequential RAM computer will no longer accurately reflect the architecture on which algorithms are being executed. In this paper we propose a model of low degree parallelism (LoPRAM) which builds upon the RAM and PRAM models yet better reflects recent advances in parallel (multi-core) architectures. This model supports a high level of abstraction that simplifies the design and analysis of parallel programs. More importantly we show that in many instances it naturally leads to work-optimal parallel algorithms via simple modifications to sequential algorithms.

Title: Algorithm design for multicore processors, Theory and applications

Peter Krusche and Alexander Tiskin, The University of Warwick.

Abstract: In this extended abstract, we outline a simple approach to designing parallel algorithms for modern multi-core processors. Our goal is to show a general method for specifying and characterizing parallel algorithms at a high level, while still being able to assess their practicality and potential for speedup. Our work is motivated by a few applications in string comparison algorithms, which is a field that benefits naturally from parallel computation since many of the algorithms run in polynomial time, but have large input sizes in practice. We show in this abstract that, even on a simple model, nontrivial theoretical results can be obtained, and point out a few open problems.

Title: A Model of Computation for Map Reduce

Howard Karloff, AT&T Labs, Siddharth Suri , and Sergei Vassilvitskii, Yahoo! Research

Abstract: In recent years the Map Reduce framework has emerged as a crucial tool for processing large amounts of data. Used daily at companies such Yahoo!, Google, and Facebook and adopted more recently by several universities, it allows easy parallelization of data intensive computations. In this work we propose a new computational class for the Map Reduce paradigm. The salient feature of the new model is a substantially sublinear limit on the total number of machines and the memory per machine. Also, we allow each machine time polynomial in the size of its input. (This is a key difference between our model and a modeling of this paradigm by Feldman et al. [2].) We define our model and show initial examples of its power.

Title: Parallel Phase Model for Manycore Parallel Machines

Zhaofang Wen, Sandia National Labs, and Junfeng Wu, Syracuse University

Title: Parallel External Memory Model

Lars Arge, University of Aarhus, Michael Goodrich, and Nodari Sitchinava, University of California – Irvine

Abstract: We would like to discuss the Parallel External Memory (PEM) model and several algorithmic results in this model. We provide solutions to the fundamental parallel problems ranging from all-prefix-sums to sorting. We also present solutions to several problems on trees and graphs, such as tree contraction and expression tree evaluation, batched lowest common ancestors, as well as finding minimum spanning tree, connected and bi-connected components, and ear decomposition of a bi-connected graph. All solutions provide asymptotically optimal linear speedup in the number of processors. They remain scalable for up to (almost) linear number processors.