

Application

Physics/Biology Apps

Commercial Apps

Social media analysis

Programming

Declarative/Expressive Lang: Python, C++, Java, Scala, Apache pig Spark SQL..

High-level API

Parallel Programming: TensorFlow, Spark, Hadoop...

Mid-level API

Distributed Processing: MPI, RPC, RDMA ...

Low-level API

Parallel Processing: Eigen, OpenMP, OpenCL, Vector-Compiler, SIMD...

Hardware


CPU/GPU/ASIC, Memory, NVM, Disks – poss. virtualised

OPTIMISATION



GPU cluster

Cluster (Amazon)



HPC



SSD



FPGA, TPU

09:45 - 10:30 Applications (High Performance Data Analytics)

09:45 - 10:15 Juha Jaykka (Univ. Cambridge COSMOS Intel Parallel Computing Centre): Advantages and disadvantages of modern parallel computing

10:15 - 10:30 Discussion

10:30 - 11:15 System Architecture I (HPC)

10:30 - 11:00 Christophe Dubach (Univ. Edinburgh): Lift: a Data-Parallel Language for High-Performance Parallel Pattern Code Generation

11:00 - 11:15 Discussion

11:15 - 11:30 Coffee Break

11:30 - 12:15 System Architecture II (Scheduling)

11:30 - 12:00 Daniel Goodman (Oracle Labs): Fine-grained parallel work scheduling in scale-up graph analytics

12:00 - 12:15 Discussion

12:15 - 13:00 Lunch

13:00 - 14:30 System Architecture III (Distributed Computing/Memory Management)

13:00 - 13:30 Aleksandar Dragojevic (Microsoft Research Cambridge): FaRM: a platform for low-latency computing

13:30 - 13:45 Sam Ainsworth (Univ. Cambridge): Graph Prefetching Using Data Structure Knowledge

13:45 - 14:15 Rajeev Raman (Univ. Leicester): In-memory memory processing of big data via succinct data structures

14:15 - 14:30 Discussion

14:30 - 14:45 Coffee Break

14:45 - 16:00 Heterogeneous Cores (Optimisation in Stream Processing and Neural Networks)

14:45 - 15:10 Alexandros Koliouisis (Imperial College London): SABER: Window-Based Hybrid Stream Processing for Heterogeneous Architectures

15:10 - 15:30 Valentin Dalibard (Univ. Cambridge): Modern Systems for Neural Networks

15:30 - 16:00 Discussion

16:00 - 16:30 Discussion of future vision of data processing stack from hardware, low-level programming, parallel programming platform, to applications

16:30 - 17:30 Closing + wine