

## R244: Paper Review Presentation Assignment

### 2017/10/17 Session 2: Data flow programming: Map/Reduce to TensorFlow

- sl829 (Stafanos) short presentation 4.** J. Dean, S. Ghemawat: [MapReduce: Simplified Data Processing on Large Clusters](#), OSDI, 2004.
- ld539 (Łukasz) 2.** M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. Franklin, S. Shenker, I. Stoica: [Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing](#), NSDI, 2013.
- ib354 (Ioana) 5.** Derek Murray, Malte Schwarzkopf, Christopher Snowton, Steven Smith, Anil Madhavapeddy and Steven Hand: [Ciel: a universal execution engine for distributed data-flow computing](#), NSDI 2011.
- jlm95 (Jesse) 6.2.** D. Murray, F. McSherry, R. Isaacs, M. Isard, et al.: [Naiad: A Timely Dataflow System](#), SOSP, 2013.
- nm583 (Nat) 8.** M. Abadi et al. [Tensorflow: A system for large-scale machine learning](#). OSDI, 2016.

### 2017/10/24 Session 3: Large-scale graph data processing: storage, processing model and parallel processing

- gaw38(George) 1.** G. Malewicz, M. Austern, A. Bik, J. Dehnert, I. Horn, N. Leiser, and G. Czajkowski: [Pregel: A System for Large-Scale Graph Processing](#), SIGMOD, 2010.
- js2173(Josh) 4.** J. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin: [Powergraph: distributed graph-parallel computation on natural graphs](#). OSDI, 2012.
- tgmp2 (TomP) 5.** J. Shun and G. Blelloch: [Ligra: A Lightweight Graph Processing Framework for Shared Memory](#), PPOPP, 2013.
- ag981(Alex) Short Presentation 11.** F. McSherry, M. Isard and D. Murray: [Scalability! But at what COST?](#) , HOTOS, 2015.
- sl715 (Stella) 9.** A. Roy, I. Mihailovic, W. Zwaenepoel: [X-Stream: Edge-Centric Graph Processing using Streaming Partitions](#), SOSP, 2013.
- tpt26 (Tudor) 14.** S. Hong, H. Chafi, E. Sedlar, K.Olukotun: [Green-Marl: A DSL for Easy and Efficient Graph Analysis](#), ASPLOS, 2012.

### 2017/11/7 Session 5: Stream Data Processing and Data/Query Model

- twb34 (TomB) 7.** T. Akidau, R. Bradshaw, C. Chambers, S. Chernyak, R. Fernandez-Moctezuma, R. Lax, S. McVeety, D. Mills, F. Perry, E. Schmidt, S. Whittle: [The Dataflow Model: A Practical Approach to Balancing Correctness, Latency, and Cost in Massive-Scale, Unbounded, Out-of-Order Data Processing](#), VLDB, 2015.

...to read the above paper, you might need to check also the following paper:

- 1.** [T. Akidau](#), A. Balikov, K. Bekiroglu, [S. Chernyak](#), J. Haberman, R. Lax, [S. McVeety](#), D. Mills, [P. Nordstrom](#), [S. Whittle](#): [MillWheel: Fault-Tolerant Stream Processing at Internet Scale](#), VLDB, 2013.
- ys485 (Yashovardhan) 9.** A. Floratou et al.: [Dhalion: self-regulating stream processing in Heron](#), VLDB, 2017.

## 2017/11/14 Session 6: Machine Learning for Optimisation of Computer Systems

### 1) Optimisation space, Multi-objective function, Manual optimisation

**tpt26(Tudor) Short Presentation 6.** B. Bodin, L. Nardi, MZ Zia et al.: [Integrating Algorithmic Parameters into Benchmarking and Design Space Exploration in 3D Scene Understanding](#), PACT, 2016.

**ag981(Alex) Short Presentation 13.** N. Lane et al.: [DeepX : A Software Accelerator for Low-Power Deep Learning Inference on Mobile Devices](#), IPSN, 2016.

### 2) Bayesian optimisation

**jlm95(Jesse) 4.** V. Dalibard, M. Schaarschmidt, and E. Yoneki: [BOAT: Building Auto-Tuners with Structured Bayesian Optimization](#), WWW, 2017.

### 3) Reinforcement Learning

**nm583(Nat) 11.** G. Tesauro et al.: [A Hybrid Reinforcement Learning Approach to Autonomic Resource Allocation](#), ICAC, 2006.

**tgmp2(ThomasP) Short Presentation 8.** V. Mnih et al.: [Playing Atari with Deep Reinforcement Learning](#), NIPS, 2013.

**js2173(Josh) 12.** A. Valadarsky et al.: [A Machine Learning Approach to Routing](#), arXiv, 2017.

## 2017/11/21 Session 7: Task scheduling, Performance, and Resource Optimisation

### 1) Scheduling/Resource Management

**ld539(Lukasz) Short Presentation 4.** I. Gog, M. Schwarzkopf, A. Gleave, R. Watson, S. Hand: [Firmament: fast, centralized cluster scheduling at scale](#), OSDI, 2016.

**ys485(Yashovardhan) Short Presentation 23.** H. Mao et al.: [Resource Management with Deep Reinforcement Learning](#), HotNets, 2016.

**sl715(Stella) 1.** A. Mirhoseini et al.: [Device Placement Optimization with Reinforcement Learning](#), ICML, 2017.

### 2) Split Computation/Trajectory

**sl829(Stefanos) 18.** Y. Kang et al.: [Neurosurgeon: Collaborative Intelligence between the Cloud and Mobile Edge](#), ASPLOS, 2017.

**twb34(ThomasB) Short Presentation 24.** M. Raghu et al.: [On the Expressive Power of Deep Neural Networks](#), PMLR, 2017.

### 3) Database Optimisation

**gaw38(George) Short Presentation 26.** A. Pavlo et al.: [Self-Driving Database Management Systems](#), CIDR, 2017.

**lb354(Ioana) 25.** D. Aken et al.: [Automatic Database Management System Tuning Through Large-scale Machine Learning](#), SIGMOD, 2017.