

## R244: Paper Review Presentation Assignment

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

**tk532 (Tejas) 5.** Derek Murray, Malte Schwarzkopf, Christopher Smowton, Steven Smith, Anil Madhavapeddy and Steven Hand: [Ciel: a universal execution engine for distributed data-flow computing](#), NSDI 2011.

**ijdo2 (Indigo) 6.2.** D. Murray, F. McSherry, R. Isaacs, M. Isard, et al.: [Naiad: A Timely Dataflow System](#), SOSP, 2013.

**acs207 (Aaron) 8.** M. Abadi et al. [Tensorflow: A system for large-scale machine learning](#). OSDI, 2016.

**You can also read: 10.** M. Abadi, M. Isard and D. Murray: [A Computational Model for TensorFlow - An Introduction](#), MAPL, 2017

**sat62 (Shyam) 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.

**dt475 (Devin) 12.** R. Nishihara, P. Moritz, et al.: [Ray: A Distributed Framework for Emerging AI Applications](#), OSDI, 2018.

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

**vrs26(Vikash) 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.

**You can read the following paper to add up the specific part of Powergraph model to Pregel presentation. 4.** J. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin: [Powergraph: distributed graph-parallel computation on natural graphs](#). OSDI, 2012.

**dk525(Dmitry) 5.** J. Shun and G. Blelloch: [Ligra: A Lightweight Graph Processing Framework for Shared Memory](#), PPOPP, 2013.

**ms2518(Marek) 9.** A. Roy, I. Mihailovic, W. Zwaenepoel: [X-Stream: Edge-Centric Graph Processing using Streaming Partitions](#), SOSP, 2013.

**cit27 (Michael) 14.** S. Hong, H. Chafi, E. Sedlar, K.Olukotun: [Green-Marl: A DSL for Easy and Efficient Graph Analysis](#), ASPLOS, 2012.

**sa894(Sami) 18.** A. Gharaibeh, E. Santos-Neto, L. Costa, M. Ripeanu: [Efficient Large-Scale Graph Processing on Hybrid CPU and GPU Systems](#), IEEE TPC, 2014.

### 2017/10/31 Session 4: Stream Data Processing and Data/Query Model

**cb2015 (Cristian) 9.** A. Floratou et al.: [Dhalion: self-regulating stream processing in Heron](#), VLDB, 2017.

**tk532 (Tejas) 11.** D. O'Keefe, T. Salonidis, and P. Pietzuch: [Frontier: Resilient Edge Processing for the Internet of Things](#), VLDB, 2018.

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

**sa894(Sami) 7.** J. Ansel et al. [Petabricks: A language and compiler for algorithmic choice](#). In Proceedings of the 2009 ACM SIGPLAN Conference on Programming Language Design and Implementation, PLDI, 2009.

**ijdo2 (Indigo) 4.** V. Dalibard, M. Schaarschmidt, and E. Yoneki: [BOAT: Building Auto-Tuners with Structured Bayesian Optimization](#), WWW, 2017.

**cit27 (Michael) 17.** T. Kraska, A. Beutel, E. Chi, and J. Dean: [The Case for Learned Index Structures](#), SIGMOD, 2018.

**dt475 (Devin) 20.** M. Jaderberg, V. Dalibard, S. Osindero, W.M. Czarnecki: [Population based training of neural networks](#), arXiv, 2017.

**vrs26(Vikash) 24.** E. Liang et al.: [RLlib: Abstractions for Distributed Reinforcement Learning](#), ICML, 2018.

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

**dk525(Dmitry) 5.** O. Alipourfard et al.: [CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics](#), NSDI, 2017.

**cb2015 (Cristian) 29.** Z. Jia, M. Zaharia, and A. Aiken: [Beyond Data and Model Parallelism for Deep Neural Networks](#), ArXiv, 2018.

**acs207 (Aaron) 31.** T. Chen, T. Moreau, Z. Jiang, L. Zheng, S. Jiao, E. Yan, H. Shen, M. Cowan, L. Wang, Y. Hu, L. Ceze, C. Guestrin, and A. Krishnamurthy: [TVM: An Automated End-to-End Optimizing Compiler for Deep Learning](#), OSDI, 2018.

**sat62 (Shyam) 1.** A. Mirhoseini et al.: [Device Placement Optimization with Reinforcement Learning](#), ICML, 2017.

**ms2518(Marek) 33.** A. Ratner, S. Bach, H. Ehrenberg, J. Fries, S. Wu, and C. Ré: [Snorkel: Rapid Training Data Creation with Weak Supervision](#), VLDB, 2017.