

R244: Paper Review Presentation Assignment

2019/10/22 Session 2: Data flow programming: Map/Reduce to TensorFlow

mcic31 (Claire) 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.

sti20 (Stefan) 6.2. D. Murray, F. McSherry, R. Isaacs, M. Isard, et al.: [Naiad: A Timely Dataflow System](#), SOSP, 2013.

htb25 (Harrison) 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

ne274 (Niall) 3. J. Gjengset, M. Schwarzkopf, J. Behrens, L. T. Araujo, M. Ek, E. Kohler, M. F. Kaashoek and R. Morris: [Noria: dynamic, partially-stateful data-flow for high-performance web applications](#), OSDI 2018.

ahb36 (Harri) 12. R. Nishihara, P. Moritz, et al.: [Ray: A Distributed Framework for Emerging AI Applications](#), OSDI, 2018.

gw412 (Grzegorz) 13. M. Schaarschmidt, S. Mika, K. Fricke, E. Yoneki: [RLgraph: Flexible Computation Graphs for Deep Reinforcement Learning](#), SysML, 2019.

2019/10/29 Session 3: Large-scale graph data processing

doaa2 (David) 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.

jsl71 (Julius) J. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin: [Powergraph: distributed graph-parallel computation on natural graphs](#). OSDI, 2012.

arb229 (Ari) J. Shun and G. Blelloch: [Ligra: A Lightweight Graph Processing Framework for Shared Memory](#), PPOPP, 2013.

tal42 (Tim) A. Roy, I. Mihailovic, W. Zwaenepoel: [X-Stream: Edge-Centric Graph Processing using Streaming Partitions](#), SOSP, 2013.

oh260 (Ollie) Z. Jia, Y. Kwon, G. Shipman, P. McCormick, M. Erez, A. Aiken: [A Distributed Multi-GPU System for Fast Graph Processing](#), VLDB, 2018.

cnfo2 (Nnaemeka) S. Hong, H. Chafi, E. Sedlar, K. Olukotun: [Green-Marl: A DSL for Easy and Efficient Graph Analysis](#), ASPLOS, 2012.

2019/11/12 Session 5: Probabilistic Programming

htb25 (Harrison) V. Dalibard, M. Schaarschmidt, and E. Yoneki: [BOAT: Building Auto-Tuners with Structured Bayesian Optimization](#), WWW, 2017.

oh260 (Ollie) J. Ai et al.: [HackPPL: a universal probabilistic programming language](#), MAPL, 2019.

2019/11/19 Session 6: Optimisation of Computer Systems and ML 1

ahb36 (Harri) 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.

arb229 (Ari) O. Alipourfard et al.: [CherryPick: Adaptively Unearthing the Best Cloud Configurations for Big Data Analytics](#), NSDI, 2017.

sti20 (Stefan) A. Klein, S. Falkner, S. Bartels, P. Hennig, F. Hutter: [Fast Bayesian Optimization of Machine Learning Hyperparameters on Large Datasets](#), AISTAS, 2017.

If you can, skim following papers:

- T. Domhan, J. T. Springenberg, F. Hutter: [Speeding up automatic hyperparameter optimization of deep neural networks by extrapolation of learning curves](#), IJCAI, 2015.

- F. Hutter et al.: [Algorithm runtime prediction: Methods&evaluation](#), Elsevier J. AI, 2014.

jsl71 (Julius) Z. Jia, M. Zaharia, and A. Aiken: [Beyond Data and Model Parallelism for Deep Neural Networks](#), SYSML, 2019.

gw412 (Grzegorz) 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.

If you can skim following papers.

- 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: End-to-End Compilation Stack for Deep Learning](#), SysML, 2017.

2019/11/26 Session 7: Optimisation of Computer Systems and ML 2

ne274 (Niall) 1. A. Mirhoseini et al.: [Device Placement Optimization with Reinforcement Learning](#), ICML, 2017.

+ A. Mirhoseini, A. Goldie, H. Pham, B. Steiner, Q. Le and J. Dean: [A Hierarchical Mode for Device Placement](#), ICLR, 2018.

cnfo2 (Nnaemeka) R. Addanki, S. B. Venkatakrishnan, S. Gupta, H. Mao, M. Alizadeh : [Placeto: Learning Generalizable Device Placement Algorithms for Distributed Machine Learning](#) , arXiv, 2019.

tal42 (Tim) Z. Jia, J. Thomas, T. Warszawski, M. Gao, M. Zaharia, A. Aiken: [Optimizing DNN Computation with Relaxed Graph Substitutions](#), SYSML, 2019.

If you can add the extended part in the following paper from 'Beyond' paper.

- Z. Jia, O. Padon, J. Thomas, T. Warszawski, M. Zaharia, A. Aiken: [TASO: Optimizing Deep Learning Computation with Automated Generation of Graph Substitutions](#): SOSP, 2019.

doaa2 (David) R. Marcus, P. Negi, Parimarjan, H. Mao, C. Zhang, M. Alizadeh, T. Kraska, O. Papaemmanouil, and N. Tatbul: [Neo: A Learned Query Optimizer](#), VLDB, 2019.

mcic31 (Claire) H. Mao et al.: [Park: An Open Platform for Learning-Augmented Computer Systems](#), OpenReview, 2019.