

R212: Paper Review Presentation Assignment

2015/10/20 Session 2: Programming in Data Centric Environment

- JVT27 (James) 1.** Yuan Yu, Michael Isard, D. Fetterly, M. Budiu, U. Erlingsson, P.K. Gunda, J. Currey: [DryadLINQ: A System for General-Purpose Distributed Data-Parallel Computing Using a High-Level Language](#), OSDI, 2008.
- CL554 (Christopher) 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.
- OW238 (Olivia) 6.2.** D. Murray, F. McSherry, R. Isaacs, M. Isard, P. Barham, M. Abadi: [Naiad: A Timely Dataflow System](#), SOSP, 2013.

2015/10/27 Session 3: Processing Models of Large-Scale Graph Data

- Muthu (dmk45) 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.
- Kenneth (wckl2) 3.** Z. Qian, X. Chen, N. Kang, M. Chen, Y. Yu, T. Moscibroda, Z. Zhang: [MadLINQ: large-scale distributed matrix computation for the cloud](#), EuroSys, 2012.
- James (jvt27) 7.** J. Gonzalez, Y. Low, H. Gu, D. Bickson, and C. Guestrin: [Powergraph: distributed graph-parallel computation on natural graphs](#). OSDI, 2012.
- Olivia (ow238) 5.** J. Shun and G. Blelloch: [Ligra: A Lightweight Graph Processing Framework for Shared Memory](#), PPOPP, 2013.

2015/11/10 Session 5: Stream Data Processing and Data/Query Model

- Kenneth (wckl2) 6.** B. Gedik, H. Andrade, K. Wu, P. Yu, and M. Doo: [SPADE: the system S Declarative Stream Processing Engine](#), SIGMOD, 2008.
- Christopher (CL554) 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 paper 7, probably you also check the following paper if you can.**
- 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.

2015/11/17 Session 6: Optimised Approaches in Data Processing

- James (jvt27) 1.** A. Kyrola and G. Blelloch: [Graphchi: Large-scale graph computation on just a PC](#), OSDI, 2012.
- Olivia (ow238) 2.** A. Roy, I. Mihailovic, W. Zwaenepoel: [X-Stream: Edge-Centric Graph Processing using Streaming Partitions](#), SOSP, 2013.
- ***** X-Stream scale-out version is Chaos. It will be great if you can include main idea of Chaos below.**
- 3.** A. Roy, L. Bindschaedler, J. Malicevic and W. Zwaenepoel: [Chaos: Scale-out Graph Processing from Secondary Storage](#), SOSP, 2015.
- Christopher (CL554) 11.** D. Proutzos, R. Manevich, K. Pingali: [Elixir: A System for Synthesizing Concurrent Graph Programs](#), OOPSLA, 2012.
- Muthu (dmk45) 10.** S. Hong, H. Chafi, E. Sedlar, K. Olukotun: [Green-Marl: A DSL for Easy and Efficient Graph Analysis](#), ASPLOS, 2012.

2015/11/24 Session 7: Machine Learning for Computer System's Optimisation

- Olivia (ow238) 1.** N.J. Yadwadkar, B. Hariharan, J. Gonzalez and R. Katz: [Faster Jobs in Distributed Data Processing using Multi-Task Learning](#), SDM, 2015.
- Kenneth (wckl2) 4.** K. LaCurts et al.: [Cicada: Introducing Predictive Guarantees for Cloud Networks](#), HOTCLOUD, 2014.
- Christopher (CL554) 5.** M. [Carvalho](#) et al.: [Long-term SLOs for reclaimed cloud computing resources](#), SOCC, 2014.
- Muthu (dmk45) 8.** X. Dutreih et al.: [Using Reinforcement Learning for Autonomic Resource Allocation in Clouds: Towards a Fully Automated Workflow](#), ICAS, 2011.