

NAIAD EGG

NAIAD

- Dataflow graph programming language
- DD implementation written in Rust
- used for graph algorithms by distributing nodes of computation across the available nodes

egg: e-graphs good

EGG

- Equivalent program representations
- Equivalent graphs collected by a class of graphs
- Allows speedy rewrites and action over equivalences
- Rust implementation



WHAT'S THE PROBLEM?

- Naiad for graph algorithms GraphLINQ, isn't adapted to skew, so can get nodes with heavy loads after partitioning
- This partition is basically a graph.
- Use Egg to find equivalent graphs that are more balanced?
- Could extend this to allowing more dynamic scheduling as loads change.
- Apply it to PageRank (well known implementation in GraphLINQ)
- No work done yet.

REFERENCES

- Murray, D.G., McSherry, F., Isaacs, R., Isard, M., Barham, P., Abadi, M. Naiad: A timely dataflow system. In Proc. SOSP (Nov. 2013), 439–455
- Derek G. Murray, Frank McSherry, Michael Isard, Rebecca Isaacs, Paul Barham, and Martin Abadi. 2016. Incremental, iterative data processing with timely dataflow. Commun. ACM 59, 10 (October 2016), 75–83. https://doi.org/10.1145/2983551
- Max Willsey, Chandrakana Nandi, Yisu Remy Wang, Oliver Flatt, Zachary Tatlock, and Pavel Panchekha. 2021. Egg: Fast and extensible equality saturation. Proc. ACM Program. Lang. 5, POPL, Article 23 (January 2021), 29 pages. https://doi.org/10.1145/3434304
- Rong Chen, Jiaxin Shi, Yanzhe Chen, Binyu Zang, Haibing Guan, and Haibo Chen. 2019. PowerLyra: Differentiated Graph Computation and Partitioning on Skewed Graphs. ACM Trans. Parallel Comput. 5, 3, Article 13 (September 2018), 39 pages. https://doi.org/10.1145/3298989