Optimizing Recursive Joins in Graph Database Management Systems

Anurag Chakraborty

David R. Cheriton School of Computer Science
University of Waterloo
a8chakra@uwaterloo.ca

1 Abstract

Recursive joins such as shortest, all shortest, and variable length path queries are a core feature provided by graph database management systems. However, these queries are computationally expensive for large datasets and frequently suffer from high execution time due to skew from a few nodes being highly connected.

Existing work on efficiently executing these queries proposes using a morsel-driven parallelism [1] (MDP) approach for bulk path finding. The MDP approach assigns fixed-size "morsels" to threads, which comprise a set of starting nodes from which the path traversal needs to be started. This does not address the skew problem since it does not enable multiple threads to work on the same recursive join computation at the same time.

This presentation will focus on efficiently executing recursive join queries by integrating a hybrid parallelization scheduler into a GDBMS [2] query pipeline. We will examine the scheduler's design within our recursive join operator, morselized workload distribution between worker threads, query execution plans, the tradeoffs between various approaches such as MS-BFS [3] and direction optimizing BFS [4] and when to trigger which approach. We will also explore the use of lock free data structures used internally by the operator to handle query cases such as returning path length or returning graph paths to support the concurrent progress of multiple recursive join computations and experimental evaluation results obtained on large graphs (LDBC, LiveJournal, Graph500).

References

[1] D. ten Wolde, T. Singh, G. Szárnyas, and P. A. Boncz, "DuckPGQ: Efficient Property Graph Queries in an Analytical RDBMS," in Proceedings of the 13th Conference on Innovative Data Systems Research (CIDR 2023)

[2] X. Feng, G. Jin, Z. Chen, C. Liu and S. Salihoğlu, "Kùzu Graph Database Management System," in Proceedings of the 13th Conference on Innovative Data Systems Research (CIDR 2023)

[3] M. Then, M. Kaufmann, F. Chirigati, T. HoangVu, K. Pham, A. Kemper, T. Neumann and Huy T. Vo, "The More the Merrier: Efficient Multi-Source Graph Traversal," VLDB 2014

[4] S. Beamer, K. Asanovic and D. Patterson, "Direction-Optimizing Breadth-First Search," SC '12: Proceedings of the International Conference on High Performance Computing, Networking, Storage and Analysis