A Unified Multi-Round Algorithm for Distributed Join Evaluation with Optimality Guarantees

Heba Aamer, and Bas Ketsman

Vrije Universiteit Brussel, Belgium
{heba.mohamed,bas.ketsman}@vub.be

Shared nothing systems and frameworks are omnipresent. A popular model within the database community to design and analyze algorithms for these systems is the massively parallel computation (MPC) model [1] in which algorithms are expressed in rounds, and each round consists of a communication and a computation phase. The overall goal of the model is to find algorithms that minimize the communication cost, which is measured as the maximum load received by any of the serves in any of the rounds. In addition, the number of rounds is restricted to either a single round or a constant number of rounds.

For the join evaluation problem, unified worst-case optimal algorithms exist for the one-round version of the MPC model [3]. Both for unconstrained as well as so-called skew-free databases, in which the degrees of values are always below a certain threshold. In the multi-round MPC model, opportunities exist to decrease the load further than is the case in the single-round MPC model. Still, optimal algorithms are only known for specific restrictive classes of join queries, like queries whose hypergraph is acyclic [2] or only consisting of binary edges [4]. The two algorithms are rather complicated and make use of techniques that rely on properties tied to the class of queries they are designed for.

In our work, we propose a unified algorithm that can be applied to any join query, independent of its structural properties. The load of our algorithm does not exceed the load guarantees of the best-known worst-case optimal algorithm for the one-round MPC model, and it is optimal for queries for which the optimal worst-case load in the constant-round MPC model is known, including queries whose hypergraph is acyclic or a graph.

References

[1] Beame, P., Koutris, P. and Suciu, D. (2014). Skew in parallel query processing. In *PODS*, 212–223.

[2] Hu, X. (2021). Cover or pack: New upper and lower bounds for massively parallel joins. In *PODS*, 181–198.

[3] Koutris, P., Salihoglu, S. and Suciu, D. (2018). Algorithmic aspects of parallel query processing. In *SIGMOD*, 1659–1664.

[4] Ketsman, B., Suciu, D. and Tao, Y. (2022). A near-optimal parallel algorithm for joining binary relations. *Logical Methods in Computer Science*.