

Worst Case Optimal Join Algorithms: Techniques, Results, and Open Problems

Hung Q Ngo

Relational AI Inc. (US)

Abstract. Worst case optimal join algorithms are the class of join algorithms whose runtime match the worst-case output size of a given join query. While the first provably worst case optimal join algorithm was discovered relatively recently, the techniques and results surrounding these algorithms grow out of decades of research from a wide range of areas, intimately connecting graph theory, algorithms, information theory, constraint satisfaction, database theory, and geometric inequalities. These ideas are not just paper ware, one such algorithm was the work-horse join algorithm of a successful commercial database and data analytics engine.

This talk aims to be a gentle introduction to worst case optimal algorithms, the intuition behind them, historical and recent developments, and connections to information theory, and several fundamental open problems.