

Beyond Heuristics: I/O-oriented Algorithms and Structures with Performance Guarantees

Yufei Tao^{1,2}

¹ Department of Computer Science and Engineering
Chinese University of Hong Kong
New Territories, Hong Kong
Email: taoyf@cse.cuhk.edu.hk

² Division of Web Science
Korea Advanced Institute of Science and Technology
Daejeon, Korea
Email: taoyf@kaist.ac.kr

Abstract

System development often results in algorithms and access methods that are heuristic in nature. They have the merits of being easy to understand, simple to implement, and reasonably effective on many real datasets. Nonetheless, common criticism is that they typically come with no non-trivial performance guarantees. For years the database community has been striving to discover methods that are small and sweet. That is, such a method should be implementable in practice (i.e., small), and in the mean time, must retain attractive efficiency even in the worst case (i.e., sweet). This talk serves as an introduction to the current progress in the research of small and sweet algorithms and data structures. We will review solutions to several classic problems including spatial join, skyline retrieval, range searching, and so on. Special discussion will be dedicated to techniques generally useful for obtaining good worst-case I/O bounds.