



香港中文大學  
The Chinese University of Hong Kong

Institute of Theoretical Computer Science and Communications

*ITCSC Seminar*

## Dynamic Indexability and the Optimality of B-trees and Hash Tables

By

**Prof. Ke Yi**

*Assistant Professor, Department of Computer Science and Engineering, HKUST*

**March 15, 2010 (Monday)**

~~4:00 pm - 5:00 pm~~ 3:00 pm - 4:00 pm

**Rm. 121, Ho Sin Hang Engineering Building, CUHK**

### Abstract:

B-trees and hash tables are two fundamental external index structures in database systems, and are used to support range queries and lookups, two basic types of queries on a table. Despite the long history and extensive studies, their optimality is yet a mystery in the dynamic setting, namely, when insertions and deletions of elements are to be supported. For external indexes, the goal is to minimize the number of block accesses (or I/Os), so a standard optimization for fast updates is to use the available internal memory to buffer the recent updates and commit the changes to disk in batches, which can significantly lower the (amortized) update cost. Unfortunately, all known techniques for doing so cause certain levels of degradation in query performance.

In this paper, we prove several inherent tradeoffs between the update cost and the query performance for any index supporting range queries and lookups, respectively, in the presence of a memory buffer. These tradeoffs match the existing simple dynamic versions of B-trees and hash tables for all reasonable parameter ranges, thus establishing the optimality of these two fundamental structures. More interestingly, we obtain our results in a unified framework named dynamic indexability. Indexability theory was initiated in 1997, but up to now it has only been used to study static indexing problems and only in two or higher dimensions. In this regard, by adding in the dynamic aspect, we have enriched the indexability theory and made it applicable to even more basic indexing problems like B-trees and hash tables.

### Biography:

Ke Yi received his B.E. from Tsinghua University and Ph.D. from Duke University, in 2001 and 2006 respectively, both in computer science. After spending one year at AT&T Labs as a researcher, he has been an Assistant Professor in the Department of Computer Science and Engineering at Hong Kong University of Science and Technology since 2007. His research interests include algorithms, data structures, and databases.

\*\*\* ALL ARE WELCOME \*\*\*