

April 2011

MADALGO seminars by Nodari Sitchinava, Karlsruhe Institute of Technology

A Parallel Buffer Tree

Abstract:

We present the parallel buffer tree, a parallel external-memory (PEM) data structure for batched search problems. Our data structure is a non-trivial extension of Arge's sequential buffer tree to a private-cache multiprocessor environment and reduces the number of I/O operations by the number of available processor cores compared to its sequential counterpart, thereby taking full advantage of multi-core parallelism.

The basic parallel buffer tree is a batched search tree data structure that supports batched parallel processing of updates to the structure and membership queries. We also show how to extend the basic parallel buffer tree to support range queries. Our data structure is able to process a sequence of N updates and queries in optimal $\mathcal{O}_p(\mathcal{V} \cdot \mathcal{P}_{\text{sort}}\{N\} + \mathcal{P}_{\text{scan}}\{K\})$ parallel I/O complexity, where K is the size of the output reported in the process and $\mathcal{P}_{\text{sort}}\{N\}$ and $\mathcal{P}_{\text{scan}}\{N\}$ are the parallel I/O complexities of, respectively, sorting and scanning N elements using P processors.