

Exercises and hand-ins

Advanced database technology

April 15, 2003

Hand-in

<p>To be handed in at the latest April 24 at 10.00 AM.</p>

This problem considers range queries in persistent B-trees. You have seen that a persistent B-tree allows search for an element x at time t using $O(\log_B N_t)$ I/Os, where N_t is the number of elements in the tree at time t and the degree of the B-tree is $\Theta(B)$. For ordinary B-trees you have seen that *range queries* for elements in an interval $[a; b]$ can be performed in $O(\log_B N_t + Z/B)$ I/Os, where Z is the number of elements in the interval. This problem considers range queries “at time t ” in a persistent B-tree.

1. What is the problem with using the same solution as in an ordinary B-tree (traversing a linked list of leaves)?
2. Argue that B-tree searches can be performed in $O(\log_B N_t + Z/B)$ I/Os in a B-tree where leaves are *not* in a linked list. Conclude that this I/O complexity can also be realized in a persistent B-tree.
3. Consider a static set of N horizontal line segments (as on page 10 of the slides). Devise a linear space data structure that allows reporting of all Z line segments crossing a given vertical line segment in $O(\log_B N + Z/B)$ I/Os.

Other exercises for discussion on April 24

No exercises this week.