

Exercises and hand-ins

Advanced database technology

March 6, 2003

Hand-in

No hand-in for March 13.

Exercises for discussion on March 13

1. GUW 13.4.3.
2. In last week's lecture, you found bad key sets (resulting in long chains) for some specific hash functions. Argue that any fixed hash function will have bad key sets if the number of possible keys is large enough. Why does it help to pick a function at random?
3. Consider linear hashing where we want to keep the space utilization $\alpha = 1/2$. Suppose that B is so large that we have no block overflows. What is the number of I/Os needed for inserting N keys in a table of initial size 1?
4. Now consider the above question for uniform rehashing, assuming that whenever the space utilization is above $\alpha = 2/3$ we rehash to a table of twice the size.