Weekplan: External Memory I

Philip Bille   Inge Li Gørtz   Eva Rotenberg

References and Reading

[1] The Input/Output Complexity of Sorting and Related Problems, A. Aggarwal and J. Vitter, CACM 1988. Set $P = 1$ when reading this.

[2] Organization and Maintenance of Large Ordered Indexes, R. Bayer, E. McCreight, Acta Inform., 1972.

[3] Introduction to Algorithms, 3rd edition, Chap. 18, T.H. Cormen, C.E. Leiserson, R.L. Rivest, C. Stein, 2009.

We recommend reading [1] and [3] in detail. [2] is the original paper introducing B-trees.

Exercises

1 [w] **Prefix Sum** Given an array A of N elements, the *prefix-sum* of A is the array P such that $P[i] = \sum_{j \leq i} A[i]$. Show how to compute the prefix sum of A efficiently in external memory

2 [w] **Memory Hierarchy** Determine the configuration of the memory hierarchy on your own computer. Also, what is the cache-inclusion policy?

3 Stacks and Queues Consider stacks and queue in external memory. Solve the following exercises.

3.1 Show how to efficiently implement a stack in external memory. What is the worst-case and amortized I/Os per operation?

3.2 Do the same for a queue.

4 Linked Lists Consider a data structure that maintains a sequence of element $L = e_1, \ldots, e_N$ under the following operations:

- insert(e, e'): Insert element e' immediately after element e in the sequential order in L (extending the length of the sequence by 1).
- delete(e): Delete the element e in L.
- traverse(): Report the elements in L in sequence.

We assume that the argument e is a pointer to the element. Show how to efficiently implement the operations in external memory. *Hint:* What are the optimal I/O bounds for these operations you can hope to achieve? Try to achieve that.

5 Dictionaries and Hashing Recall the standard dictionaries based on hashing from your previous algorithms studies. What are the I/O complexities for these solutions? Consider the *exact* number of I/Os.

6 Range Reporting Suppose we want to extend B-trees to support the following range reporting operation:

- report(i, j): Report all elements with keys k, such that $i \leq k \leq j$.

Show how to efficiently implement report on B-trees. Your solution should have a good dependency on the size of the output.

7 Insertions in B-tree Consider the following B-tree of order 4. The capital letters represent subtrees. Show the the tree after inserting 59.

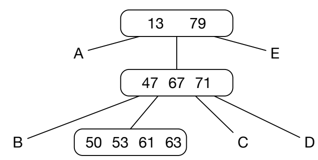

8 B-tree Construction Show how to efficiently construct a B-tree from an array of N elements.

9 Optimality of B-trees Suppose that we want to search among N keys. Furthermore, suppose that the only way of accessing disk blocks is by following pointers. Show that a search takes at least $\Omega(\log_B N/M)$ I/Os in the worst case. *Hint:* Consider the size C_t of the set of blocks that can be accessed in at most t I/Os. Assume that our memory initially is full of pointers.

10 Dynamic Programming Let S and T be strings of length N and consider the classic $O(N^2)$ time solution for computing the longest common subsequence of S and T. Show how to implement the algorithm efficiently in external memory.