Seminar on Advanced topics in data structures

Fall 2001/2002

Haim Kaplan
phone: 03-6408814
email: haimk@math.tau.ac.il

October 27, 2001

We shall focus on the following topics in data structures

1. Persistent data structures
2. Perfect hashing
3. Integer priority queues
4. Cache oblivious algorithms

Depending upon the number of students participating we may touch other topics as well.

How to get in: To participate in the seminar you must choose a paper and register it with me by Thursday (25/10) of this week. You can get copies of the papers to examine either by coming to my office or at the library. Once you decide that you want to participate and picked a paper please come to register the paper at my office (its also possible to do that by phone or email if you cannot come in person). Please note that as soon as you register your chances to get in with the paper that you want are higher, in particular if there are a lot of students in class. Papers in the list below are ordered more or less in the order in which I would like them to be presented so you can estimate the approximate week of your presentation.

Once registered you must do the presentation you commited to!

Your grade: You'd have to attend all meetings and actively participate in the discussion so as to make sure you understand the material. You are encouraged to ask questions and you get credit for it. To encourage you to ask everything you like I guarantee not to take any credit from the speaker even in case that he or she does not answer well. So there is absolutely no way you can hurt your friend by asking, and you may hurt yourself if you refrain from asking.

Participation worth up to 10% of your final mark. The other 90% will be determined by the quality of your presentation.

Guidelines for preparing your presentation:

1. All the papers on the list are theoretical and you may have to spend quite a bit of time in order to understand your paper. Good understanding of the paper is the key for a good presentation. So start working early and try to gain a good understanding of the results.

2. Some of the papers on the list are related to others. So attending the other lectures would certainly help you when you work on your paper.

3. You are not restricted to use any particular media for your presentation. Clear conventional blackboard talk certainly can gain maximum credit. Remember that sometimes a picture worth 1000 words. Media such as slides or computer slides is particularly useful when you want to show pictures and animate operations since doing that on the board may take too much time. Your time is limited so you have to use it carefully.

4. Its likely that in the time you have you will not be able to cover everything that is written in your paper. Depending on how crowded the seminar would be, it may not be possible to give you extra time. So plan the talk carefully, you need not cover everything in the paper but what you cover should be transfered clearly. Don't cover stuff that is in your paper but has already been covered by previous speakers.

References

[1] N. Sarnak and R. E. Tarjan. Planar point location using persistent search trees. *Communications of the ACM*, 29(7):669–679, 1986.

[2] M. L. Fredman, J. Komlós, and E. Szemerédi. Storing a sparse table with $o(1)$ worst case access time. *Journal of the ACM*, 31:538–544, 1984.

[3] P. van Emde Boaz, R. Kaas, and E. Zijlstra. Design and implementation of an efficient priority queue. *Mathematical Systems Theory*, 10:99–127, 1977.

[4] D. Knuth. Notes on the van Emde Boaz construction of priority deques: An instructive use of recursion. manuscript, 1977.

[5] P. van Emde Boaz. Preserving order in a forest in less than logarithmic time. *Information Processing Letters*, 6(3):80–82, 1977.

[6] D. E. Willard. New trie data structures which support very fast search operations. *J. Computer and Systems Sciences*, 28:379–394, 1984.

[7] D. E. Willard. Log-logarithmic worst-case range queries are possible in space $\Theta(n)$. *Information Processing Letters*, 17:81–84, 1983.

[8] K. Mehlhorn and S. Naher. Dynamic fractional cascading. *Algorithmica*, 5:215–241, 1990.

[9] P. Bozanis, N. Kitsios, C. Makris, and A. Tsakalidis. The space-optimal version of a known rectangle enclosure reporting algorithm. *Information Processing Letters*, 61:37–41, 1997.

[10] A. Itai, A. G. Konheim, and M. Rodeh. A sparse table implementation of priority queues. In *Proceedings of the 8th Colloquium on Automata, Languages, and Programming*, pages 417–431, 1981.

[11] M. Bender, E. Demaine, and M. Farach-Colton. Cache-oblivious B-trees. In *In Proceedings of the 41st IEEE FOCS, pages 399–409, Nov*, 2000.

[12] W. Goddard, C. Kenyon, V. King, and L. J. Schulman. Optimal randomized algorithms for local sorting and set-maxima. *SIAM Journal on Computing*, 22(2):272–283, 1993.

[13] Mikkel Thorup. Parallel shortcutting of rooted trees. *Journal of Algorithms*, 23(1):139–159, 1997.