

Exam curriculum for Advanced database technology, spring 2006

The course book

- G UW = Database Systems: The Complete Book by Hector Garcia-Molina, Jeff Ullman, and Jennifer Widom, 2002.
 - Section 11.1-11.7.3, 12, 13.0-13.4, 14.0-14.1, 15.0-15.8, 16.2-16.7, 17.0-17.1, 17.2.4, 17.2.5, 17.4, 18.0-18.8, 19.3, 20.6.
 - Note that examples are not curriculum, meaning that the exam will not require knowledge of any particular example. (On the other hand it may be useful to read some of the examples to help you understand the text.)

Papers and other materials

All material, except [CLRS01], is available at <http://www.itu.dk/people/pagh/ADBT06/>. From outside ITU you can find the papers at <http://www.itu.dk/people/pagh/ADBT06/Intranet/> (password protected with the login and password supplied to course participants by e-mail).

- [AgrawalSrikant94] = Rakesh Agrawal, Ramakrishnan Srikant. Fast Algorithms for Mining Association Rules. Proceeding of the 20th International Conference on Very Large Data Bases, 1994. The whole paper, except section 3.3.
- [Arge01] = Lars Arge. External memory data structures. Section 1, 2.1 (persistent B-trees), 3-4 (static versions only), 4.1, 9.
- [AGMS99] = N Alon, P Gibbons, Y Matias, M Szegedy. Tracking join and self-join sizes in limited storage. PODS '99. Section 1, 2, 4, 5, except proofs.
- [BrinPage98] = Sergey Brin and Lawrence Page. The Anatomy of a Large-Scale Hypertextual Web Search Engine. Section 1, 2, 4.2.
- [BrodalFagerberg03] = Gerth Brodal and Rolf Fagerberg. Lower Bounds for External Memory Dictionaries, SODA 2003. Section 1 and 3.4.
- [CLRS01] = Cormen, Leiserson, Rivest, and Stein. Introduction to Algorithms, 2nd edition. Page 405-409. (Or [CLR90] = Cormen, Leiserson, and Rivest. Introduction to Algorithms, Page 356-360.)
- [KarkkainenRao03] = Juha KarkkainenRao and S. Srinivasa Rao. Full-text Indexes in External Memory. Section 7.1-7.4.
- [MaheshwariZeh03] = Anil Maheshwari and Norbert Zeh. A Survey of Techniques for Designing I/O-Efficient Algorithms. Section 3.1-3.2.
- [MethaAgrawalRissanen96] = M Mehta, R Agrawal, and J Rissanen. SLIQ: A fast scalable classifier for data mining. Proceedings of EDBT '96. The whole paper.
- [RG] = Ramakrishnan and Gehrke. Database Management Systems, 3rd ed. Section 16.6.3 (handed out).

- [Sanders03] = Peter Sanders. Memory Hierarchies - Models and Lower bounds. Section 1.3-1.5.
- [Pagh03] = Rasmus Pagh. Basic External Memory Data Structures. Section 1, 3.0-3.2, 4.
- [Ullman00] = Jeff Ullman. Association rules and frequent item sets. Note for CS345, Stanford University, 2000. The entire note.
- Slides from the lectures.
Available at <http://www.itu.dk/people/pagh/ADBT06/> .