Ph.D. Written Examination Syllabus

August 11, 2016

1 Analysis of Algorithms

Syllabus

Philosophy: We are testing analytical ability: how well students think on their feet, rather than details of some data structure, etc. S there will be less reliance on memory. The questions will be more of a general nature (and will require more analytical/creative ability) than a typical final exam.

All sections marked with a * in CLRS (4.4, 5.4, 11.5, 12.4, 16.4, 16.5, 21.4, 26.4, 26.5) are excluded.

1. Preliminaries (CLRS Chapters 3 and 4):

   - pages 21, CLRS Ch 3 Growth of Functions
   - pages 29, CLRS Ch 4 Recurrences
   - pages 30, CLRS Ch 5 Probabilistic Analysis and Randomized Algorithms
   - pages 6, CLRS Ch C.3 Discrete random variables

   - Algorithm growth, Big-Oh and similar notations.
   - Recurrence equations, and their solution. Applications of recurrences.
   - Randomized algorithms, expected runtime analysis

2. Sorting (CLRS Chapters 6-9):

   - pages 18, CLRS Ch 6 Heapsort
   - pages 20, CLRS Ch 7 Quicksort
   - pages 18, CLRS Ch 8 Sorting in Linear Time
   - pages 13, CLRS Ch 9 Medians and Order Statistics

   - Algorithms (Heapsort, quicksort, mergesort, radixsort).
   - Proof that sorting (by comparisons) takes Omega(n log(n)) time.

- Medians and order statistics.

3. Divide and conquer techniques (CLRS Section 2.3, Chapter 7, and Section 28.2):

   - pages 10, CLRS Section 2.3 Designing Algorithms
   - pages 20, CLRS Ch 7 Quicksort
   - pages 7, CLRS Section 28.2 Strassen's algorithm

   - For mergesort.
   - For quicksort.
   - Strassen's algorithm for matrix multiplication.

4. Trees and hashing (CLRS Chapters 10–14 and 18):

   - pages 21, CLRS Ch 10 Elementary Data Structures
   - pages 42, CLRS Ch 11 Hash Tables
   - pages 20, CLRS Ch 12 Binary Search Trees
   - pages 29, CLRS Ch 13 Red-Black Trees
   - pages 21, CLRS Ch 18 B-trees
   - pages 17, CLRS Ch 14 Augmenting Data Structures

   - Stacks, queues, linked lists.
   - Hashing techniques and functions.
   - Binary search trees.
   - Balanced search trees.
   - B-trees.
   - Dynamic order statistics, interval trees

5. Dynamic programming (CLRS Chapter 15):

   - pages 47, CLRS Ch 15 Dynamic Programming

   - Examples and principles.

6. Greedy algorithms (CLRS Chapters 16 and 23);

   - pages 25, CLRS Ch 16 Greedy Algorithms
   - pages 19, CLRS Ch 23 Minimum Spanning Trees

   - Huffman codes.

- Minimal spanning tree.

7. Graph algorithms (CLRS Chapters 22-26):

   - pages 34, CLRS 22 Elementary Graph Algorithms
   - pages 19, CLRS 23 Minimum Spanning Trees
   - pages 40, CLRS 24 Single-Source Shortest Paths
   - pages 23, CLRS 25 All-pairs Shortest Paths
   - pages 56, CLRS 26 Maximum Flow

   - Graph representations.
   - Graph traversals.
   - Spanning trees or forests.
   - Shortest paths.
   - Maximum flow.

8. Amortized Analysis (CLRS Ch 17)

   - pages 25, CLRS 17 Amortized Analysis

   - Analyzing a sequence of operations
   - Aggregate, accounting, and potential methods

9. Data Structures for Disjoint Sets (CLRS Ch 21, excluding 21.4)

   - pages 12, CLRS 21 (excluding 21.4) Data Structures for Disjoint Sets

   - Union/Find data structures
   - Linked list representations, disjoint forest representations
   - Union by rank, path compression
   - Runtime, but not the proof, of union by rank with path compression

10. NP-completeness (CLRS Chapter 34):

    - pages 56, CLRS Ch 34 NP-Completeness

    - We won't ask students to prove how NP Turing Machines reduce to satisfiability, but they should be able to do simple reductions.
    - The classes P, NP, NP-complete, NP-hard.
    - An initial NP-complete problem.
    - Reductions and proofs that other problems are NP-complete, e.g., Hamiltonian Circuit, Traveling Salesman, and Vertex Cover.

References

- Algorithms, by Cormen, Leiserson, Rivest, and Stein (CLRS) (the "green" book).

2 Operating Systems

Syllabus

1. Basic resource management:

   - pages 95-128 SGG Ch 4 Processes
   - pages 207-249 CDK Ch 6 Operating Systems Support
   - pages 129-149 SGG Ch 5 Threads
   - pages 151-188 SGG Ch 6 Process scheduling
   - pages 273-316 SGG Ch 9 Memory Management
   - pages 317-370 SGG Ch 10 Virtual Memory
   - pages 371-409 SGG Ch 11 File System Interface
   - pages 411-451 SGG Ch 12 File System Implementation

   Total: 314 pages
   Main Topics:

   - Processes and process organization
   - Threads versus processes
   - Process scheduling
   - Memory management and virtual memory
   - Disk and storage management
   - File systems

2. Synchronization:

   - pages 189-241 SGG Ch 7 Process Synchronization
   - pages 243-270 SGG Ch 8 Deadlocks

   Total pages: 82 pages
   Main Topics:

   - Mutual exclusion
   - Mutex locks
   - Semaphores
   - Condition variables
   - Classical synchronization problems
   - Monitors
   - Concurrency

- Deadlocks

3. Communication:

   - pages 65-124, CDK Ch 3 Networking
   - pages 125-164, CDK Ch 4 Interprocess Communication
   - pages 165-203, CDK Ch 5 Distributed Objects and Remote Invocation
   - pages 669-695, CDK Ch 17 Corba

   Total Pages: 166
   Main Topics:

   - Connectionless versus connection-oriented protocols
   - Blocking versus non-blocking
   - Client-server model
   - RPCs
   - Remote object invocation
   - CORBA
   - Multicast
   - Group communication
   - Message passing

4. Naming, authentication, and protection

   - pages 251-308, CDK Ch 7 Security
   - pages 629-656, SGG Ch 18 Protection
   - pages 353-384, CDK Ch 9 Name Services

   Total: 118 pages
   Main Topics:

   - Basic security principles
   - Authentication
   - Digital signatures
   - Public Key versus Shared Key
   - Needham-Schroeder, Kerberos, SSL
   - Domains
   - Capabilities
   - Name services and spaces
   - DNS

- Directory services
- Discovery services

5. Distributed file systems

   - pages 309-352, CDK CH 8 Distributed File Systems

   Total: 44 pages
   Main Topics:

   - Design principles of distributed file systems
   - NFS
   - Andrew
   - Caching
   - Consistency
   - Persistent stores

Operating Systems total: 724 pages

References

- Operating System Concepts, 6th ed by Silberschatz, Galvin and Gagne (SGG)
- Distributed Systems, 3rd ed. by Coulouris, Dollimore, and Kindberg (CDK)