

ADVANCE COURSE IN DISTRIBUTED COMPUTING

Spring 2012 (Academic year 2011/2012). Prof. Y. Afek (afek at tau.ac...).

Course Home Page www.cs.tau.ac.il/~afek/dc12.html

A graduate course exploring current topics from the literature in distributed computing, focusing on theoretical issues: models, upper and lower bounds, and proof methods. Two major topics:

1. Distributed algorithms for data communication networks (Message Passing)
2. synchronization algorithms for *asynchronous* shared memory parallel machines.

In addition we will discuss the connections and relations between these two models.

Schedule /* Subject for Change */:

| Date | Topic and references used | Additional references |
|------------------------|--|--|
| Lecture 1 | Models, broadcast and echo [1, 2]. | [1, 3, 4, 5, 6, 2]. |
| Lecture 2 | Termination Detection, Snapshot [2, 7] Synchronizers [8]. | [9, 10, 11, 12, 2], [8, 13]. |
| Lecture 3 | Electing a Leader in a Ring Network, Chapter 3 in [?] | []. |
| Lecture 4 | Electing a leader in general graphs [14, 15]. | [14, 16, 17, 18, 19, 5, 20, 21, 22, 23]. |
| Lecture 5 | Computing the maximal independent set, rings and general graphs, upper and lower bounds [24]. | [25, 26, 27, 28, 29]. |
| Lecture 6 | Data link protocols, the sequence transmission problem, and End-to-end protocols [30, 31, 32]. | [31, 33, 34, 35, 36, 37, 38, 32] |
| Lecture 7 | The Consensus problem. Algorithms and lower bounds [39, 40, 41]. | |
| Lecture 8 | The shared memory model [42]. | [42, 43, 44, 45] |
| Lecture 9 | The consensus problem, and its impossibility in asynchronous networks with one faulty processor [40] or in the shared memory model [42]. | [40, 46, 41, 47, 48, 49]. |
| Lecture 10 | Wait-free synchronization, the shared memory hierarchy, and universal constructions [42]. | [48, ?, ?] |
| Lecture 11 | Atomic snap-shots of shared memory [50], Immediate snap-shots [51] . | [52] |
| Lecture 12 | Mutual exclusion, Fast Mutual Exclusion, Adaptive Algorithms Taubenfeld Paper Moir Anderson, Lamport-87 | |
| Lecture 13 | Simulating Shared memory in message passing, Randomized Consensus. | |
| Time permitting | Lower bound techniques [53, 24]. | [53, 20, 54, 5, 34, 55, 56, 24] |

Grade: Homeworks 35%, Take home exam 65%.

(This outline will be subject to modification as the semester progresses.)

References

- [1] A. Segall. Distributed network protocols. *IEEE Trans. on Information Theory*, IT-29(1):23–35, January 1983.
- [2] E. Gafni. Perspective on distributed network protocols: A case for building blocks. In *Proc. of MILCOM-86*, Oct 5-9 1986.
- [3] R. G. Gallager. Distributed minimum hop algorithms. Technical Report LIDS-P-1175, M.I.T. Lab for Information and Decision Systems, January 1982.
- [4] Y. Afek, B. Awerbuch, and E. Gafni. Applying static network protocols to dynamic networks. In *Proc. of the 28th IEEE Ann. Symp. on Foundation of Computer Science*, pages 358–370, October 1987.
- [5] J. E. Burns. A formal model for message passing systems. Technical Report TR-91, Indiana Univ., May 1980.
- [6] L. Lamport. Time, clocks, and the ordering of events in a distributed system. *CACM*, 21(7):558–565, July 1978.
- [7] K. M. Chandy and L. Lamport. Distributed snapshots: Determining global states of distributed systems. *ACM Trans. on Computer Systems*, 3(1):63–75, January 1985.
- [8] B. Awerbuch. Complexity of network synchronization. *Journal of the ACM*, 32(4):804–823, October 1985.
- [9] W. Dijkstra and C. S. Scholten. Termination detection for diffusing computations. *Information Processing Letters*, 11-1:1–4, August 1980.
- [10] N. Francez. Distributed termination. *TOPLAS*, 2,1:42–55, January 1980.
- [11] J. Misra. Detecting termination of distributed computations using markers. In *Proc. of the 2nd ACM Symp. PODC*, pages 290–295, August 1983.
- [12] K. M. Chandy and J. Misra. Termination detection of diffusing computations in csp. *ACM TOPLAS*, 4:37–42, 1982.
- [13] B. Awerbuch and D. Peleg. Cost-sensitive analysis of communication protocols. In *Proc. of the Ninth ACM Symp. on Principles of Distributed Computing (PODC)*, pages 177–188, August 1990.
- [14] R. G. Gallager, P. A. Humblet, and P. M. Spira. A distributed algorithm for minimum weight spanning trees. *ACM Trans. Program. Lang. Syst.*, 5:66–77, January 1983.
- [15] C. Chou and E. Gafni. Understanding and verifying distributed algorithms using stratified decomposition. In *Proc. of the 7th ACM Symp. on Principles of Distributed Computing*, pages 44–65, 1988.

- [16] R. G. Gallager. Finding a leader in a network with $O(e) + O(n \log n)$ messages. Unpublished Note, 1977.
- [17] G. N. Frederickson and N. A. Lynch. Electing a leader in a synchronous ring. *Journal of The ACM*, 34(1):98–115, January 1987.
- [18] G. L. Peterson. An $O(n \log n)$ unidirectional algorithm for the circular extrema problem. *ACM Trans. Program. Lang. Syst.*, 4(4):758–762, October 1982.
- [19] D. Dolev, M. Klawe, and M. Rodeh. An $O(n \log n)$ unidirectional algorithm for extrema finding in a circle. *Journal of Algorithm*, 3:245–260, 1982.
- [20] Y. Afek and E. Gafni. Time and message bounds for election in synchronous and asynchronous complete networks. In *Proc. of the ACM Symp. on Principles of Distributed Computing*, August 1985. to appear in *Siam J on Computing*.
- [21] G. L. Peterson. Efficient algorithms for elections in meshes and complete networks. Technical report, Dep. of Computer Science, Univ. of Rochester, August 1984.
- [22] B. Awerbuch. Linear time distributed algorithms for minimum spanning trees, leader election, counting and related problems. In *Proc. of the 19th Ann. ACM Symp. on Theory of Computing*, pages 230–240, May 1987.
- [23] G. N. Frederickson and N. A. Lynch. A general lower bound for electing a leader in a ring. Technical report, Purdue University, March 1985. CSD-TR-512.
- [24] N. Linial. Distributive graph algorithms - global solutions from local data. In *Proc. of the 28th IEEE Ann. Symp. on Foundation of Computer Science*, pages 331–335, October 1987.
- [25] R. Cole and U. Vishkin. Deterministic coin tossing and accelerating cascades: micro and macro techniques for designing parallel algorithms. In *Proc. of the 16th Ann. ACM Symp. on Theory of Computing*, pages 206–219, May 1986.
- [26] A. V. Goldberg, S. A. Plotkin, and G. E. Shannon. Parallel symmetry-breaking in sparse graphs. *SIAM J. Desc. Math.*, 1:434–446, 1989.
- [27] A. Goldberg and S. Plotkin. Parallel $(\delta + 1)$ coloring of constant-degree graphs. *Information Processing Letters*, 25(4):241–245, June 1987.
- [28] A. V. Goldberg, Serge A. Plotkin, and Gregory E. Shannon. Parallel symmetry-breaking in sparse graphs. In *Proc. of the 19th Ann. ACM Symp. on Theory of Computing*, pages 315–324, May 1987.
- [29] N. Linial and M. Saks. Finding low-diameter graph decompositions distributively. In *Proc. of 2nd Ann. ACM-SIAM Symp. on Discrete Algorithms*, December 1990.
- [30] Y. Afek and E. Gafni. End-to-end communication in unreliable networks. In *Proc. of the 7th ACM Symp. on Principles of Distributed Computing*, pages 131–148, August 1988.

- [31] Y. Afek, H. Attiya, A. Fekete, M. J. Fischer, N. Lynch, Y. Mansour, D. Wang, and L. D. Zuck. Reliable communication over unreliable channels. Manuscript, 1990.
- [32] Y. Afek, E. Gafni, and A. Rosen. The slide mechanism with applications in dynamic networks. In *Proc. 11th ACM Symp. on Principles of Distributed Computing*, pages 35–46, August 1992.
- [33] A. V. Aho, J. D. Ullman, and M. Yannakakis. Modeling communication protocols by automata. In *Proc. of the 20th IEEE Ann. Symp. on Foundation of Computer Science*, pages 267–273, 1979.
- [34] A. V. Aho, J. D. Ullman, A. D. Wyner, and M. Yannakakis. Bounds on the size and transmission rate of communication protocols. *Comp. & Maths. with Appls.*, 8:3:205–214, 1982. detailed version of [AUY-79].
- [35] J. Y. Halpern and L. D. Zuck. A little knowledge goes a long way: Simple knowledge-based derivations and correctness proofs for a family of protocols. In *Proc. of the ACM Symp. on Principles of Distributed Computing*, August 1987.
- [36] A. E. Baratz and A. Segall. Reliable link initialization procedures. In H. Rudin and C.H. West, editors, *IFIP 3rd Workshop on Protocol Specification, Testing and Verification, III*, May 1983.
- [37] K. A. Bartlett, R. A. Scantlebury, and P. T. Wilkinson. A note on reliable full-duplex transmission over half-duplex links. *Comm. of the ACM*, 12:260–261, 1969.
- [38] M. V. Stenning. A data transfer protocol. *Comput. Networks*, 1:99–110, 1976.
- [39] P. Berman, J. A. Garay, and K. J. Perry. Towards optimal distributed consensus. In *Proc. of the 30th IEEE Ann. Symp. on Foundation of Computer Science*, pages 410–415, October 1989.
- [40] M. Fischer, N. Lynch, and M. Paterson. Impossibility of distributed consensus with one faulty process. *Journal of the ACM*, 32:374–382, April 1985.
- [41] M. Fischer, N. Lynch, and M. Merritt. Easy impossibility proofs for distributed consensus problems. *Distributed Computing*, 1(1):26–39, 1985.
- [42] M. Herlihy. Wait-free synchronization. *ACM Trans. on Programming Languages and Systems*, 13(1):124–149, January 1991.
- [43] M. Herlihy and J. M. Wing. Axioms for concurrent objects. In *Proc. 14th ACM Symp. on Principles of Programming Languages*, pages 13–26, January 1987.
- [44] M. P. Herlihy. Impossibility and universality results for wait-free synchronization. In *Proc. of the 7th ACM Symp. on Principles of Distributed Computing*, pages 291–302, 1988.
- [45] L. Lamport. On interprocess communication, parts I and II. *Distributed Computing*, 1:77–101, 1986.

- [46] O. Biran, S. Moran, and S. Zaks. A combinatorial characterization of the distributed tasks which are solvable in the presence of one faulty processor. In *Proc. of the 7th ACM Symp. on Principles of Distributed Computing*, pages 263–275, 1988.
- [47] D. Dolev, C. Dwork, and L. Stockmeyer. On the minimal synchronism needed for distributed consensus. *Journal of the ACM*, 34(1):77–97, January 1987.
- [48] M. C. Loui and H. H. Abu-Amara. Memory requirements for agreement among unreliable asynchronous processes. *Advances in Computing Research, JAI Press*, 4:163–183, 1987.
- [49] J. E. Burns and Gary L. Peterson. The ambiguity of choosing. In *Proc. of the 8th ACM Symp. on Principles of Distributed Computing*, pages 145–158, Edmonton, Alberta, Canada, August 1989.
- [50] Y. Afek, H. Attiya, D. Dolev, E. Gafni, M. Merritt, and N. Shavit. Atomic snapshots of shared memory. In *Proc. of the Ninth Ann. ACM Symp. on Principles of Distributed Computing*, pages 1–13, August 1990. Also *Journal of the ACM*, to appear.
- [51] E. Borowsky and E. Gafni. Immediate atomic snapshots and fast renaming. In *Proc. 12th ACM Symp. on Principles of Distributed Computing*, pages 41–51, August 1993.
- [52] J. H. Anderson. Composite registers. In *Proc. of the Ninth Ann. ACM Symp. on Principles of Distributed Computing*, pages 15–29, August 1990.
- [53] G. N. Frederickson and N. A. Lynch. The impact of synchronous communication on the problem of electing a leader in a ring. In *Proc. of the 16th Ann. ACM Symp. on Theory of Computing*, pages 493–503, 1984.
- [54] J. Pahl, E. Korach, and D. Rotem. A technique for proving lower bounds for distributed maximum-finding algorithms. In *Proc. of the 14th Ann. ACM Symp. on Theory of Computing*, pages 378–382, 1982.
- [55] O. Goldreich and L. Shrira. On the complexity of computation in the presence of link failures: the case of a ring. *Distributed Computing*, 5, 1991.
- [56] O. Goldreich and D. Snih. On the complexity of global computation in the presence of link failures: the case of uni-directional faults. In *Proc. 11th ACM Symp. on Principles of Distributed Computing*, pages 103–111, August 1992.