

July, 2010

Haim KAPLAN, Ph.D.**LIST OF PUBLICATIONS****ARTICLES**

1. H. Kaplan and R. Shamir.
The domatic number problem on some perfect graph families.
Information Processing Let. 49 (1994), 51–56.
2. M.C. Golumbic, H. Kaplan, and R. Shamir.
On the complexity of DNA physical mapping.
Advances in Applied Mathematics 15 (1994), 251–261.
3. P.W. Goldberg, M.C. Golumbic, H. Kaplan, and R. Shamir.
Four strikes against physical mapping of DNA.
Journal of Computational Biology 2:1 (1995), 139–152.
4. M.C. Golumbic, H. Kaplan, and R. Shamir.
Graph sandwich problems.
Journal of Algorithms 19 (1995), 449–473 (See also Proc. No. 1).
5. H. Kaplan and R. Shamir.
Pathwidth, bandwidth, and completion problems to proper interval graphs with small cliques.
SIAM J. Comput. 25:3 (1996), 540–561 (See also Proc. No. 2).
6. A.L. Buchsbaum, H. Kaplan, A. Rogers, and J.R. Westbrook.
A new, simpler linear-time dominators algorithm.
ACM Transactions on Programming Languages and Systems (TOPLAS) 20:6 (1998), 1265–1296 (See also Proc. No. 8).
7. H. Kaplan and R. Shamir.
Bounded degree interval sandwich problems.
Algorithmica 24:2 (1999), 96–104 (See also Proc. No. 4).
8. H. Kaplan, R. Shamir, and R.E. Tarjan.
Tractability of parameterized completion problems on chordal, strongly chordal, and proper interval graphs.
SIAM J. Comput. 28:5 (1999), 1906–1922 (See also Proc. No. 2).
9. H. Kaplan, R. Shamir, and R.E. Tarjan.
A faster and simpler algorithm for sorting signed permutations by reversals.
SIAM J. Comput. 29:3 (1999), 880–892 (See also Proc. No. 7).
10. E. Cohen, H. Kaplan, and J.D. Oldham.
Managing TCP connections under persistent HTTP.
Computer Networks 31:11–16 (1999), 1709–1723 (See also Proc. No. 15).

11. H. Kaplan and R. E. Tarjan.
Purely functional, real-time dequeues with catenation.
Journal of the ACM 46:5 (1999), 577–603 (See also Proc. No. 3).
12. H. Kaplan, C. Okasaki, and R. E. Tarjan.
Simple confluent persistent catenable lists.
SIAM J. Comput. 30:3 (2000), 965–977 (See also Proc. No. 9).
13. H. N. Gabow, H. Kaplan, and R. E. Tarjan.
Unique maximum matching algorithms.
Journal of Algorithms 40 (2001), 159–183 (See also Proc. No. 18).
14. E. Cohen and H. Kaplan.
Caching documents with variable sizes and fetching costs: An LP-based approach.
Algorithmica 32:3 (2002), 459–466 (See also Proc. No. 14).
15. E. Cohen, H. Kaplan.
Exploiting regularities in Web traffic patterns for cache replacement.
Algorithmica 33:3 (2002), 300–334 (See also Proc. No. 16).
16. S. Arora, A. Frieze, and H. Kaplan.
A new rounding procedure for the assignment problem with applications to dense graph arrangement problems.
Mathematical Programming (series A) 92:1 (2002), 1–36 (See also Proc. No. 6).
17. E. Cohen and H. Kaplan.
Refreshment policies for Web content caches.
Computer Networks 38:6 (2002), 795–808 (See also Proc. No. 27).
18. E. Cohen and H. Kaplan.
Prefetching the means for document transfer: A new approach for reducing Web latency.
Computer Networks 39:4 (2002), 437–455 (See also Proc. No. 19).
19. E. Cohen, H. Kaplan, and U. Zwick.
Competitive analysis of the LRFU paging algorithm.
Algorithmica 33:4 (2002), 511–516 (See also Proc. No. 29).
20. N. Alon, H. Kaplan, M. Krivelevich, D. Malkhi, and J. Stern.
Scalable secure storage when half the system is faulty.
Information and Computation 174:2 (2002), 203–213 (See also Proc. No. 21).
21. Y. Afek, A. Bremner-Barr, H. Kaplan, E. Cohen, and M. Merritt.
Restoration by path concatenation: Fast recovery of MPLS paths.
Distributed Computing 15:4 (2002), 273–283, (See also Proc. No. 31). (Special issue of invited papers from PODC 2001.)

22. E. Cohen and H. Kaplan.
Proactive caching of DNS records: Addressing a performance bottleneck.
Computer Networks 41:6 (2003), 707–726 (See also Proc. No. 22).
23. E. Cohen, H. Kaplan, and U. Zwick.
Connection caching: Model and algorithms.
Journal of Computer and System Sciences (JCSS) 67:1 (2003), 92–126. (See also Proc. No. 17 and Proc. No. 20).
24. E. Cohen, E. Halperin, H. Kaplan, and U. Zwick.
Reachability and distance queries via 2-hop labels.
SIAM J. Comput. 32:5 (2003), 1338–1355 (See also Proc. No. 35).
25. A. Fiat and H. Kaplan
Making data structures confluent persistent.
Journal of Algorithms 48:1 (2003), 16–58 (See also Proc. No. 23). (Special issue of invited papers from SODA 2001)
26. A. Bremler-Barr, E. Cohen, H. Kaplan and Y. Mansour
Predicting and bypassing end-to-end internet service degradations.
IEEE Journal on Selected Areas in Communications 21:6 (2003), 961–978 (See also Proc. No. 41). (Special issue on internet and WWW measurement, mapping, and modeling.)
27. Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke.
Optimal oblivious routing in polynomial time.
Journal of Computer and System Sciences (JCSS) 69:3 (2004), 383–394 (See also Proc. No. 46). (Special issue of invited papers from STOC 2003.)
28. E. Cohen and H. Kaplan.
Balanced-replication algorithms for distribution trees.
SIAM J. Comput. 34:1 (2004), 227–247 (See also Proc. No. 40).
29. S. Alstrup, C. Gavoille, H. Kaplan, and T. Rauhe.
Nearest common ancestors: A survey and a new algorithm for a distributed environment.
Theory of Computing Systems 37:3 (2004), 441–456 (See also Proc. No. 38). (Special issue of papers from SPAA 2002.)
30. H. Kaplan and N. Shafrir.
The greedy algorithm for shortest superstrings.
Information Processing Letters 93:1 (2005), 13–17.
31. E. Cohen, E. Halperin, and H. Kaplan.
Performance aspects of distributed caches using TTL-based consistency.
Theoretical Computer Science 331:1 (2005), 73–96 (See also Proc. No. 28). (special issue of invited papers from ICALP 2001.)

32. M.C. Golumbic, H. Kaplan and E. Verbin.
On the complexity of cell flipping in permutation diagrams, and multiprocessor scheduling problems.
Discrete Mathematics 296:1 (2005), 25–41 (See also Proc. No. 11).
33. H. Kaplan and E. Verbin.
Sorting signed permutations by reversals, revisited.
Journal of Computer and System Sciences 70:3 (2005), 321–341 (see also Proc. No. 47). (Special Issue on Bioinformatics II)
34. H. Kaplan, M. Lewenstein, N. Shafrir, and M. Sviridenko.
Approximation algorithms for asymmetric TSP by decomposing directed regular multi-graphs.
Journal of the ACM 52:4 (2005), 602–626 (see also Proc. No. 48).
35. A.L. Buchsbaum, H. Kaplan, A. Rogers, and J.R. Westbrook.
Corrigendum: a new, simpler linear-time dominators algorithm.
ACM Transactions on Programming Languages and Systems (TOPLAS) 27:3 (2005), 383–387 (See also Journal No. 6).
36. A. Barkan and H. Kaplan.
Partial alphabetic trees.
Journal of algorithms 58:2 (2006), 81–103 (See also Proc. No. 39).
37. H. Kaplan and N. Shafrir.
The greedy algorithm for edit distance with moves.
Information Processing Letters 97:1 (2006), 23–27.
38. S. Abiteboul, S. Alstrup, H. Kaplan, T. Milo, and T. Rauhe.
Compact labeling schemes for ancestor queries.
SIAM J. Comput. 35:6 (2006), 1295–1309 (See also Proc. No. 24).
39. E. Cohen and H. Kaplan
Spatially-decaying aggregation over a network.
Journal of Computer and System Sciences 73:3 (2007), 265–288, special issue on database theory (See also Proc. No. 50).
40. G. Alexandron, H. Kaplan and M. Sharir.
Kinetic and dynamic data structures for convex hulls and upper envelopes.
Computational Geometry: Theory and Applications 36:2 (2007), 144–158 (See also Proc. No. 53).
41. K. Chen, A. Fiat, H. Kaplan, M. Levy, J. Matousek, E. Mossel, J. Pach, M. Sharir, S. Smorodinsky, U. Wagner, and E. Welzl.
Online conflict-free coloring for intervals.
SIAM J. Comput. 36:5 (2006), 1342–1359.

42. H. Kaplan, T. Milo, and R. Shabo.
Compact labeling scheme for XML ancestor queries.
Theory of Computing Systems 40:1 (2007), 55–99 (See also Proc. No. 34).
43. E. Cohen, A. Fiat, and H. Kaplan
Associative search in peer to peer networks: Harnessing latent semantics.
Computer Networks 51:8 (2007), 1861–1881 (See also Proc. No. 42-44).
44. N. Alon, H. Kaplan, M. Krivelevich, D. Malkhi, and J. Stern.
Addendum to "scalable secure storage when half the system is faulty".
Information and Computation 205:7 (2007) 1114–1116 (See also Journal No. 20)
45. H. Kaplan, S. Landau, and E. Verbin.
A simpler analysis of Burrows-Wheeler based compression.
Theoretical Computer Science 387:3 (2007) 220–235, special issue on the Burrow-Wheeler Transform (See also Proc. No. 56).
46. H. Kaplan and R.E. Tarjan.
Thin heaps, thick heaps.
ACM Transactions on Algorithms 4:1 (2008), article 3.
47. P. K. Agarwal, H. Kaplan, and M. Sharir.
Kinetic and dynamic data structures for closest pairs and all nearest neighbors.
ACM Transactions on Algorithms 5:1 (2008), article 4.
48. H. Kaplan, N. Rubin, M. Sharir, and E. Verbin.
Counting colors in boxes.
SIAM J. Comput. 38:3 (2008), 982–1011 (See also Proc. No. 64).
49. N. Alon, H. Kaplan, G. Nivasch, M. Sharir, and S. Smorodinsky.
Weak ϵ -nets and interval chains.
Journal of the ACM 55:6 (2008), article 28 (See also Proc. No. 73).
50. E. Cohen, N. Grossuag, and H. Kaplan.
Processing top-k queries from samples.
Computer Networks 55:14 (2008), 2605–2622 (See also Proc. No. 62).
51. A.L. Buchsbaum, L. Georgiadis, H. Kaplan, A. Rogers, R.E. Tarjan, and J.R. Westbrook.
Linear-time pointer-machine algorithms for dominators and other path-evaluation problems.
SIAM J. Comput. 38:4 (2008), 1533–1573 (See also Proc. No. 8).
52. E. Cohen and H. Kaplan.
Tighter estimation using bottom-k sketches.
Proceedings of the VLDB Endowment 1:1 (2008), 213–224.

53. K. Chen, H. Kaplan, and M. Sharir.
Online conflict free coloring for halfplanes, congruent disks, and axis-parallel rectangles.
ACM transactions on algorithms 5:2 (2009), article 16.
54. Y. Giora and H. Kaplan.
Optimal dynamic vertical ray shooting in rectilinear planar subdivisions.
ACM transactions on algorithms 5:3 (2009), article 28, special issue of invited papers from SODA 2007 (See also Proc. No. 63).
55. H. Kaplan, N. Rubin, and M. Sharir.
Linear data structures for fast ray-shooting amidst convex polyhedra.
Algorithmica 55:2 (2009), 283–310, special issue of invited papers from ESA 2007 (See also Proc. No. 71).
56. H. Kaplan and Y. Nussbaum.
Certifying algorithms for recognizing proper circular-arc graphs and unit circular-arc graphs.
Discrete Applied Mathematics 157:15 (2009) 3216–3230 (See also Proc. No. 59).
57. E. Cohen, N.G. Duffield, H. Kaplan, C. Lund, and M. Thorup.
Composable, scalable, and accurate weight summarization of unaggregated data sets.
Proceedings of the VLDB Endowment 2:1 (2009), 431–442.
58. E. Cohen, H. Kaplan, and S. Sen.
Coordinated weighted sampling for estimating aggregates over multiple weight assignments.
Proceedings of the VLDB Endowment 2:1 (2009), 646–657.
59. E. Cohen, H. Kaplan, and T. Milo.
Labeling dynamic XML trees.
SIAM J. Comput. 39:5 (2010), 2048–2074 (See also Proc. No. 37).
60. P. K. Agarwal, S. Bereg, O. Daescu, H. Kaplan, S. Ntafos, M. Sharir, and B. Zhu.
Guarding a terrain by two watchtowers.
Algorithmica 58:2 (2010), 352–390 (See also Proc. No. 51).

Accepted for publication

1. H. Kaplan, M. Sharir, and E. Shustin.
On lines and joints.
Discrete and Computational Geometry.
2. H. Kaplan and Y. Nussbaum.
A simpler linear-time recognition of circular-arc graphs.
Algorithmica (See also Proc. No. 61).

3. H. Kaplan, N. Rubin, and M. Sharir.
Line transversals of convex polyhedra in R^3 .
SIAM J. Comput. (See also Proc. No. 79).
4. H. Kaplan and Y. Nussbaum.
Maximum flow in directed planar graphs with vertex capacities.
Algorithmica (See also Proc. No. 82).
5. L. Georgiadis, H. Kaplan, N. Shafrir, R. E. Tarjan, and R. Werneck.
Data structures for mergeable trees.
ACM transactions on algorithms.

Submitted for publication

1. H. Kaplan, E. Ramos, and M. Sharir.
Range minima queries with respect to a random permutation, and approximate range counting.
Discrete and Computational Geometry (October 2007) (See also Proc. No. 54).
2. H. Kaplan, E. Ramos, and M. Sharir.
The overlay of minimization diagrams in a randomized incremental construction.
Discrete and Computational Geometry (October 2007) (See also Proc. No. 54).
3. E. Cohen, N. Duffield, H. Kaplan, C. Lund and M. Thorup.
Algorithms and estimators for summarization of unaggregated data streams
Journal of Computer and System Sciences (JCSS) (September 2007) (See also Proc. No. 68, and Proc. No. 72).
4. M. Hershcovitch and H. Kaplan
I/O efficient dynamic data structures for longest prefix queries.
Algorithmica (October 2008) (See also Proc. No. 74).
5. G. Elekes and H. Kaplan and M. Sharir.
On lines, joints, and incidences in three dimensions.
Discrete and Computational Geometry (May 2009).
6. H. Kaplan, N. Rubin, and M. Sharir.
A kinetic triangulation scheme for moving points in the plane.
Computational Geometry: Theory and Applications (May 2010).
7. H. Kaplan and N. Shafrir.
Approximate repeats
Algorithmica (June 2010).

CHAPTERS IN BOOKS (COLLECTION OF ARTICLES)

1. H. Kaplan.
Persistent data structures.
Handbook of Data Structures and Applications, chapter 31.
D.P. Mehta and S. Sahni, editors, Chapman & Hall/CRC , 2005.
2. E. Cohen and H. Kaplan.
The Time-To-Live based consistency mechanism: Understanding performance issues and their impact.
Web Content Delivery, Chapter 3.
X. Tang, J. Xu, and S. T. Chanson, editors, Springer, 2005
3. A.V. Goldberg, H. Kaplan, and R. Werneck.
Reach for A*: Shortest path algorithms with preprocessing.
DIMACS: Series in Discrete Mathematics and Theoretical Computer Science, 93–140.
C. Demetrescu, A.V. Goldberg, D.S. Johnson, editors, AMS, 2009 (See also Proc. No. 55 and Proc. No. 66).

BOOKS (edited)

1. H. Kaplan.
Proceedings of the 12th scandinavian symposium and workshops on algorithm theory, LNCS 6139, 2010.

PUBLICATIONS IN PROCEEDINGS OF CONFERENCES

1. M.C. Golumbic, H. Kaplan, and R. Shamir.
Algorithms and complexity of sandwich problems in graphs.
19th International Workshop on Graph Theoretic Concepts in Computer Science (WG), Springer Verlag, LNCS 790 (1994), 57–69.
2. H. Kaplan, R. Shamir, and R.E. Tarjan.
Tractability of parameterized completion problems on chordal and interval graphs: Minimum fill-in and physical mapping.
35th IEEE Annual Symposium on Foundations of Computer Science (FOCS), IEEE (1994), 780–791.
3. H. Kaplan and R.E. Tarjan
Persistent lists with catenation via recursive slow-down.
27th Annual ACM Symposium on Theory of Computing (STOC), ACM (1995), 93–102.
4. H. Kaplan and R. Shamir
Physical maps and interval sandwich problems: Bounded degrees help.
of the 5th Israeli Symposium on the Theory of Computing and Systems (ISTCS) (1996), 195–201.

5. H. Kaplan and R.E. Tarjan.
Purely functional representations of catenable sorted lists.
28th Annual ACM Symposium on Theory of Computing (STOC), ACM (1996), 202–211.
6. S. Arora, A. Frieze, and H. Kaplan.
A new rounding procedure for the assignment problem with applications to dense graph arrangement problems.
37th IEEE Annual Symposium on Foundations of Computer Science (FOCS), IEEE (1996), 21–30.
7. H. Kaplan, R. Shamir, and R.E. Tarjan.
Faster and simpler algorithm for sorting signed permutations by reversals.
8th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (1997), 344–351.
Also in Proc. RECOMB 97, p. 163.
8. A.L. Buchsbaum, H. Kaplan, A. Rogers, and J.R. Westbrook.
Linear-time pointer-machine algorithms for least common ancestors, MST verification, and dominators.
30th Annual ACM Symposium on Theory of Computing (STOC) ACM (1998), 279–288.
9. H. Kaplan, C. Okasaki, and R.E. Tarjan.
Simple confluent persistent catenable lists.
6th Skandinavian Workshop on Algorithm Theory (SWAT), Springer Verlag, LNCS 1432 (1998), 119–130.
10. H. Kaplan, M. Strauss, and M. Szegedy.
Just the fax–differentiating voice and fax phone lines using call billing data.
Proc. Symposium on Quantitative Analysis for Decision Making, AT&T (1998), 102–111 (See also Proc. No. 12).
11. M.C. Golumbic and H. Kaplan.
Cell flipping in permutation diagrams.
15th annual Symposium on Theoretical Aspects of Computer Science (STACS), Springer Verlag, LNCS 1373 (1998), 577–586.
12. H. Kaplan, M. Strauss, and M. Szegedy.
Just the fax–differentiating voice and fax phone lines using call billing data (abstract).
10th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (1999), S935–S936.
13. H. Kaplan and M. Szegedy.
On-line complexity of monotone set systems.
10th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (1999), 507–516.

14. E. Cohen and H. Kaplan.
LP-based analysis of greedy-dual-size.
10th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (1999), S879–S880.
15. E. Cohen, H. Kaplan, and J.D. Oldham.
Managing TCP connections under persistent HTTP.
8th International World Wide Web Conference (WWW8), Elsevier Science (1999) (Special issue of the Computer Networks journal, See also Jour. no. 10.).
16. E. Cohen and H. Kaplan.
Exploiting regularities in Web traffic patterns for cache replacement.
31st Annual ACM Symposium on Theory of Computing (STOC), ACM (1999), 109–118.
17. E. Cohen, H. Kaplan, and U. Zwick.
Connection caching.
31st Annual ACM Symposium on Theory of Computing (STOC), ACM (1999), 612–621.
18. H.N. Gabow, H. Kaplan, and R.E. Tarjan.
Unique maximum matching algorithms.
31st Annual ACM Symposium on Theory of Computing (STOC), ACM (1999), 70–78.
19. E. Cohen and H. Kaplan.
Prefetching the means for document transfer: A new approach for reducing web latency.
19th annual conference on Computer Communications (INFOCOM), IEEE (2000), 854–863.
20. E. Cohen, H. Kaplan, and U. Zwick.
Connection caching under various models of communication.
12th ACM Symposium on Parallel Algorithms and Architectures (SPAA), ACM (2000), 54–63.
21. N. Alon, H. Kaplan, M. Krivelevich, D. Malkhi, and J. Stern.
Scalable secure storage when half the system is faulty.
27th International Colloquium on Automata, Languages and Programming (ICALP), Springer (2000), LNCS 1853, 576–587.
22. E. Cohen and H. Kaplan.
Proactive caching of DNS records: Addressing a performance bottleneck.
First Symp. Applications and the Internet (SAINT), IEEE (2001), 85–94.
23. A. Fiat and H. Kaplan.
Making data structures confluent persistent.
12th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2001), 537–546.

24. S. Abiteboul, H. Kaplan and T. Milo.
Compact labeling schemes for ancestor queries.
12th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2001), 547–556.
25. H. Kaplan, R.E. Tarjan, and K. Tsioutsoulis.
Faster kinetic heaps and their use in broadcast scheduling.
12th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2001), 836–844.
26. E. Cohen and H. Kaplan.
The age penalty and its effect on cache performance.
3rd USENIX Symposium on Internet Technologies and Systems, March 2001.
27. E. Cohen and H. Kaplan.
Refreshment policies for Web content caches.
20th annual conference on Computer Communications (INFOCOM), IEEE (2001), vol. 1, 1398–1406.
28. E. Cohen, E. Halperin, and H. Kaplan.
Performance aspects of distributed caches using TTL-based consistency.
28th International Colloquium on Automata, Languages and Programming (ICALP), Springer Verlag, LNCS 2076 (2001), 744–756.
29. E. Cohen, H. Kaplan, and U. Zwick.
Competitive analysis of the LRFU paging algorithm.
7th International Workshop on Algorithms and Data Structures (WADS), Springer Verlag, LNCS 2125 (2001), 148–154.
30. H. Kaplan and T. Milo.
Short and simple labels for small distances and other functions.
7th International Workshop on Algorithms and Data Structures (WADS), Springer Verlag, LNCS 2125 (2001), 246–257.
31. A. Bremner-Barr, Y. Afek, H. Kaplan, E. Cohen, and M. Merritt.
Restoration by path concatenation: Fast recovery of MPLS paths.
20th ACM Symposium on Principles of Distributed Computing (PODC), ACM (2001), 43–52.
32. E. Cohen and H. Kaplan.
Aging through cascaded caches: Performance issues in the distribution of Web content.
ACM SIGCOMM’01, 2001, Special issue of the journal Computer Communication Review 41:4 (2001), 41–53.
33. H. Kaplan, N. Shafrir, and R. E. Tarjan.
Union-find with deletions.
13th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2002), 19–28.

34. H. Kaplan, T. Milo, and R. Shabo.
A comparison of labeling schemes for ancestor queries.
13th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2002), 954–963.
35. E. Cohen, E. Halperin, H. Kaplan, and U. Zwick.
Reachability and distance queries via 2-hop labels.
13th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2002), 937–946.
36. H. Kaplan, N. Shafrir, and R. E. Tarjan.
Meldable heaps and boolean union find.
34th Annual ACM Symposium on Theory of Computing (STOC), ACM (2002), 573–582.
37. E. Cohen, H. Kaplan, and T. Milo.
Labeling dynamic XML trees.
21st ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS), ACM (2002), 271–281.
38. S. Alstrup, C. Gavoille, H. Kaplan, and T. Rauhe.
Nearest common ancestors: A survey and a new distributed algorithm
14th Annual Symposium on Parallel Algorithms & Architectures (SPAA), ACM 2002, 258–264.
39. A. Barkan and H. Kaplan.
Partial alphabetic trees.
10th Annual European Symposium (ESA), Springer Verlag, LNCS 2461, 114–125, (2002).
40. E. Cohen and H. Kaplan.
Balanced-replication algorithms for distribution trees.
10th Annual European Symposium (ESA), Springer Verlag, LNCS 2461, 297–309, (2002).
41. A. Bremler-Barr, E. Cohen, H. Kaplan and Y. Mansour.
Predicting and bypassing end-to-end internet service degradations.
2nd Internet Measurement Workshop (IMW), 2002.
Online at www.icir.org/vern/imw-2002
42. E. Cohen, A. Fiat, and H. Kaplan.
A case for associative peer to peer overlays.
First Workshop on Hot Topics in Networks (HotNets-I), 2002.
Online at www.cs.washington.edu/hotnets (Appeared as a special issue of the Journal Computer Communication Review 33(1), 95–100 2003.)
43. E. Cohen, A. Fiat, and H. Kaplan.
Efficient sequences of trials.

- 14th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2003), 737 - 746.
44. E. Cohen, A. Fiat, and H. Kaplan.
Associative search in peer to peer networks: Harnessing latent semantics.
22nd annual conference on Computer Communications (INFOCOM), IEEE (2003),
available online from <http://www.ieee-infocom.org/2003/>.
45. H. Kaplan, E. Molad, and R. E. Tarjan.
Dynamic rectangular intersection with priorities.
35th Annual ACM Symposium on Theory of Computing (STOC), ACM (2003), 639-648.
46. Y. Azar, E. Cohen, A. Fiat, H. Kaplan, and H. Räcke.
Optimal oblivious routing in polynomial time.
35th Annual ACM Symposium on Theory of Computing (STOC), ACM (2003), 383-388.
47. H. Kaplan and E. Verbin.
Efficient data structures and a new randomized approach for sorting signed permutations by reversals.
14th Annual Symposium on Combinatorial Pattern Matching (CPM), Springer Verlag, LNCS 2676, 170-185, (2003).
48. H. Kaplan, M. Lewenstein, N. Shafrir, and M. Sviridenko.
Approximation algorithms for asymmetric TSP by decomposing directed regular multigraphs.
44th Annual IEEE Symposium on Foundations of Computer Science (FOCS), IEEE (2003), 56-65.
49. E. Cohen and H. Kaplan.
Efficient estimation algorithms for neighborhood variance and other moments.
15th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM (2004), 157 - 166.
50. E. Cohen and H. Kaplan.
Spatially-decaying aggregation over a network: Model and algorithms.
ACM SIGMOD International Conference on Management of Data, ACM (2004), 707-718.
51. P. K. Agarwal, S. Bereg, O. Daescu, H. Kaplan, S. Ntafos and B. Zhu.
Guarding a terrain by two watchtowers.
21st Annual ACM Symposium on Computational Geometry (SoCG), ACM (2005), 346-355
52. H. Kaplan, E. Kushilevitz, and Y. Mansour.
Learning with attribute costs.

- 37th Annual ACM Symposium on Theory of Computing (STOC), ACM (2005), 356 – 365.
53. G. Alexandron, H. Kaplan, and M. Sharir.
Kinetic and dynamic data structures for convex hulls and upper envelopes.
9th Workshop on Algorithms and Data Structures (WADS), Springer Verlag, LNCS 3608 (2005), 269–281.
54. H. Kaplan and M. Sharir.
Randomized incremental constructions of three-dimensional convex hulls and planar Voronoi diagrams, and approximate range counting.
17th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2006), 484–493.
55. A.V. Goldberg, H. Kaplan, and R. Werneck.
Reach for A*: Efficient point-to-point shortest path algorithms.
8th Workshop on Algorithm Engineering and Experiments (ALENEX) and the 3rd Workshop on Analytic Algorithmics and Combinatorics (ANALCO), SIAM (2006), 129–143.
56. H. Kaplan, S. Landau, and E. Verbin.
A simpler analysis of Burrows-Wheeler based compression.
17th Annual Symposium on Combinatorial Pattern Matching (CPM), Springer Verlag (2006), LNCS 4009, 282–293. (**Best paper award**)
57. A. Fiat, H. Kaplan, M. Levy, S. Olonetsky, and R. Shabo.
On the price of stability for designing undirected networks with fair cost allocations.
33rd International Colloquium on Automata, Languages and Programming (ICALP), Springer (2006), LNCS 4051, 608–618.
58. H. Kaplan, M. Sharir, and E. Verbin.
Colored intersection searching via sparse rectangular matrix multiplication.
22nd Annual ACM Symposium on Computational Geometry (SOCG), ACM (2006), 52–60.
59. H. Kaplan and Y. Nussbaum.
Certifying algorithms for recognizing proper circular-arc graphs and unit circular-arc graphs.
32nd International Workshop on Graph Theoretic Concepts in Computer Science (WG), Springer Verlag, LNCS 4271 (2006), 289–300.
60. H. Kaplan, E. Porat, and N. Shafrir.
Finding the position of the k-mismatch and approximate tandem repeats.
10th Scandinavian Workshop on Algorithm Theory (SWAT), Springer Verlag, LNCS 4059 (2006), 90–101.
61. H. Kaplan and Y. Nussbaum.
A simpler linear-time recognition of circular-arc graphs.

- 10th Scandinavian Workshop on Algorithm Theory (SWAT), Springer Verlag, LNCS 4059 (2006), 41–52.
62. E. Cohen, N. Grossaug, and H. Kaplan.
Processing top-k queries from samples.
2nd Conference on Future Networking Technologies (CoNEXT), ACM (2006), 76–87.
63. Y. Gijora and H. Kaplan.
Optimal dynamic vertical ray shooting in rectilinear planar subdivisions.
18th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2007), 19–28
64. H. Kaplan, N. Rubin, M. Sharir, and E. Verbin.
Counting colors in boxes.
18th ACM-SIAM Symposium on Discrete Algorithms (SODA), ACM-SIAM (2007), 785–794
65. P.K. Agarwal, H. Kaplan, and M. Sharir.
Computing the volume of the union of cubes.
23rd Annual ACM Symposium on Computational Geometry (SOCG), ACM (2007), 294–301.
66. A.V. Goldberg, H. Kaplan, and R. Werneck.
Better landmarks within reach.
6th Workshop on Experimental Algorithms (WEA), 2007. Springer Verlag, LNCS 4525 (2007), 38–51.
67. A. Fiat, H. Kaplan, M. Levy, S. Olonetsky.
Strong price of anarchy for machine load balancing.
34th International Colloquium on Automata, Languages and Programming (ICALP), Springer (2007), LNCS 4596, 583–594.
68. E. Cohen, N. Duffield, H. Kaplan, C. Lund and M. Thorup
Sketching unaggregated data streams for subpopulation-size queries
26th ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS), ACM (2007), 253–262.
69. H. Kaplan and E. Verbin.
Most Burrows-Wheeler based compressors are not optimal.
18th Annual Symposium on Combinatorial Pattern Matching (CPM), Springer (2007), LNCS 4580, 107–118.
70. E. Cohen and H. Kaplan.
Summarizing data using bottom-k sketches.
26th Annual ACM Symposium on Principles of Distributed Computing (PODC), ACM (2007), 225–234.

71. H. Kaplan, N. Rubin, and M. Sharir.
Linear data structures for fast ray-shooting amidst convex polyhedra.
15th Annual European Symposium on Algorithms (ESA), Springer (2007), LNCS 4698, 287–298.
72. E. Cohen, N. Duffield, H. Kaplan, C. Lund and M. Thorup.
Algorithms and estimators for accurate summarization of internet traffic.
7th ACM/SIGCOMM Internet Measurement Conference (IMC), ACM (2007), 265–278.
73. N. Alon, H. Kaplan, G. Nivasch, M. Sharir, and S. Smorodinsky.
Weak ϵ -nets and interval chains.
19th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM (2008), 1194–1203.
74. M. Hershcovitch and H. Kaplan.
I/O efficient dynamic data structures for longest prefix queries.
11th Skandinavian Workshop on Algorithm Theory (SWAT), Springer, LNCS 5124 (2008), 29–40.
75. H. Kaplan and N. Shafrir.
Finding path minima in incremental unrooted trees.
16th Annual European Symposium on Algorithms (ESA), Springer (2008), LNCS 5193, 565–576.
76. E. Cohen and H. Kaplan.
Estimating aggregates over multiple subsets.
8th IEEE International Conference on Data Mining (ICDM), IEEE (2008), 761–766.
77. H. Kaplan and U. Zwick.
A simpler implementation and analysis of Chazelle’s soft heaps.
20th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM (2009), 477–485.
78. E. Cohen, N. Duffield, H. Kaplan, C. Lund and M. Thorup.
Stream sampling for variance-optimal estimation of subset sums.
20th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM (2009), 1255–1264.
79. H. Kaplan, N. Rubin, and M. Sharir.
Line transversals of convex polyhedra in R^3 .
20th ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM (2009), 170–179.
80. E. Cohen and H. Kaplan.
Leveraging discarded samples for tighter estimation of multiple-set aggregates.
11th International Joint Conference on Measurement and Modeling of Computer Systems (SIGMETRICS), ACM (2009), 251–262.

81. D. Feldman, A. Fiat, H. Kaplan, and K. Nissim.
Private coresets.
41st Annual ACM Symposium on Theory of Computing (STOC), ACM (2009), 361–370.
82. H. Kaplan and Y. Nussbaum.
Maximum flow in directed planar graphs with vertex capacities.
17th Annual European Symposium on Algorithms (ESA), Springer (2009), LNCS 5757, 397–407.
83. E. Cohen, M. Feldman, A. Fiat, H. Kaplan, S. Olonetsky.
Envy-free makespan approximation: extended abstract.
11th ACM Conference on Electronic Commerce (EC), ACM (2010), 159–166.
84. P.K. Agarwal, J. Gao, L. J. Guibas, H. Kaplan, V. Koltun, N. Rubin, and M. Sharir.
Kinetic stable delaunay graphs.
26th Annual ACM Symposium on Computational Geometry (SOCG), ACM (2010), 127–136.
85. H. Kaplan, N. Rubin, and M. Sharir.
A kinetic triangulation scheme for moving points in the plane.
26th Annual ACM Symposium on Computational Geometry (SOCG), ACM (2010), 137–146.
86. R. Boim, H. Kaplan, T. Milo, and R. Rubinfeld.
Improved recommendations via (more) collaboration.
13th International Workshop on the Web and Databases (WebDB), ACM (2010).

Accepted to conferences

1. H. Kaplan, N. Rubin, and M. Sharir.
Improved bounds for geometric permutations.
The 51st Annual Symposium on Foundations of Computer Science (FOCS2010).

PATENTS FILED

1. E. Cohen, H. Kaplan, and Uri Zwick.
Method for preconnecting to a server on a network.
Filed October, 1998, issued August 12, 2003, US6606645.
2. E. Cohen and H. Kaplan.
Method for prewarming a connecting to a server on a network.
Filed November, 1999.
3. E. Cohen and H. Kaplan.
Method and Apparatus for Proactive Caching of Domain Name Server Records.
Filed November, 1999.

4. E. Cohen, A. Fiat, and H. Kaplan.
Associative search in peer-to-peer networks.
Filed May, 2003.

TECHNICAL REPORTS

1. H. Kaplan, M. Strauss, and M. Szegedy.
A study of facsimile capability by geographical regions.
Technical Report 98.19.1, AT&T, 1998.
2. E. Cohen, M. Feldman, A. Fiat, H. Kaplan, and S. Olonetsky.
Envy-Free Makespan Approximation.
CoRR, abs/0909.1072, 2009.