

BIBLIOGRAPHY

Books

1. (with J. E. Hopcroft) *Formal Languages and Their Relation to Automata*, Addison–Wesley, Reading, Mass., 1969. Japanese translation, 1971; Czech translation, 1978; unauthorized Chinese translation, 1979.
2. (with A. V. Aho) *The Theory of Parsing, Translation, and Compiling, Vol. I: Parsing*, Prentice–Hall, Englewood Cliffs, N. J., 1972. Japanese translation, 1977; unauthorized Russian translation, 1978.
3. (with A. V. Aho) *The Theory of Parsing, Translation, and Compiling, Vol. II: Compiling*, Prentice–Hall, Englewood Cliffs, N. J., 1973. Japanese translation, 1977; unauthorized Russian translation, 1978.
4. (with A. V. Aho and J. E. Hopcroft) *The Design and Analysis of Computer Algorithms*, Addison–Wesley, Reading, Mass., 1974. Japanese translation, 1977; Russian translation, 1979; Polish translation, 1980; Hungarian translation, 1983.
5. *Fundamental Concepts of Programming Systems*, Addison–Wesley, Reading, Mass., 1976. Japanese translation, 1981.
6. (with A. V. Aho) *Principles of Compiler Design*, Addison–Wesley, Reading, Mass., 1977. Japanese translation, 1979; World student edition, 1982; Spanish translation, 1983.
7. (with J. E. Hopcroft) *Introduction to Automata Theory, Languages, and Computation*, Addison–Wesley, Reading, Mass., 1979. Japanese translation, 1983; Spanish translation, 1983; Slovene translation, 1986; German translation, 1989; Polish translation, 1994;
8. *Principles of Database Systems*, Computer Science Press, New York, N. Y., 1980. Revised edition, 1982. Russian translation of first edition, 1983; unauthorized Chinese translation of first edition, 1985; Japanese translation of second edition, 1985; Indian edition, 1985; Polish translation of second edition, 1988.
9. (with J. E. Hopcroft and A. V. Aho) *Data Structures and Algorithms*, Addison–Wesley, Reading, Mass., 1983. Japanese translation, 1983; Spanish translation, 1985; French translation, 1988; Russian translation, 2000.
10. *Computational Aspects of VLSI*, Computer Science Press, New York, N. Y., 1984. Japanese translation, 1990; Russian translation, 1990.
11. (with A. V. Aho and R. Sethi) *Compiler Design: Principles, Tools, and Techniques*, Addison–Wesley, Reading, Mass., 1986. World student edition, 1987; Japanese translation, 1988; German translation, 1989; Chinese translation, 2001; Russian translation, 2001.
12. *Principles of Database and Knowledge-Base Systems* Computer Science Press, New York, N. Y.: Volume I: *Classical Database Systems*, 1988; Volume II: *The New Technologies*, 1989. Italian translation, 1991.
13. (with A. V. Aho) *Foundations of Computer Science* Computer Science Press, New York, N. Y., 1992. Italian translation, 1993; French translation, 1993; Chinese translation, 1993; Spanish translation, 1994; German translation, 1994; C Edition, 1994.
14. *Elements of ML Programming* Prentice–Hall, Englewood Cliffs, N. J., 1994. Japanese translation, 1996; second edition (ML97), 1998.
15. (with J. Widom) *A First Course in Database Systems* Prentice–Hall, Englewood Cliffs, N. J., 1997, 2002, 2008. Hungarian translation, 1999; Chinese translation 2000; Korean translation 2000; Polish translation 2000; Spanish translation, 2000; Italian translation, 2001; Hungarian translation of the third edition, 2009.

16. (with H. Garcia-Molina and J. Widom) *Database System Implementation* Prentice–Hall, Englewood Cliffs, N. J., 1999. Chinese translation, 2001.
17. (with J. E. Hopcroft and R. Motwani) *Introduction to Automata Theory, Languages, and Computation*, Addison–Wesley, Reading, Mass., 2001. Russian translation, 2002; Italian translation, 2003; Spanish translation, 2003; Polish translation, 2003; German translation, 2011.
18. (with H. Garcia-Molina and J. Widom) *The Complete Book of Database Systems* Prentice–Hall, Englewood Cliffs, N. J., 2002, 2009. Russian translation, 2003; Polish translation, 2003; World student edition, 2003.
19. (with A. V. Aho, M. Lam, and R. Sethi) *Compilers: Principles, Techniques, and Tools, Second Edition* Addison–Wesley, Reading, Mass. 2006.

Chapters of Books

20. (with A. V. Aho and R. Sethi) Code optimization and finite Church-Rosser systems, in *Design and Optimization of Compilers* (R. Rustin, ed.), Prentice–Hall, Englewood Cliffs, N. J., 1972.
21. Applications of language theory to compiler design, in *Currents in the Theory of Computing* (A. V. Aho, ed.), Prentice–Hall, Englewood Cliffs, N. J., 1973. Japanese translation, 1976.
22. (with M. R. Garey and R. L. Graham) An analysis of some packing algorithms, in *Combinatorial Algorithms* (R. Rustin, ed.), Algorithmics Press, New York, 1973.
23. (with J. E. Hopcroft) Fast algorithms for set manipulation, in *Combinatorial Algorithms* (R. Rustin, ed.), Algorithmics Press, New York, 1973.
24. Complexity of sequencing problems, in *Computer and Job Shop Scheduling Theory* (E. G. Coffman, ed.), John Wiley, New York, 1976.
25. The complexity of code generation, in *Algorithms and Complexity—New Directions and Recent Results* (J. Traub, ed.), Academic Press, New York, 1976.
26. An approach to processing queries in a logic-based query language, *Large Scale Knowledge Base and Reasoning Systems*, (M. Brodie and J. Mylopoulos, eds.), pp. 147–164, Springer-Verlag, 1986.
27. Logic and database systems, *ibid.*, pp. 121–124.
28. (with R. Fagin, G. M. Kuper, and M. Y. Vardi) Updating logical databases, *Advances in Computing Research* 3 (P. Kanellakis, ed.), pp. 1–18, JAI Press, 1986.
29. (with A. C. Hartmann) Model categories for theories of parallel systems, *Parallel Computing: Theory and Comparisons* (G. J. Lipovski and M. Malek, eds.) pp. 369–381, Wiley, New York, 1987.
30. Fast algorithms for deductive database queries, *Proc. NEC Conf. on Algorithms and Architectures*, SIAM Press, Phila., 1992.
31. The interface between language theory and database theory, in *Theoretical Studies in Computer Science*, a volume dedicated to Seymour Ginsburg on his 2⁸th birthday, pp. 133–152, Academic Press, Cambridge MA, 1992.
32. New frontiers in database system research, in *Future Tendencies in Computer Science, Control, and Applied Mathematics* (A. Bensoussan and J.-P. Verjus, eds.), pp. 87–104, Springer-Verlag, Berlin, 1992.
33. The parallel complexity of logical inference, in *Synthesis of Parallel Algorithms* (J. Reif, ed.), Morgan-Kaufmann, San Mateo, 1993.
34. Assigning an appropriate meaning to database logic with negation, in *Computers as Our Better Partners* (H. Yamada, Y. Kambayashi, and S. Ohta, eds.), pp. 216–225, World Scientific Press, March, 1994.

Journal Papers

35. Decoding of cyclic codes using position invariant functions, *IBM J.* **9** (July, 1965), pp. 233-240.
36. Near optimal single synchronization error correcting code, *IEEE Trans. on Information Theory* **IT-12** (Oct., 1966), pp. 418-424.
37. On the capabilities of codes to correct synchronization errors, *IEEE Trans. on Information Theory* **IT-13** (Jan., 1967), pp. 95-107.
38. (with J. E. Hopcroft) Nonerasing stack automata, *J. Computer and System Sciences* **1** (Aug., 1967), pp. 166-188.
39. (with J. E. Hopcroft) An approach to a unified theory of automata, *BSTJ* **46** (Oct., 1967), pp. 1793-1829.
40. (with J. E. Hopcroft) Decidable and undecidable questions about automata, *J. ACM* **15** (Apr., 1968), pp. 317-324.
41. (with J. E. Hopcroft) Deterministic stack automata and the quotient operator, *J. Computer and System Sciences* **2** (June, 1968), pp. 1-12.
42. (with A. V. Aho) The theory of languages, *si Math. Syst. Th.* **2** (June, 1968), pp. 97-125.
43. (with J. E. Hopcroft) Relations between time and tape complexities, *J. ACM* **15** (July, 1968), pp. 414-429.
44. (with J. E. Hopcroft) Sets accepted by one way stack automata are context sensitive, *Information and Control* **13** (Aug., 1968), pp. 114-133.
45. (with A. V. Aho and J. E. Hopcroft) Time and tape complexity of pushdown automaton languages, *Information and Control* **13** (Sept., 1968), pp. 186-206.
46. (with J. E. Hopcroft) Some results on tape bounded Turing machines, *J. ACM* **16** (Jan., 1969), pp. 168-177.
47. (with A. V. Aho) Syntax directed translation and the pushdown assembler, *J. Computer and System Sciences* **3** (Feb., 1969), pp. 37-56.
48. (with P. Weiner) Uniform synthesis of sequential machines, *BSTJ* **48** (May, 1969), pp. 1115-1127.
49. (with A. D. Friedman and R. L. Graham) Universal single transition time asynchronous state assignments, *IEEE Trans. on Computers* **C-18** (June, 1969), pp. 541-547.
50. (with A. V. Aho) Properties of syntax directed translations, *J. Computer and System Sciences* **3** (Aug., 1969), pp. 319-334.
51. (with A. V. Aho and J. E. Hopcroft) A general theory of translation, *Math. Syst. Th.* **3** (Sept., 1969), pp. 193-221.
52. Halting stack automata, *J. ACM* **16** (Oct., 1969), pp. 550-563.
53. (with A. V. Aho and J. E. Hopcroft) On the computational power of pushdown automata, *J. Computer and System Sciences* **4** (Apr., 1970), pp. 129-136.
54. (with R. Sethi) The generation of optimal code for arithmetic expressions, *J. ACM* **17** (Oct., 1970), pp. 715-728. Reprinted in *Compiler Techniques* (B. W. Pollack, ed.), Auerbach Press, New York.
55. (with A. V. Aho) A characterization of two way deterministic classes of languages, *J. Computer and System Sciences* **4** (Dec., 1970), pp. 523-538.
56. (with A. V. Aho and P. J. Denning) Principles of optimal page replacement, *J. ACM* **18** (Jan., 1971), pp. 80-93.

57. (with A. V. Aho) Characterizations and extensions of pushdown translations, *Math. Syst. Th.* **5** (June, 1971), pp. 172–192.
58. (with A. V. Aho) Translations on a context free grammar, *Information and Control* **19** (Dec., 1971), pp. 439–475.
59. (with A. V. Aho) Optimization of straight line programs, *SIAM J. Computing* **1** (Mar., 1972), pp. 1–19.
60. (with A. V. Aho and P. J. Denning) Weak and mixed strategy precedence parsing, *J. ACM* **19** (Apr., 1972), pp. 225–243.
61. (with A. V. Aho) Equivalence of programs with structured variables, *J. Computer and System Sciences* **6** (Apr., 1972), pp. 125–137.
62. (with M. S. Hecht) Flow graph reducibility, *SIAM J. Computing* **1** (June, 1972), pp. 188–202.
63. (with A. V. Aho and M. R. Garey) The transitive reduction of a directed graph, *SIAM J. Computing* **1** (June, 1972), pp. 131–137.
64. A note on the efficiency of hashing functions, *J. ACM* **19** (July, 1972), pp. 569–575.
65. (with P. Weiner) Modular decomposition and nondeterministic sequential machines, *IEEE Trans. on Computers* **C-21** (Oct., 1972), pp. 1124–1128.
66. (with A. V. Aho) Optimization of LR(k) parsers, *J. Computer and System Sciences* **6** (Dec., 1972), pp. 573–602.
67. (with A. V. Aho) Linear precedence functions for weak precedence grammars, *Intl. J. Comp. Math.* **3** (Dec., 1972), pp. 149–155.
68. (with A. V. Aho) Error detection in precedence parsers, *Math. Syst. Th.* **7** (June, 1973), pp. 97–113.
69. (with A. V. Aho) A technique for speeding up LR(k) parsers, *SIAM J. Computing* **2** (June, 1973), pp. 106–127.
70. (with A. Birman) Parsing algorithms with backtrack, *Information and Control* **23** (Aug., 1973), pp. 1–34.
71. (with J. E. Hopcroft) Set merging algorithms, *SIAM J. Computing* **2** (Dec., 1973), pp. 294–303.
72. Fast algorithms for the elimination of common subexpressions, *Acta Inf.* **2** (Dec. 1973), pp. 191–213.
73. (with A. V. Aho) Dynamic memories with fast sequential access, *IEEE Trans. on Computers* **C-23** (Mar., 1974), pp. 272–276.
74. (with M. S. Hecht) Characterizations of reducible flow graphs, *J. ACM* **21** (July, 1974), pp. 367–375.
75. (with D. S. Johnson, A. Demers, M. R. Garey, and R. L. Graham) Worst-case performance bounds for simple one-dimensional packing algorithms, *SIAM J. Computing* **4** (Dec., 1974), pp. 299–325. Reprinted in *Best Computer Papers of 1974*, Auerbach Press, 1975.
76. \mathcal{NP} -complete scheduling problems, *J. Computer and System Sciences* **10** (June, 1975), pp. 384–393.
77. (with A. V. Aho and S. C. Johnson) Deterministic parsing of ambiguous grammars, *Comm. ACM* **18** (Aug., 1975), pp. 441–452.
78. (with H. B. Hunt III and T. G. Szymanski) On the complexity of LR(k) testing, *Comm. ACM* **18** (Dec., 1975), pp. 707–715.
79. (with M. S. Hecht) A simple algorithm for global flow analysis problems, *SIAM J. Computing* **4** (Dec., 1975), pp. 519–532.
80. (with A. V. Aho and K. Steiglitz) Evaluating polynomials at fixed sets of points, *SIAM J. Computing* **4** (Dec., 1975), pp. 533–539.

81. (with J. Kam) Global optimization problems and iterative algorithms, *J. ACM* **23** (Jan., 1976), pp. 158–171.
82. (with A. V. Aho and D. S. Hirschberg) Bounds on the complexity of the longest common subsequence problem, *J. ACM* **23** (Jan., 1976), pp. 1–12.
83. (with A. V. Aho and J. E. Hopcroft) On finding lowest common ancestors in trees, *SIAM J. Computing* **5** (Mar., 1976), pp. 115–132.
84. (with M. A. Harrison and W. L. Ruzzo) Protection in operating systems, *Comm. ACM* **19** (Aug., 1976), pp. 461–470.
85. (with A. V. Aho) Node listings for reducible flow graphs, *J. Computer and System Sciences* **13** (Dec., 1976), pp. 286–299.
86. (with A. V. Aho and S. C. Johnson) Code generation for expressions with common subexpressions, *J. ACM* **24** (Jan., 1977), pp. 146–160.
87. (with H. B. Hunt III and T. G. Szymanski) Operations on sparse relations, *Comm. ACM* **20** (Mar., 1977), pp. 171–176.
88. (with J. Kam) Monotone data flow analysis frameworks, *Acta Inf.* **7** (Mar., 1977), pp. 305–317.
89. (with T. G. Szymanski) Evaluating relational expressions with sparse and dense arguments, *SIAM J. Computing* **6** (Mar., 1977), pp. 109–122.
90. (with J. Kam) A model of statistical databases and their security, *ACM Trans. on Database Systems* **2** (Mar., 1977), pp. 1–10.
91. (with M. M. Geller, H. B. Hunt III, and T. G. Szymanski) Economy of description by parsers, DPDA's and PDA's, *Theoretical Computer Science* **4** (Apr., 1977), pp. 143–154.
92. (with A. C. Fong) Finding the depth of a flow graph, *J. Computer and System Sciences* **15** (Dec., 1977), pp. 300–309.
93. (with A. V. Aho and Y. Sagiv) Equivalences among relational expressions, *SIAM J. Computing* **8** (May, 1979), pp. 218–246.
94. (with A. V. Aho) Optimal partial match retrieval when fields are independently specified, *ACM Trans. on Database Systems* **4** (June, 1979), pp. 168–179.
95. (with A. V. Aho and C. Beeri) The theory of joins in relational databases, *ACM Trans. on Database Systems* **4** (Sept., 1979), pp. 297–314.
96. (with A. V. Aho and Y. Sagiv) Efficient optimization of a class of relational expressions, *ACM Trans. on Database Systems* **4** (Dec., 1979), pp. 435–454.
97. (with M. A. Kaplan) A scheme for the automatic inference of variable types, *J. ACM* **27** (Jan., 1980), pp. 128–145.
98. (with D. Maier, A. O. Mendelzon, and F. Sadri) Adequacy of decompositions of relational databases, *J. Computer and System Sciences* **21** (Dec., 1980), pp. 368–379.
99. (with C. Beeri, A. O. Mendelzon, and Y. Sagiv) Equivalence of relational database schemes, *SIAM J. Computing* **10** (May, 1981), pp. 352–370.
100. (with A. V. Aho, T. G. Szymanski, and Y. Sagiv) Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions, *SIAM J. Computing* **10** (Aug., 1981), pp. 405–421.
101. (with S. Toueg) Deadlock-free packet switching networks, *SIAM J. Computing* **10** (Aug., 1981), pp. 594–611.

102. (with F. Sadri) The theory of functional and template dependencies, *Theoretical Computer Science* **17** (March, 1982), pp. 317–332.
103. (with F. Sadri) Template dependencies: a large class of dependencies in relational databases and its complete axiomatization, *J. ACM* **29** (April, 1982), pp. 363–372.
104. (with R. W. Floyd) The compilation of regular expressions into integrated circuits, *J. ACM* **29** (July, 1982), pp. 603–622.
105. (with A. V. Aho, A. Wyner, and M. Yannakakis) Bounds on the size and transmission rate of communication protocols, *Computers and Mathematics with Applications* **8** (July, 1982), pp. 205–214.
106. (with R. Fagin and A. O. Mendelzon) A simplified universal relation assumption and its properties, *ACM Trans. on Database Systems* **7** (Sept., 1982), pp. 343–360.
107. (with R. Fagin, D. Maier, and M. Yannakakis) Tools for template dependencies, *SIAM J. Computing* **12** (Feb., 1983), pp. 36–59.
108. (with D. Maier) Maximal objects and the semantics of universal relation database systems, *ACM Trans. on Database Systems* **8** (March, 1983), pp. 1–14.
109. On Kent’s “Consequences of assuming a universal relation,” *ACM Trans. on Database Systems* **8** (Dec., 1983), pp. 637–643.
110. (with D. Maier and M. Y. Vardi) On the foundations of the universal relation model, *ACM Trans. on Database Systems* **9** (June, 1984), pp. 283–308.
111. (with D. Maier) Connections in acyclic hypergraphs, *Theoretical Computer Science* **32** (July, 1984), pp. 185–200.
112. (with H. F. Korth, G. M. Kuper, J. Feigenbaum, and A. van Gelder) System/U: a database system based on the universal relation assumption, *ACM Trans. on Database Systems* **9** (Sept., 1984), pp. 331–347.
113. Flux, sorting, and supercomputation for AI applications, *J. Parallel and Distributed Computing* **1** (Nov., 1984), pp. 133–151.
114. Implementation of logical query languages for database systems, *ACM Trans. on Database Systems* **10** (Sept., 1985), pp. 289–321. (Invited paper from 1985 SIGMOD Conf.)
115. (with C. H. Papadimitriou) A communication-time tradeoff, *SIAM J. Computing* **16** (Aug., 1987), pp. 639–647.
116. (with A. R. Karlin and H. W. Trickey) Algorithms for the compilation of regular expressions into PLA’s, *Algorithmica* **2** (Oct., 1987), pp. 283–324.
117. (with A. Van Gelder) Parallel complexity of logical query programs, *Algorithmica* **3** (Jan., 1988), pp. 5–42.
118. (with A. Van Gelder) Testing applicability of top-down capture rules, *J. ACM* **35** (April, 1988), pp. 345–373.
119. (with F. Afrati, C. H. Papadimitriou, G. Papageorgiou, A. Roussou, and Y. Sagiv) On the convergence of query evaluation, *J. Computer and System Sciences* **38** (April, 1989), pp. 341–359.
120. (with D. Peleg) An optimal synchronizer for the hypercube, *SIAM J. Computing* **18** (Aug. 1989), pp. 740–747.
121. (with M. Yannakakis) High-probability transitive closure algorithms, *SIAM J. Computing* **20** (Feb., 1991), pp. 100–125.
122. (with M. Yannakakis) The I/O complexity of transitive closure, *J. Mathematics and AI* **3** (July, 1991), pp. 331–360.

123. (with A. Silberschatz and M. Stonebraker, eds.) Database systems: achievements and opportunities, *Comm. ACM* **34** (Oct., 1991), pp. 110–120. Also in *SIGMOD Record* **19**:4, pp. 6–22, Dec., 1990.
124. (with R. Ramakrishnan, Y. Sagiv, and M. Vardi) Logical query optimization by proof-tree transformation, *J. Computer and System Sciences* **47** (Aug., 1993), pp. 222–248.
125. (with R. Ramakrishnan) A survey of research in deductive database systems, *J. Logic Programming* **23**:2 (May, 1995), pp. 125–149.
126. (with J. F. Naughton, R. Ramakrishnan, and Y. Sagiv) Argument reduction through factoring, *Theoretical Computer Science* **146** (July, 1995), pp. 269–310.
127. (with H. Garcia-Molina, Y. Papakonstantinou, D. Quass, A. Rajaraman, Y. Sagiv, V. Vassalos, and J. Widom) The TSIMMIS approach to mediation: data models and languages, *J. Intelligent Information Systems* **8**:2 (March, 1997), pp. 117–132.
128. (with D. Quass, A. Rajaraman, Y. Sagiv, and J. Widom) Querying semistructured heterogeneous information, *J. Systems Integration* **7**:3–4 (Sept., 1997), pp. 381–407.
129. (with A. Levy and A. Rajaraman) Answering queries using limited external processors, *J. Computer and System Sciences* **58** (Feb., 1999), pp. 69–82.
130. (with R. Motwani, E. Cohen, M. Datar, S. Fujiwara, A. Gionis, P. Indyk, and C. Yang) Finding interesting associations without support pruning, *IEEE Trans. on Knowledge and Data Engineering* **13**:1 (Jan., 2001), pp. 64–78.
131. (with A. Rajaraman) Querying websites using compact skeletons, *J. CSS* **66**:4, pp. 809–851, 2003.
132. (with F. N. Afrati and C. Li) Using views to generate efficient evaluation plans for queries, *J. CSS* **73**:5, pp. 703–724, 2007.
133. Advising students for success, *Comm. ACM* **52**:3, pp. 34–37, 2009.

Conference Papers and Technical Reports (Not Otherwise Published)

134. (with C. V. Freiman) Optimal multiplier trees of carry save adders with mixed delay, *Proc. IFIP 65* (Aug., 1965), pp. 632–633.
135. (with J. E. Hopcroft) A survey of formal language theory, *Proc. First Annl. Princeton Conf. on Inf. Sciences and Systems* (Mar., 1967), pp. 68–75.
136. (with A. V. Aho and R. Sethi) A formal approach to code optimization, *Proc. SIGPLAN Conf. on Compiler Optimization* (July, 1970), pp. 86–100.
137. (with J. E. Hopcroft) An $n \log n$ algorithm for detecting reducible flow graphs, *Proc. Sixth Annual Princeton Conf. on Information Sciences and Systems* (March, 1972), pp. 119–122.
138. Applications of language theory to compiler design, *Proc. SJCC* (May, 1972), pp. 235–240.
139. (with A. C. Fong and J. Kam) Applications of lattice algebra to loop optimization, *Proc. Second ACM Conf. on Principles of Programming Languages* (Jan., 1975), pp. 1–9.
140. (with F. Y. Chin) Asymptotic complexity of partial fraction expansion, TR-172, Computer Science Lab., Dept. of EECS, Princeton Univ., Princeton, N. J.
141. Data flow analysis, *Proc. USA–Japan Computer Conf.* (Aug., 1975), pp. 335–342.
142. (with A. C. Fong) Induction variables for very high level languages, *Proc. Third ACM Conf. on Principles of Programming Languages* (Jan., 1976), pp. 104–112.
143. (with A. V. Aho and S. C. Johnson) Code generation for machines with multiregister operations, *Proc. Fourth ACM Conf. on Principles of Programming Languages* (Jan., 1977), pp. 21–28.

144. (with A. V. Aho) Universality of a data retrieval language, *Proc. Sixth ACM Conf. on Principles of Programming Languages* (Jan., 1979), pp. 110–120.
145. (with H. Korth) SYSTEM/U: a database system based on the universal relation assumption, *Proc. XPI Conference* (June, 1980), Stonybrook, NY.
146. (with D. Dolev, K. Karplus, A. Siegel, and A. Strong) Optimal wiring between rectangles, *Proc. Thirteenth ACM Symposium on Theory of Computing* (May, 1981), pp. 312–317.
147. (with C. Beeri, R. Fagin, D. Maier, A. O. Mendelzon, and M. Yannakakis) Properties of acyclic database schemes, *Proc. Thirteenth ACM Symposium on Theory of Computing* (May, 1981), pp. 355–362.
148. A view of directions in relational database theory (invited paper), *Automata, Languages, and Programming* (S. Even and O. Kariv, eds.), Springer-Verlag, Berlin (July, 1981), pp. 165–176.
149. (with H. W. Trickey) A regular expression compiler, *IEEE COMPCON* (Feb., 1982), pp. 345–348.
150. (with D. Dolev, K. Karplus, A. Siegel, and A. Strong) Optimal algorithms for structural assembly, *VLSI Design*, pp. 38–43, March, 1982.
151. The U. R. strikes back, *Proc. First ACM Symp. on Principles of Database Systems* (March, 1982), pp. 10–22.
152. Combining state machines with regular expressions for automatic synthesis of VLSI circuits, STAN-CS-82-927, Computer Science Dept., Stanford Univ., Sept., 1982.
153. Query processing in universal relation systems, *Database Engineering* 5 (Sept., 1982), pp. 6–10.
154. (with R. Fagin and M. Y. Vardi) On the semantics of updates in databases, *Proc. Second ACM Symp. on Principles of Database Systems* (March, 1983), pp. 352–365.
155. (with A. V. Aho and M. Yannakakis) On notions of information transfer in VLSI circuits, *Proc. Fifteenth ACM Symposium on Theory of Computing* (May, 1983), pp. 133–139.
156. (with D. Maier) Fragments of relations, *Proc. ACM Symp. on Management of Data* (June, 1983).
157. Universal relation interfaces for database systems, *Proc. IFIP 83* (Sept., 1983), pp. 243–252. Invited paper.
158. (with A. R. Karlin and H. W. Trickey) Experience with a regular expression compiler, *Proc. IEEE Conf. on Computer Design/VLSI in Computers*, pp. 656–665, Oct., 1983.
159. Some thoughts on supercomputer organization, *IEEE COMPCON*, pp. 424–432, Feb., 1984.
160. (with G. M. Kuper and M. Y. Vardi) On the equivalence of logical databases, *Proc. Third Annual ACM Symp. on Principles of Database Systems*, April, 1984, pp. 221–228.
161. VLSI complexity and supercomputer organization, *Proc. Fourth Jerusalem Conf. on Information Technology*, May, 1984, pp. 646–649.
162. (with D. Dolev, D. Maier, and H. Mairson) Error-correcting codes for write-once memory, *Proc. Sixteenth ACM Symp. on Theory of Computing*, May, 1984, pp. 225–229.
163. (with A. Keller) On complementary and independent mappings in databases, *Proc. ACM Conf. on Management of Data*, June, 1984, pp. 143–148.
164. (with Y. Sagiv) Analysis of a top-down capture rule, STAN-CS-84-1009, July, 1984.
165. (with F. Bancilhon, D. Maier, and Y. Sagiv) Magic sets and other strange ways to implement logic programs, *Proc. Fifth ACM Symp. on Principles of Database Systems* (March, 1986), pp. 1–15.
166. (with K. Morris and A. Van Gelder) Design overview of the NAIL! system, (invited paper), *Proc. Third Intl. Symp. on Logic Programming*, pp. 554–568, July, 1986.

167. Database theory: past and future, (invited paper), *Proc. Sixth ACM Symp. on Principles of Database Systems* (March, 1987), pp. 1–10.
168. (with K. Morris, J. F. Naughton, Y. Saraiya, and A. Van Gelder) An overview of the NAIL! system, *Database Engineering*, Dec., 1987, pp. 28–43.
169. (with M. Y. Vardi) The complexity of ordering subgoals, *Proc. Seventh ACM Symp. on Principles of Database Systems* (March, 1988), pp. 74–81.
170. Review of *Mathematics of Information Processing* (M. Anshel and W. Gewirtz, eds.) in *SIAM Reviews* **30**:1, pp. 165–168, March, 1988.
171. Bottom-up beats top-down for datalog, *Proc. Eighth ACM Symp. on Principles of Database Systems* (March, 1989), pp. 140–149.
172. (with J. F. Naughton, R. Ramakrishnan, and Y. Sagiv) Efficient evaluation of right-, left-, and multilinear rules, *Proc. ACM Symp. on Management of Data* (June, 1989), pp. 235–242.
173. The theory of deductive database systems, (invited paper), *IEEE Compcon*, pp. 496–502, Feb., 1990.
174. (with C. Zaniolo) Deductive databases: achievements and future directions, *SIGMOD Record* **19**:4, Dec., 1990.
175. A comparison of deductive and object-oriented database systems, (invited paper) *Proc. 2nd Intl. Conf. on Deductive and Object-Oriented Databases*, pp. 263–277, Springer Verlag, Berlin, Dec., 1991.
176. (with A. Gupta) Generalizing conjunctive query containment for view maintenance and integrity constraint verification, *Joint Intl. Conf. on Logic Programming, Workshop on Deductive Databases*, Aug., 1992.
177. (with H. Craig Howard, K. Law, P. Teicholz, A. Keller, S. Tiwari, A. Gupta, and K. Krishnamurthy) Integrated data exchange and concurrent design for engineered facilities, *AAAS Workshop on Advances in Data Management for the Scientist and Engineer*, Feb., 1993, pp. 47–56.
178. (with A. Gupta, Y. Sagiv, and J. Widom) Constraint checking with partial information, *Proc. Thirteenth ACM Symp. on Principles of Database Systems* (May, 1994), pp. 45–55.
179. (with H. Craig Howard, K. Law, P. Teicholz, A. Keller, S. Tiwari, A. Gupta, and K. Krishnamurthy) Versions, configurations, and constraints in CEDB, unpublished memorandum, July, 1994.
180. The role of theory today, unpublished memorandum, Jan., 1995, to appear in *Computing Surveys*.
181. (with S. Chawathe, H. Garcia-Molina, J. Hammer, K. Ireland, Y. Papakonstantinou, and J. Widom) The TSIMMIS project: integration of heterogeneous information sources, *AAAI Spring Symp. on Information Gathering*, Palo Alto, March, 1995.
182. (with A. Keller) A version numbering scheme with a useful lexicographic order, *Intl. Conf. on Data Engg.*, March, 1995.
183. (with A. Rajaraman and Y. Sagiv) Query optimization using templates with binding patterns, *Proc. Fourteenth ACM Symp. on Principles of Database Systems* (May, 1995), pp. 105–112.
184. (with Y. Papakonstantinou, A. Gupta, and H. Garcia-Molina) A query translation scheme for rapid implementation of wrappers, Fourth DOOD, Singapore, Dec., 1995.
185. (with Y. Papakonstantinou and H. Garcia-Molina) Medmaker: a mediation system based on declarative specifications, ICDE, Feb., 1996.
186. (with A. Silberschatz and M. Stonebraker, eds.) Database research; achievements and opportunities into the 21st century, *SIGMOD Record* **25**:1 (March, 1996), pp. 52–63.
187. (with V. Harinarayan and A. Rajaraman) Implementing data cubes efficiently, *ACM SIGMOD Intl. Symp. on Management of Data* (June, 1996), pp. 205–216. Winner: Best Paper Award.

188. (with A. Rajaraman) Information integration by outerjoins and full disjunctions, *Proc. Fifteenth ACM Symp. on Principles of Database Systems* (June, 1996), pp. 238–248.
189. Efficient implementation of data cubes via materialized views, *Proc. 2nd Intl. Symp. on Knowledge Discovery and Data Mining*, Aug., 1996, pp. 386–388.
190. The database approach to knowledge representation, *Proc. 13th Natl. Conf. on AI*, Aug., 1996.
191. Information integration using logical views, *Intl. Conf. on Database Theory*, Jan., 1997.
192. (with S. Nestorov, S. Chawathe, and J. Wiener) Representative objects: concise representation of hierarchical objects, *ICDE*, May, 1997.
193. (with H. Gupta, V. Harinarayan, and A. Rajaraman) Index Selection for OLAP, *ICDE*, May, 1997.
194. (with S. Brin, R. Motwani, and S. Tsur) Dynamic itemset counting and implication rules for market-basket data, *SIGMOD*, May, 1997.
195. (with S. Tsur, S. Abiteboul, C. Clifton, R. Motwani, S. Nestorov, and A. Rosenthal) Query flocks: a generalization of association-rule mining, *SIGMOD*, June 1998, pp. 1–12.
196. (with C. Li, R. Yerneni, V. Vassilos, H. Garcia-Molina, Y. Papakonstantinou, and M. Valiveti) Capability-based mediation in TSIMMIS, *SIGMOD*, June, 1998, pp. 564–566.
197. (with M. Fang, N. Shivakumar, H. Garcia-Molina, and R. Motwani) Computing iceberg queries efficiently, *VLDB*, Aug., 1998, pp. 299–310.
198. (with C. Silverstein, S. Brin, and R. Motwani) Scalable techniques for mining causal structures, *VLDB*, Aug., 1998, pp. 594–605.
199. (with R. Yerneni, C. Li, and H. Garcia-Molina) Optimizing large join queries in mediator systems, *ICDT* (Jan., 1999), pp. 348–364.
200. (with R. Yerneni, C. Li, and H. Garcia-Molina) Computing capabilities of mediators, *SIGMOD*, June, 1999, pp. 443–454.
201. (with S. Fujiwara and R. Motwani) Dynamic miss-counting algorithms: finding implication and similarity rules with confidence pruning, *ICDE*, Feb, 2000.
202. (with C. Li and M. Bawa) Minimizing view sets without losing query-answering power, *ICDT*, Jan., 2001.
203. Improving the efficiency of database-system teaching, *SIGMOD*, June 2003, pp. 1–3, 2003.
204. Gradiance on-line accelerated learning, *ACSC*, pp. 3–6, 2005.
205. (with A. das Sarma and J. Widom) Schema design for uncertain databases, *Alberto Mendelzon Workshop*, 2009.
206. (with F. N. Afrati) Optimizing joins in a map-reduce environment, *EDBT*, March, 2010.