

# RONALD FAGIN

IBM Fellow  
IBM Research – Almaden  
650 Harry Road  
San Jose, California 95120-6099

Phone: 408-927-1726

Fax: 845-491-2916

Email: fagin@us.ibm.com

Home page: <http://researcher.ibm.com/person/us-fagin>

## Education

Ph.D. in Mathematics, University of California at Berkeley, 1973

Thesis: *Contributions to the Model Theory of Finite Structures*

Advisor: Prof. Robert L. Vaught

National Science Foundation Graduate Fellowship 1967-72

Research Assistantship 1972-73

Passed Ph.D. Qualifying Exams “With Distinction” (top 5%).

Bachelor’s Degree in Mathematics, Dartmouth College, 1967

Held National Merit Scholarship, 1963-67

Elected to Phi Beta Kappa, 1966

Graduated Summa cum laude

Graduated With Highest Distinction in Mathematics.

## Professional Experience

IBM Fellow, 2012–present

IBM Research – Almaden (formerly IBM Research Laboratory), San Jose, California, 1975–present

- Manager, Foundations of Computer Science, 1979–2012
- Acting Senior Manager, Principles and Methodologies, 2004
- Acting Senior Manager, Mathematics and Related Computer Science, 1995

Research Fellow, IBM Haifa Research Laboratory, 1996-1997 (and was visitor to IBM Haifa during summers of 1998, 1999, and 2000)

Visiting Professor, Pontifícia Universidade Católica do Rio de Janeiro, 1981

IBM Watson Research Center, Yorktown Heights, N.Y., 1973–1975

- Served as Technical Assistant to Director of Computer Science Dept.

Computer programmer for Kiewit Computation Center, Dartmouth College, Hanover, N.H., 1967

Research Assistant to John G. Kemeny, Chairman of Mathematics Dept., Dartmouth College, Hanover, N.H., 1965–1967

## Honors

1. “Ron Fagin Special Event” was held at the 2016 ACM SIGMOD/PODS Conference “in honor of Ron Fagin’s influential and long-lasting research contributions to the principles of database systems”.
2. Presented one of two “Gems of PODS” talks (for research on score aggregation) at the 2016 ACM SIGMOD/PODS Conference.
3. The paper “Data exchange: semantics and query answering” (coauthors Phokion Kolaitis, Renée J. Miller, and Lucian Popa) was named the second most highly cited paper of the decade 2005-2014 for the journal *Theoretical Computer Science*.
4. Won IBM Corporate Award for “Pioneering Research in Rank and Score Aggregation”, 2015.
5. Won Best Paper Award, 2015 International Conference on Database Theory (ICDT), for the paper “A Declarative Framework for Linking Entities” (co-authors Douglas Burdick, Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan).
6. Had research on “Rank and Score Aggregation” upgraded to an “Outstanding” IBM Research Division Accomplishment for 2014.
7. Elected to the American Academy of Arts and Sciences, 2014.
8. Elected to the US National Academy of Engineering, 2014, for “contributions to theory and practice of data management”.
9. Won Alberto O. Mendelzon Test-of-Time award of the 2014 ACM Symposium on Principles of Database Systems (PODS) for 2004 PODS paper “Composing schema mappings: second-order dependencies to the rescue” (coauthors Phokion Kolaitis, Lucian Popa and Wang-Chiew Tan).
10. Won 2014 Gödel Prize for 2003 paper “Optimal aggregation algorithms for middleware” (coauthors Amnon Lotem and Moni Naor).
11. Won Test-of-Time award of ICDT 2013 for 2003 ICDT paper “Data exchange: semantics and query answering” (coauthors Phokion Kolaitis, Renée J. Miller, and Lucian Popa)
12. Named an IBM Fellow, 2012
13. Won IEEE Computer Society W. Wallace McDowell Award “for fundamental and lasting contributions to the theory of databases”, 2012.

14. Won Alberto O. Mendelzon Test-of-Time award of PODS 2011 for 2001 PODS paper “Optimal aggregation algorithms for middleware” (coauthors Amnon Lotem and Moni Naor).
15. Won IEEE Computer Society Technical Achievement Award “for pioneering contributions to the theory of rank and score aggregation”, 2011.
16. Won IBM Corporate Award for “Advances in Database Theory”, 2010.
17. Won Best Paper Award, 2010 ICDT, for the paper “Composition with target constraints” (coauthors Marcelo Arenas and Alan Nash).
18. Had research on Tivoli Storage Manager (in particular differential backup) upgraded to an “Extraordinary” IBM Research Division Accomplishment for 2009.
19. Received IBM Outstanding Innovation Award, 2009, for “Advances in database theory”.
20. Had research on “Schema mappings research: theory and practice” upgraded to an “Outstanding” IBM Research Division Accomplishment for 2007, and awarded an IBM Outstanding Technical Achievement Award for this work in 2008.
21. Elected to IBM Academy of Technology, 2007, for “fundamental contributions to computer science theory and its application to IBM products”.
22. The paper “Functional dependencies in a relational database and propositional logic” was selected as a significant paper published in the first 50 years of the IBM journals, 2007.
23. Received IBM supplemental Patent Issue Award, given for key IBM patents, for encryption keys patent, 2007.
24. Named Fellow of the American Association for the Advancement of Science, 2006, for “fundamental contributions to computational complexity theory, database theory, and the theory of multi-agent systems”.
25. Received IBM Outstanding Innovation Award, 2006, for foundations of schema mappings.
26. Had research on “Foundations of schema mappings” named as IBM Research Division Accomplishment for 2005, jointly with Phokion Kolaitis and Lucian Popa.
27. Won Pat Goldberg Memorial Best Paper Award for one of the three Best IBM Research Papers in Computer Science, Electrical Engineering and Math published in 2004 for the paper “Composing schema mappings: second-order dependencies to the rescue” (coauthors Phokion Kolaitis, Lucian Popa and Wang-Chiew Tan).
28. Received 2004 ACM SIGMOD Edgar F. Codd Innovations Award for “fundamental contributions to database theory”.
29. Received IBM supplemental Patent Issue Award, given for key IBM patents, for differential backup patent, 2003.

30. Received IBM Outstanding Innovation Award, 2002, for differential backup for Tivoli Storage Manager.
31. Won award for one of the three Best IBM Research Papers in Computer Science, Electrical Engineering and Math published in 2001 for the paper “Optimal aggregation algorithms for middleware” (coauthors Amnon Lotem and Moni Naor).
32. Had research on “Aggregation algorithms for middleware” named as IBM Research Division Accomplishment for 2001.
33. Received IBM Outstanding Innovation Award, 2001, for research on aggregation algorithms for middleware.
34. Named Docteur Honoris Causa, University of Paris, 2001.
35. Won Best Paper Award, 2001 PODS, for the paper “Optimal aggregation algorithms for middleware” (coauthors Amnon Lotem and Moni Naor).
36. Named Fellow of the ACM (Association for Computing Machinery), 2000, for “creating the field of finite model theory, and for fundamental research in relational database theory and in reasoning about knowledge”.
37. Had research on “Differential backup for Tivoli Storage Manager” named as IBM Research Division Accomplishment for 2000, jointly with Randal Burns, Darrell Long, Larry Stockmeyer, Miklos Ajtai, David Pease, Bernie Lopez, Robert Morris, and Norm Pass.
38. Named Fellow of the IEEE (Institute of Electrical and Electronic Engineers), 1997, for “contributions to finite-model theory and to relational database theory”.
39. Received IBM Outstanding Innovation Award (jointly with Moshe Y. Vardi), 1992, for zero-one laws.
40. Received IBM Outstanding Innovation Award (jointly with Joseph Y. Halpern and Moshe Y. Vardi), 1987, for research on reasoning about knowledge.
41. Won MIT Press Publisher’s Prize for Best Paper, 1985 International Joint Conference on Artificial Intelligence, for the paper “Belief, Awareness, and Limited Reasoning” (coauthor Joseph Y. Halpern).
42. Received IBM Outstanding Innovation Award (jointly with Nicholas J. Pippenger and H. Raymond Strong), 1981, for developing extendible hashing, a fast access method for dynamic files.
43. Received IBM Outstanding Innovation Award, 1981, for fundamental contributions to relational database theory.

## **Keynote Talks**

1. Israel Database Day, Technion, Haifa, 2017
2. Big Data and Computational Intelligence Workshop, Beijing, 2016
3. Hot Topics in High-Performance Distributed Computing Workshop, 2015
4. ACM International Conference on Information and Knowledge Management, 2013
5. Australasian Computer Science Week, 2009
6. Milner Lecture, Univ. of Edinburgh, 2007
7. International Workshop on Exchange and Integration of Data, 2006
8. IEEE Symposium on Logic in Computer Science, 2000
9. Association for Symbolic Logic Annual Meeting, 1999
10. ACM Symposium on Principles of Database Systems, 1998
11. International Conference on Database Theory, 1990
12. International Colloquium on Automata, Languages, and Programming, 1984
13. Colloquium on Trees in Algebra and Programming, 1983
14. IBM Symposium on Mathematical Foundations of Computer Science, Japan, 1982

## **Program Committee Chairmanships**

1. International Conference on Database Theory, 2009
2. ACM Symposium on Theory of Computing, 2005
3. Theoretical Aspects of Reasoning about Knowledge, 1994
4. ACM Symposium on Principles of Database Systems, 1984

## **Conference committees**

1. Program committee, ACM Symposium on Principles of Database Systems, 2018
2. Invited papers committee, STOC Theory Fest, 2017
3. Program committee, WebDB, 2016
4. Program committee, Foundations of Information and Knowledge Systems, 2016
5. Program committee, Scalable Uncertainty Management, 2015

6. Program committee, Scalable Uncertainty Management, 2014
7. Chair, International Conference on Database Theory Test-of-Time Award committee, 2014
8. Program committee, Foundations of Information and Knowledge Systems, 2014
9. Program committee, International Conference on Database Theory, 2014
10. Program committee, SIGMOD/PODS PhD Symposium, 2013
11. Program committee, Foundations of Information and Knowledge Systems, 2012
12. Program committee, International Conference on Database Theory, 2012
13. Program committee, Foundations of Information and Knowledge Systems, 2010
14. Program committee, Workshop on Ranking in Databases, 2009
15. Program committee, Alberto Mendelzon International Workshop on Foundations of Data Management, 2009
16. Program committee chair, International Conference on Database Theory, 2009
17. Program committee, ACM Symposium on Theory of Computing, 2008
18. Program committee, ACM SIGMOD Conference, 2008
19. Program committee, Very Large Data Bases Conference, 2008
20. Program committee, Scalable Uncertainty Management, 2007
21. Program committee, Flexible Query Answering Systems, 2006
22. Program committee, International World Wide Web Conference, 2006
23. Program committee chair, ACM Symposium on Theory of Computing, 2005
24. Program committee, International World Wide Web Conference, 2004
25. Program committee, Flexible Query Answering Systems, 2004
26. Program committee, International Conference on Database Theory, 2003
27. Program committee, Workshop on Distributed Data and Structures, 2002
28. Program committee, ACM Symposium on Principles of Database Systems, 2002
29. Program committee, BISC International Workshop on Fuzzy Logic and the Internet, 2001
30. Program committee, ACM Symposium on Theory of Computing, 2001
31. Program committee, ACM Symposium on Theory of Computing, 1999

32. Program committee, Next Generation Information Technologies and Systems, Israel, 1999
33. Program committee, International Conference on Database Theory, 1999
34. Program committee chair, Theoretical Aspects of Reasoning about Knowledge, 1994
35. Program committee, IEEE Symposium on Foundations of Computer Science, 1993
36. Conference chair, Theoretical Aspects of Reasoning about Knowledge, 1992
37. Program committee, Theoretical Aspects of Reasoning about Knowledge, 1992
38. Program committee, Mathematical Fundamentals of Database and Knowledge Base Systems, 1991
39. Program committee, ACM Symposium on Theory of Computing, 1990
40. Member, Board of Directors, Theoretical Aspects of Reasoning about Knowledge (now Theoretical Aspects of Rationality and Knowledge), 1990–2004
41. Program committee, Theoretical Aspects of Reasoning about Knowledge, 1990
42. Program committee, Third Logical Biennial (Kleene '90), 1990
43. Program committee, ACM Symposium on Theory of Computing, 1988
44. Program committee, International Symposium on Model Theoretic Methods in Computer Science, 1984
45. Program committee chair, ACM Symposium on Principles of Database Systems, 1984
46. General chair, ACM Symposium on Principles of Database Systems, 1983
47. Program committee, ACM Symposium on Theory of Computing, 1983
48. Program committee, ACM Symposium on Principles of Database Systems, 1982
49. Program committee, ACM SIGMOD Conference, 1979

## Editorships

1. Editor, Journal of Computer and System Sciences, 2007–present. Associate Editor, 1984–2007.
2. Member, Editorial Board, Chicago Journal of Theoretical Computer Science, 1994–present
3. Member, Editorial Board, Foundations and Trends in Databases, 2006–2014
4. Member, Editorial Board, Methods of Logic in Computer Science, 1990–1996
5. Editor of special issue of Theory of Computing Systems for ICDT '09
6. Editor of special issue of SIAM J. Computing for STOC '05

7. Editor of special issue of Journal of Computer and System Sciences for STOC '99
8. Editor of special issue of Journal of Computer and System Sciences for PODS '92 (took over on the death of Paris Kanellakis)

## Patents

### Patents issued:

1. Method and apparatus for incorporating weights into data combinational rules  
01/11/2000 Issued as Patent 6014664 in U.S.  
Co-inventor: E.L. Wimmers
2. System and method for differential compression of data from a plurality of binary sources  
04/16/2002 Issued as Patent 6374250 in U.S.  
Co-inventors: M. Ajtai, R.C. Burns, L.J. Stockmeyer
3. Digital pen using speckle tracking  
02/03/2004 Issued as Patent 6686579 in U.S.  
Co-inventors: H.J. Rosen, N. Megiddo, R.J. Morris, T.G. Zimmerman, S. Rajagopalan
4. Static index pruning for information retrieval systems  
02/02/2005 Issued as Patent ZL0182128.8 in Canada  
02/15/2008 Issued as Patent 4080878 in Japan  
09/19/2012 Issued as Patent 60147138 in Denmark  
09/19/2012 Issued as Patent 1346296 in France  
09/19/2012 Issued as Patent 1346296 in United Kingdom  
Co-inventors: D. Carmel, D. Cohen, E.D. Farchi, Y.S. Maarek, M. Herscovici, A.Y. Soffer
5. Method for assigning encryption keys  
09/20/2005 Issued as Patent 6947563 in U.S.  
Co-inventors: M. Naor, D. Naor, J.B. Lottspiech, N. Megiddo
6. Method for schema mapping and data transformation  
12/12/2006 Issued as Patent 7149746 in U.S.  
Co-inventors: R.E. Miller, M.A. Hernández, L. Haas, L. Popa, F.G. Naumann
7. System, method, and service for ranking search results using a modular scoring system  
08/14/2007 Issued as Patent 7257277 in U.S.  
Co-inventors: J.A. Tomlin, K. Mc Curley, D. Sivakumar, J. Novak, D.P. Williamson, R. Kumar
8. Lossy index compression  
02/15/2008 issued as Patent 4080878 in Japan  
04/08/2008 Issued as Patent 7356527 in U.S.  
Co-inventors: D. Carmel, D. Cohen, E.D. Farchi, Y.S. Maarek, M. Herscovici, A.Y. Soffer



9. System and Method for Performing a High-Level Multi-Dimensional Query on a Multi-Structural Database  
04/14/2009 Issued as Patent 7519582 in U.S.  
Co-inventors: R. Guha, Ph. Kolaitis, R. Kumar, J. Novak, D. Sivakumar, and A. Tomkins.
10. System and Method of Performing an Inverse Schema Mapping  
03/09/2010 Issued as Patent 7676484 in U.S.
11. Sequential Composition of Schema Mappings  
03/30/2010 Issued as Patent 7689581 in U.S.  
Co-inventors: P. Kolaitis, L. Popa, W-C. Tan, and C. Yu.
12. Quasi-inverses of Schema Mappings  
08/24/2010 Issued as Patent 7783680 in U.S.  
Co-inventors: P. Kolaitis, L. Popa, and W-C. Tan.

#### Patents filed for by IBM:

1. Optimal approximate approach to aggregating information  
05/21/2002 Filed as Docket ARC920020019US1 in U.S.
2. System and method for translating data from a source schema to a target schema  
04/01/2003 Filed as Docket ARC920030001US1 in U.S.  
Co-inventors: R.E. Miller, M.A. Hernández, L. Popa.
3. Efficient similarity search and classification via rank aggregation  
06/09/2003 Filed as Docket ARC920030029US1 in U.S.  
Co-inventors: D. Sivakumar, R. Kumar
4. Efficient semi-supervised concept organization accelerated via inequality process  
8/1/17 Filed as Docket ARC920170010US1 in U.S.  
Co-inventors: A. Alba, K. Clarkson, C. Drews, D. F. Gruhl, N. R. Lewis, P. N.. Mendes, M. Nagaraiian, C. Ramakrishnan

#### Publications

1. Representation theory for a class of denumerable Markov chains. *J. Math. Analysis and Applications* 23, 1968, pp. 500–530. (Senior Thesis, Dartmouth College.)
2. *Contributions to the model theory of finite structures*. Doctoral dissertation, Univ. of California, Berkeley, June 1973.
3. Generalized first-order spectra and polynomial-time recognizable sets. *Complexity of Computation*, ed. R. Karp, SIAM-AMS Proceedings 7, 1974, pp. 43–73.
4.  $L_{\kappa,\lambda}$ -equivalence of ordinals and a compactness result. *Amer. Math. Soc. Notices* 21, 2, 1974, p. A-322.
5. Monadic generalized spectra. *Zeitschr. f. math. Logik und Grundlagen d. Math.* 21, 1975, pp. 89–96.

6. A two-cardinal characterization of double spectra. *Zeitschr. f. math. Logik und Grundlagen d. Math.* 21, 1975, pp. 121–122.
7. A spectrum hierarchy. *Zeitschr. f. math. Logik und Grundlagen d. Math.* 21, 1975, pp. 123–134.
8. The independence of miss ratio on page size, with Malcolm C. Easton. *J. ACM* 23, 1, Jan. 1976, pp. 128–146.
9. A counterintuitive example of computer paging. *Comm. ACM* 19, 2, Feb. 1976, pp. 96–97. Corrigendum: *Comm. ACM* 19, 4, April 1976, p. 187.
10. Probabilities on finite models. *J. Symbolic Logic* 41, 1, March 1976, pp. 50–58. Abstract appeared in *Notices Amer. Math. Soc.*, 1972, A714 (Abstract no. 72T-E90).
11. Asymptotic miss ratios over independent references. *J. Computer and System Sciences* 14, 2, April 1977, pp. 222–250.
12. A complete axiomatization for functional and multivalued dependencies in database relations, with Catriel Beeri and John H. Howard, Jr. *Proc. 1977 ACM SIGMOD Symposium*, ed. D. C. P. Smith, Toronto, pp. 47–61.
13. The number of finite relational structures. *Discrete Math.* 19, 1, July 1977, pp. 17–21.
14. Multivalued dependencies and a new normal form for relational databases. *ACM Trans. on Database Systems* 2, 3, Sept. 1977, pp. 262–278.
15. The decomposition versus the synthetic approach to relational database design. *Proc. 1977 Very Large Data Bases Conference*, Tokyo, pp. 441–446. Reprinted in: *Tutorial: Data Base Management in the 1980's*, ed. J. A. Larson and H. A. Freeman, IEEE Computer Society, NY, 1981, pp. 269–274.
16. Functional dependencies in a relational database and propositional logic. *IBM J. Research and Development* 21, 6, Nov. 1977, pp. 534–544.
17. Efficient calculation of expected miss ratios in the independent reference model, with Thomas G. Price. *SIAM J. Computing* 7, 3, Aug. 1978, pp. 288–297.
18. On an authorization mechanism. *ACM Trans. on Database Systems* 3, 3, Sept. 1978, pp. 310–319.
19. Cold-start vs. warm-start miss ratios, with Malcolm C. Easton. *Comm. ACM* 21, 10, Oct. 1978, pp. 866–871.
20. Normal forms and relational database operators. *Proc. 1979 ACM SIGMOD Conference*, ed. P. A. Bernstein, pp. 153–160.
21. Extendible hashing—a fast access method for dynamic files, with Jurg Nievergelt, Nicholas J. Pippenger, and H. Raymond Strong. *ACM Trans. on Database Systems* 4, 3, Sept. 1979, pp. 315–344.

22. An equivalence between relational database dependencies and a fragment of propositional logic, with Yehoshua Sagiv, Claude Delobel, and D. Stott Parker, Jr. *J. ACM* 28, 3, July 1981, pp. 435–453. Corrigendum: *J. ACM* 34, 4, Oct. 1987, pp. 1016–1018.
23. A note on the existence of continuous functionals, with J. Lawrence Carter. *Theoretical Computer Science* 16, 1981, pp. 231–235.
24. A normal form for relational databases that is based on domains and keys. *ACM Trans. on Database Systems* 6, 3, Sept. 1981, pp. 387–415.
25. Properties of acyclic database schemes, with Catriel Beeri, David Maier, Alberto Mendelzon, Jeffrey D. Ullman, and Mihalis Yannakakis. *Proc. 13th ACM Symposium on the Theory of Computing*, 1981, pp. 355–362.
26. Armstrong databases. **Invited paper**, *Proc. 7th IBM Symposium on Mathematical Foundations of Computer Science*, Kanagawa, Japan, May 1982.
27. A simplified universal relation assumption and its properties, with Alberto Mendelzon and Jeffrey D. Ullman. *ACM Trans. on Database Systems* 7, 3, Sept. 1982, pp. 343–360.
28. Horn clauses and database dependencies. *J. ACM* 29, 4, Oct. 1982, pp. 952–985. Preliminary version appeared in *Proc. 12th ACM Symposium on the Theory of Computing*, 1980, pp. 123–134.
29. Armstrong databases for functional and inclusion dependencies, with Moshe Y. Vardi. *Information Processing Letters* 16, Jan. 1983, pp. 13–19.
30. Tools for template dependencies, with David Maier, Jeffrey D. Ullman, and Mihalis Yannakakis. *SIAM J. Computing* 12, 1, Feb. 1983, pp. 36–59.
31. On the semantics of updates for databases, with Jeffrey D. Ullman and Moshe Y. Vardi. *Proc. 2nd ACM Symposium on Principles of Database Systems*, Atlanta, 1983, pp. 352–365.
32. A fair carpool scheduling algorithm, with John H. Williams. *IBM J. Research and Development* 27, 2, March 1983, pp. 133–139.
33. Acyclic database schemes of various degrees: a painless introduction. **Invited paper**, *Proc. CAAP83 8th Colloquium on Trees in Algebra and Programming*, Springer-Verlag Lecture Notes in Computer Science 159, 1983, ed. G. Ausiello and M. Protasi, pp. 65–89.
34. Degrees of acyclicity for hypergraphs and relational database schemes. *J. ACM* 30, 3, July 1983, pp. 514–550.
35. On the desirability of acyclic database schemes, with Catriel Beeri, David Maier, and Mihalis Yannakakis. *J. ACM* 30, 3, July 1983, pp. 479–513.
36. On the structure of Armstrong relations for functional dependencies, with Catriel Beeri, Martin Dowd, and Richard Statman. *J. ACM* 31, 1, Jan. 1984, pp. 30–46.

37. Inclusion dependencies and their interaction with functional dependencies, with Marco Casanova and Christos Papadimitriou. *J. Computer and System Sciences* 28, 1, Feb. 1984, pp. 29–59. Preliminary version appeared in *Proc. 1st ACM Symposium on Principles of Database Systems*, Los Angeles, March 1982, pp. 171–176.
38. The theory of data dependencies: an overview, with Moshe Y. Vardi. **Invited paper**, *Proc. 11th International Colloquium on Automata, Languages, and Programming*, Antwerp, July 1984. Appeared in Springer–Verlag Lecture Notes in Computer Science 172, 1984, ed. J. Paradaens, pp. 1–22.
39. Decreasing the nesting depth of expressions involving square roots, with Allan Borodin, John Hopcroft, and Martin Tompa. *Journal of Symbolic Computation* 1, 1985, pp. 169–188.
40. An internal semantics for modal logic, with Moshe Y. Vardi. *Proc. 17th ACM Symposium on Theory of Computing*, Providence, RI, May 1985, pp. 305–315.
41. Bounded depth, polynomial–size circuits for symmetric functions, with Maria M. Klawe, Nicholas J. Pippenger, and Larry J. Stockmeyer. *Theoretical Computer Science* 36, 2–3, April 1985, pp. 239–250.
42. Knowledge and implicit knowledge in a distributed environment, with Moshe Y. Vardi. *Conf. on Theoretical Aspects of Reasoning about Knowledge*, ed. J. Y. Halpern, Morgan Kaufmann, 1986, pp. 187–206.
43. Updating logical databases, with Gabriel Kuper, Jeffrey D. Ullman, and Moshe Y. Vardi. *Advances in Computing Research*, vol. 3, The Theory of Databases, ed. P. Kanellakis and F. P. Preparata, JAI Press, Greenwich, Conn., 1986, pp. 1–18.
44. The theory of data dependencies: a survey, with Moshe Y. Vardi. *Mathematics of Information Processing*, Proceedings of Symposia in Applied Mathematics, American Mathematical Society, 1986, vol. 34, pp. 19–72.
45. A simple characterization of database dependency implication, with Yoshito Hanatani. *Information Processing Letters* 22, 6, May 1986, 281–283.
46. Belief, awareness, and limited reasoning, with Joseph Y. Halpern. *Artificial Intelligence* 34, 1988, pp. 39–76. Preliminary version appeared in *International Joint Conference on Artificial Intelligence (IJCAI–85)*, Aug. 1985, pp. 491–501.
47. Modelling knowledge and action in distributed systems, with Joseph Y. Halpern. *Distributed Computing* 3, 4, 1989, pp. 159–179. Preliminary version appeared under the title “A formal model of knowledge, action, and communication in distributed systems” in *ACM Symposium on Principles of Distributed Computing*, 1985, pp. 224–236. Abridged version appeared in *Proc. Concurrency* 88, 1988, ed. F. H. Vogt, pp. 18–32.
48. I’m OK if you’re OK: On the notion of trusting communication, with Joseph Y. Halpern, *J. of Philosophical Logic* 17, 4, Nov. 1988, 329–354. Reprinted in *Philosophical Logic and Artificial Intelligence*, ed. R. Thomason, Kluwer, 1989, pp. 9–34. Preliminary version

- appeared in *Proc. IEEE Symposium on Logic in Computer Science*, June 1987, pp. 280–292.
49. Reachability is harder for directed than for undirected finite graphs, with Miklos Ajtai. *J. Symbolic Logic* 55, 1, March 1990, pp. 113–150. Preliminary version appeared in *Proc. 29th IEEE Symposium on Foundations of Computer Science*, 1988, pp. 358–367.
  50. A logic for reasoning about probabilities, with Joseph Y. Halpern and Nimrod Megiddo. *Information and Computation* 87, July/Aug. 1990, pp. 78–128. (Special issue for selected papers from the 1988 IEEE Symposium on Logic in Computer Science Conference).
  51. Uncertainty, belief, and probability, with Joseph Y. Halpern. *Computational Intelligence* 7, 1991, pp. 160–173. Preliminary version appeared in *International Joint Conference on Artificial Intelligence (IJCAI-89)*, pp. 1161–1167.
  52. A model-theoretic analysis of knowledge, with Joseph Y. Halpern and Moshe Y. Vardi. *J. ACM* 91, 2, April 1991, pp. 382–428. Preliminary version appeared in *Proc. 25th IEEE Symposium on Foundations of Computer Science*, West Palm Beach, Florida, Oct. 1984, pp. 268–278.
  53. Database systems: achievements and opportunities, *Comm. ACM* 34, 10, Oct. 1991, pp. 110–120, ed. A. Silberschatz, M. Stonebraker, and J. Ullman. The report was drafted by the NSF Invitational Workshop on Future Directions in DBMS Research, Feb. 22–23, 1990, of which I was a member (the other invitees were Michael Brodie, Peter Buneman, Michael Carey, Ashok Chandra, Hector Garcia-Molina, Jim Gray, David Lomet, David Maier, Marie Ann Niemat, Avi Silberschatz, Michael Stonebraker, Irving Traiger, Jeffrey Ullman, Gio Wiederhold, Carlo Zaniolo, and Maria Zemankova).
  54. A new approach to updating beliefs, with Joseph Y. Halpern. In *Uncertainty in Artificial Intelligence: Volume VI*, ed. P. P. Bonissone, M. Henrion, L. N. Kanal and J. Lemmer, Elsevier, 1991, pp. 347–374. Preliminary version appeared in *Conference on Uncertainty in AI*, 1990, pp. 317–324.
  55. What can machines know? On the properties of knowledge in distributed systems, with Joseph Y. Halpern and Moshe Y. Vardi. *J. ACM* 39, 2, 1992, pp. 328–376. Preliminary version appeared under the name “What can machines know? On the epistemic properties of machines” in *National Conference on Artificial Intelligence (AAAI-86)*, 1986, pp. 428–434.
  56. Two views of belief: Belief as generalized probability and belief as evidence, with Joseph Y. Halpern. *Artificial Intelligence* 54, 1992, pp. 275–317. Preliminary version appeared in *National Conference on Artificial Intelligence (AAAI-90)*, Aug. 1990, pp. 112–119.
  57. What is an inference rule?, with Joseph Y. Halpern and Moshe Y. Vardi. *J. Symbolic Logic* 57, 3, Sept. 1992, pp. 1018–1045.
  58. Simple conditions for guaranteeing higher normal forms for relational databases, with C. J. Date. *ACM Trans. on Database Systems* 17, 3, Sept. 1992, 465–476. Reprinted in

- “Relational database, writings 1989 – 1991” by C. J. Date with Hugh Darwen, Addison Wesley, 1992.
59. Response to “Remarks on two new theorems of Date and Fagin”, with C. J. Date. *SIGMOD RECORD*, March 1993, pp. 57–58.
  60. Finite model theory—a personal perspective. *Theoretical Computer Science* 116, 1993, pp. 3–31. Preliminary version appeared as **Invited paper**, *3rd International Conference on Database Theory*, Dec. 1990, Springer–Verlag Lecture Notes in Computer Science 470, pp. 3–24.
  61. An operational semantics for knowledge bases, with Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *National Conference on Artificial Intelligence (AAAI-94)*, 1994, pp. 1142–1147. This is a later version of “Modeling knowledge bases”, *Proc. IJCAI Workshop on Hybrid Reasoning*, 1993.
  62. A quantitative analysis of modal logic. *J. Symbolic Logic* 59, 1, March 1994, pp. 209–252.
  63. Reasoning about knowledge and probability, with Joseph Y. Halpern. *J. ACM* 41, 2, 1994, pp. 340–367. Preliminary version appeared in *Second Conf. on Theoretical Aspects of Reasoning about Knowledge*, ed. M. Y. Vardi, Morgan Kaufmann, 1988, pp. 277–293. Corrigendum: *J. ACM* 45, 1, Jan. 1998, p. 214.
  64. Towards heterogeneous multimedia information systems: the Garlic approach, with Michael J. Carey, Laura M. Haas, Peter M. Schwarz, Manish Arya, William F. Cody, Myron Flickner, Allen W. Luniewski, Wayne Niblack, Dragutin Petkovic, John Thomas, John H. Williams, and Edward L. Wimmers. *RIDE-DOM '95 (5th Int'l Workshop on Research Issues in Data Engineering: Distributed Object Management)*, 1995, pp. 124–131.
  65. Querying multimedia data from multiple repositories by content: the Garlic project, with William F. Cody, Laura M. Haas, Wayne Niblack, Manish Arya, Michael J. Carey, Myron Flickner, Denis S. Lee, Dragutin Petkovic, Peter M. Schwarz, John Thomas, Mary Tork Roth, John H. Williams, and Edward L. Wimmers. *IFIP 2.6 3rd Working Conference on Visual Database Systems (VDB-3)*, 1995.
  66. *Reasoning about Knowledge*, with Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. **Book**, MIT Press, 1995. Paperback edition, 2003.
  67. On monadic NP vs. monadic co-NP, with Larry Stockmeyer and Moshe Vardi. *Information and Computation* 120, 1, July 1995, pp. 78–92. Preliminary version appeared in *1993 IEEE Structure in Complexity Theory Conference*, pp. 19–30.
  68. Comparing the power of monadic NP games. *Logic and Computational Complexity*, Springer–Verlag Lecture Notes in Computer Science 960, 1995, ed. D. Leivant, pp. 414–425.
  69. A nonstandard approach to the logical omniscience problem, with Joseph Y. Halpern and Moshe Y. Vardi. *Artificial Intelligence* 79, 2, Dec. 1995, pp. 203–240. Preliminary version appeared in *Third Conf. on Theoretical Aspects of Reasoning about Knowledge*, ed. R. Parikh, Morgan Kaufmann, 1990, pp. 41–55.

70. Comparing information without leaking it, with Moni Naor and Peter Winkler. *Comm. ACM* 39, 5, May 1996, pp. 77–85.
71. Efficiently extendible mappings for balanced data distribution, with David M. Choy and Larry Stockmeyer. *Algorithmica* 16, 1996, pp. 215–232.
72. Common knowledge: Now you have it, now you don't, with Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. In *Intelligent Systems: A Semiotics Perspective*, Proc. 1996 Int'l Multidisciplinary Conf., Vol. I, Oct. 1996, pp. 177–183.
73. The Garlic project, with Mary Tork Roth, Manish Arya, Laura Haas, Michael J. Carey, William Cody, Peter Schwarz, John Thomas, and Edward L. Wimmes. *Proc. 1996 ACM SIGMOD Symposium*, Montreal, page 557.
74. On winning strategies in Ehrenfeucht–Fraïssé games, with Sanjeev Arora. *Theoretical Computer Science* 174, 1997, pp. 97–121.
75. Easier ways to win logical games. **Invited paper**, *Descriptive Complexity and Finite Models*, DIMACS Series in Discrete Mathematics and Theoretical Computer Science, American Mathematical Society, Volume 31, 1997, ed. N. Immerman and Ph. Kolaitis, pp. 1–32.
76. Comparing the power of games on graphs. *Mathematical Logic Quarterly* 43, 4, 1997, pp. 431–455. This is an expanded version of “Comparing the power of monadic NP games”, *Logic and Computational Complexity*, Springer–Verlag Lecture Notes in Computer Science 960, 1995, ed. D. Leivant, pp. 414–425.
77. Reasoning about knowledge: a response by the authors, with Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Minds and Machines* 7, 1997, p. 113.
78. Knowledge-based programs, with Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Distributed Computing* 10, 4, 1997, pp. 199–225. (Special issue for selected papers from the 1995 ACM Symposium on Principles of Distributed Computing).
79. Spectra with only unary function symbols, with Arnaud Durand and Bernd Loescher. *1997 Annual Conference of the European Association for Computer Science Logic (CSL'97)*. In Springer–Verlag Lecture Notes in Computer Science 1414, 1998, ed. M. Nielsen and W. Thomas, pp. 189–202.
80. Fuzzy queries in multimedia database systems. **Invited paper**, *Proc. 17th ACM Symposium on Principles of Database Systems*, Seattle, 1998, pp. 1–10.
81. Relaxing the triangle inequality in pattern matching, with Larry Stockmeyer. *International Journal of Computer Vision* 30, 3, 1998, pp. 219–231.
82. Common knowledge revisited, with Joseph Y. Halpern, Yoram Moses, and Moshe Y. Vardi. *Annals of Pure and Applied Logic* 96, 1999, pp. 89–105. Preliminary version appeared in *Sixth Conf. on Theoretical Aspects of Rationality and Knowledge*, ed. Y. Shoham, Morgan Kaufmann, 1996, pp. 283–298. Reprinted in *Knowledge Contributors* (V.F. Henricks, K.F. Jorgensen, and S.A. Pedersen, eds.), Kluwer, 2003, pp. 87–104.

83. Combining fuzzy information from multiple systems. *J. Computer and System Sciences* 58, 1999, pp. 83–99. (Special issue for selected papers from the 1996 ACM Symposium on Principles of Database Systems).
84. The hierarchical approach to modeling knowledge and common knowledge, with John Geanakoplos, Joseph Y. Halpern, and Moshe Y. Vardi. *International Journal of Game Theory* 28, 3, 1999, pp. 331–365. Preliminary version appeared as “The expressive power of the hierarchical approach to modeling knowledge and common knowledge”, *Fourth Conf. on Theoretical Aspects of Reasoning about Knowledge*, ed. Y. Moses, Morgan Kaufmann, 1992, pp. 229–244.
85. Allowing users to weight search terms, with Yoelle Maarek. RIAO (Recherche d’Informations Assistee par Ordinateur = Computer-Assisted Information Retrieval) 2000, pp. 682–700.
86. Ephemeral document clustering for web applications, with Yoelle Maarek, Israel Ben-Shaul, and Dan Pelleg. IBM Research Report RJ10186, April 2000.
87. The closure of monadic NP, with Miklos Ajtai and Larry Stockmeyer. *J. Computer and System Sciences* 60, June 2000, pp. 660–716. (Special issue for selected papers from the 1998 ACM Symposium on Theory of Computing).
88. Logic, complexity, and games. **Invited paper**, *Proc. 15th IEEE Symposium on Logic in Computer Science*, Santa Barbara, 2000, p. 3.
89. A formula for incorporating weights into scoring rules, with Edward L. Wimmers. *Theoretical Computer Science* 239, 2000, pp. 309–338. (Special issue for selected papers from the 1997 International Conference on Database Theory). Preliminary version appeared under the title “Incorporating user preferences in multimedia queries”, in *Proc. 6th International Conference on Database Theory*, Jan. 1997, Springer–Verlag Lecture Notes in Computer Science 1186, ed. F. Afrati and Ph. Kolaitis, Delphi, pp. 247–261.
90. Static index pruning for information retrieval systems, with David Carmel, Doron Cohen, Eitan Farchi, Michael Herscovici, Yoelle Maarek, and Aya Soffer. *Proc. 24th ACM SIGIR Conference on Research and Development in Information Retrieval*, New Orleans, Sept. 2001, pp. 43–50.
91. The Clio project: managing heterogeneity, with Renée J. Miller, Mauricio A. Hernández, Laura M. Haas, Lingling Yan, C. T. Howard Ho, and Lucian Popa. *ACM SIGMOD Record* 30, 1, March 2001, pp. 78–83.
92. Data-driven understanding and refinement of schema mappings, with Ling Ling Yan, Renée J. Miller, and Laura M. Haas. *Proc. 2001 ACM SIGMOD Symposium*, Santa Barbara, pp. 485–496.
93. Random walks with “back buttons”, with Anna Karlin, Jon Kleinberg, Prabhakar Raghavan, Sridhar Rajagopalan, Ronitt Rubinfeld, Madhu Sudan, and Andrew Tomkins, *Annals of Applied Probability* 11, 3, 2001, pp. 810–862. Preliminary version appeared in *Proc. 32nd ACM Symposium on Theory of Computing*, 2000, pp. 484–493.



94. Compactly encoding unstructured inputs with differential compression, with Miklos Ajtai, Randal Burns, Larry Stockmeyer, and Darrell Long. *J. ACM* 49, 3, May 2002, pp. 318–367.
95. Combining fuzzy information: an overview. *SIGMOD Record* 31,2, June 2002, pp. 109–118.
96. Query strategies for priced information, with Moses Charikar, Venkatesan Guruswami, Jon Kleinberg, Prabhakar Raghavan, and Amit Sahai, *J. Computer and System Sciences*, 64, 4, June 2002, pp. 785–819. (Special issue for selected papers from the 2000 ACM Symposium on Theory of Computing).
97. Translating web data, with Mauricio Hernández, Renée J. Miller, Lucian Popa, and Yannis Velegrakis. *Proc. 2002 Very Large Data Bases Conference*, Hong Kong, pp. 598–609.
98. Schema management, with Periklis Andritsos, Ariel Fuxman, Laura M. Haas, Mauricio A. Hernández, Howard Ho, Anastasios Kementsietsidis, Renée J. Miller, Felix Nauman, Lucian Popa, Yannis Velegrakis, Charlotte Vilarem and Ling-Ling Yan, *IEEE Data Engineering Bulletin* 25, 3, 2002, pp. 33–39.
99. Comparing top  $k$  lists, with Ravi Kumar and D. Sivakumar. *SIAM J. Discrete Mathematics* 17, 1 (2003), pp. 134–160. Extended abstract appeared in *Proc. 2003 ACM Symposium on Discrete Algorithms (SODA '03)*, pp. 28–36.
100. Searching the workplace web, with Ravi Kumar, Kevin S. McCurley, Jasmine Novak, D. Sivakumar, John A. Tomlin, and David P. Williamson. *Proc. 2003 International World Wide Web Conference (WWW2003)*, pp. 366–375.
101. Efficient similarity search and classification via rank aggregation, with Ravi Kumar and D. Sivakumar. *Proc. 2003 ACM SIGMOD Symposium*, San Diego, pp. 301–312. Corrigendum: 2008 ACM SIGMOD Symposium, Vancouver, pp. 1375–1376.
102. Optimal aggregation algorithms for middleware, with Amnon Lotem and Moni Naor. *J. Computer and System Sciences* 66, 2003, pp. 614–656. (Special issue for selected papers from the 2001 ACM Symposium on Principles of Database Systems).
103. Comparing and aggregating rankings with ties, with Ravi Kumar, Mohammad Mahdian, D. Sivakumar, and Erik Vee. *Proc. 23rd ACM Symposium on Principles of Database Systems*, Paris, 2004, pp. 47–58.
104. Data exchange: semantics and query answering, with Phokion Kolaitis, Renée J. Miller, and Lucian Popa. *Theoretical Computer Science* 336 (2005), pp. 89–124. (Special issue for selected papers from the 2003 International Conference on Database Theory).
105. Data exchange: getting to the core, with Phokion Kolaitis and Lucian Popa. *ACM Trans. on Database Systems* 30,1, Mar. 2005, pp. 174–210. (Special issue for selected papers from the 2003 ACM Symposium on Principles of Database Systems.)

106. Composing schema mappings: second-order dependencies to the rescue, with Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan. *ACM Trans. on Database Systems* 30, 4, Dec. 2005, pp. 994–1005. (Special issue for selected papers from the 2004 ACM SIGMOD/PODS Conference.)
107. Multi-structural databases, with R. Guha, Ravi Kumar, Jasmine Novak, D. Sivakumar, and Andrew Tomkins. *Proc. 24th ACM Symposium on Principles of Database Systems*, Baltimore, 2005, pp. 184–195.
108. Efficient Implementation of Large-Scale Multi-Structural Databases, with Phokion Kolaitis, Ravi Kumar, Jasmine Novak, D. Sivakumar, and Andrew Tomkins. *Proc. 2005 Very Large Data Bases Conference*, pp. 958–969.
109. Comparing partial rankings, with Ravi Kumar, Mohammad Mahdian, D. Sivakumar, and Erik Vee. *SIAM J. Discrete Mathematics* 20, 3, 2006, pp. 628–648.
110. Inverting schema mappings. *ACM Trans. on Database Systems* 32, 4, Nov. 2007. (Special issue for selected papers from the 2006 ACM Symposium on Principles of Database Systems.)
111. Towards a theory of schema mapping optimization, with Phokion Kolaitis, Alan Nash, and Lucian Popa. *Proc. 27th ACM Symposium on Principles of Database Systems*, Vancouver, 2008, pp. 33–42.
112. Quasi-inverses of schema mappings, with Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan. *ACM Trans. on Database Systems* 33, 2, June 2008, Article No. 11. (Special issue for selected papers from the 2007 ACM Symposium on Principles of Database Systems.)
113. Clio: Schema mapping creation and data exchange, with Laura M. Haas, Mauricio A. Hernández, Renée J. Miller, Lucian Popa, and Yannis Velegarakis. In A. T. Borgida, V. K. Chaudhri, P. Giorgini, and E. S. Yu, editors, *Conceptual Modeling: Foundations and Applications, Essays in Honor of John Mylopoulos*, LNCS 5600, Springer, 2009, pp. 198–236.
114. Understanding queries in a search database system, with Benny Kimelfeld, Yunyao Li, Sriram Raghavan, and Shivakumar Vaithyanathan. *Proc. 29th ACM Symposium on Principles of Database Systems*, Indianapolis, 2010, pp. 273–284.
115. The structure of inverses in schema mappings, with Alan Nash. *J. ACM* 57, 6, Oct. 2010, Article 31.
116. Epistemic privacy, with Alexandre Evfimievski and David Woodruff. *J. ACM* 58, 1 (December 2010), Article 2. (Invited from the 2008 ACM Symposium on Principles of Database Systems.)
117. Schema mapping evolution through composition and inversion, with Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan. In Z. Bellahsene, A. Bonifati, and E. Rahm, editors, *Schema Matching and Mapping*, Springer, 2011, pp. 191–222.

118. Reverse data exchange: coping with nulls, with Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan. *ACM Trans. on Database Systems* 36, 2 (May 2011), Article 11. (Invited from the 2009 ACM Symposium on Principles of Database Systems.)
119. Rewrite rules for search database systems, with Benny Kimelfeld, Yunyao Li, Sriram Raghavan and Shivakumar Vaithyanathan. *Proc. 29th ACM Symposium on Principles of Database Systems*, Athens, 2011, pp. 271–282.
120. Probabilistic data exchange, with Benny Kimelfeld and Phokion G. Kolaitis. *J. ACM* 58, 4 (July 2011), Article 15. Preliminary version appeared in *Proc. 2010 International Conference on Database Theory*, pp. 76–88.
121. Composition with target constraints, with Marcelo Arenas and Alan Nash. *Logical Methods in Computer Science* 7, 3:13 (2011), pp. 1–38. (Special issue for selected papers from the 2010 International Conference on Database Theory.)
122. A normal form for preventing redundant tuples in relational databases, with Hugh Darwen and C. J. Date. *Proc. 2012 International Conference on Database Theory*, pp. 114–126.
123. Local transformations and conjunctive-query equivalence, with Phokion G. Kolaitis. *Proc. 30th ACM Symposium on Principles of Database Systems*, Scottsdale, 2012, pp. 179–190.
124. Solutions and query rewriting in data exchange, with Marcelo Arenas, Pablo Barceló, and Leonid Libkin. *Information and Computation*, 2013, pp. 28–61. Preliminary version appeared under the name “Locally consistent transformations and query answering in data exchange”, in *Proc. 23rd ACM Symposium on Principles of Database Systems*, Paris, 2004, pp. 229–240.
125. Document spanners: a formal approach to information extraction with Benny Kimelfeld, Frederick Reiss, and Stijn Vansummeren. *J. ACM* 62, 2 (April 2015), Article 12. Preliminary version appeared in *Proc. 31st ACM Symposium on Principles of Database Systems*, New York, 2013, pp. 37–48, under the title “Spanners: a formal framework for information extraction”.
126. Dichotomies in the complexity of preferred repairs, with Benny Kimelfeld and Phokion Kolaitis. *Proc. 33rd ACM Symposium on Principles of Database Systems*, Melbourne, 2015, pp. 3–15.
127. Declarative cleaning of inconsistencies in information extraction, with Benny Kimelfeld, Frederick Reiss, and Stijn Vansummeren. *ACM Trans. on Database Systems* 41,1, Article 6 (February 2016), 44 pages. Preliminary version appeared in *Proc. 32nd ACM Symposium on Principles of Database Systems*, Snowbird, 2014, pp. 164–175, under the title “Cleaning inconsistencies in information extraction via prioritized repairs”.
128. A relational framework for information extraction, with Benny Kimelfeld, Frederick Reiss, and Stijn Vansummeren. *SIGMOD RECORD*, December 2015, pp. 5–16.
129. A declarative framework for linking entities, with Douglas Burdick, Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan. *ACM Trans. on Database Systems* 41,3, Article 17

(June 2016), 38 pages. Preliminary version appeared in *Proc. International Conference on Database Theory*, 2015, pp. 25–43.

130. Optimal score aggregation algorithms. *Proc. ACM Symposium on Principles of Database Systems (PODS '16)*, p. 55.
131. An algorithmic view of voting, with Ravi Kumar, Mohammad Mahdian, D. Sivakumar, and Erik Vee. *Siam J. Discrete Math.* 30,4 (2016), pp. 1978–1996.
132. Expressive power of entity-linking frameworks, with Douglas Burdick, Phokion Kolaitis, Lucian Popa, and Wang-Chiew Tan. *Proc. International Conference on Database Theory*, 2017, pp. 10:1–10:18.