

Data Management for Data Analytics

material for presentations and projects

L. Libkin

The references are taken from DBLP (<http://dblp.uni-trier.de>), the main bibliographical source for computer science research. You can search DBLP by authors' names, to find those papers. Once they are found, clicking on the electronic edition icon next to the paper gives you access to the source as long as the university carries a subscription. Otherwise, try entering the title, in quotes, as well as pdf in a google search, and you should find copies on authors' webpages.

1 Background: Conjunctive queries, complexity, expressiveness, reasoning

1. Ashok K. Chandra, Philip M. Merlin: Optimal Implementation of Conjunctive Queries in Relational Data Bases. *STOC 1977*: 77-90
2. Mihalis Yannakakis: Algorithms for Acyclic Database Schemes. *VLDB 1981*: 82-94
3. Georg Gottlob, Nicola Leone, Francesco Scarcello: The complexity of acyclic conjunctive queries. *Journal of the ACM* 48(3):431-498 (2001)
4. Georg Gottlob, Nicola Leone, Francesco Scarcello: Hypertree Decompositions and Tractable Queries. *J. Comput. Syst. Sci.* 64(3):579-627 (2002)
5. Martin Grohe, Thomas Schwentick, Luc Segoufin: When is the evaluation of conjunctive queries tractable? *STOC 2001*: 657-666
6. Moshe Y. Vardi: The Complexity of Relational Query Languages (Extended Abstract) .*STOC 1982*: 137-146
7. Christos H. Papadimitriou, Mihalis Yannakakis: On the Complexity of Database Queries. *J. Comput. Syst. Sci.* 58(3): 407-427 (1999)
8. Neil Immerman: Languages that Capture Complexity Classes. *SIAM J. Comput.* 16(4): 760-778 (1987)
9. Martin Grohe: Fixed-point definability and polynomial time on graphs with excluded minors. *Journal of the ACM* 59(5): 27 (2012)
10. Phokion G. Kolaitis, Moshe Y. Vardi: Conjunctive-Query Containment and Constraint Satisfaction. *J. Comput. Syst. Sci.* 61(2): 302-332 (2000)

11. Martin Grohe: From polynomial time queries to graph structure theory. *Commun. ACM* 54(6): 104-112 (2011)
12. Leonid Libkin: The finite model theory toolbox of a database theoretician. *PODS 2009*: 65-76
13. Leonid Libkin: Expressive power of SQL. *Theor. Comput. Sci.* 296(3): 379-404 (2003)
14. David Maier, Alberto O. Mendelzon, Yehoshua Sagiv: Testing Implications of Data Dependencies. *ACM Trans. Database Syst.* 4(4): 455-469 (1979)
15. Alin Deutsch, Alan Nash, Jeffrey B. Remmel: The chase revisited. *PODS 2008*: 149-158
16. Paolo Guagliardo and Leonid Libkin: A Formal Semantics of SQL Queries, Its Validation, and Applications, *PVLDB*, September 2017.

2 Volume: Scalability of queries, join processing, approximations, efficiency via syntactic restrictions

17. Michael Armbrust, Kristal Curtis, Tim Kraska, Armando Fox, Michael J. Franklin, David A. Patterson: PIQL: Success-Tolerant Query Processing in the Cloud. *PVLDB* 5(3):181-192 (2011)
18. Michael Armbrust, Armando Fox, David A. Patterson, Nick Lanham, Beth Trushkowsky, Jesse Trutna, Haruki Oh: SCADS: Scale-Independent Storage for Social Computing Applications. *CIDR 2009*
19. Michael Armbrust, Eric Liang, Tim Kraska, Armando Fox, Michael J. Franklin, David A. Patterson: Generalized scale independence through incremental precomputation. *SIGMOD 2013*:625-636
20. Wenfei Fan, Floris Geerts, Frank Neven: Making Queries Tractable on Big Data with Preprocessing. *PVLDB* 6(9): 685-696 (2013)
21. Wenfei Fan, Floris Geerts, Leonid Libkin: On scale independence for querying big data. *PODS 2014*:51-62
22. Yang Cao, Wenfei Fan, Tianyu Wo, Wenyuan Yu: Bounded Conjunctive Queries. *PVLDB* 7(12): 1231-1242 (2014)
23. Foto N. Afrati, Jeffrey D. Ullman: Transitive closure and recursive Datalog implemented on clusters. *EDBT 2012*: 132-143
24. Foto N. Afrati, Jeffrey D. Ullman: Optimizing Multiway Joins in a Map-Reduce Environment. *IEEE Trans. Knowl. Data Eng.* 23(9): 1282-1298 (2011)
25. Mahmoud Abo Khamis, Hung Q. Ngo, Christopher Ré, Atri Rudra: Joins via Geometric Resolutions: Worst-case and Beyond. *PODS 2015*: 213-228
26. Hung Q. Ngo, Christopher Ré, Atri Rudra: Skew strikes back: new developments in the theory of join algorithms. *SIGMOD Record* 42(4): 5-16 (2013)

27. Hung Q. Ngo, Ely Porat, Christopher Ré, Atri Rudra: Worst-case optimal join algorithms. PODS 2012: 37-48
28. Yannis E. Ioannidis: Approximations in Database Systems. ICDT 2003: 16-30
29. Minos N. Garofalakis, Phillip B. Gibbons: Approximate Query Processing: Taming the TeraBytes. VLDB 2001
30. Pablo Barcelo, Leonid Libkin, Miguel Romero: Efficient Approximations of Conjunctive Queries. SIAM J. Comput. 43(3): 1085-1130 (2014)
31. Markus Frick, Martin Grohe: Deciding first-order properties of locally tree-decomposable structures. Journal of the ACM 48(6): 1184-1206 (2001)
32. Joerg Flum, Martin Grohe: Fixed-Parameter Tractability, Definability, and Model-Checking. SIAM J. Comput. 31(1): 113-145 (2001)
33. Joerg Flum, Markus Frick, Martin Grohe: Query evaluation via tree-decompositions. Journal of the ACM 49(6): 716-752 (2002)
34. Albert Atserias, Martin Grohe, Daniel Marx: Size Bounds and Query Plans for Relational Joins. SIAM J. Comput. 42(4): 1737-1767 (2013)
35. Paul Beame, Paraschos Koutris, Dan Suciu: Communication steps for parallel query processing. PODS 2013: 273-284
36. Tom J. Ameloot, Gaetano Geck, Bas Ketsman, Frank Neven, Thomas Schwentick: Parallel-Correctness and Transferability for Conjunctive Queries. PODS 2015: 47-58

3 Variety: semi-structured data

37. Frank Neven: Automata, Logic, and XML. CSL 2002: 2-26
38. Leonid Libkin: Logics for Unranked Trees: An Overview. Logical Methods in Computer Science 2(3) (2006)
39. Georg Gottlob, Christoph Koch, Reinhard Pichler: Efficient algorithms for processing XPath queries. ACM Trans. Database Syst. 30(2): 444-491 (2005)
40. Georg Gottlob, Christoph Koch, Reinhard Pichler, Luc Segoufin: The complexity of XPath query evaluation and XML typing. Journal of the ACM 52(2): 284-335 (2005)
41. Frank Neven, Thomas Schwentick: Query automata over finite trees. Theor. Comput. Sci. 275(1-2): 633-674 (2002)
42. Georg Gottlob, Christoph Koch: Monadic datalog and the expressive power of languages for Web information extraction. Journal of the ACM 51(1): 74-113 (2004)
43. Pablo Barcelo, Leonid Libkin: Temporal logics over unranked trees. LICS 2005: 31-40

44. Leonid Libkin, Cristina Sirangelo: Reasoning about XML with temporal logics and automata. *J. Applied Logic* 8(2): 210-232 (2010)
45. Thomas Schwentick: XPath query containment. *SIGMOD Record* 33(1): 101-109 (2004)
46. Wenfei Fan, Leonid Libkin: On XML integrity constraints in the presence of DTDs. *Journal of the ACM* 49(3): 368-406 (2002)
47. Marcelo Arenas, Wenfei Fan, Leonid Libkin: On the Complexity of Verifying Consistency of XML Specifications. *SIAM J. Comput.* 38(3): 841-880 (2008)
48. Wim Martens, Frank Neven, Thomas Schwentick: Simple off the shelf abstractions for XML schema. *SIGMOD Record* 36(3): 15-22 (2007)
49. Wim Martens, Frank Neven, Thomas Schwentick, Geert Jan Bex: Expressiveness and complexity of XML Schema. *ACM Trans. Database Syst.* 31(3): 770-813 (2006)
50. Mikolaj Bojanczyk, Anca Muscholl, Thomas Schwentick, Luc Segoufin: Two-variable logic on data trees and XML reasoning. *Journal of the ACM* 56(3) (2009)
51. Tony Tan: Extending two-variable logic on data trees with order on data values and its automata. *ACM Trans. Comput. Log.* 15(1): 8 (2014)
52. Henrik Bjrklund, Wim Martens, Thomas Schwentick: Conjunctive query containment over trees. *J. Comput. Syst. Sci.* 77(3): 450-472 (2011)
53. Wojciech Czerwinski, Wim Martens, Pawel Parys, Marcin Przybylko: The (Almost) Complete Guide to Tree Pattern Containment. *PODS 2015*: 117-130
54. Maarten Marx: Conditional XPath. *ACM Trans. Database Syst.* 30(4): 929-959 (2005)
55. Balder ten Cate, Maarten Marx: Navigational XPath: calculus and algebra. *SIGMOD Record* 36(2): 19-26 (2007)
56. Loredana Afanasiev, Maarten Marx: An analysis of XQuery benchmarks. *Inf. Syst.* 33(2): 155-181 (2008)
57. Claire David, Leonid Libkin, Tony Tan: Efficient reasoning about data trees via integer linear programming. *ACM Trans. Database Syst.* 37(3): 19 (2012)
58. Luc Segoufin, Cristina Sirangelo: Constant-Memory Validation of Streaming XML Documents Against DTDs. *ICDT 2007*: 299-313
59. Wim Martens, Frank Neven, Thomas Schwentick, Geert Jan Bex: Expressiveness and complexity of XML Schema. *ACM Trans. Database Syst.* 31(3): 770-813 (2006)

4 Variety: graph data

60. Isabel F. Cruz, Alberto O. Mendelzon, Peter T. Wood: A Graphical Query Language Supporting Recursion. *SIGMOD Conference 1987*: 323-330

61. Mariano P. Consens, Alberto O. Mendelzon: Low Complexity Aggregation in GraphLog and Datalog. *Theor. Comput. Sci.* 116(1): 95-116 (1993)
62. Pablo Barcelo: Querying graph databases. *PODS 2013*: 175-188
63. Renzo Angles, Claudio Gutiérrez: Survey of graph database models. *ACM Comput. Surv.* 40(1) (2008)
64. Peter T. Wood: Query languages for graph databases. *SIGMOD Record* 41(1): 50-60 (2012)
65. Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Moshe Y. Vardi: Rewriting of Regular Expressions and Regular Path Queries. *J. Comput. Syst. Sci.* 64(3): 443-465 (2002)
66. Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Moshe Y. Vardi: Reasoning on regular path queries. *SIGMOD Record* 32(4): 83-92 (2003)
67. Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Moshe Y. Vardi: Containment of Conjunctive Regular Path Queries with Inverse. *KR 2000*: 176-185
68. Leonid Libkin, Wim Martens, Domagoj Vrgoc: Querying graph databases with XPath. *ICDT 2013*: 129-140
69. Pablo Barcelo, Leonid Libkin, Anthony Widjaja Lin, Peter T. Wood: Expressive Languages for Path Queries over Graph-Structured Data. *ACM Trans. Database Syst.* 37(4): 31 (2012)
70. Pablo Barcelo, Diego Figueira, Leonid Libkin: Graph Logics with Rational Relations. *Logical Methods in Computer Science* 9(3) (2013)
71. Dominik D. Freydenberger, Nicole Schweikardt: Expressiveness and Static Analysis of Extended Conjunctive Regular Path Queries. *AMW 2011*
72. Jelle Hellings, Bart Kuijpers, Jan Van den Bussche, Xiaowang Zhang: Walk logic as a framework for path query languages on graph databases. *ICDT 2013*: 117-128
73. Pablo Barcelo, Leonid Libkin, Juan L. Reutter: Querying Regular Graph Patterns. *Journal of the ACM* 61(1): 8:1-8:54 (2014)
74. Jorge Perez, Marcelo Arenas, Claudio Gutierrez: Semantics and complexity of SPARQL. *ACM Trans. Database Syst.* 34(3) (2009)
75. Marcelo Arenas, Jorge Perez: Querying semantic web data with SPARQL. *PODS 2011*: 305-316
76. Marcelo Arenas, Georg Gottlob, Andreas Pieris: Expressive languages for querying the semantic web. *PODS 2014*: 14-26
77. Leonid Libkin, Wim Martens, Domagoj Vrgoc: Querying Graphs with Data. *J. ACM* 63(2): 14:1-14:53 (2016)
78. Leonid Libkin, Juan L. Reutter, Adrian Soto, Domagoj Vrgoc: TriAL: A Navigational Algebra for RDF Triplestores. *ACM Trans. Database Syst.* 43(1): 5:1-5:46 (2018)
79. Wenfei Fan, Xin Wang, Yinghui Wu: Querying big graphs within bounded resources. *SIGMOD Conference 2014*: 301-312

80. Wenfei Fan: Graph pattern matching revised for social network analysis. ICDT 2012: 8-21
81. Jorge Perez, Marcelo Arenas, Claudio Gutierrez: nSPARQL: A navigational language for RDF. J. Web Sem. 8(4): 255-270 (2010)
82. Marcelo Arenas, Sebastian Conca, Jorge Perez: Counting beyond a Yottabyte, or how SPARQL 1.1 property paths will prevent adoption of the standard. WWW 2012: 629-638
83. Katja Losemann, Wim Martens: The complexity of regular expressions and property paths in SPARQL. ACM Trans. Database Syst. 38(4): 24 (2013)
84. Wim Martens, Tina Trautner: Evaluation and Enumeration Problems for Regular Path Queries. ICDT 2018: 19:1-19:21
85. Angela Bonifati, Wim Martens, Thomas Timm: An Analytical Study of Large SPARQL Query Logs. PVLDB 11(2): 149-161 (2017)
86. Oskar van Rest, Sungpack Hong, Jinha Kim, Xuming Meng, Hassan Chafi: PGQL: a property graph query language. GRADES 2016.
87. Nadime Francis, Alastair Green, Paolo Guagliardo, Leonid Libkin, Tobias Lindaaker, Victor Marsault, Stefan Plantikow, Mats Rydberg, Petra Selmer, Andres Taylor: Cypher: An Evolving Query Language for Property Graphs. SIGMOD Conference 2018: 1433-1445
88. Renzo Angles, Marcelo Arenas, Pablo Barcelo, Peter A. Boncz, George H. L. Fletcher, Claudio Gutierrez, Tobias Lindaaker, Marcus Paradies, Stefan Plantikow, Juan F. Sequeda, Oskar van Rest, Hannes Voigt: G-CORE: A Core for Future Graph Query Languages. SIGMOD Conference 2018: 1421-1432

5 Interoperability: data integration

89. Alon Y. Levy, Alberto O. Mendelzon, Yehoshua Sagiv, Divesh Srivastava: Answering Queries Using Views. PODS 1995: 95-104
90. Rachel Pottinger, Alon Y. Halevy: MiniCon: A scalable algorithm for answering queries using views. VLDB J. 10(2-3): 182-198 (2001)
91. Oliver M. Duschka, Michael R. Genesereth: Answering Recursive Queries Using Views. PODS 1997: 109-116
92. Alon Y. Levy, Anand Rajaraman, Joann J. Ordille: Query-Answering Algorithms for Information Agents. AAAI/IAAI, Vol. 1 1996: 40-47
93. Alon Y. Halevy: Theory of Answering Queries Using Views. SIGMOD Record 29(4): 40-47 (2000)
94. Xin Luna Dong, Alon Y. Halevy, Cong Yu: Data integration with uncertainty. VLDB J. 18(2): 469-500 (2009)

95. Serge Abiteboul, Oliver M. Duschka: Complexity of Answering Queries Using Materialized Views. *PODS 1998*: 254-263
96. Serge Abiteboul, Omar Benjelloun, Ioana Manolescu, Tova Milo, Roger Weber: Active XML: Peer-to-Peer Data and Web Services Integration. *VLDB 2002*: 1087-1090
97. Alin Deutsch, Yannis Katsis, Yannis Papakonstantinou: Determining source contribution in integration systems. *PODS 2005*: 304-315
98. Andrea Cali, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini: Data integration under integrity constraints. *Inf. Syst.* 29(2): 147-163 (2004)
99. Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, Riccardo Rosati: Logical Foundations of Peer-To-Peer Data Integration. *PODS 2004*: 241-251
100. Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini: Description Logics for Information Integration. *Computational Logic: Logic Programming and Beyond 2002*: 41-60
101. Andrea Cali, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini: On the Expressive Power of Data Integration Systems. *ER 2002*: 338-350
102. Maurizio Lenzerini: Data Integration Needs Reasoning. *LPNMR 2001*: 54-61
103. Jayant Madhavan, Shirley Cohen, Xin Luna Dong, Alon Y. Halevy, Shawn R. Jeffery, David Ko, Cong Yu: Web-Scale Data Integration: You can afford to Pay as You Go. *CIDR 2007*: 342-350
104. Shawn R. Jeffery, Michael J. Franklin, Alon Y. Halevy: Pay-as-you-go user feedback for data-space systems. *SIGMOD Conference 2008*: 847-860
105. Khalid Belhajjame, Norman W. Paton, Alvaro A. A. Fernandes, Cornelia Hedeler, Suzanne M. Embury: User Feedback as a First Class Citizen in Information Integration Systems. *CIDR 2011*: 175-183

6 Interoperability: data exchange

106. Ronald Fagin, Phokion G. Kolaitis, Lucian Popa: Data exchange: getting to the core. *ACM Trans. Database Syst.* 30(1): 174-210 (2005)
107. Georg Gottlob, Alan Nash: Efficient core computation in data exchange. *Journal of the ACM* 55(2) (2008)
108. Marcelo Arenas, Pablo Barcelo, Ronald Fagin, Leonid Libkin: Solutions and query rewriting in data exchange. *Inf. Comput.* 228: 28-61 (2013)
109. Marcelo Arenas, Pablo Barcelo, Juan L. Reutter: Query Languages for Data Exchange: Beyond Unions of Conjunctive Queries. *Theory Comput. Syst.* 49(2): 489-564 (2011)
110. Ronald Fagin, Phokion G. Kolaitis, Lucian Popa, Wang Chiew Tan: Composing schema mappings: Second-order dependencies to the rescue. *ACM Trans. Database Syst.* 30(4): 994-1055 (2005)

111. Marcelo Arenas, Ronald Fagin, Alan Nash: Composition with Target Constraints .Logical Methods in Computer Science 7(3) (2011)
112. Philip A. Bernstein, Todd J. Green, Sergey Melnik, Alan Nash: Implementing mapping composition. VLDB J. 17(2): 333-353 (2008)
113. Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Riccardo Rosati: On reconciling data exchange, data integration, and peer data management. PODS 2007: 133-142
114. Ronald Fagin: Inverting schema mappings. ACM Trans. Database Syst. 32(4) (2007)
115. Patricia C. Arocena, Boris Glavic, Radu Ciucanu, Rene J. Miller: The iBench Integration Metadata Generator. PVLDB 9(3): 108-119 (2015)
116. Bogdan Alexe, Wang Chiew Tan, Yannis Velegrakis: STBenchmark: towards a benchmark for mapping systems. PVLDB 1(1): 230-244 (2008)
117. Marcelo Arenas, Jorge Pérez, Juan L. Reutter: Data exchange beyond complete data. Journal of the ACM 60(4): 28 (2013)
118. Andre Hernich: Answering Non-Monotonic Queries in Relational Data Exchange .Logical Methods in Computer Science 7(3) (2011)
119. Gosta Grahne, Adrian Onet: Representation systems for data exchange. ICDT 2012: 208-221
120. Gosta Grahne, Ali Moallemi, Adrian Onet: Intuitionistic Data Exchange. AMW 2015
121. Andre Hernich, Leonid Libkin, Nicole Schweikardt: Closed world data exchange. ACM Trans. Database Syst. 36(2): 14 (2011)
122. Leonid Libkin, Cristina Sirangelo: Data exchange and schema mappings in open and closed worlds. J. Comput. Syst. Sci. 77(3): 542-571 (2011)
123. Marcelo Arenas, Leonid Libkin: XML data exchange: Consistency and query answering. Journal of the ACM 55(2) (2008)
124. Shun'ichi Amano, Claire David, Leonid Libkin, Filip Murlak: XML Schema Mappings: Data Exchange and Metadata Management. Journal of the ACM 61(2): 12:1-12:48 (2014)
125. Rada Chirkova, Leonid Libkin, Juan L. Reutter: Tractable XML data exchange via relations. Frontiers of Computer Science 6(3): 243-263 (2012)
126. Pablo Barcelo, Jorge Perez, Juan L. Reutter: Schema mappings and data exchange for graph databases. ICDT 2013:189-200
127. Ronald Fagin, Benny Kimelfeld, Phokion G. Kolaitis: Probabilistic data exchange. Journal of the ACM 58(4): 15 (2011)

7 Veracity: Incomplete information

128. Tomasz Imielinski, Witold Lipski Jr.: Incomplete Information in Relational Databases. *Journal of the ACM* 31(4): 761-791 (1984)
129. Serge Abiteboul, Paris C. Kanellakis, Gosta Grahne: On the Representation and Querying of Sets of Possible Worlds. *Theor. Comput. Sci.* 78(1): 158-187 (1991)
130. Raymond Reiter: What Should a Database Know? *J. Log. Program.* 14(1&2): 127-153 (1992)
131. Raymond Reiter: A sound and sometimes complete query evaluation algorithm for relational databases with null values. *Journal of the ACM* 33(2): 349-370 (1986)
132. Tomasz Imielinski: Incomplete Deductive Databases. *Ann. Math. Artif. Intell.* 3(2-4): 259-293 (1991)
133. Jan Chomicki, Tomasz Imielinski: Finite Representation of Infinite Query Answers. *ACM Trans. Database Syst.* 18(2): 181-223 (1993)
134. Tomasz Imielinski, Ron van der Meyden, Kumar V. Vadaparty: Complexity Tailored Design: A New Design Methodology for Databases With Incomplete Information. *J. Comput. Syst. Sci.* 51(3): 405-432 (1995)
135. Leonid Libkin: Certain answers as objects and knowledge. *Artificial Intelligence*, Volume 232, March 2016, Pages 119
136. Leonid Libkin: SQL's Three-Valued Logic and Certain Answers. *ICDT 2015*: 94-109
137. Amélie Gheerbrant, Leonid Libkin, Cristina Sirangelo: Naive Evaluation of Queries over Incomplete Databases. *ACM Trans. Database Syst.* 39(4): 31:1-31:42 (2014)
138. Amélie Gheerbrant, Leonid Libkin: Certain Answers over Incomplete XML Documents: Extending Tractability Boundary. *Theory Comput. Syst.* 57(4): 892-926 (2015)
139. Pablo Barcelo, Leonid Libkin, Antonella Poggi, Cristina Sirangelo: XML with incomplete information. *Journal of the ACM* 58(1): 4 (2010)
140. Dan Olteanu, Christoph Koch, Lyublena Antova: World-set decompositions: Expressiveness and efficient algorithms. *Theor. Comput. Sci.* 403(2-3):265-284 (2008)
141. Lyublena Antova, Christoph Koch, Dan Olteanu: 10^{10^6} worlds and beyond: efficient representation and processing of incomplete information. *VLDB J.* 18(5):1021-1040 (2009)
142. Robert Fink, Dan Olteanu, Swaroop Rath: Providing support for full relational algebra in probabilistic databases. *ICDE 2011*: 315-326
143. Serge Abiteboul, Benny Kimelfeld, Yehoshua Sagiv, Pierre Senellart: On the expressiveness of probabilistic XML models. *VLDB J.* 18(5): 1041-1064 (2009)
144. Nilesh N. Dalvi, Christopher Re, Dan Suciu: Queries and materialized views on probabilistic databases. *J. Comput. Syst. Sci.* 77(3): 473-490 (2011)
145. Paolo Guagliardo, Leonid Libkin, Making SQL queries correct on incomplete databases: a feasibility study, *PODS 2016*: 211–223.

8 Veracity: Inconsistent data; data cleaning

146. Marcelo Arenas, Leopoldo E. Bertossi, Jan Chomicki: Consistent Query Answers in Inconsistent Databases. *PODS 1999*: 68-79
147. Marcelo Arenas, Leopoldo E. Bertossi, Jan Chomicki: Scalar Aggregation in FD-Inconsistent Databases. *ICDT 2001*: 39-53
148. Pablo Barcelo, Leopoldo E. Bertossi: Logic Programs for Querying Inconsistent Databases. *PADL 2003*: 208-222
149. Leopoldo E. Bertossi, Loreto Bravo: Consistent Query Answers in Virtual Data Integration Systems. *Inconsistency Tolerance 2005*: 42-83
150. Jef Wijsen: Database repairing using updates. *ACM Trans. Database Syst.* 30(3): 722-768 (2005)
151. Jef Wijsen: On the consistent rewriting of conjunctive queries under primary key constraints. *Inf. Syst.* 34(7): 578-601 (2009)
152. Jef Wijsen: Certain conjunctive query answering in first-order logic. *ACM Trans. Database Syst.* 37(2): 9 (2012)
153. Gaëlle Fontaine: Why Is It Hard to Obtain a Dichotomy for Consistent Query Answering? *ACM Trans. Comput. Log.* 16(1): 7:1-7:24 (2015)
154. Wenfei Fan, Floris Geerts, Jef Wijsen: Determining the currency of data. *PODS 2011*: 71-82
155. Wenfei Fan, Floris Geerts, Xibei Jia, Anastasios Kementsidis: Conditional functional dependencies for capturing data inconsistencies. *ACM Trans. Database Syst.* 33(2) (2008) 9
156. Wenfei Fan, Floris Geerts, Jianzhong Li, Ming Xiong: Discovering Conditional Functional Dependencies. *IEEE Trans. Knowl. Data Eng.* 23(5): 683-698 (2011)
157. Shuai Ma, Wenfei Fan, Loreto Bravo: Extending inclusion dependencies with conditions. *Theor. Comput. Sci.* 515: 64-95 (2014)
158. Leopoldo E. Bertossi, Solmaz Kolahi, Laks V. S. Lakshmanan: Data cleaning and query answering with matching dependencies and matching functions. *ICDT 2011*: 268-279

9 Veracity: Ontology mediated querying

159. Shqiponja Ahmetaj, Magdalena Ortiz, Mantas Simkus: Polynomial Datalog Rewritings for Expressive Description Logics with Closed Predicates. *IJCAI 2016*: 878-885
160. Meghyn Bienvenu, Peter Hansen, Carsten Lutz, Frank Wolter: First Order-Rewritability and Containment of Conjunctive Queries in Horn Description Logics. *IJCAI 2016*: 965-971

161. Georg Gottlob, Stanislav Kikot, Roman Kontchakov, Vladimir V. Podolskii, Thomas Schwentick, Michael Zakharyashev: The price of query rewriting in ontology-based data access. *Artif. Intell.* 213: 42-59 (2014)
162. Georg Gottlob, Marco Manna, Andreas Pieris: Polynomial Rewritings for Linear Existential Rules. *IJCAI 2015*: 2992-2998
163. Georg Gottlob, Giorgio Orsi, Andreas: Query Rewriting and Optimization for Ontological Databases. *ACM Trans. Database Syst.* 39(3): 25:1-25:46 (2014)
164. Georg Gottlob, Marco Manna, Andreas Pieris: Polynomial Combined Rewritings for Existential Rules. *KR 2014*
165. Peter Hansen, Carsten Lutz, Inan Seylan, Frank Wolter: Efficient Query Rewriting in the Description Logic EL and Beyond. *IJCAI 2015*: 3034-3040
166. Melanie Koenig, Michel Leclere, Marie-Laure Mugnier, Michael Thomazo: Sound, complete and minimal UCQ-rewriting for existential rules. *Semantic Web* 6(5): 451-475 (2015)
167. Roman Kontchakov, Carsten Lutz, David Toman, Frank Wolter, Michael Zakharyashev: The Combined Approach to Ontology-Based Data Access. *IJCAI 2011*: 2656-2661
168. Riccardo Rosati, Alessandro Almatelli: Improving Query Answering over DL-Lite Ontologies. *KR 2010*