

References

- Abadi, D. J., Carney, D., Çetintemel, U., Cherniack, M., Convey, C., Lee, S., Stonebraker, M., Tatbul, N., and Zdonik, S. (2003). Aurora: a new model and architecture for data stream management. *VLDB J.*, 12(2):120–139. 498, 499
- Abadi, D. J., Ahmad, Y., Balazinska, M., Çetintemel, U., Cherniack, M., Hwang, J.-H., Lindner, W., Maskey, A., Rasin, A., Ryvkina, E., Tatbul, N., Xing, Y., and Zdonik, S. B. (2005). The design of the Borealis stream processing engine. In *Proc. 2nd Biennial Conf. on Innovative Data Systems Research*, pages 277–289. 499
- Abadi, D. J., Marcus, A., Madden, S. R., and Hollenbach, K. (2007). Scalable semantic web data management using vertical partitioning. In *Proc. 33rd Int. Conf. on Very Large Data Bases*, pages 411–422. 593
- Abadi, D. J., Marcus, A., Madden, S., and Hollenbach, K. (2009). SW-Store: a vertically partitioned DBMS for semantic web data management. *VLDB J.*, 18(2): 385–406. 593
- Aberer, K. (2001). P-grid: A self-organizing access structure for P2P information systems. In *Proc. Int. Conf. on Cooperative Inf. Syst.*, pages 179–194. 433
- Aberer, K. (2003). Guest editor’s introduction. *ACM SIGMOD Rec.*, 32(3):21–22. 433
- Aberer, K., Cudré-Mauroux, P., Datta, A., Despotovic, Z., Hauswirth, M., Puceva, M., and Schmidt, R. (2003a). P-grid: a self-organizing structured P2P system. *ACM SIGMOD Rec.*, 32(3):29–33. 433, 434
- Aberer, K., Cudré-Mauroux, P., and Hauswirth, M. (2003b). Start making sense: The chatty web approach for global semantic agreements. *J. Web Semantics*, 1(1): 89–114. 433
- Abiteboul, S., Quass, D., McHugh, J., Widom, J., and Wiener, J. (1997). The Lorel query language for semistructured data. *Int. J. Digit. Libr.*, 1(1):68–88. 593
- Abiteboul, S., Buneman, P., and Suciu, D. (1999). *Data on the Web: From Relations to Semistructured Data and XML*. Morgan Kaufmann. 592
- Abiteboul, S., Manolescu, I., Rigaux, P., Rousset, M.-C., and Senellart, P. (2011). *Web Data Management*. Cambridge University Press. 592
- Abou-Rjeili, A. and Karypis, G. (2006). Multilevel algorithms for partitioning power-law graphs. In *Proc. 20th IEEE Int. Parallel & Distributed Processing Symp.*, pages 124–124. 499

- Abouzeid, A., Bajda-Pawlikowski, K., Abadi, D., Silberschatz, A., and Rasin, A. (2009). HadoopDB: an architectural hybrid of MapReduce and DBMS technologies for analytical workloads. *Proc. VLDB Endowment*, 2(1):922–933. 536
- Abramson, N. (1973). The ALOHA system. In Abramson, N. and Kuo, F. F., editors, *Computer Communication Networks*. Prentice-Hall. D-6
- Adali, S., Candan, K. S., Papakonstantinou, Y., and Subrahmanian, V. S. (1996a). Query caching and optimization in distributed mediator systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 137–148. 326
- Adali, S., Candan, K. S., Papakonstantinou, Y., and Subrahmanian, V. S. (1996b). Query caching and optimization in distributed mediator systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 137–148. 328
- Adamic, L. and Huberman, B. (2000). The nature of markets in the world wide web. *Quart. J. Electron. Comm.*, 1:5–12. 498
- Adiba, M. (1981). Derived relations: A unified mechanism for views, snapshots and distributed data. In *Proc. 7th Int. Conf. on Very Data Bases*, pages 293–305. 123
- Adiba, M. and Lindsay, B. (1980). Database snapshots. In *Proc. 6th Int. Conf. on Very Data Bases*, pages 86–91. 123
- Adler, M. and Mitzenmacher, M. (2001). Towards compressing web graphs. In *Proc. Data Compression Conf.*, pages 203–212. 592
- Aggarwal, C. C., editor. (2007). *Data Streams: Models and Algorithms*. Springer. 498
- Agichtein, E., Lawrence, S., and Gravano, L. (2004). Learning to find answers to questions on the web. *ACM Trans. Internet Tech.*, 4(3):129–162. 593
- Agrawal, D. and El Abbadi, A. (1990). Locks with constrained sharing. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 85–93. C-32
- Agrawal, D. and Sengupta, S. (1993). Modular synchronization in distributed, multiversion databases: Version control and concurrency control. *IEEE Trans. Knowl. and Data Eng.*, 5(1):126–137. 237
- Agrawal, D., Das, S., and El Abbadi, A. (2012). *Data Management in the Cloud: Challenges and Opportunities*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers. 31, 32
- Agrawal, R. and DeWitt, D. J. (1985). Integrated concurrency control and recovery mechanisms. *ACM Trans. Database Syst.*, 10(4):529–564. C-44
- Agrawal, S., Narasayya, V., and Yang, B. (2004). Integrating vertical and horizontal partitioning into automated physical database design. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 85
- Akal, F., Böhm, K., and Schek, H.-J. (2002). Olap query evaluation in a database cluster: A performance study on intra-query parallelism. In *Proc. 6th East European Conf. Advances in Databases and Information Systems*, pages 218–231. 382
- Akal, F., Türker, C., Schek, H.-J., Breitbart, Y., Grabs, T., and Veen, L. (2005). Fine-grained replication and scheduling with freshness and correctness guarantees. In *Proc. 31st Int. Conf. on Very Large Data Bases*, pages 565–576. 273
- Akbarinia, R. and Martins, V. (2007). Data management in the APPA system. *J. Grid Comp.*, 5(3):303–317. 433
- Akbarinia, R., Martins, V., Pacitti, E., and Valduriez, P. (2006). Design and implementation of Atlas P2P architecture. In Baldoni, R., Cortese, G., and Davide, F., editors, *Global Data Management*, pages 98–123. IOS Press. 433, 434

- Akbarinia, R., Pacitti, E., and Valduriez, P. (2007a). Processing top-k queries in distributed hash tables. In *Proc. 13th Int. Euro-Par Conf.*, pages 489–502. 434
- Akbarinia, R., Pacitti, E., and Valduriez, P. (2007b). Query processing in P2P systems. Technical Report 6112, INRIA, Rennes, France. 434
- Akbarinia, R., Pacitti, E., and Valduriez, P. (2007c). Best position algorithms for top-k queries. In *Proc. 33rd Int. Conf. on Very Large Data Bases*, pages 495–506. 434
- Akbarinia, R., Pacitti, E., and Valduriez, P. (2007d). Data currency in replicated dhts. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 211–222. 434
- Akidau, T., Balikov, A., Bekiroglu, K., Chernyak, S., Haberman, J., Lax, R., McVeety, S., Mills, D., Nordstrom, P., and Whittle, S. (2013). MillWheel: Fault-tolerant stream processing at internet scale. *Proc. VLDB Endowment*, 6(11):1033–1044. 499
- Alagiannis, I., Borovica, R., Branco, M., Idreos, S., and Ailamaki, A. (2012). NoDB: efficient query execution on raw data files. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 241–252. 86
- Alagiannis, I., Idreos, S., and Ailamaki, A. (2014). H2O: A hands-free adaptive store. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1103–1114. 86
- Alamoudi, A. A., Grover, R., Carey, M. J., and Borkar, V. R. (2015). External data access and indexing in AsterixDB. In *Proc. 24th ACM Int. Conf. on Information and Knowledge Management*, pages 3–12. 535
- Albutiu, M.-C., Kemper, A., and Neumann, T. (2012). Massively parallel sort-merge joins in main memory multi-core database systems. *Proc. VLDB Endowment*, 5(10):1064–1075. 381
- Allard, T., Hébrail, G., Masegla, F., and Pacitti, E. (2015). Chiaroscuro: Transparency and privacy for massive personal time-series clustering. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 779–794. 434
- Alomari, M., Cahill, M., Fekete, A., and Rohm, U. (2008). The cost of serializability on platforms that use snapshot isolation. In *Proc. 24th Int. Conf. on Data Engineering*, pages 576–585. 238
- Alomari, M., Fekete, A., and Rohm, U. (2009). A robust technique to ensure serializable executions with snapshot isolation DBMS. In *Proc. 25th Int. Conf. on Data Engineering*, pages 341–352. 238
- Alsberg, P. A. and Day, J. D. (1976). A principle for resilient sharing of distributed resources. In *Proc. 2nd Int. Conf. on Software Engineering*, pages 562–570. 237
- Alsubaiee, S., Altowim, Y., Altwajry, H., Behm, A., Borkar, V. R., Bu, Y., Carey, M. J., Cetindil, I., Cheelangi, M., Faraaz, K., Gabrielova, E., Grover, R., Heilbron, Z., Kim, Y., Li, C., Li, G., Ok, J. M., Onose, N., Pirzadeh, P., Tsotras, V. J., Vernica, R., Wen, J., and Westmann, T. (2014). AsterixDB: A scalable, open source DBMS. *Proc. VLDB Endowment*, 7(14):1905–1916. 535
- Altingövde, I. S. and Ulusoy, Ö. (2004). Exploiting interclass rules for focused crawling. *IEEE Intelligent Systems*, 19(6):66–73. 593
- Aluç, G. (2015). *Workload Matters: A Robust Approach to Physical RDF Database Design*. PhD thesis, University of Waterloo. 593
- Alvarez, V., Schuhknecht, F. M., Dittrich, J., and Richter, S. (2014). Main memory adaptive indexing for multi-core systems. In *Proc. 10th Workshop on Data Management on New Hardware*, pages 3:1—3:10. 86

- Amdahl, G. M. (1967). Validity of the single processor approach to achieving large scale computing capabilities. In *Proc. Spring Joint Computer Conf.*, pages 483–485. 381
- Amsaleg, L., Franklin, M. J., Tomasic, A., and Urhan, T. (1996). Scrambling query plans to cope with unexpected delays. In *Proc. 4th Int. Conf. on Parallel and Distributed Information Systems*, pages 208–219. 177, 328
- Andreev, K. and Racke, H. (2006). Balanced graph partitioning. *Theor. Comp. Sci.*, 39(6):929–939. 499
- Angles, R. and Gutierrez, C. (2008). The expressive power of SPARQL. In *Proc. 7th Int. Semantic Web Conf.*, pages 114–129. 593
- ANSI. (1992). *Database Language SQL*, ANSI X3.135-1992 edition. C-12
- Antoniou, G. and Plexousakis, D. (2018). Semantic web. In Liu, L. and Özsu, M. T., editors, *Encyclopedia of Database Systems*, pages 3425–3429. Springer New York, New York, NY. 571, 593
- Apache. (2016). Apache Giraph. <http://giraph.apache.org>. Last accessed June 2019. 500
- Apers, P., van den Berg, C., Flokstra, J., Grefen, P., Kersten, M., and Wilschut, A. (1992). Prisma/DB: a parallel main-memory relational DBMS. *IEEE Trans. Knowl. and Data Eng.*, 4:541–554. 381
- Apers, P. M. G. (1981). Redundant allocation of relations in a communication network. In *Proc. 5th Berkeley Workshop on Distributed Data Management and Computer Networks*, pages 245–258. 85
- Arasu, A. and Widom, J. (2004). A denotational semantics for continuous queries over streams and relations. *ACM SIGMOD Rec.*, 33(3):6–11. 498
- Arasu, A., Cho, J., Garcia-Molina, H., Paepcke, A., and Raghavan, S. (2001). Searching the web. *ACM Trans. Internet Tech.*, 1(1):2–43. 542, 592, 593
- Arasu, A., Babu, S., and Widom, J. (2006). The CQL continuous query language: Semantic foundations and query execution. *VLDB J.*, 15(2):121–142. 498, 499
- Armbrust, M., Xin, R. S., Lian, C., Huai, Y., Liu, D., Bradley, J. K., Meng, X., Kaftan, T., Franklin, M. J., Ghodsi, A., and Zaharia, M. (2015). Spark SQL: Relational data processing in Spark. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1383–1394. 498, 536
- Arocena, G. and Mendelzon, A. (1998). WebOQL: Restructuring documents, databases and webs. In *Proc. 14th Int. Conf. on Data Engineering*, pages 24–33. 593
- Asad, O. and Kemme, B. (2016). Adaptcache: Adaptive data partitioning and migration for distributed object caches. In *Proc. ACM/IFIP/USENIX 17th Int. Middleware Conf.*, pages 7:1–7:13. 85
- Aspnes, J. and Shah, G. (2003). Skip graphs. In *Proc. 14th Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 384–393. 433
- Astrahan, M. M., Blasgen, M. W., Chamberlin, D. D., Eswaran, K. P., Gray, J. N., Griffiths, P. P., King, W. F., Lorie, R. A., McJones, P. R., Mehl, J. W., Putzolu, G. R., Traiger, I. L., Wade, B. W., and Watson, V. (1976). System r: A relational database management system. *ACM Trans. Database Syst.*, 1(2):97–137. B-27, C-44
- Avnur, R. and Hellerstein, J. (2000). Eddies: Continuously adaptive query processing. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 261–272. 177, 328

- Ayad, A. and Naughton, J. (2004). Static optimization of conjunctive queries with sliding windows over unbounded streaming information sources. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 419–430. 498
- Azar, Y., Broder, A. Z., Karlin, A. R., and Upfal, E. (1999). Balanced allocations. *SIAM J. on Comput.*, 29(1):180–200. 499
- Babb, E. (1979). Implementing a relational database by means of specialized hardware. *ACM Trans. Database Syst.*, 4(1):1–29. 380
- Babcock, B., Babu, S., Datar, M., Motwani, R., and Widom, J. (2002). Models and issues in data stream systems. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 1–16. 498
- Balazinska, M., Kwon, Y., Kuchta, N., and Lee, D. (2007). Moirae: History-enhanced monitoring. In *Proc. 3rd Biennial Conf. on Innovative Data Systems Research*, pages 375–386. 498
- Balke, W.-T., Nejd, W., Siberski, W., and Thaden, U. (2005). Progressive distributed top-k retrieval in peer-to-peer networks. In *Proc. 21st Int. Conf. on Data Engineering*, pages 174–185. 434
- Ball, M. O. and Hardie, F. (1967). Effects and detection of intermittent failures in digital systems. Technical Report Internal Report 67-825-2137, IBM. Cited in [Siewiorek and Swarz, 1982]. C-38
- Balter, R., Berard, P., and Decitre, P. (1982). Why control of concurrency level in distributed systems is more important than deadlock management. In *Proc. ACM SIGACT-SIGOPS 1st Symp. on the Principles of Distributed Computing*, pages 183–193. C-22
- Bancilhon, F. and Spyrtatos, N. (1981). Update semantics of relational views. *ACM Trans. Database Syst.*, 6(4):557–575. 123
- Barbara, D., Garcia-Molina, H., and Spauster, A. (1986). Policies for dynamic vote reassignment. In *Proc. 6th IEEE Int. Conf. on Distributed Computing Systems*, pages 37–44. 273
- Barbara, D., Molina, H. G., and Spauster, A. (1989). Increasing availability under mutual exclusion constraints with dynamic voting reassignment. *ACM Trans. Comp. Syst.*, 7(4):394–426. 273
- Barthels, C., Loesing, S., Alonso, G., and Kossmann, D. (2015). Rack-scale in-memory join processing using RDMA. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1463–1475. 381
- Batini, C. and Lenzirini, M. (1984). A methodology for data schema integration in entity-relationship model. *IEEE Trans. Softw. Eng.*, SE-10(6):650–654. 327
- Batini, C., Lenzirini, M., and Navathe, S. B. (1986). A comparative analysis of methodologies for database schema integration. *ACM Comput. Surv.*, 18(4):323–364. 326, 327
- Beeri, C., Bernstein, P. A., and Goodman, N. (1989). A model for concurrency in nested transaction systems. *J. ACM*, 36(2):230–269. 238
- Bell, D. and Grimson, J. (1992). *Distributed Database Systems*. Addison Wesley. Reading. 31
- Bell, D. and Lapuda, L. (1976). Secure computer systems: Unified exposition and Multics interpretation. Technical Report MTR-2997 Rev.1, MITRE Corp, Bedford, MA. 124
- Berenson, H., Bernstein, P., Gray, J., Melton, J., O’Neil, E., and O’Neil, P. (1995). A critique of ansi sql isolation levels. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1–10. 237, 238, C-12, C-13, C-27

- Bergamaschi, S. (2001). Semantic integration of heterogeneous information sources. *Data & Knowl. Eng.*, 36(3):215–249. 327
- Bergman, M. K. (2001). The deep web: Surfacing hidden value. *J. Electronic Publishing*, 7(1). 592
- Bergsten, B., Couprie, M., and Valduriez, P. (1991). Prototyping DBS3, a shared-memory parallel database system. In *Proc. Int. Conf. on Parallel and Distributed Information Systems*, pages 226–234. 381
- Bergsten, B., Couprie, M., and Valduriez, P. (1993). Overview of parallel architectures for databases. *The Comp. J.*, 36(8):734–739. 381
- Berkholz, C., Keppeler, J., and Schweikardt, N. (2017). Answering conjunctive queries under updates. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 303–318. 124
- Berlin, J. and Motro, A. (2001). Autoplex: Automated discovery of content for virtual databases. In *Proc. Int. Conf. on Cooperative Inf. Syst.*, pages 108–122. 327
- Berners-Lee, T. (2006). Linked data. Accessible at <https://www.w3.org/DesignIssues/LinkedData.html>. Last accessed June 2019. 593
- Bernstein, P. and Blaustein, B. (1982). Fast methods for testing quantified relational calculus assertions. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 39–50. 124
- Bernstein, P. and Melnik, S. (2007). Model management: 2.0: Manipulating richer mappings. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1–12. 324, 326, 329
- Bernstein, P., Blaustein, B., and Clarke, E. M. (1980a). Fast maintenance of semantic integrity assertions using redundant aggregate data. In *Proc. 6th Int. Conf. on Very Data Bases*, pages 126–136. 124
- Bernstein, P., Shipman, P., and Rothnie, J. B. (1980b). Concurrency control in a system for distributed databases (SDD-1). *ACM Trans. Database Syst.*, 5(1):18–51. 237
- Bernstein, P. A. and Chiu, D. M. (1981). Using semi-joins to solve relational queries. *J. ACM*, 28(1):25–40. 176
- Bernstein, P. A. and Goodman, N. (1981). Concurrency control in distributed database systems. *ACM Comput. Surv.*, 13(2):185–222. 31, 237, C-28
- Bernstein, P. A. and Goodman, N. (1983). Multiversion concurrency control — theory and algorithms. *ACM Trans. Database Syst.*, 8(4):465–483. 237
- Bernstein, P. A. and Goodman, N. (1984). An algorithm for concurrency control and recovery in replicated distributed databases. *ACM Trans. Database Syst.*, 9(4): 596–615. 273
- Bernstein, P. A. and Newcomer, E. (1997). *Principles of Transaction Processing for the Systems Professional*. Morgan Kaufmann. 237
- Bernstein, P. A., Goodman, N., Wong, E., Reeve, C. L., and Jr, J. B. R. (1981). Query processing in a system for distributed databases (SDD-1). *ACM Trans. Database Syst.*, 6(4):602–625. 176
- Bernstein, P. A., Hadzilacos, V., and Goodman, N. (1987). *Concurrency Control and Recovery in Database Systems*. Addison Wesley. 31, 237, 273, C-37, C-46, C-47, C-48
- Bernstein, P. A., Giunchiglia, F., Kementsietsidis, A., Mylopoulos, J., Serafini, L., and Zaihrayeu, I. (2002). Data management for peer-to-peer computing : A vision.

- In *Proc. 5th Int. Workshop on the World Wide Web and Databases*, pages 89–94. 433
- Bernstein, P. A., Fekete, A., Guo, H., Ramakrishnan, R., and Tamma, P. (2006). Relaxed concurrency serializability for middle-tier caching and replication. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 599–610. 273
- Beyer, K. S., Ercegovac, V., Krishnamurthy, R., Raghavan, S., Rao, J., Reiss, F., Shekita, E. J., Simmen, D. E., Tata, S., Vaithyanathan, S., and Zhu, H. (2009). Towards a scalable enterprise content analytics platform. *Q. Bull. IEEE TC on Data Eng.*, 32(1):28–35. 498
- Bharat, K. and Broder, A. (1998). A technique for measuring the relative size and overlap of public web search engines. *Comp. Networks and ISDN Syst.*, 30:379 – 388. (Proc. 7th Int. World Wide Web Conf.). 592
- Bhowmick, S. S., Madria, S. K., and Ng, W. K. (2004). *Web Data Management*. Springer. 592
- Bifet, A., Gavaldà, R., Holmes, G., and Pfahringer, B. (2018). *Machine Learning for Data Streams: with Practical Examples in MOA*. MIT Press. 498
- Binnig, C., Hildenbrand, S., Färber, F., Kossmann, D., Lee, J., and May, N. (2014). Distributed snapshot isolation: global transactions pay globally, local transactions pay locally. *VLDB J.*, 23:987–1011. 237, 238
- Biscondi, N., Brunie, L., Flory, A., and Kosch, H. (1996). Encapsulation of intra-operation parallelism in a parallel match operator. In *Proc. ACPC Conf.*, volume 1127 of *Lecture Notes in Computer Science*, pages 124–135. 382
- Bitton, D., Boral, H., DeWitt, D. J., and Wilkinson, W. K. (1983). Parallel algorithms for the execution of relational database operations. *ACM Trans. Database Syst.*, 8(3):324–353. 381
- Bitton, D., DeWitt, D. J., Hsiao, D. K., and Menon, J. (1984). A taxonomy of parallel sorting. *ACM Comput. Surv.*, 16(3):287–318. 381
- Bizer, C., Vidal, M.-E., and Skaf-Molli, H. (2018). Linked open data. In Liu, L. and Özsu, M. T., editors, *Encyclopedia of Database Systems*, pages 2096–2101. Springer New York, New York, NY. 593
- Blanas, S., Patel, J. M., Ercegovac, V., Rao, J., Shekita, E. J., and Tian, Y. (2010). A comparison of join algorithms for log processing in MapReduce. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 975–986. 498
- Blaustein, B. (1981). *Enforcing Database Assertions: Techniques and Applications*. PhD thesis, Harvard University, Cambridge, Mass. 124
- Bleiholder, J. and Naumann, F. (2009). Data fusion. *ACM Comput. Surv.*, 41(1): 1:1–1:41. 588, 594
- Bonato, A. (2008). *A Course on the Web Graph*. American Mathematical Society. 592
- Bondiombouy, C. and Valduriez, P. (2016). Query processing in multistore systems: an overview. *Int. J. Cloud Computing*, 5(4):309–346. 536
- Bondiombouy, C., Kolev, B., Levchenko, O., and Valduriez, P. (2016). Multistore big data integration with CloudMdsQL. *Trans. Large-Scale Data- and Knowledge-Centered Syst.*, 28:48–74. 536
- Bonifati, A., Summa, G., Pacitti, E., and Draïdi, F. (2014). Query reformulation in PDMS based on social relevance. *Trans. Large-Scale Data- and Knowledge-Centered Syst.*, 13:59–90. 434
- Bonnet, P., Gehrke, J., and Seshadri, P. (2001). Towards sensor database systems. In *Proc. 2nd Int. Conf. on Mobile Data Management*, pages 3–14. 499

- Bönström, V., Hinze, A., and Schweppe, H. (2003). Storing RDF as a graph. In *Proc. 1st Latin American Web Congress*, pages 27 – 36. 593
- Boral, H. and DeWitt, D. (1983). Database machines: An idea whose time has passed? A critique of the future of database machines. In *Proc. 3rd Int. Workshop on Database Machines*, pages 166–187. 380
- Boral, H., Alexander, W., Clay, L., Copeland, G., Danforth, S., Franklin, M., Hart, B., Smith, M., and Valduriez, P. (1990). Prototyping bubba, a highly parallel database system. *IEEE Trans. Knowl. and Data Eng.*, 2(1):4–24. 381
- Borkar, D., Mayuram, R., Sangudi, G., and Carey, M. J. (2016). Have your data and query it too: From key-value caching to big data management. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 239–251. 535
- Bornea, M. A., Dolby, J., Kementsietsidis, A., Srinivas, K., Dantressangle, P., Udrea, O., and Bhattacharjee, B. (2013). Building an efficient RDF store over a relational database. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 121–132. 593
- Borr, A. (1988). High performance SQL through low-level system integration. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 342–349. 237
- Bouganim, L., Florescu, D., and Valduriez, P. (1996). Dynamic load balancing in hierarchical parallel database systems. In *Proc. 22th Int. Conf. on Very Large Data Bases*, pages 436–447. 382
- Bouganim, L., Florescu, D., and Valduriez, P. (1999). Multi-join query execution with skew in NUMA multiprocessors. *Distrib. Parall. Databases*, 7(1). in press. 382
- Bratbergsengen, K. (1984). Hashing methods and relational algebra operations. In *Proc. 10th Int. Conf. on Very Large Data Bases*, pages 323–333. B-14
- Breitbart, Y. and Korth, H. F. (1997). Replication and consistency: Being lazy helps sometimes. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 173–184. 273
- Breitbart, Y. and Silberschatz, A. (1988). Multidatabase update issues. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 135–142. 381
- Breitbart, Y., Olson, P. L., and Thompson, G. R. (1986). Database integration in a distributed heterogeneous database system. In *Proc. 2nd Int. Conf. on Data Engineering*, pages 301–310. 326
- Brewer, E., Ying, L., Greenfield, L., Cypher, R., and T’so, T. (2016). Disks for data centers. Technical report, Google. 497
- Brewer, E. A. (2000). Towards robust distributed systems (abstract). In *Proc. ACM SIGACT-SIGOPS 19th Symp. on the Principles of Distributed Computing*, page 7. 535
- Bright, M. W., Hurson, A. R., and Pakzad, S. H. (1994). Automated resolution of semantic heterogeneity in multidatabases. *ACM Trans. Database Syst.*, 19(2): 212–253. 327
- Brill, D., Templeton, M., and Yu, C. (1984). Distributed query processing strategies in MERMAID: A front-end to data management systems. In *Proc. 1st Int. Conf. on Data Engineering*, pages 211–218. 328
- Brin, S. and Page, L. (1998). The anatomy of a large-scale hypertextual web search engine. *Comp. Netw.*, 30(1-7):107 – 117. 592
- Broder, A., Kumar, R., Maghoul, F., Raghavan, P., Rajagopalan, S., Stata, R., Tomkins, A., and Wiener, J. (2000). Graph structure in the web. *Comp. Netw.*, 33(1-6): 309–320. 592

- Broekstra, J., Kampman, A., and van Harmelen, F. (2002). Sesame: A generic architecture for storing and querying RDF and RDF schema. In *Proc. 1st Int. Semantic Web Conf.*, pages 54–68. 593
- Bruno, N. and Chaudhuri, S. (2002). Exploiting statistics on query expressions for optimization. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 263–274. B-23
- Bu, Y., Howe, B., Balazinska, M., and Ernst, M. D. (2010). HaLoop: efficient iterative data processing on large clusters. *Proc. VLDB Endowment*, 3(1):285–296. 499
- Bu, Y., Howe, B., Balazinska, M., and Ernst, M. D. (2012). The HaLoop approach to large-scale iterative data analysis. *VLDB J.*, 21(2):169–190. 499
- Bu, Y., Borkar, V. R., Jia, J., Carey, M. J., and Condie, T. (2014). Pregelix: Bigger graph analytics on a dataflow engine. *Proc. VLDB Endowment*, 8(2):161–172. 500
- Buchmann, A., Özsu, M. T., Hornick, M., Georgakopoulos, D., and Manola, F. (1992). A transaction model for active distributed object systems. In Elmagarmid, A. K., editor, *Transaction Models for Advanced Database Applications*, pages 123–158. Morgan Kaufmann. C-18, C-19
- Bugiotti, F., Bursztyn, D., Deutsch, A., Ileana, I., and Manolescu, I. (2015). Invisible glue: Scalable self-tuning multi-stores. In *Proc. 7th Biennial Conf. on Innovative Data Systems Research*. 536
- Buneman, P., Davidson, S., Hillebrand, G. G., and Suciu, D. (1996). A query language and optimization techniques for unstructured data. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 505–516. 593
- Cahill, M. J., Röhm, U., and Fekete, A. D. (2009). Serializable isolation for snapshot databases. *ACM Trans. Database Syst.*, 34(4):Article 20. 238
- Calì, A. and Calvanese, D. (2002). Optimized querying of integrated data over the web. In *Engineering Information Systems in the Internet Context*, pages 285–301. 326
- Callan, J. P. and Connell, M. E. (2001). Query-based sampling of text databases. *ACM Trans. Information Syst.*, 19(2):97–130. 592
- Callan, J. P., Connell, M. E., and Du, A. (1999). Automatic discovery of language models for text databases. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 479–490. 592
- Cammert, M., Krämer, J., Seeger, B., and S.Vaupel. (2006). An approach to adaptive memory management in data stream systems. In *Proc. 22nd Int. Conf. on Data Engineering*, page 137. 498
- Canaday, R. H., Harrison, R. D., Ivie, E. L., Rydery, J. L., and Wehr, L. A. (1974). A back-end computer for data base management. *Commun. ACM*, 17(10):575–582. 380
- Cao, P. and Wang, Z. (2004). Query processing issues in image (multimedia) databases. In *Proc. ACM SIGACT-SIGOPS 23rd Symp. on the Principles of Distributed Computing*, pages 206–215. 434
- Carbone, P., Katsifodimos, A., Ewen, S., Markl, V., Haridi, S., and Tzoumas, K. (2015). Apache Flink™: Stream and batch processing in a single engine. *Q. Bull. IEEE TC on Data Eng.*, 38(4):28–38. 499
- Carey, M. and Lu, H. (1986). Load balancing in a locally distributed database system. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 108–119. 177
- Castano, S. and Antonellis, V. D. (1999). A schema analysis and reconciliation tool environment for heterogeneous databases. In *Proc. 3rd Int. Conf. on Database Eng. and Applications*, pages 53 – 62. 327

- Castano, S., Fugini, M. G., Martella, G., and Samarati, P. (1995). *Database Security*. Addison Wesley. 124
- Castro, M. and Liskov, B. (1999). Practical byzantine fault tolerance. In *Proc. 3rd USENIX Symp. on Operating System Design and Implementation*, pages 173–186. 239
- Cellary, W., Gelenbe, E., and Morzy, T. (1988). *Concurrency Control in Distributed Database Systems*. North-Holland. 237, 238
- Ceri, S. and Owicki, S. (1982). On the use of optimistic methods for concurrency control in distributed databases. In *Proc. 6th Berkeley Workshop on Distributed Data Management and Computer Networks*, pages 117–130. 237
- Ceri, S. and Pelagatti, G. (1983). Correctness of query execution strategies in distributed databases. *ACM Trans. Database Syst.*, 8(4):577–607. 31, 176
- Ceri, S. and Pernici, B. (1985). DATAID–D: Methodology for distributed database design. In Albano, V. d. A. and di Leva, A., editors, *Computer-Aided Database Design*, pages 157–183. North-Holland. 84
- Ceri, S. and Widom, J. (1993). Managing semantic heterogeneity with production rules and persistent queues. In *Proc. 19th Int. Conf. on Very Large Data Bases*, pages 108–119. 327
- Ceri, S., Martella, G., and Pelagatti, G. (1982a). Optimal file allocation in a computer network: A solution method based on the knapsack problem. *Comp. Netw.*, 6: 345–357. 85
- Ceri, S., Negri, M., and Pelagatti, G. (1982b). Horizontal data partitioning in database design. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 128–136. 85
- Ceri, S., Navathe, S. B., and Wiederhold, G. (1983). Distribution design of logical database schemes. *IEEE Trans. Softw. Eng.*, SE-9(4):487–503. 84, 85
- Ceri, S., Gottlob, G., and Pelagatti, G. (1986). Taxonomy and formal properties of distributed joins. *Inf. Syst.*, 11(1):25–40. 176
- Ceri, S., Pernici, B., and Wiederhold, G. (1987). Distributed database design methodologies. *Proc. IEEE*, 75(5):533–546. 31, 84
- Chairunnanda, P., Daudjee, K., and Özsu, M. T. (2014). ConfluxDB: multi-master replication for partitioned snapshot isolation databases. *Proc. VLDB Endowment*, 7(11):947–958. 237, 238
- Chakrabarti, K., Keogh, E., Mehrotra, S., and Pazzani, M. (2002). Locally adaptive dimensionality reduction for indexing large time series databases. *ACM Trans. Database Syst.*, 27. 593
- Chakrabarti, S., Dom, B., and Indyk, P. (1998). Enhanced hypertext classification using hyperlinks. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 307 – 318. 592
- Chamberlin, D. (2018). *SQL++ For SQL Users: A Tutorial*. CouchBase Inc. 535
- Chamberlin, D., Gray, J., and Traiger, I. (1975). Views, authorization and locking in a relational database system. In *Proc. National Computer Conf.*, pages 425–430. 123
- Chamberlin, D. D., Astrahan, M. M., King, W. F., Lorie, R. A., Mehl, J. W., Price, T. G., Schkolnick, M., Selinger, P. G., Slutz, D. R., Wade, B. W., and Yost, R. A. (1981). Support for repetitive transactions and ad hoc queries in System R. *ACM Trans. Database Syst.*, 6(1):70–94. B-32
- Chambers, C., Raniwala, A., Perry, F., Adams, S., Henry, R. R., Bradshaw, R., and Weizenbaum, N. (2010). FlumeJava: easy, efficient data-parallel pipelines.

- In *Proc. ACM SIGPLAN 2010 Conf. on Programming Language Design and Implementation*, pages 363–375. 498
- Chandra, T. D., Griesemer, R., and Redstone, J. (2007). Paxos made live: An engineering perspective. In *Proc. ACM SIGACT-SIGOPS 26th Symp. on the Principles of Distributed Computing*, pages 398–407. 239
- Chandrasekaran, S., Cooper, O., Deshpande, A., Franklin, M. J., Hellerstein, J. M., Hong, W., Krishnamurthy, S., Madden, S., Raman, V., Reiss, F., and Shah, M. A. (2003). TelegraphCQ: Continuous dataflow processing for an uncertain world. In *Proc. 1st Biennial Conf. on Innovative Data Systems Research*. 498, 499
- Chang, F., Dean, J., Ghemawat, S., Hsieh, W., Wallach, D., Burrows, M., Chandra, T., Fikes, A., and Gruber, R. (2008). Bigtable: A distributed storage system for structured data. *ACM Trans. Comp. Syst.*, 26(2):Article 4. 535
- Chang, S. K. and Liu, A. C. (1982). File allocation in a distributed database. *Int. J. Comput. Inf. Sci.*, 11(5):325–340. 85
- Chattopadhyay, B., Lin, L., Liu, W., Mittal, S., Aragona, P., Lychagina, V., Kwon, Y., and Wong, M. (2011). Tenzing: A SQL implementation on the MapReduce framework. *Proc. VLDB Endowment*, 4(12):1318–1327. 498
- Chaudhuri, S., Ganjam, K., Ganti, V., and Motwani, R. (2003). Robust and efficient fuzzy match for online data cleaning. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 313–324. 327
- Chen, R., Shi, J., Chen, Y., and Chen, H. (2015). PowerLyra: Differentiated graph computation and partitioning on skewed graphs. In *Proc. 10th ACM SIGOPS/EuroSys European Conf. on Comp. Syst.*, pages 1:1–1:15. 499
- Chiu, D. M. and Ho, Y. C. (1980). A methodology for interpreting tree queries into optimal semi-join expressions. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 169–178. 176, 177
- Cho, J. and Garcia-Molina, H. (2000). The evolution of the web and implications for an incremental crawler. In *Proc. 26th Int. Conf. on Very Large Data Bases*. 593
- Cho, J. and Ntoulas, A. (2002). Effective change detection using sampling. In *Proc. 28th Int. Conf. on Very Large Data Bases*. 593
- Cho, J., Garcia-Molina, H., and Page, L. (1998). Efficient crawling through URL ordering. *Comp. Networks and ISDN Syst.*, 30(1-7):161–172. 593
- Chockler, G., Keidar, I., and Vitenberg, R. (2001). Group communication specifications: a comprehensive study. *ACM Comput. Surv.*, 33(4):427–469. 273
- Chong, E., Das, S., Eadon, G., and Srinivasan, J. (2005). An efficient SQL-based RDF querying scheme. In *Proc. 31st Int. Conf. on Very Large Data Bases*, pages 1216–1227. 593
- Chu, W. W. (1969). Optimal file allocation in a multiple computer system. *IEEE Trans. Comput.*, C-18(10):885–889. 85
- Chu, W. W. (1973). Optimal file allocation in a computer network. In Abramson, N. and Kuo, F. F., editors, *Computer Communication Networks*, pages 82–94. 85
- Chu, W. W. (1976). Performance of file directory systems for data bases in star and distributed networks. In *Proc. National Computer Conf.*, volume 45, pages 577–587. 31
- Chu, W. W. and Nahouraii, E. E. (1975). File directory design considerations for distributed databases. In *Proc. 1st Int. Conf. on Very Data Bases*, pages 543–545. 31
- Chu, X., Ilyas, I. F., and Papotti, P. (2013). Discovering Denial Constraints. *Proc. VLDB Endowment*, 6(13):1498–1509. 594

- Chundi, P., Rosenkrantz, D. J., and Ravi, S. S. (1996). Deferred updates and data placement in distributed databases. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 469–476. 273
- Civelek, F. N., Dogac, A., and Spaccapietra, S. (1988). An expert system approach to view definition and integration. In *Proc. 7th Int'l. Conf. on Entity-Relationship Approach*, pages 229–249. 124
- Codd, E. F. (1970). A relational model for large shared data banks. *Commun. ACM*, 13(6):377–387. A-7, A-16
- Codd, E. F. (1972). Relational completeness of data base sublanguages. In Rustin, R., editor, *Relational Databases*, pages 65–98. Prentice-Hall, Englewood Cliffs, N.J. A-7
- Codd, E. F. (1974). Recent investigations in relational data base systems. *Proceedings of IFIP Congress, Information Processing 74*, pages 1017–1021. A-4
- Codd, E. F. (1979). Extending the database relational model to capture more meaning. *ACM Trans. Database Syst.*, 4(4):397–434. A-2
- Cohen, J. (2009). Graph twiddling in a MapReduce world. *Computing in Science & Engineering*, 11(4):29–41. 499
- Cole, R. L. and Graefe, G. (1994). Optimization of dynamic query evaluation plans. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 150–160. 177, B-32, B-33
- Coletta, R., Castanier, E., Valduriez, P., Frisch, C., Ngo, D., and Bellahsene, Z. (2012). Public data integration with websmatch. In *Proc. Int. Workshop on Open Data*, pages 5–12. 500
- Copeland, G., Alexander, W., Boughter, E., and Keller, T. (1988). Data placement in bubba. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 99–108. 381
- Corbett, J. C., Dean, J., Epstein, M., Fikes, A., Frost, C., Furman, J. J., Ghemawat, S., Gubarev, A., Heiser, C., Hochschild, P., Hsieh, W., Kanthak, S., Kogan, E., Li, H., Lloyd, A., Melnik, S., Mwaura, D., Nagle, D., Quinlan, S., Rao, R., Rolig, L., Saito, Y., Szymaniak, M., Taylor, C., Wang, R., and Woodford, D. (2013). Spanner: Google's globally distributed database. *ACM Trans. Comp. Syst.*, 31(3):8:1–8:22. 238, 500
- Crainiceanu, A., Linga, P., Gehrke, J., and Shanmugasundaram, J. (2004). Querying peer-to-peer networks using p-trees. In *Proc. 7th Int. Workshop on the World Wide Web and Databases*, pages 25–30. 433
- Cranor, C., Johnson, T., Spatscheck, O., and Shkapenyuk, V. (2003). Gigascope: High performance network monitoring with an SQL interface. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 647–651. 498, 499
- Crespo, A. and Garcia-Molina, H. (2002). Routing indices for peer-to-peer systems. In *Proc. 22nd IEEE Int. Conf. on Distributed Computing Systems*, pages 23–33. 433
- Cuenca-Acuna, F., Peery, C., Martin, R., and Nguyen, T. (2003). PlanetP: using gossiping to build content addressable peer-to-peer information sharing communities. In *IEEE Int. Symp. on High Performance Distributed Computing*, pages 236–249. 434
- Curino, C., Jones, E., Zhang, Y., and Madden, S. (2010). Schism: a workload-driven approach to database replication and partitioning. *Proc. VLDB Endowment*, 3(1): 48–57. 85

- Curino, C., Jones, E. P. C., Madden, S., and Balakrishnan, H. (2011). Workload-aware database monitoring and consolidation. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 313–324. 31
- Cusumano, M. A. (2010). Cloud computing and SaaS as new computing platforms. *Commun. ACM*, 53(4):27–29. 31
- Dasgupta, S., Coakley, K., and Gupta, A. (2016). Analytics-driven data ingestion and derivation in the AWESOME polystore. In *Proc. 2016 IEEE Int. Conf. on Big Data*, pages 2555–2564. 536
- Daswani, N., Garcia-Molina, H., and Yang, B. (2003). Open problems in data-sharing peer-to-peer systems. In *Proc. 9th Int. Conf. on Database Theory*, pages 1–15. 433
- Date, C. J. (1987). *A Guide to the SQL Standard*. Addison Wesley. A-17
- Daudjee, K. and Salem, K. (2004). Lazy database replication with ordering guarantees. In *Proc. 20th Int. Conf. on Data Engineering*, pages 424–435. 273
- Daudjee, K. and Salem, K. (2006). Lazy database replication with snapshot isolation. In *Proc. 32nd Int. Conf. on Very Large Data Bases*, pages 715–726. 273
- Davenport, R. A. (1981). Design of distributed data base systems. *Comp. J.*, 24(1): 31–41. 84
- Davidson, S. B. (1984). Optimism and consistency in partitioned distributed database systems. *ACM Trans. Database Syst.*, 9(3):456–481. 273
- Davidson, S. B., Garcia-Molina, H., and Skeen, D. (1985). Consistency in partitioned networks. *ACM Comput. Surv.*, 17(3):341–370. 237, 272
- Dawson, J. L. (1980). A user demand model for distributed database design. In *Digest of Papers – COMPCON*, pages 211–216. 84
- Dayal, U. and Bernstein, P. (1978). On the updatability of relational views. In *Proc. 4th Int. Conf. on Very Data Bases*, pages 368–377. 123
- Dayal, U. and Hwang, H. (1984). View definition and generalization for database integration in MULTIBASE: A system for heterogeneous distributed database. *IEEE Trans. Softw. Eng.*, SE-10(6):628–644. 326, 327, 328
- Dayal, U., M.Hsu, and Ladin, R. (1991). A transactional model for long-running activities. In *Proc. 17th Int. Conf. on Very Large Data Bases*, pages 113–122. C-18, C-19
- Dean, J. and Ghemawat, S. (2004). MapReduce: Simplified data processing on large clusters. In *Proc. 6th USENIX Symp. on Operating System Design and Implementation*, pages 137–149. 497
- Dean, J. and Ghemawat, S. (2010). MapReduce: a flexible data processing tool. *Commun. ACM*, 53(1):72–77. 497
- DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., Sivasubramanian, S., Vosshall, P., and Vogels, W. (2007). Dynamo: Amazon’s highly available key-value store. In *Proc. 21st ACM Symp. on Operating System Principles*, pages 205–220. 535
- Demers, A. J., Greene, D. H., Hauser, C., Irish, W., Larson, J., Shenker, S., Sturgis, H. E., Swinehart, D. C., and Terry, D. B. (1987). Epidemic algorithms for replicated database maintenance. In *Proc. ACM SIGACT-SIGOPS 6th Symp. on the Principles of Distributed Computing*, pages 1–12. 433
- Denning, P. J. (1968). The working set model for program behavior. *Commun. ACM*, 11(5):323–333. C-40
- Denning, P. J. (1980). Working sets: Past and present. *IEEE Trans. Softw. Eng.*, SE-6(1):64–84. C-40

- Deshpande, A. and Gupta, A. (2018). *Principles of Graph Data Management and Analytics*. ACM Books. Forthcoming. 499
- Devine, R. (1993). Design and implementation of DDH: A distributed dynamic hashing algorithm. In *Proc. 4th Int. Conf. on Foundations of Data Organization and Algorithms*, pages 101–114. 433
- Dewitt, D. and Stonebraker, M. (2009). MapReduce: A major step backwards. https://homes.cs.washington.edu/~billhowe/mapreduce_a_major_step_backwards.html. 497
- DeWitt, D., Naughton, J., Schneider, D., and Seshadri, S. (1992). Practical skew handling in parallel joins. In *Proc. 22th Int. Conf. on Very Large Data Bases*, pages 27–40. 381, 382
- DeWitt, D. J. and Gerber, R. (1985). Multi processor hash-based join algorithms. In *Proc. 11th Int. Conf. on Very Large Data Bases*, pages 151–164. 381
- DeWitt, D. J. and Gray, J. (1992). Parallel database systems: The future of high performance database systems. *Commun. ACM*, 35(6):85–98. 85, 381
- DeWitt, D. J., Katz, R., Olken, F., Shapiro, L., Stonebraker, M., and Wood, D. (1984). Implementation techniques for main memory database systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1–8. 381
- DeWitt, D. J., Gerber, R. H., Graek, G., Heytens, M. L., Kumar, K. B., and Muralikrishna, M. (1986). Gamma: A high performance dataflow database machine. In *Proc. 12th Int. Conf. on Very Large Data Bases*, pages 228–237. 381
- DeWitt, D. J., Paulson, E., Robinson, E., Naughton, J., Royalty, J., Shankar, S., and Krioukov, A. (2008). Clustera: an integrated computation and data management system. *Proc. VLDB Endowment*, 1:28–41. 497
- DeWitt, D. J., Halverson, A., Nehme, R. V., Shankar, S., Aguilar-Saborit, J., Avanes, A., Flasz, M., and Gramling, J. (2013). Split query processing in Polybase. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1255–1266. 536
- Dhamankar, R., Lee, Y., Doan, A., Halevy, A. Y., and Domingos, P. (2004). iMAP: Discovering complex mappings between database schemas. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 383–394. 327
- Ding, L. and Rundensteiner, E. (2004). Evaluating window joins over punctuated streams. In *Proc. 13th ACM Int. Conf. on Information and Knowledge Management*, pages 98–107. 498
- Ding, L., Mehta, N., Rundensteiner, E., and Heineman, G. (2004). Joining punctuated streams. In *Advances in Database Technology, Proc. 9th Int. Conf. on Extending Database Technology*, pages 587–604. 498
- Dinh, T. T. A., Liu, R., Zhang, M., Chen, G., Ooi, B. C., and Wang, J. (2018). Untangling blockchain: A data processing view of blockchain systems. *IEEE Trans. Knowl. and Data Eng.*, 30(7):1366–1385. 434
- Do, H. and Rahm, E. (2002). COMA: a system for flexible combination of schema matching approaches. In *Proc. 28th Int. Conf. on Very Large Data Bases*, pages 610–621. 327
- Doan, A. and Halevy, A. Y. (2005). Semantic integration research in the database community: A brief survey. *AI Magazine*, 26(1):83–94. 327
- Doan, A., Domingos, P., and Halevy, A. Y. (2001). Reconciling schemas of disparate data sources: A machine-learning approach. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 509–520. 327
- Doan, A., Domingos, P., and Halevy, A. (2003a). Learning to match the schemas of data sources: A multistrategy approach. *Machine Learning*, 50(3):279–301. 327

- Doan, A., Madhavan, J., Dhamankar, R., Domingos, P., and Halevy, A. (2003b). Learning to match ontologies on the semantic web. *VLDB J.*, 12(4):303–319. 433
- Doan, A., Halevy, A., and Ives, Z. (2012). *Principles of Data Integration*. Morgan Kaufmann. 326
- Dogac, A., Kalinichenko, L., Özsu, M. T., and Sheth, A., editors. (1998). *Advances in Workflow Systems and Interoperability*. Springer. 238, C-18
- Dong, X. L. and Naumann, F. (2009). Data fusion: resolving data conflicts for integration. *Proc. VLDB Endowment*, 2(2):1654–1655. 594
- Dong, X. L. and Srivastava, D. (2015). *Big Data Integration*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers. 500, 594
- Dong, X. L., Berti-Equille, L., and Srivastava, D. (2009a). Truth discovery and copying detection in a dynamic world. *Proc. VLDB Endowment*, 2(1):562–573. 594
- Dong, X. L., Berti-Equille, L., and Srivastava, D. (2009b). Integrating conflicting data: the role of source dependence. *Proc. VLDB Endowment*, 2(1):550–561. 594
- Dowdy, L. W. and Foster, D. V. (1982). Comparative models of the file assignment problem. *ACM Comput. Surv.*, 14(2):287–313. 31, 85
- Du, W., Krishnamurthy, R., and Shan, M. (1992). Query optimization in a heterogeneous DBMS. In *Proc. 18th Int. Conf. on Very Large Data Bases*, pages 277–291. 328
- Du, W., Shan, M., and Dayal, U. (1995). Reducing multidatabase query response time by tree balancing. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 293–303. 177, 328
- Duggan, J., Elmore, A. J., Stonebraker, M., Balazinska, M., Howe, B., Kepner, J., Madden, S., Maier, D., Mattson, T., and Zdonik, S. B. (2015). The BigDAWG polystore system. *ACM SIGMOD Rec.*, 44(2):11–16. 536
- Duschka, O. M. and Genesereth, M. R. (1997). Answering recursive queries using views. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 109–116. 326, 328
- Eager, D. L. and Sevcik, K. C. (1983). Achieving robustness in distributed database systems. *ACM Trans. Database Syst.*, 8(3):354–381. 273
- Edwards, J., McCurley, K., and Tomlin, J. (2001). An adaptive model for optimizing performance of an incremental web crawler. In *Proc. 10th Int. World Wide Web Conf.* 593
- Effelsberg, W. and Härder, T. (1984). Principles of database buffer management. *ACM Trans. Database Syst.*, 9(4):560–595. C-40
- El Abbadi, A., Skeen, D., and Cristian, F. (1985). An efficient, fault-tolerant protocol for replicated data management. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 215–229. 273
- Elbushra, M. M. and Lindström, J. (2015). Causal consistent databases. *Open Journal of Databases*, 2(1):17–35. 535
- Elmagarmid, A., Rusinkiewicz, M., and Sheth, A., editors. (1999). *Management of Heterogeneous and Autonomous Database Systems*. Morgan Kaufmann. 326
- Elmagarmid, A. K. (1986). A survey of distributed deadlock detection algorithms. *ACM SIGMOD Rec.*, 15(3):37–45. 31, 238
- Elmagarmid, A. K., editor. (1992). *Transaction Models for Advanced Database Applications*. Morgan Kaufmann. 238

- Elmagarmid, A. K., Soundararajan, N., and Liu, M. T. (1988). A distributed deadlock detection and resolution algorithm and its correctness proof. *IEEE Trans. Softw. Eng.*, 14(10):1443–1452. 238
- Elmagarmid, A. K., Leu, Y., Litwin, W., and Rusinkiewicz, M. (1990). A multidatabase transaction model for InterBase. In *Proc. 16th Int. Conf. on Very Large Data Bases*, pages 507–518. C-18
- Elmasri, R., Larson, J., and Navathe, S. B. (1987). Integration algorithms for database and logical database design. Technical report, Honeywell Corporate Research Center, Golden Valley, Minn. 327
- Elmore, A. J., Arora, V., Taft, R., Pavlo, A., Agrawal, D., and El Abbadi, A. (2015). Squall: Fine-grained live reconfiguration for partitioned main memory databases. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 299–313. 86, 239
- Elseidy, M., Elguindy, A., Vitorovic, A., and Koch, C. (2014). Scalable and adaptive online joins. *Proc. VLDB Endowment*, 7(6):441–452. 499
- Embley, D. W., Jackman, D., and Xu, L. (2001). Multifaceted exploitation of metadata for attribute match discovery in information integration. In *Proc. Workshop on Information Integration on the Web*, pages 110–117. 327
- Embley, D. W., Jackman, D., and Xu, L. (2002). Attribute match discovery in information integration: exploiting multiple facets of metadata. *Journal of the Brazilian Computing Society*, 8(2):32–43. 327
- Epstein, R., Stonebraker, M., and Wong, E. (1978). Query processing in a distributed relational database system. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 169–180. 177, B-21
- Eswaran, K. P. (1974). Placement of records in a file and file allocation in a computer network. In *Information Processing '74*, pages 304–307. 85
- Eswaran, K. P., Gray, J. N., Lorie, R. A., and Traiger, I. L. (1976). The notions of consistency and predicate locks in a database system. *Commun. ACM*, 19(11):624–633. C-31
- Etzion, O. and Niblett, P. (2010). *Event Processing in Action*. Manning. 498
- Evrendilek, C., Dogac, A., Nural, S., and Ozcan, F. (1997). Multidatabase query optimization. *Distrib. Parall. Databases*, 5(1):77–114. 177, 328
- Eyal, I., Gencer, A. E., Siner, E. G., and van Renesse, R. (2016). Bitcoin-NG: A scalable blockchain protocol. In *Proc. 13th USENIX Symp. on Networked Systems Design & Implementation*, pages 45–59. 434
- Fagin, R. (1977). Multivalued dependencies and a new normal form for relational databases. *ACM Trans. Database Syst.*, 2(3):262–278. A-4
- Fagin, R. (1979). Normal forms and relational database operators. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 153–160. A-4
- Fagin, R. (2002). Combining fuzzy information: an overview. *ACM SIGMOD Rec.*, 31(2):109–118. 327
- Fagin, R., Lotem, J., and Naor, M. (2003). Optimal aggregation algorithms for middleware. *Journal of Computer and System Sciences*, 66(4):614–656. 434
- Fagin, R., Kolaitis, P. G., Miller, R. J., and Popa, L. (2005). Data exchange: semantics and query answering. *Theor. Comp. Sci.*, 336(1):89–124. 325
- Faleiro, J. M. and Abadi, D. J. (2015). Rethinking serializable multiversion concurrency control. *Proc. VLDB Endowment*, 8(11):1190–1201. 237

- Faloutsos, C. and Christodoulakis, S. (1984). Signature files: an access method for documents and its analytical performance evaluation. *ACM Trans. Information Syst.*, 2(4):267–288. 593
- Farid, M. H., Roatis, A., Ilyas, I. F., Hoffmann, H., and Chu, X. (2016). CLAMS: bringing quality to data lakes. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 2089–2092. 594
- Farrag, A. A. and Özsu, M. T. (1989). Using semantic knowledge of transactions to increase concurrency. *ACM Trans. Database Syst.*, 14(4):503–525. 238
- Fekete, A., Lynch, N., Merritt, M., and Wehl, W. (1987a). Nested transactions and read/write locking. Technical Memo MIT/LCS/TM–324, Massachusetts Institute of Technology, Cambridge, Mass. 238
- Fekete, A., Lynch, N., Merritt, M., and Wehl, W. (1987b). Nested transactions, conflict-based locking, and dynamic atomicity. Technical Memo MIT/LCS/TM–340, Massachusetts Institute of Technology, Cambridge, Mass. 238
- Fekete, A., Lynch, N., Merritt, M., and Wehl, W. (1989). Commutativity-based locking for nested transactions. Technical Memo MIT/LCS/TM–370b, Massachusetts Institute of Technology, Cambridge, Mass. 238
- Fernandez, M., Florescu, D., and Levy, A. (1997). A query language for a web-site management system. *ACM SIGMOD Rec.*, 26(3):4–11. 593
- Fernandez, R. C., Migliavacca, M., Kalyvianaki, E., and Pietzuch, P. (2013). Integrating scale out and fault tolerance in stream processing using operator state management. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 725–736. 499
- Fernández-Moctezuma, R., Tufte, K., and Li, J. (2009). Inter-operator feedback in data stream management systems via punctuation. In *Proc. 4th Biennial Conf. on Innovative Data Systems Research*. 498
- Ferraiolo, D. and Kuhn, R. (1992). Role-based access control. In *Proc. National Computer Conf.*, pages 554–563. 124
- Fisher, M. K. and Hochbaum, D. S. (1980). Database location in computer networks. *J. ACM*, 27(4):718–735. 85
- Fisher, P. S., Hollist, P., and Slonim, J. (1980). A design methodology for distributed data bases. In *Digest of Papers – COMPCON*, pages 199–202. 84
- Florentin, J. J. (1974). Consistency auditing of databases. *Comp. J.*, 17(1):52–58. 124
- Florescu, D., Levy, A., and Mendelzon, A. (1998). Database techniques for the World-Wide Web: a survey. *ACM SIGMOD Rec.*, 27(3):59–74. 592, 593
- Friedman, M., Levy, A. Y., and Millstein, T. D. (1999). Navigational plans for data integration. In *Proc. 16th National Conf. on Artificial Intelligence and 11th Innovative Applications of Artificial Intelligence Conf.*, pages 67–73. 326
- Fu, Y., Ong, K. W., Papakonstantinou, Y., and Zamora, E. (2014). FORWARD: data-centric UIs using declarative templates that efficiently wrap third-party JavaScript components. *Proc. VLDB Endowment*, 7(13):1649–1652. 536
- Furtado, C., Lima, A. A. B., Pacitti, E., Valduriez, P., and Mattoso, M. (2008). Adaptive hybrid partitioning for OLAP query processing in a database cluster. *Int. Journal of High Performance Computing and Networking*, 5(4):251–262. 382
- Fushimi, S., Kitsuregawa, M., and Tanaka, H. (1986). An overview of the system software of a parallel relational database machine GRACE. In *Proc. 12th Int. Conf. on Very Large Data Bases*, pages 209–219. 381

- Gadepally, V., Chen, P., Duggan, J., Elmore, A. J., Haynes, B., Kepner, J., Madden, S., Mattson, T., and Stonebraker, M. (2016). The BigDAWG polystore system and architecture. In *Proc. IEEE High Performance Extreme Computing Conf.*, pages 1–6. 536
- Galhardas, H., Florescu, D., Shasha, D., Simon, E., and Saita, C.-A. (2001). Declarative data cleaning: Language, model, and algorithms. In *Proc. 27th Int. Conf. on Very Large Data Bases*, pages 371–380. 327
- Gallaire, H., Minker, J., and Nicolas, J.-M. (1984). Logic and databases: A deductive approach. *ACM Comput. Surv.*, 16(2):153–186. A-7
- Gançarski, S., Naacke, H., Pacitti, E., and Valduriez, P. (2007). The leganet system: Freshness-aware transaction routing in a database cluster. *Inf. Syst.*, 32(7):320–343. 382
- Ganesan, P., Yang, B., and Garcia-Molina, H. (2004). One torus to rule them all: Multidimensional queries in P2P systems. In *Proc. 7th Int. Workshop on the World Wide Web and Databases*, pages 19–24. 433
- Gankidi, V. R., Teletia, N., Patel, J. M., Halverson, A., and DeWitt, D. J. (2014). Indexing HDFS data in PDW: splitting the data from the index. *Proc. VLDB Endowment*, 7(13):1520–1528. 536
- Garcia-Molina, H. (1982). Elections in distributed computing systems. *IEEE Trans. Comput.*, C-31(1):48–59. 237
- Garcia-Molina, H. (1983). Using semantic knowledge for transaction processing in a distributed database. *ACM Trans. Database Syst.*, 8(2):186–213. 238, C-17
- Garcia-Molina, H. and Salem, K. (1987). Sagas. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 249–259. 238, C-15, C-17
- Garcia-Molina, H. and Wiederhold, G. (1982). Read-only transactions in a distributed database. *ACM Trans. Database Syst.*, 7(2):209–234. 238
- Garcia-Molina, H., Gawlick, D., Klein, J., Kleissner, K., and Salem, K. (1990). Coordinating multi-transaction activities. Technical Report CS-TR-247-90, Department of Computer Science, Princeton University. 238, C-17
- Garcia-Molina, H., Papakonstantinou, Y., Quass, D., Rajaraman, A., Sagiv, Y., Ullman, J. D., Vassalos, V., and Widom, J. (1997). The TSIMMIS approach to mediation: Data models and languages. *J. Intell. Information Syst.*, 8(2):117–132. 326
- Garofalakis, M. N. and Ioannidis, Y. E. (1996). Multi-dimensional resource scheduling for parallel queries. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 365–376. 381
- Gavish, B. and Pirkul, H. (1986). Computer and database location in distributed computer systems. *IEEE Trans. Comput.*, C-35(7):583–590. 85
- GE. (1976). *MADMAN User Manual*. General Electric Company, Schenectady, N.Y. C-36
- Gedik, B. (2014). Partitioning functions for stateful data parallelism in stream processing. *VLDB J.*, 23:517–539. 499
- Gelenbe, E. and Gardy, D. (1982). The size of projections of relations satisfying a functional dependency. In *Proc. 8th Int. Conf. on Very Data Bases*, pages 325–333. B-21
- Georgakopoulos, D., Hornick, M., and Sheth, A. (1995). An overview of workflow management: From process modeling to workflow automation infrastructure. *Distrib. Parall. Databases*, 3:119–153. 238, C-18
- Ghemawat, S., Gobioff, H., and Leung, S. (2003). The Google file system. In *Proc. 19th ACM Symp. on Operating System Principles*, pages 29–43. 497

- Ghoting, A., Krishnamurthy, R., Pednault, E. P. D., Reinwald, B., Sindhwani, V., Tatkonda, S., Tian, Y., and Vaithyanathan, S. (2011). Systemml: Declarative machine learning on MapReduce. In *Proc. 27th Int. Conf. on Data Engineering*, pages 231–242. 498
- Gifford, D. K. (1979). Weighted voting for replicated data. In *Proc. 7th ACM Symp. on Operating System Principles*, pages 50–159. 273
- Gilbert, S. and Lynch, N. A. (2002). Brewer’s conjecture and the feasibility of consistent, available, partition-tolerant web services. *SIGACT News*, 33(2):51–59. 535
- Glasbergen, B., Abebe, M., Daudjee, K., Foggo, S., and Pacaci. (2018). Apollo: Learning query correlations for predictive caching in geo-distributed systems. In *Proc. 21st Int. Conf. on Extending Database Technology*, pages 253–264. 86
- Golab, L. and Özsu, M. T. (2003). Processing sliding window multi-joins in continuous queries over data streams. In *Proc. 29th Int. Conf. on Very Large Data Bases*, pages 500–511. 498
- Golab, L. and Özsu, M. T. (2010). *Data Stream Systems*. Synthesis Lectures on Data Management. Morgan & Claypool. 498
- Goldman, K. J. (1987). Data replication in nested transaction systems. Technical Report MIT/LCS/TR-390, Massachusetts Institute of Technology, Cambridge, Mass. 238
- Goldman, R. and Widom, J. (1997). Dataguides: Enabling query formulation and optimization in semistructured databases. In *Proc. 23th Int. Conf. on Very Large Data Bases*, pages 436–445. 593
- Gonzalez, J. E., Low, Y., Gu, H., Bickson, D., and Guestrin, C. (2012). PowerGraph: Distributed graph-parallel computation on natural graphs. In *Proc. 10th USENIX Symp. on Operating System Design and Implementation*, pages 17–30. 500
- Gonzalez, J. E., Xin, R. S., Dave, A., Crankshaw, D., Franklin, M. J., and Stoica, I. (2014). GraphX: graph processing in a distributed dataflow framework graph processing in a distributed dataflow framework. In *Proc. 11th USENIX Symp. on Operating System Design and Implementation*, pages 599–613. 498, 499
- Goodman, J. R. and Woest, P. J. (1988). The Wisconsin multicube: A new large-scale cache-coherent multiprocessor. Technical Report TR766, University of Wisconsin-Madison. 381
- Gounaris, A., Paton, N. W., Fernandes, A. A. A., and Sakellariou, R. (2002). Adaptive query processing: A survey. In *Proc. British National Conf. on Databases*, pages 11–25. 177, 328
- Graefe, G. (1990). Encapsulation of parallelism in the Volcano query processing systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 102–111. 381
- Graefe, G. (1993). Query evaluation techniques for large databases. *ACM Comput. Surv.*, 25(2):73–170. 176, 381
- Graefe, G. (1994). Volcano - an extensible and parallel query evaluation system. *IEEE Trans. Knowl. and Data Eng.*, 6(1):120–135. 176, B-33
- Graefe, G. and Kuno, H. (2010a). Self-selecting, self-tuning, incrementally optimized indexes. In *Proc. 13th Int. Conf. on Extending Database Technology*, pages 371–381. 86
- Graefe, G. and Kuno, H. (2010b). Adaptive indexing for relational keys. In *Proc. Workshops of 26th Int. Conf. on Data Engineering*, pages 69–74. 86

- Graefe, G., Idreos, S., Kuno, H., and Manegold, S. (2010). Benchmarking adaptive indexing. In *Proc. TPC Technology Conference on Performance Evaluation, Measurement and Characterization of Complex Systems*, pages 169–184. 86
- Graefe, G., Halim, F., Idreos, S., Kuno, H., and Manegold, S. (2012). Concurrency control for adaptive indexing. *Proc. VLDB Endowment*, 5(7):656–667. 86
- Graefe, G., Halim, F., Idreos, S., Kuno, H. A., Manegold, S., and Seeger, B. (2014). Transactional support for adaptive indexing. *VLDB J.*, 23(2):303–328. 86
- Grapa, E. and Belford, G. G. (1977). Some theorems to aid in solving the file allocation problem. *Commun. ACM*, 20(11):878–882. 85
- Gravano, L., Garcia-Molina, H., and Tomasic, A. (1999). Gloss: Text-source discovery over the internet. *ACM Trans. Database Syst.*, 24(2):229–264. 592
- Gray, J. (1979). Notes on database operating systems. In Bayer, R., Graham, R., and Seegmüller, G., editors, *Operating Systems – An Advanced Course*, pages 393–481. Springer, New York. 31, 237, 239, C-43, C-44, C-48, C-49
- Gray, J. (1981). The transaction concept: Virtues and limitations. In *Proc. 7th Int. Conf. on Very Data Bases*, pages 144–154. C-2
- Gray, J. (1987). Why do computers stop and what can be done about it. In *CIPS (Canadian Information Processing Society) Edmonton '87 Conf. Tutorial Notes*. C-14, C-38
- Gray, J. and Lamport, L. (2006). Consensus on transaction commit. *ACM Trans. Database Syst.*, 31(1):133–160. 239
- Gray, J. and Reuter, A. (1993). *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann. 237
- Gray, J., Helland, P., O’Neil, P. E., and Shasha, D. (1996). The dangers of replication and a solution. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 173–182. 272
- Gray, J. N., Lorie, R. A., Putzolu, G. R., and Traiger, I. L. (1976). Granularity of locks and degrees of consistency in a shared data base. In Nijssen, G. M., editor, *Modelling in Data Base Management Systems*, pages 365–394. C-10
- Gray, J. N., McJones, P., Blasgen, M., Lindsay, B., Lorie, R., Price, T., Putzolu, F., and Traiger, I. (1981). The recovery manager of the System R database manager. *ACM Comput. Surv.*, 13(2):223–242. 239, C-38, C-44, C-49
- Grefen, P. and Widom, J. (1997). Protocols for integrity constraint checking in federated databases. *Distrib. Parall. Databases*, 5(4):327–355. 124
- Griffiths, P. P. and Wade, B. W. (1976). An authorization mechanism for a relational database system. *ACM Trans. Database Syst.*, 1(3):242–255. 124
- Grossman, R. L. and Gu, Y. (2009). On the varieties of clouds for data intensive computing. *Q. Bull. IEEE TC on Data Eng.*, 32(1):44–50. 27
- Guha, S. and McGregor, A. (2006). Approximate quantiles and the order of the stream. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 273–279. 498
- Gulisano, V., Jiménez-Peris, R., Patino-Martinez, M., and Valduriez, P. (2010). StreamCloud: A large scale data streaming system. In *Proc. 30th IEEE Int. Conf. on Distributed Computing Systems*. 498
- Gulisano, V., Jiménez-Peris, R., Patino-Martinez, M., and Valduriez, P. (2012). StreamCloud: An elastic and scalable data streaming system. *IEEE Trans. Parall. Dist. Sys.*, 23(12):2351–2365. 498
- Gulli, A. and Signorini, A. (2005). The indexable web is more than 11.5 billion pages. In *Proc. 14th Int. World Wide Web Conf.*, pages 902–903. 592

- Gummadi, P. K., Gummadi, R., Gribble, S. D., Ratnasamy, S., Shenker, S., and Stoica, I. (2003). The impact of DHT routing geometry on resilience and proximity. In *Proc. Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 381–394. 433
- Güntzer, U., Kießling, W., and Balke, W.-T. (2000). Optimizing multi-feature queries for image databases. In *Proc. 26th Int. Conf. on Very Large Data Bases*, pages 419–428. 434
- Gupta, A. and Mumick, I. S., editors. (1999). *Materialized Views: Techniques, Implementations, and Applications*. M.I.T. Press. 123
- Gupta, A., Mumick, I. S., and Subrahmanian, V. S. (1993). Maintaining views incrementally. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 157–166. 123
- Gupta, A., Jagadish, H., and Mumick, I. S. (1996). Data integration using self-maintainable views. In *Advances in Database Technology, Proc. 5th Int. Conf. on Extending Database Technology*, pages 140–144. 123
- Gupta, A., Agrawal, D., and El Abbadi, A. (2003). Approximate range selection queries in peer-to-peer systems. In *Proc. 1st Biennial Conf. on Innovative Data Systems Research*, pages 141–151. 434
- Haas, L. (2007). Beauty and the beast: The theory and practice of information integration. In *Proc. 11th Int. Conf. on Database Theory*, pages 28–43. 326
- Haas, L., Kossmann, D., Wimmers, E., and Yang, J. (1997a). Optimizing queries across diverse data sources. In *Proc. 23th Int. Conf. on Very Large Data Bases*, pages 276–285. 177, 328
- Haas, L. M., Kossmann, D., Wimmers, E. L., and Yang, J. (1997b). Optimizing queries across diverse data sources. In *Proc. 23th Int. Conf. on Very Large Data Bases*, pages 276–285. 326
- Haas, P. and Hellerstein, J. (1999a). Ripple joins for online aggregation. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 287–298. 498
- Haas, P. J. and Hellerstein, J. M. (1999b). Ripple joins for online aggregation. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 287–298. 177, 328, 381
- Hacıgümüş, H., Sankaranarayanan, J., Tatemura, J., LeFevre, J., and Polyzotis, N. (2013). Odyssey: A multi-store system for evolutionary analytics. *Proc. VLDB Endowment*, 6(11):1180–1181. 536
- Haderle, C. M. D., Lindsay, B., Pirahesh, H., and Schwarz, P. (1992). Aries: A transaction recovery method supporting fine-granularity locking and partial rollbacks using write-ahead logging. *ACM Trans. Database Syst.*, 17(1):94–162. 238, C-43
- Hadzilacos, V. (1988). A theory of reliability in database systems. *J. ACM*, 35(1): 121–145. 237, 238
- Halevy, A., Rajaraman, A., and Ordille, J. (2006). Data integration: the teenage years. In *Proc. 32nd Int. Conf. on Very Large Data Bases*, pages 9–16. 326
- Halevy, A. Y. (2001). Answering queries using views: A survey. *VLDB J.*, 10(4): 270–294. 326, 328
- Halevy, A. Y., Etzioni, O., Doan, A., Ives, Z. G., Madhavan, J., McDowell, L., and Tatarinov, I. (2003). Crossing the structure chasm. In *Proc. 1st Biennial Conf. on Innovative Data Systems Research*. 329
- Halici, U. and Dogac, A. (1989). Concurrency control in distributed databases through time intervals and short-term locks. *IEEE Trans. Softw. Eng.*, 15(8):994–995. 238

- Halim, F., Idreos, S., Karras, P., and Yap, R. H. C. (2012). Stochastic database cracking: Towards robust adaptive indexing in main-memory column-stores. *Proc. VLDB Endowment*, 5(6):502–513. 86
- Hammad, M., Aref, W., and Elmagarmid, A. (2003a). Stream window join: Tracking moving objects in sensor-network databases. In *Proc. 15th Int. Conf. on Scientific and Statistical Database Management*, pages 75–84. 498
- Hammad, M., Aref, W., Franklin, M., Mokbel, M., and Elmagarmid, A. (2003b). Efficient execution of sliding window queries over data streams. Technical Report CSD TR 03-035, Purdue University. 498
- Hammad, M., Mokbel, M., Ali, M., Aref, W., Catlin, A., Elmagarmid, A., Eltabakh, M., Elfeky, M., Ghanem, T., Gwadera, R., Ilyas, I., Marzouk, M., and Xiong, X. (2004). Nile: a query processing engine for data streams. In *Proc. 20th Int. Conf. on Data Engineering*, page 851. 498
- Hammad, M., Aref, W., and Elmagarmid, A. (2005). Optimizing in-order execution of continuous queries over streamed sensor data. In *Proc. 17th Int. Conf. on Scientific and Statistical Database Management*, pages 143–146. 498
- Hammer, M. and Niamir, B. (1979). A heuristic approach to attribute partitioning. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 93–101. 85
- Hammer, M. and Shipman, D. W. (1980). Reliability mechanisms for SDD-1: A system for distributed databases. *ACM Trans. Database Syst.*, 5(4):431–466. 237, 239
- Han, M. (2015). On improving distributed Pregel-like graph processing systems. Master’s thesis, University of Waterloo, David R. Cheriton School of Computer Science. URL <http://hdl.handle.net/10012/9484>. 500
- Han, M. and Daudjee, K. (2015). Giraph unchained: Barrierless asynchronous parallel execution in Pregel-like graph processing systems. *Proc. VLDB Endowment*, 8(9): 950–961. 500
- Härder, T. and Reuter, A. (1983). Principles of transaction-oriented database recovery. *ACM Comput. Surv.*, 15(4):287–317. 31, 239, C-38, C-45, C-46, C-47, C-48
- Hartig, O. (2012). SPARQL for a web of linked data: Semantics and computability. In *Proc. 9th Extended Semantic Web Conf.*, pages 8–23. 575
- Hartig, O. (2013a). An overview on execution strategies for linked data queries. *Datenbank-Spektrum*, 13(2):89–99. 594
- Hartig, O. (2013b). SQUIN: a traversal based query execution system for the web of linked data. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1081–1084. 594
- Harvey, N. J. A., Jones, M. B., Saroiu, S., Theimer, M., and Wolman, A. (2003). SkipNet: A scalable overlay network with practical locality properties. In *Proc. 4th USENIX Symp. on Internet Tech. and Systems*. 433
- He, B., Chang, K. C.-C., and Han, J. (2004). Mining complex matchings across web query interfaces. In *Proc. ACM SIGMOD Workshop on Research Issues in Data Mining and Knowledge Discovery*, pages 3–10. 327
- He, Q. and Ling, T. W. (2006). An ontology-based approach to the integration of entity-relationship schemas. *Data & Knowl. Eng.*, 58(3):299–326. 327
- Hedley, Y. L., Younas, M., James, A., and Sanderson, M. (2004a). A two-phase sampling technique for information extraction from hidden web databases. In *WIDM04*, pages 1–8. 592

- Hedley, Y.-L., Younas, M., James, A. E., and Sanderson, M. (2004b). Query-related data extraction of hidden web documents. In *Proc. 27th Annual Int. ACM SIGIR Conf. on Research and Development in Information Retrieval*, pages 558–559. 592
- Heinze, T., Pappalardo, V., Jerzak, Z., and Fetzer, C. (2014). Auto-scaling techniques for elastic data stream processing. In *Proc. 8th Int. Conf. Distributed Event-Based Systems*, pages 318–321. 499
- Heinze, T., Roediger, L., Meister, A., Ji, Y., Jerzak, Z., and Fetzer, C. (2015). Online parameter optimization for elastic data stream processing. In *Proc. 6th ACM Symp. on Cloud Computing*, pages 276–287. 499
- Helal, A. A., Heddaya, A. A., and Bhargava, B. B. (1997). *Replication Techniques in Distributed Systems*. Kluwer Academic Publishers. 272
- Hellerstein, J. M., Haas, P., and Wang, H. (1997). Online aggregation. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 171–182. 498
- Hellerstein, J. M., Franklin, M. J., Chandrasekaran, S., Deshpande, A., Hildrum, K., Madden, S., Raman, V., and Shah, M. A. (2000). Adaptive query processing: Technology in evolution. *Q. Bull. IEEE TC on Data Eng.*, 23(2):7–18. 177, 328
- Herlihy, M. (1987). Concurrency versus availability: Atomicity mechanisms for replicated data. *ACM Trans. Comp. Syst.*, 5(3):249–274. 273
- Herman, D. and Verjus, J. P. (1979). An algorithm for maintaining the consistency of multiple copies. In *Proc. 1st IEEE Int. Conf. on Distributed Computing Systems*, pages 625–631. 237
- Hersh, W. (2001). Managing gigabytes - compressing and indexing documents and images (second edition). *Inf. Retr.*, 4(1):79–80. 593
- Hevner, A. R. and Schneider, G. M. (1980). An integrated design system for distributed database networks. In *Digest of Papers - COMPCON*, pages 459–465. 84
- Hevner, A. R. and Yao, S. B. (1979). Query processing in distributed database systems. *IEEE Trans. Softw. Eng.*, 5(3):177–182. B-22
- Hirate, Y., Kato, S., and Yamana, H. (2006). Web structure in 2005. In *Proc. 4th Int. Workshop on Algorithms and Models for the Web-Graph*, pages 36 – 46. 592
- Hoffer, H. A. and Severance, D. G. (1975). The use of cluster analysis in physical data base design. In *Proc. 1st Int. Conf. on Very Data Bases*, pages 69–86. 85
- Hoffer, J. A. (1975). *A Clustering Approach to the Generation of Subfiles for the Design of a Computer Data Base*. PhD thesis, Department of Operations Research, Cornell University, Ithaca, N.Y. 85
- Hoffman, J. L. (1977). *Model Methods for Computer Security and Privacy*. Prentice-Hall. 124
- Holze, M. and Ritter, N. (2008). Autonomic databases: Detection of workload shifts with n-gram-models. In *Proc. 12th East European Conf. Advances in Databases and Information Systems*, pages 127–142. 86
- Hong, W. (1992). Exploiting inter-operation parallelism in XPRS. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 19–28. 381
- Hong, W. and Stonebraker, M. (1993). Optimization of parallel query execution plans in XPRS. *Distrib. Parall. Databases*, 1(1):9–32. 381
- Hoque, I. and Gupta, I. (2013). LFGraph: simple and fast distributed graph analytics. In *Proc. 1st ACM SIGOPS Conf. on Timely Results in Operating Syst.*, pages 9:1–9:17. 500
- Hortonworks. (2014). White paper: A modern data architecture with Apache Hadoop: the journey to the data lake. Technical report, Hortonworks. Last accessed August 2018. 500

- Hsiao, H. I. and DeWitt, D. (1991). A performance study of three high-availability data replication strategies. In *Proc. Int. Conf. on Parallel and Distributed Information Systems*, pages 18–28. 381
- Hsu, M., editor. (1993). *IEEE Quart. Bull. Data Eng., Special Issue on Workflow and Extended Transaction Systems*, volume 16. IEEE Computer Society. C-18
- Huang, Z. and He, Y. (2018). Auto-detect: Data-driven error detection in tables. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1377–1392. 587, 594
- Huebsch, R., Hellerstein, J., Lanham, N., Loo, B. T., Shenker, S., and Stoica, I. (2003). Querying the internet with pier. In *Proc. 29th Int. Conf. on Very Large Data Bases*, pages 321–332. 434
- Hull, R. (1997). Managing semantic heterogeneity in databases: A theoretical perspective. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 51–61. 327
- Hwang, J., Balazinska, M., Rasin, A., Cetintemel, U., Stonebraker, M., and Zdonik, S. (2005). High-availability algorithms for distributed stream processing. In *Proc. 21st Int. Conf. on Data Engineering*, pages 779–790. 499
- Idreos, S. (2010). *Database Cracking: Towards Auto-tuning Database Kernels*. PhD thesis, University of Amsterdam. 86
- Idreos, S., Kersten, M. L., and Manegold, S. (2007a). Updating a cracked database. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 413–424. 86
- Idreos, S., Kersten, M. L., and Manegold, S. (2007b). Database cracking. In *Proc. 3rd Biennial Conf. on Innovative Data Systems Research*, pages 68–78. 86
- Idreos, S., Kersten, M. L., and Manegold, S. (2009). Self-organizing tuple reconstruction in column-stores. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 297–308. 86
- Idreos, S., Alagiannis, I., Johnson, R., and Ailamaki, A. (2011). Here are my data files. here are my queries. where are my results? In *Proc. 5th Biennial Conf. on Innovative Data Systems Research*, pages 57–68. 86
- Ilyas, I. and Chu, X. (2019). *Principles of Data Cleaning*. ACM Books. 327, 594
- Ilyas, I. F. and Chu, X. (2015). Trends in cleaning relational data: Consistency and deduplication. *Foundations and Trends in Databases*, 5(4):281–393. 594
- Ilyas, I. F., Beskales, G., and Soliman, M. A. (2008). A survey of top-k query processing techniques in relational database systems. *ACM Comput. Surv.*, 40(4): 1–58. 434
- Ioannidis, Y. and Wong, E. (1987). Query optimization by simulated annealing. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 9–22. 176, 381, B-2, B-17
- Ipeirotis, P. G. and Gravano, L. (2002). Distributed search over the hidden web: Hierarchical database sampling and selection. In *Proc. 28th Int. Conf. on Very Large Data Bases*, pages 394–405. 592
- Irani, K. B. and Khabbaz, N. G. (1982). A methodology for the design of communication networks and the distribution of data in distributed computer systems. *IEEE Trans. Comput.*, C-31(5):419–434. 85
- Isloor, S. and Marsland, T. (1980). The deadlock problem : An overview. *Computer*, 13(9):58–78. 31, 238
- Ito, J., Narula, N., and Ali, R. (2017). The blockchain will do to the financial system what the internet did to media. Accessible at <https://hbr.org/2017/03/the->

- blockchain-will-do-to-banks-and-law-firms-what-the-internet-did-to-media/. Last accessed February 2019. 434
- Jagadish, H. V., Ooi, B. C., and Vu, Q. H. (2005). BATON: A balanced tree structure for peer-to-peer networks. In *Proc. 31st Int. Conf. on Very Large Data Bases*, pages 661–672. 433, 434
- Jagadish, H. V., Ooi, B. C., Tan, K.-L., Vu, Q. H., and Zhang, R. (2006). Speeding up search in peer-to-peer networks with a multi-way tree structure. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1–12. 433
- Jajodia, S. and Mutchler, D. (1987). Dynamic voting. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 227–238. 273
- Jajodia, S. and Sandhu, R. S. (1991). Towards a multilevel secure relational data model. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 50–59. 124
- Jajodia, S., Atluri, V., Keefe, T. F., McCollum, C. D., and Mulkamala, R. (2001). Multilevel security transaction processing. *J. Computer Security*, 9(3):165–195. 124
- Jhingran, A. D., Mattos, N., and Pirahesh, H. (2002). Information integration: A research agenda. *IBM Systems J.*, 41(4):555–562. 326
- Jimenez-Peris, R. and Patiño Martínez, M. (2011). System and method for highly scalable decentralized and low contention transactional processing. US Patent 9,760,597 B2, EU Patent 2780832. 238, 536
- Jiménez-Peris, R., Patiño-Martínez, M., and Alonso, G. (2002). Non-intrusive, parallel recovery of replicated data. In *Proc. 21st Symp. on Reliable Distributed Systems*, pages 150–159. 382
- Jiménez-Peris, R., Patiño-Martínez, M., Kemme, B., and Alonso, G. (2002). Improving the scalability of fault-tolerant database clusters. In *Proc. 22nd IEEE Int. Conf. on Distributed Computing Systems*, pages 477–484. 273, 382
- Jiménez-Peris, R., Patiño-Martínez, M., Alonso, G., and Kemme, B. (2003). Are quorums an alternative for data replication? *ACM Trans. Database Syst.*, 28(3): 257–294. 273, 382
- Johnson, T., Muthukrishnan, S., Shkapenyuk, V., and Spatscheck, O. (2005). A heartbeat mechanism and its application in Gigascope. In *Proc. 31st Int. Conf. on Very Large Data Bases*, pages 1079–1088. 498
- Johnson, T., Muthukrishnan, S. M., Shkapenyuk, V., and Spatscheck, O. (2008). Query-aware partitioning for monitoring massive network data streams. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1135–1146. 499
- Kaelbling, L. P., Littman, M. L., and Moore, A. P. (1996). Reinforcement learning: A survey. *J. Autom. Reasoning*, 4:237–285. 593
- Kalogeraki, V., Gunopulos, D., and Zeinalipour-Yazti, D. (2002). A local search mechanism for peer-to-peer networks. In *Proc. 11th Int. Conf. on Information and Knowledge Management*, pages 300–307. 433
- Kambayashi, Y., Yoshikawa, M., and Yajima, S. (1982). Query processing for distributed databases using generalized semi-joins. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 151–160. 176
- Kang, J., Naughton, J., and Viglas, S. (2003). Evaluating window joins over unbounded streams. In *Proc. 19th Int. Conf. on Data Engineering*, pages 341–352. 498
- Kaoudi, Z. and Manolescu, I. (2015). RDF in the clouds: A survey. *VLDB J.*, 24: 67–91. 593

- Kara, A., Ngo, H. Q., Nikolic, M., Olteanu, D., and Zhang, H. (2019). Counting triangles under updates in worst-case optimal time. In *Proc. 22nd Int. Conf. on Database Theory*, pages 1:1–1:18. 124
- Karlapalem, K. and Navathe, S. B. (1994). Materialization of redesigned distributed relational databases. Technical Report HKUST-CS94-14, Hong Kong University of Science and Technology, Department of Computer Science. 85
- Karlapalem, K., Navathe, S. B., and Ammar, M. (1996). Optimal redesign policies to support dynamic processing of applications on a distributed relational database system. *Inf. Syst.*, 21(4):353–367. 85
- Karypis, G. and Kumar, V. (1995). Multilevel graph partitioning schemes. In *Proc. 1995 Int. Conf. on Parallel Processing*, pages 113–122. 499
- Kashyap, V. and Sheth, A. P. (1996). Semantic and schematic similarities between database objects: A context-based approach. *VLDB J.*, 5(4):276–304. 327
- Katz, B. and Lin, J. (2002). Annotating the world wide web using natural language. In *Proc. 2nd Workshop on NLP and XML*, pages 1–8. 593
- Kazerouni, L. and Karlapalem, K. (1997). Stepwise redesign of distributed relational databases. Technical Report HKUST-CS97-12, Hong Kong University of Science and Technology, Department of Computer Science. 85
- Keeton, K., Patterson, D., and Hellerstein, J. M. (1998). A case for intelligent disks (idisks). *ACM SIGMOD Rec.*, 27(3):42–52. 380
- Keller, A. M. (1982). Update to relational databases through views involving joins. In *Proc. 2nd Int. Conf. on Databases: Improving Usability and Responsiveness*, pages 363–384. 123
- Kementsietsidis, A., Arenas, M., and Miller, R. J. (2003). Managing data mappings in the hyperion project. In *Proc. 19th Int. Conf. on Data Engineering*, pages 732–734. 433
- Kemme, B. and Alonso, G. (2000a). A new approach to developing and implementing eager database replication protocols. *ACM Trans. Database Syst.*, 25(3):333–379. 273, 382
- Kemme, B. and Alonso, G. (2000b). Don't be lazy, be consistent: Postgres-R, a new way to implement database replication. In *Proc. 26th Int. Conf. on Very Large Data Bases*, pages 134–143. 273, 382
- Kemme, B., Bartoli, A., and Babaoglu, O. (2001). Online reconfiguration in replicated databases based on group communication. In *Proc. Int. Conf. on Dependable Systems and Networks*, pages 117–130. 382
- Kemme, B., Peris, R. J., and Patino-Martinez, M. (2010). *Database Replication*. Morgan & Claypool. 272
- Kemper, A. and Neumann, T. (2011). HyPer: A hybrid OLTP&OLAP main memory database system based on virtual memory snapshots. In *Proc. 27th Int. Conf. on Data Engineering*, pages 195–206. 239
- Kermarrec, A.-M. and van Steen, M. (2007). Gossiping in distributed systems. *Operating Systems Rev.*, 41(5):2–7. 433
- Kermarrec, A.-M., Rowstron, A., Shapiro, M., and Druschel, P. (2001). The icecube approach to the reconciliation of diverging replicas. In *Proc. ACM SIGACT-SIGOPS 20th Symp. on the Principles of Distributed Computing*, pages 210–218. 434
- Khayyat, Z., Awara, K., Alonazi, A., Jamjoom, H., Williams, D., and Kalnis, P. (2013). Mizan: A system for dynamic load balancing in large-scale graph processing. In *Proc. 8th ACM SIGOPS/EuroSys European Conf. on Comp. Syst.*, pages 169–182. 500

- Khoshafian, S. and Valduriez, P. (1987). Sharing persistence and object-orientation: A database perspective. In *Int. Workshop on Database Programming Languages*, pages 181–205. 176, 381
- Kim, W. and Seo, J. (1991). Classifying schematic and data heterogeneity in multidatabase systems. *Computer*, 24(12):12–18. 326
- Kirsch, J. and Amir, Y. (2008). Paxos for system builders: An overview. In *Proc. 2nd Workshop on Large-Scale Distributed Systems and Middleware*, pages 3:1–3:6. 239
- Kitsuregawa, M. and Ogawa, Y. (1990). Bucket spreading parallel hash: A new, robust, parallel hash join method for data skew in the super database computer. In *Proc. 16th Int. Conf. on Very Large Data Bases*, pages 210–221. 381
- Kitsuregawa, M., Tanaka, H., and Moto-Oka, T. (1983). Application of hash to data base machine and its architecture. *New Generation Computing*, 1(1):63–74. 381
- Kiveris, R., Lattanzi, S., Mirrokni, V., Rastogi, V., and Vassilvitskii, S. (2014). Connected components in MapReduce and beyond. In *Proc. 5th ACM Symp. on Cloud Computing*, pages 18:1–18:13. 499
- Kleinberg, J., Kumar, R., Raghavan, P., Rajagopalan, S., and Tomkins, A. (1999). The web as a graph: Measurements, models, and methods. In *Proc. 5th Annual Int. Conf. Computing and Combinatorics*, pages 1–17. 592
- Kleinberg, J. M. (1999). Authoritative sources in a hyperlinked environment. *J. ACM*, 46(5):604–632. 592, 593
- Knapp, E. (1987). Deadlock detection in distributed databases. *ACM Comput. Surv.*, 19(4):303–328. 31, 238
- Knuth, D. E. (1973). *The Art of Computer Programming, Volume III: Sorting and Searching*. Addison-Wesley. 382
- Koch, C. (2001). *Data Integration against Multiple Evolving Autonomous Schemata*. Ph.D. thesis, Technical University of Vienna. 280, 326
- Koch, C. (2010). Incremental query evaluation in a ring of databases. In *Proc. 29th ACM SIGACT-SIGMOD-SIGART Symp. on Principles of Database Systems*, pages 87–98. 124
- Koch, C., Ahmad, Y., Kennedy, O., Nikolic, M., Nötzli, A., Lupei, D., and Shaikhha, A. (2014). DBToaster: higher-order delta processing for dynamic, frequently fresh views. *VLDB J.*, 23(2):253–278. 124
- Kohler, W. H. (1981). A survey of techniques for synchronization and recovery in decentralized computer systems. *ACM Comput. Surv.*, 13(2):149–183. 238
- Kolev, B., Bondiombouy, C., Valduriez, P., Jiménez-Peris, R., Pau, R., and Pereira, J. (2016a). The cloudmssql multistore system. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 2113–2116. 536
- Kolev, B., Valduriez, P., Bondiombouy, C., Jiménez-Peris, R., Pau, R., and Pereira, J. (2016b). CloudMdsQL: querying heterogeneous cloud data stores with a common language. *Distrib. Parall. Databases*, 34(4):463–503. 536
- Kolev, B., Levchenko, O., Pacitti, E., Valduriez, P., Vilaça, R., Gonçalves, R. C., Jiménez-Peris, R., and Kranas, P. (2018). Parallel polyglot query processing on heterogeneous cloud data stores with LeanXscale. In *Proc. 2018 IEEE Int. Conf. on Big Data*, pages 1757–1766. 536
- Koller, D. and Friedman, N. (2009). *Probabilistic Graphical Models: Principles and Techniques*. The MIT Press. 594
- Kollias, J. G. and Hatzopoulos, M. (1981). Criteria to aid in solving the problem of allocating copies of a file in a computer network. *Comp. J.*, 24(1):29–30. 85

- Konopnicki, D. and Shmueli, O. (1995). W3QS: A query system for the World Wide Web. In *Proc. 21th Int. Conf. on Very Large Data Bases*, pages 54–65. 593
- Korth, H., Levy, E., and Silberschatz, A. (1990). Compensating transactions: A new recovery paradigm. In *Proc. 16th Int. Conf. on Very Large Data Bases*, pages 95–106. C-17
- Kossmann, D. (2000). The state of the art in distributed query processing. *ACM Comput. Surv.*, 32(4):422–469. 176, 328
- Krishnamurthy, R., Litwin, W., and Kent, W. (1991). Language features for interoperability of databases with schematic discrepancies. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 40–49. 326
- Kshemkalyani, A. and Singhal, M. (1994). On characterization and correctness of distributed deadlocks. *J. Parall. and Distrib. Comput.*, 22(1):44–59. 238
- Kubiatowicz, J., Bindel, D., Chen, Y., Czerwinski, S., Eaton, P., Geels, D., Gummadi, R., Rhea, S., Weatherspoon, H., Weimer, W., Wells, C., and Zhao, B. (2000). Oceanstore: an architecture for global-scale persistent storage. In *ACM Int. Conf. on Architectural Support for Programming Languages and Operating Systems (ASPLOS)*, pages 190–201. 434
- Kulkarni, S., Bhagat, N., Fu, M., Kedigehalli, V., Kellogg, C., Mittal, S., Patel, J. M., Ramasamy, K., and Taneja, S. (2015). Twitter heron: Stream processing at scale. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 239–250. 499
- Kumar, A. and Segev, A. (1993). Cost and availability tradeoffs in replicated data concurrency control. *ACM Trans. Database Syst.*, 18(1):102–131. 273
- Kumar, R., Raghavan, P., Rajagopalan, S., Sivakumar, D., Tomkins, A., and Upfal, E. (2000). The Web as a graph. In *Proc. 19th ACM SIGACT-SIGMOD-SIGART Symp. on Principles of Database Systems*, pages 1–10. 541, 592
- Kumar, V., editor. (1996). *Performance of Concurrency Control Mechanisms in Centralized Database Systems*. Prentice-Hall. 238
- Kung, H. and Robinson, J. (1981). On optimistic methods for concurrency control. *ACM Trans. Database Syst.*, 6(2):213–226. 237
- Kung, H. T. and Papadimitriou, C. H. (1979). An optimality theory of concurrency control for databases. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 116–125. C-14
- Kwok, C. C. T., Etzioni, O., and Weld, D. S. (2001). Scaling question answering to the web. In *Proc. 10th Int. World Wide Web Conf.*, pages 150–161. 593
- Lacroix, M. and Pirotte, A. (1977). Domain-oriented relational languages. In *Proc. 3rd Int. Conf. on Very Data Bases*, pages 370–378. A-17
- Ladwig, G. and Tran, T. (2010). Linked data query processing strategies. In *Proc. 9th Int. Semantic Web Conf.*, pages 453–469. 594
- Ladwig, G. and Tran, T. (2011). SIHJoin: Querying remote and local linked data. In *Proc. 8th Extended Semantic Web Conf.*, pages 139–153. 594
- Lage, J. P., da Silva, A. S., Golgher, P. B., and Laender, A. H. F. (2002). Collecting hidden web pages for data extraction. In *Proc. 4th Int. Workshop on Web Information and Data Management*, pages 69–75. 592
- Lakshmanan, L. V. S., Sadri, F., and Subramanian, I. N. (1996). A declarative language for querying and restructuring the web. In *Proc. 6th Int. Workshop on Research Issues on Data Eng.*, pages 12–21. 593
- Lam, S. S. and Özsu, M. T. (2002). Querying web data – the WebQA approach. In *Proc. 3rd Int. Conf. on Web Information Systems Eng.*, pages 139–148. 593

- Lamport, L. (1998). The part-time parliament. *ACM Trans. Comp. Syst.*, 16(2): 133–169. 239
- Lamport, L. (2001). Paxos made simple. *ACM SIGACT News*, 32(4):51–58. 239
- Lampson, B. and Sturgis, H. (1976). Crash recovery in distributed data storage system. Technical report, Xerox Palo Alto Research Center, Palo Alto, Calif. 237, C-22
- Landers, T. and Rosenberg, R. L. (1982). An overview of multibase. In Schneider, H.-J., editor, *Distributed Data Bases*, pages 153–184. North-Holland, Amsterdam. 328
- Langville, A. N. and Meyer, C. D. (2006). *Google's PageRank and Beyond*. Princeton University Press. 593
- Lanzelotte, R., Valduriez, P., and Zaït, M. (1993). On the effectiveness of optimization search strategies for parallel execution spaces. In *Proc. 19th Int. Conf. on Very Large Data Bases*, pages 493–504. B-17, B-18
- Lanzelotte, R., Valduriez, P., Zaït, M., and Ziane, M. (1994). Industrial-strength parallel query optimization: issues and lessons. *Inf. Syst.*, 19(4):311–330. 381
- Larriba-Pey, J. L., Martínez-Bazán, N., and Domínguez-Sal, D. (2014). Introduction to graph databases. In Koubarakis, M., Stamou, G., Stoilos, G., Horrocks, I., Kolaitis, P., Lausen, G., and Weikum, G., editors, *Reasoning Web: Reasoning on the Web in the Big Data Era*, pages 171–194. Springer. 499
- Larson, P.-Å., Blanas, S., Diaconu, C., Freedman, C., Patel, J. M., and Zwillig, M. (2011). High-performance concurrency control mechanisms for main-memory databases. *Proc. VLDB Endowment*, 5(4):298–309. 239
- Law, Y.-N., Wang, H., and Zaniolo, C. (2004). Query languages and data models for database sequences and data streams. In *Proc. 30th Int. Conf. on Very Large Data Bases*, pages 492–503. 498
- Lawrence, S. and Giles, C. L. (1998). Searching the world wide web. *Science*, 280(5360):98–100. 592
- Lawrence, S. and Giles, C. L. (1999). Accessibility of information on the web. *Nature*, 400(6740):107–9. 592
- Lee, K.-H., Lee, Y.-J., Choi, H., Chung, Y. D., and Moon, B. (2012). Parallel data processing with mapreduce: A survey. *ACM SIGMOD Rec.*, 40(4):11–20. 497
- LeFevre, J., Sankaranarayanan, J., Hacigumus, H., Tatemura, J., Polyzotis, N., and Carey, M. J. (2014). MISO: Souping up big data query processing with a multistore system. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1591–1602. 536
- Leis, V., Boncz, P. A., Kemper, A., and Neumann, T. (2014). Morsel-driven parallelism: a NUMA-aware query evaluation framework for the many-core age. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 743–754. 381
- Lenoski, D., Laudon, J., Gharachorloo, K., Weber, W. D., Gupta, A., Henessy, J., Horowitz, M., and Lam, M. S. (1992). The Stanford Dash multiprocessor. *Computer*, 25(3):63–79. 381
- Lenzerini, M. (2002). Data integration: a theoretical perspective. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 233–246. 326
- Leung, J. Y. and Lai, E. K. (1979). On minimum cost recovery from system deadlock. *IEEE Trans. Comput.*, 28(9):671–677. C-37
- Levandovski, J. J., Larson, P. Å., and Stoica, R. (2013). Identifying hot and cold data in main-memory databases. In *Proc. 29th Int. Conf. on Data Engineering*, pages 26–37. 86

- Levin, K. D. and Morgan, H. L. (1975). Optimizing distributed data bases: A framework for research. In *Proc. National Computer Conf.*, pages 473–478. 31, 84, 85
- Levy, A. Y., Mendelzon, A. O., Sagiv, Y., and Srivastava, D. (1995). Answering queries using views. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 95–104. 328
- Levy, A. Y., Rajaraman, A., and Ordille, J. J. (1996a). The world wide web as a collection of views: Query processing in the information manifold. In *Proc. Workshop on Materialized Views: Techniques and Applications*, pages 43–55. 326
- Levy, A. Y., Rajaraman, A., and Ordille, J. J. (1996b). Querying heterogeneous information sources using source descriptions. In *Proc. 22th Int. Conf. on Very Large Data Bases*, pages 251–262. 326, 328
- Li, F., Ooi, B. C., Özsu, M. T., and Wu, S. (2014). Distributed data management using MapReduce. *ACM Comput. Surv.*, 46(3):Article No. 31. 497
- Li, H.-G., Chen, S., Tatemura, J., Agrawal, D., Candan, K. S., and Hsiung, W.-P. (2006). Safety guarantee of continuous join queries over punctuated data streams. In *Proc. 32nd Int. Conf. on Very Large Data Bases*, pages 19–30. 498
- Li, J., Maier, D., Tufte, K., Papadimos, V., and Tucker, P. a. (2005). Semantics and evaluation techniques for window aggregates in data streams. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 311–322. 498
- Li, W.-S. and Clifton, C. (2000). Semint: A tool for identifying attribute correspondences in heterogeneous databases using neural networks. *Data & Knowl. Eng.*, 33(1):49–84. 329
- Li, W.-S., Clifton, C., and Liu, S.-Y. (2000). Database integration using neural networks: Implementation and experiences. *Knowl. and Information Syst.*, 2(1): 73–96. 329
- Lim, L., Wang, M., Padmanabhan, S., Vitter, J. S., and Agarwal, R. (2003). Dynamic maintenance of web indexes using landmarks. In *Proc. 12th Int. World Wide Web Conf.*, pages 102–111. 593
- Lima, A., Mattoso, M., and Valduriez, P. (2004). OLAP query processing in a database cluster. In *Proc. 10th Int. Euro-Par Conf.*, pages 355–362. 382
- Lin, Q., Chang, P., Chen, G., Ooi, B. C., Tan, K., and Wang, Z. (2016). Towards a Non-2PC transaction management in distributed database systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1659–1674. 85
- Lin, Y., Kemme, B., Patiño Martínez, M., and Jiménez-Peris, R. (2005). Middleware based data replication providing snapshot isolation. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 419–430. 273
- Lindsay, B. (1979). Notes on distributed databases. Technical Report RJ 2517, IBM San Jose Research Laboratory, San Jose, Calif. C-49
- Litwin, W., Neimat, M.-A., and Schneider, D. A. (1993). LH* – linear hashing for distributed files. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 327–336. 433
- Liu, B., Zhu, Y., and Rundensteiner, E. (2006). Run-time operator state spilling for memory intensive long running queries. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 347–358. 498
- Livny, M., Khoshafian, S., and Boral, H. (1987). Multi-disk management. In *Proc. ACM SIGMETRICS Conf. on Measurement and Modeling of Computer Systems*, pages 69–77. 381

- Lohman, G., Mohan, C., Haas, L., Daniels, D., Lindsay, B., Selinger, P., and Wilms, P. (1985). Query processing in R*. In Kim, W., Reiner, D. S., and Batory, D. S., editors, *Query Processing in Database Systems*, pages 31–47. Springer. 176, 177
- Lomet, D., Feket, A., Wang, R., and Ward, P. (2012). Multi-version concurrency via timestamp range conflict management. In *Proc. 28th Int. Conf. on Data Engineering*, pages 714–725. 237
- Low, Y., Gonzalez, J., Kyrola, A., Bickson, D., and Guestrin, C. (2010). GraphLab: new framework for parallel machine learning. In *Proc. 26th Conf. on Uncertainty in Artificial Intelligence*, pages 340–349. 500
- Low, Y., Gonzalez, J., Kyrola, A., Bickson, D., Guestrin, C., and Hellerstein, J. M. (2012). Distributed graphlab: A framework for machine learning in the cloud. *Proc. VLDB Endowment*, 5(8):716–727. 500
- Lu, H., Shan, M.-C., and Tan, K.-L. (1991). Optimization of multi-way join queries for parallel execution. In *Proc. 17th Int. Conf. on Very Large Data Bases*, pages 549–560. 381
- Lu, H., Ooi, B., and Goh, C. (1992). On global multidatabase query optimization. *ACM SIGMOD Rec.*, 21(4):6–11. 328
- Lu, H., Ooi, B., and Goh, C. (1993). Multidatabase query optimization: Issues and solutions. In *Proc. 3rd Int. Workshop on Res. Issues in Data Eng.*, pages 137–143. 328
- Lugowski, A., Alber, D., Buluç, A., Gilbert, J. R., Reinhardt, S., Teng, Y., and Waranis, A. (2012). A flexible open-source toolbox for scalable complex graph analysis. In *Proc. 2012 SIAM Int. Conf. on Data Mining*, pages 930–941. 500
- Lumsdaine, A., Gregor, D., Hendrickson, B., and Berry, J. (2007). Challenges in parallel graph processing. *Parallel Processing Letters*, 17(01):5–20. 499
- Lunt, T. F. and Fernández, E. B. (1990). Database security. *ACM SIGMOD Rec.*, 19(4):90–97. 124
- Lv, Q., Cao, P., Cohen, E., Li, K., and Shenker, S. (2002). Search and replication in unstructured peer-to-peer networks. In *Proc. 16th Annual Int. Conf. on Supercomputing*, pages 84–95. 433
- Lynch, N. (1983a). Multilevel atomicity: A new correctness criterion for database concurrency control. *ACM Trans. Database Syst.*, 8(4):484–502. 238
- Lynch, N. (1983b). Concurrency control for resilient nested transactions. In *Proc. 2nd ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 166–181. 238
- Lynch, N. and Merritt, M. (1986). Introduction to the theory of nested transactions. Technical Report MIT/LCS/TR-367, Massachusetts Institute of Technology, Cambridge, Mass. 238
- Lynch, N., Merritt, M., Weihl, W. E., and Fekete, A. (1993). *Atomic Transactions in Concurrent Distributed Systems*. Morgan Kaufmann. 238
- Mackert, L. and Lohman, G. (1986a). R* optimizer validation and performance evaluation for distributed queries. In *Proc. 12th Int. Conf. on Very Large Data Bases*, pages 149–159. 177
- Mackert, L. F. and Lohman, G. (1986b). R* optimizer validation and performance evaluation for local queries. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 84–95. 177, B-31
- Madden, S. and Franklin, M. J. (2002). Fjording the stream: An architecture for queries over streaming sensor data. In *Proc. 18th Int. Conf. on Data Engineering*, pages 555–566. 498

- Madden, S., Shah, M., Hellerstein, J., and Raman, V. (2002a). Continuously adaptive continuous queries over streams. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 49–60. 498
- Madden, S., Shah, M. A., Hellerstein, J. M., and Raman, V. (2002b). Continuously adaptive continuous queries over streams. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 49–60. 328
- Madhavan, J., Bernstein, P., and Rahm, E. (2001). Generic schema matching with Cupid. In *Proc. 27th Int. Conf. on Very Large Data Bases*, pages 49–58. 327
- Mahmoud, . A. and Riordon, J. S. (1976). Optimal allocation of resources in distributed information networks. *ACM Trans. Database Syst.*, 1(1):66–78. 85
- Maiyya, S., Zakhary, V., Agrawal, D., and El Abbadi, A. (2018). Database and distributed computing fundamentals for scalable, fault-tolerant, and consistent maintenance of blockchains. *Proc. VLDB Endowment*, 11(12):2098–2101. 434
- Malewicz, G., Austern, M. H., Bik, A. J. C., Dehnert, J. C., Horn, I., Leiser, N., and Czajkowski, G. (2010). Pregel: a system for large-scale graph processing. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 135–146. 500
- Manber, U. and Myers, G. (1990). Suffix arrays: a new method for on-line string searches. In *Proc. 1st Annual ACM-SIAM Symp. on Discrete Algorithms*, pages 319–327. 593
- Manegold, S., Boncz, P. A., and Kersten, M. L. (2002). Optimizing main-memory join on modern hardware. *IEEE Trans. Knowl. and Data Eng.*, 14(4):709–730. 381
- Manolescu, I., Florescu, D., and Kossmann, D. (2001). Answering XML queries on heterogeneous data sources. In *Proc. 27th Int. Conf. on Very Large Data Bases*, pages 241–250. 326
- Martins, V. and Pacitti, E. (2006). Dynamic and distributed reconciliation in p2p-dht networks. In *European Conf. on Parallel Computing (Euro-Par)*, pages 337–349. 434
- Martins, V., Akbarinia, R., Pacitti, E., and Valduriez, P. (2006a). Reconciliation in the APPA P2P system. In *Proc. IEEE Int. Conf. on Parallel and Distributed Systems*, pages 401–410. 434
- Martins, V., Pacitti, E., and Valduriez, P. (2006b). Survey of data replication in P2P systems. Technical Report 6083, INRIA, Rennes, France. 434
- Martins, V., Pacitti, E., Dick, M. E., and Jimenez-Peris, R. (2008). Scalable and topology-aware reconciliation on P2P networks. *Distrib. Parall. Databases*, 24(1–3):1–43. 434
- McBrien, P. and Poulouvasilis, A. (2003). Defining peer-to-peer data integration using both as view rules. In *Proc. 1st Int. Workshop on Databases, Information Systems and Peer-to-Peer Computing*, pages 91–107. 433
- McCallum, A., Nigam, K., Rennie, J., and Seymore, K. (1999). A machine learning approach to building domain-specific search engines. In *Proc. 16th Int. Joint Conf. on AI*. 593
- McCann, R., AlShebli, B., Le, Q., Nguyen, H., Vu, L., and Doan, A. (2005). Mapping maintenance for data integration systems. In *Proc. 31st Int. Conf. on Very Large Data Bases*, pages 1018–1029. 329
- McCormick, W. T., Schweitzer, P. J., and White, T. W. (1972). Problem decomposition and data reorganization by a clustering technique. *Oper. Res.*, 20(5):993–1009. 85
- McCune, R. R., Weninger, T., and Madey, G. (2015). Thinking like a vertex: A survey of vertex-centric frameworks for large-scale distributed graph processing. *ACM Comput. Surv.*, 48(2):25:1–25:39. 499

- Medina-Mora, R., Wong, H., and Flores, P. (1993). Action workflow as the enterprise integration technology. *Q. Bull. IEEE TC on Data Eng.*, 16(2):49–52. C-18
- Mehta, M. and DeWitt, D. (1995). Managing intra-operator parallelism in parallel database systems. In *Proc. 21th Int. Conf. on Very Large Data Bases*. 381, 382
- Melnik, S., Raghavan, S., Yang, B., and Garcia-Molina, H. (2001). Building a distributed full-text index for the web. In *Proc. 10th Int. World Wide Web Conf.*, pages 396–406. 593
- Melnik, S., Garcia-Molina, H., and Rahm, E. (2002). Similarity flooding: A versatile graph matching algorithm and its application to schema matching. In *Proc. 18th Int. Conf. on Data Engineering*, pages 117–128. 327
- Menasce, D. A. and Muntz, R. R. (1979). Locking and deadlock detection in distributed databases. *IEEE Trans. Softw. Eng.*, SE-5(3):195–202. 237
- Mendelzon, A. O., Mihaila, G. A., and Milo, T. (1997). Querying the World Wide Web. *Int. J. Digit. Libr.*, 1(1):54–67. 593
- Meng, W., Yu, C., Kim, W., Wang, G., Phan, T., and Dao, S. (1993). Construction of relational front-end for object-oriented database systems. In *Proc. 9th Int. Conf. on Data Engineering*, pages 476–483. 328
- Milán-Franco, J. M., Jiménez-Peris, R., Patiño-Martínez, M., and Kemme, B. (2004). Adaptive middleware for data replication. In *Proc. ACM/IFIP/USENIX 5th Int. Middleware Conf.*, pages 175–194. 382
- Miller, R. J., Haas, L. M., and Hernández, M. A. (2000). Schema mapping as query discovery. In *Proc. 26th Int. Conf. on Very Large Data Bases*, pages 77–88. 327
- Miller, R. J., Hernández, M. A., Haas, L. M., Yan, L., Ho, C. T. H., Fagin, R., and Popa, L. (2001). The Clio project: Managing heterogeneity. *ACM SIGMOD Rec.*, 31(1):78–83. 327
- Milo, T. and Zohar, S. (1998). Using schema matching to simplify heterogeneous data translation. In *Proc. 24th Int. Conf. on Very Large Data Bases*, pages 122–133. 327
- Minoura, T. and Wiederhold, G. (1982). Resilient extended true-copy token scheme for a distributed database system. *IEEE Trans. Softw. Eng.*, SE-8(3):173–189. 273
- Mitchell, T. (1997). *Machine Learning*. McGraw-Hill. 593
- Mitzenmacher, M. (2001). The power of two choices in randomized load balancing. *IEEE Trans. Paralle. Dist. Sys.*, 12(10):1094–1104. 499
- Mohan, C. (1979). Data base design in the distributed environment. Working Paper WP-7902, Department of Computer Sciences, University of Texas at Austin. 84
- Mohan, C. and Lindsay, B. (1983). Efficient commit protocols for the tree of processes model of distributed transactions. In *Proc. ACM SIGACT-SIGOPS 2nd Symp. on the Principles of Distributed Computing*, pages 76–88. 237, 239
- Mohan, C. and Yeh, R. T. (1978). *Distributed Data Base Systems: A Framework for Data Base Design*. In *Distributed Data Bases, Infotech State-of-the-Art Report*. Infotech. 31
- Mohan, C., Lindsay, B., and Obermarck, R. (1986). Transaction management in the r* distributed database management system. *ACM Trans. Database Syst.*, 11(4): 378–396. 237
- Morgan, H. L. and Levin, K. D. (1977). Optimal program and data location in computer networks. *Commun. ACM*, 20(5):315–322. 85
- Moss, E. (1985). *Nested Transactions*. M.I.T. Press. 238, C-15, C-16
- Muthukrishnan, S. (2005). Data streams: Algorithms and applications. *Foundations and Trends in Theoretical Computer Science*, 1(2):117–236. 498

- Naacke, H., Tomasic, A., and Valduriez, P. (1999). Validating mediator cost models with Disco. *Networking and Information Systems Journal*, 2(5):639–663. 328
- Najork, M. and Wiener, J. L. (2001). Breadth-first crawling yields high-quality pages. In *Proc. 10th Int. World Wide Web Conf.*, pages 114–118. 593
- Nakamoto, S. (2008). Bitcoin: A peer-to-peer electronic cash system. Accessible at <https://bitcoin.org/bitcoin.pdf/>. Last accessed February 2019. 434
- Nasir, M. A. U., Morales, G. D. F., García-Soriano, D., Kourtellis, N., and Serafini, M. (2015). The power of both choices: Practical load balancing for distributed stream processing engines. In *Proc. 31st Int. Conf. on Data Engineering*, pages 137–148. 499
- Nasir, M. A. U., Morales, G. D. F., Kourtellis, N., and Serafini, M. (2016). When two choices are not enough: Balancing at scale in distributed stream processing. In *Proc. 32nd Int. Conf. on Data Engineering*, pages 589–600. 499
- Naumann, F., Ho, C.-T., Tian, X., Haas, L. M., and Megiddo, N. (2002). Attribute classification using feature analysis. In *Proc. 18th Int. Conf. on Data Engineering*, page 271. 327
- Navathe, S. B., Ceri, S., Wiederhold, G., and Dou, J. (1984). Vertical partitioning of algorithms for database design. *ACM Trans. Database Syst.*, 9(4):680–710. 85
- Nejdl, W., Siberski, W., and Sintek, M. (2003). Design issues and challenges for rdf- and schema-based peer-to-peer systems. *ACM SIGMOD Rec.*, 32(3):41–46. 433, 434
- Nepal, S. and Ramakrishna, M. (1999). Query processing issues in image (multimedia) databases. In *Proc. 15th Int. Conf. on Data Engineering*, pages 22–29. 434
- Neumann, T. and Weikum, G. (2008). RDF-3X: a RISC-style engine for RDF. *Proc. VLDB Endowment*, 1(1):647–659. 593
- Neumann, T. and Weikum, G. (2009). The RDF-3X engine for scalable management of RDF data. *VLDB J.*, 19(1):91–113. 593
- Newman, M. E. J., Watts, D. J., and Strogatz, S. H. (2002). Random graph models of social networks. In *(Sackler NAS Colloquium) Self-Organized Complexity in the Physical, Biological, and Social Sciences*, pages 2566–2573. National Academy of Sciences. 499
- Niamir, B. (1978). Attribute partitioning in a self-adaptive relational database system. Technical Report 192, Laboratory for Computer Science, Massachusetts Institute of Technology, Cambridge, Mass. 85
- Nicolas, J. M. (1982). Logic for improving integrity checking in relational data bases. *Acta Informatica*, 18:227–253. 124
- Nikolic, M. and Olteanu, D. (2018). Incremental view maintenance with triple lock factorization benefits. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 365–380. 124
- Nodine, M. and Zdonik, S. (1990). Cooperative transaction hierarchies: A transaction model to support design applications. In *Proc. 16th Int. Conf. on Very Large Data Bases*, pages 83–94. C-18
- Novakovic, S., Daglis, A., Bugnion, E., Falsafi, B., and Grot, B. (2014). Scale-out NUMA. In *Architectural Support for Programming Languages and Operating Systems, ASPLOS*, pages 3–18. 381
- Obermack, R. (1982). Distributed deadlock detection algorithm. *ACM Trans. Database Syst.*, 7(2):187–208. 31, 237, 238
- Okcan, A. and Riedewald, M. (2011). Processing theta-joins using MapReduce. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 949–960. 498

- Olston, C., Reed, B., Srivastava, U., Kumar, R., and Tomkins, A. (2008). Pig latin: a not-so-foreign language for data processing. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1099–1110. 498
- Ong, K. W., Papakonstantinou, Y., and Vernoux, R. (2014). The SQL++ semi-structured data model and query language: A capabilities survey of SQL-on-Hadoop, NoSQL and NewSQL databases. CoRR/abs/1405.3631. 535
- Ongaro, D. and Ousterhout, J. (2014). In search of an understandable consensus algorithm. In *Proc. USENIX 2014 Annual Technical Conf.*, pages 305–320. 239, 536
- Ooi, B., Shu, Y., and Tan, K.-L. (2003). Relational data sharing in peer-based data management systems. *ACM SIGMOD Rec.*, 32(3):59–64. 433
- Ouksel, A. M. and Sheth, A. P. (1999). Semantic interoperability in global information systems: A brief introduction to the research area and the special section. *ACM SIGMOD Rec.*, 28(1):5–12. 327
- Özsoyoglu, Z. M. and Zhou, N. (1987). Distributed query processing in broadcasting local area networks. In *Proc. 20th Hawaii Int. Conf. on System Sciences*, pages 419–429. 177
- Özsu, M. T. (2016). A survey of RDF data management systems. *Front. Comput. Sci.*, 10(3):418–432. 593
- Pacaci, A. and Özsu, M. T. (2018). Distribution-aware stream partitioning for distributed stream processing systems. In *Proc. 5th ACM SIGMOD Workshop on Algorithms and Systems for MapReduce and Beyond*, pages 6:1–6:10. 499
- Pacitti, E. and Simon, E. (2000). Update propagation strategies to improve freshness in lazy master replicated databases. *VLDB J.*, 8(3-4):305–318. 273
- Pacitti, E., Simon, E., and de Melo, R. (1998). Improving data freshness in lazy master schemes. In *Proc. 18th IEEE Int. Conf. on Distributed Computing Systems*, pages 164–171. 272
- Pacitti, E., Minet, P., and Simon, E. (1999). Fast algorithms for maintaining replica consistency in lazy master replicated databases. In *Proc. 25th Int. Conf. on Very Large Data Bases*, pages 126–137. 273
- Pacitti, E., Coulon, C., Valduriez, P., and Özsu, M. T. (2005). Preventive replication in a database cluster. *Distrib. Parall. Databases*, 18(3):223–251. 382
- Pacitti, E., Valduriez, P., and Mattoso, M. (2007). Grid data management: open problems and new issues. *Journal of Grid Computing*, 5(3):273–281. 434
- Pacitti, E., Akbarinia, R., and Dick, M. E. (2012). *P2P Techniques for Decentralized Applications*. Synthesis Lectures on Data Management. Morgan & Claypool Publishers. 433
- Page, L., Brin, S., Motwani, R., and Winograd, T. (1998). The pagerank citation ranking: Bringing order to the web. Technical report, Stanford University. 593
- Palopoli, L. (2003). Experiences using DIKE, a system for supporting cooperative information system and data warehouse design. *Inf. Syst.*, 28(7):835–865. 327
- Palopoli, L., Saccà, D., and Ursino, D. (1998). Semi-automatic semantic discovery of properties from database schemas. In *Proc. 2nd Int. Conf. on Database Eng. and Applications*, pages 244–253. 327
- Palopoli, L., Saccà, D., Terracina, G., and Ursino, D. (1999). A unified graph-based framework for deriving nominal interscheme properties, type conflicts and object cluster similarities. In *Proc. Int. Conf. on Cooperative Inf. Syst.*, pages 34–45. 327

- Palopoli, L., Saccà, D., Terracina, G., and Ursino, D. (2003). Uniform techniques for deriving similarities of objects and subschemes in heterogeneous databases. *IEEE Trans. Knowl. and Data Eng.*, 15(2):271–294. 327
- Papadimitriou, C. H. (1979). Serializability of concurrent database updates. *J. ACM*, 26(4):631–653. C-14
- Papadimitriou, C. H. (1986). *The Theory of Concurrency Control*. Computer Science Press. 237
- Papakonstantinou, Y., Garcia-Molina, H., and Widom, J. (1995). Object exchange across heterogeneous information sources. In *Proc. 11th Int. Conf. on Data Engineering*, pages 251–260. 593
- Pape, C. L., Gançarski, S., and Valduriez, P. (2004). Refresco: Improving query performance through freshness control in a database cluster. In *Proc. Confederated Int. Conf. DOA, CoopIS and ODBASE*, Lecture Notes in Computer Science 3290, pages 174–193. 273
- Paris, J. F. (1986). Voting with witnesses: A consistency scheme for replicated files. In *Proc. 6th IEEE Int. Conf. on Distributed Computing Systems*, pages 606–612. 273
- Pasetto, D. and Akhriev, A. (2011). A comparative study of parallel sort algorithms. In *Proc. 26th ACM SIGPLAN Conf. on Object-Oriented Programming Systems, Languages & Applications*, pages 203–204. 381
- Passerini, A., Frasconi, P., and Soda, G. (2001). Evaluation methods for focused crawling. In *Proc. 7th Congress of the Italian Association for Artificial Intelligence*, pages 33–39. 593
- Pasupuleti, P. and Purra, B. S. (2015). *Data Lake Development with Big Data*. Packt Books. 500
- Patiño-Martínez, M., Jiménez-Peris, R., Kemme, B., and Alonso, G. (2000). Scalable replication in database clusters. In *Proc. 14th Int. Symp. on Distributed Computing*, pages 315–329. 273, 382
- Pavlo, A., Paulson, E., Rasin, A., Abadi, D. J., DeWitt, D. J., Madden, S., and Stonebraker, M. (2009). A comparison of approaches to large-scale data analysis. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 165–178. 497
- Pavlo, A., Curino, C., and Zdonik, S. B. (2012). Skew-aware automatic database partitioning in shared-nothing, parallel OLTP systems. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 61–72. 85
- Perez-Sorrosal, F., Vuckovic, J., Patiño-Martínez, M., and Jiménez-Peris, R. (2006). Highly available long running transactions and activities for J2EE. In *Proc. 26th IEEE Int. Conf. on Distributed Computing Systems*, page 2. 382
- Petraki, E., Idreos, S., and Manegold, S. (2015). Holistic indexing in main-memory column-stores. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1153–1166. 86
- Piatetsky-Shapiro, G. and Connell, C. (1984). Accurate estimation of the number of tuples satisfying a condition. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 256–276. B-19
- Pike, R., Dorward, S., Griesemer, R., and Quinlan, S. (2005). Interpreting the data: Parallel analysis with sawzall. *Sci. Program.*, 13(4):277–298. 498
- Pirahesh, H., Mohan, C., Cheng, J. M., Liu, T. S., and Selinger, P. G. (1990). Parallelism in rdbms : Architectural issues and design. In *Proc. 2nd Int. Symp. on Databases in Distributed and Parallel Systems*, pages 4–29. 381

- Pirk, H., Manegold, S., and Kersten, M. (2014). Waste not . . . efficient co-processing of relational data. In *Proc. 30th Int. Conf. on Data Engineering*, pages 508–519. 86
- Plattner, C. and Alonso, G. (2004). Ganymed: Scalable replication for transactional web applications. In *Proc. ACM/IFIP/USENIX 5th Int. Middleware Conf.*, pages 155–174. 273
- Plugge, E., Membrey, P., and Hawkins, T. (2010). *The Definitive Guide to MongoDB: The NoSQL Database for Cloud and Desktop Computing*. Apress. 535
- Poosala, V., Ioannidis, Y., Haas, P., and Shekita, E. (1996). Improved histograms for selectivity estimation of range predicates. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 294–305. B-23
- Popa, L., Velegrakis, Y., Miller, R. J., Hernandez, M. A., and Fagin, R. (2002). Translating web data. In *Proc. 28th Int. Conf. on Very Large Data Bases*. 327
- Porto, F., Laber, E. S., and Valduriez, P. (2003). Cherry picking: A semantic query processing strategy for the evaluation of expensive predicates. In *Proc. Brazilian Symposium on Databases*, pages 356–370. 177, 328
- Ports, D. R. K. and Grittner, K. (2012). Serializable snapshot isolation in postgresql. *Proc. VLDB Endowment*, 5(12):1850–1861. 238
- Pottinger, R. and Levy, A. Y. (2000). A scalable algorithm for answering queries using views. In *Proc. 26th Int. Conf. on Very Large Data Bases*, pages 484–495. 328
- Pu, C. (1988). Superdatabases for composition of heterogeneous databases. In *Proc. 4th Int. Conf. on Data Engineering*, pages 548–555. 238, 327, C-17
- Pu, C. and Leff, A. (1991). Replica control in distributed systems: An asynchronous approach. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 377–386. 272
- Qian, Z., He, Y., Su, C., Wu, Z., Zhu, H., Zhang, T., Zhou, L., Yu, Y., and Zhang, Z. (2013). Timestream: Reliable stream computation in the cloud. In *Proc. 8th ACM SIGOPS/EuroSys European Conf. on Comp. Syst.*, pages 1–14. 499
- Qin, L., Yu, J. X., Chang, L., Cheng, H., Zhang, C., and Lin, X. (2014). Scalable big graph processing in mapreduce. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 827–838. 499
- Quamar, A., Kumar, K. A., and Deshpande, A. (2013). Sword: Scalable workload-aware data placement for transactional workloads. In *Proc. 16th Int. Conf. on Extending Database Technology*, pages 430–441. 85
- Raghavan, S. and Garcia-Molina, H. (2001). Crawling the hidden web. In *Proc. 27th Int. Conf. on Very Large Data Bases*. 592
- Raghavan, S. and Garcia-Molina, H. (2003). Representing web graphs. In *Proc. 19th Int. Conf. on Data Engineering*, pages 405–416. 592, 593
- Rahal, A., Zhu, Q., and Larson, P.-Å. (2004). Evolutionary techniques for updating query cost models in a dynamic multidatabase environment. *VLDB J.*, 13(2): 162–176. 328
- Rahimi, S. (1987). Reference architecture for distributed database management systems. In *Proc. 3th Int. Conf. on Data Engineering*. Tutorial Notes. 32
- Rahimi, S. K. and Haug, F. S. (2010). *Distributed Database Management Systems – A Practical Approach*. Wiley. 31
- Rahm, E. and Bernstein, P. a. (2001). A survey of approaches to automatic schema matching. *VLDB J.*, 10(4):334–350. 327

- Rahm, E. and Do, H. H. (2000). Data cleaning: Problems and current approaches. *Q. Bull. IEEE TC on Data Eng.*, 23(4):3–13. 327, 594
- Rahm, E. and Marek, R. (1995). Dynamic multi-resource load balancing in parallel database systems. In *Proc. 21th Int. Conf. on Very Large Data Bases*, pages 395–406. 381
- Ramabhadran, S., Ratnasamy, S., Hellerstein, J. M., and Shenker, S. (2004). Brief announcement: prefix hash tree. In *Proc. ACM SIGACT-SIGOPS 23rd Symp. on the Principles of Distributed Computing*, page 368. 433, 434
- Ramamoorthy, C. V. and Wah, B. W. (1983). The isomorphism of simple file allocation. *IEEE Trans. Comput.*, 32:221–223. 85
- Ramamritham, K. and Pu, C. (1995). A formal characterization of epsilon serializability. *IEEE Trans. Knowl. and Data Eng.*, 7(6):997–1007. 238, 272
- Raman, V. and Hellerstein, J. M. (2001). Potter’s wheel: An interactive data cleaning system. In *Proc. 27th Int. Conf. on Very Large Data Bases*, pages 381–390. 327
- Raman, V., Deshpande, A., and Hellerstein, J. M. (2003). Using state modules for adaptive query processing. In *Proc. 19th Int. Conf. on Data Engineering*, pages 353–365. 177, 328
- Rao, J., Zhang, C., Megiddo, N., and Lohman, G. (2002). Automating physical database design in a parallel database. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*. 85
- Rastogi, V., Machanavajjhala, A., Chitnis, L., and Sarma, A. D. (2013). Finding connected components in map-reduce in logarithmic rounds. In *Proc. 29th Int. Conf. on Data Engineering*, pages 50–61. 499
- Ratnasamy, S., Francis, P., Handley, M., Karp, R., and Schenker, S. (2001). A scalable content-addressable network. *ACM SIGCOMM Computer Communication Review*, 31(4):161–172. 433
- Ray, I., Mancini, L. V., Jajodia, S., and Bertino, E. (2000). Asep: A secure and flexible commit protocol for mls distributed database systems. *IEEE Trans. Knowl. and Data Eng.*, 12(6):880–899. 124
- Redmond, E. and Wilson, J. R. (2012). *Seven Databases in Seven Weeks: A Guide to Modern Databases and the NoSQL Movement*. The Pragmatic Programmers. 535
- Reed, D. P. (1978). *Naming and Synchronization in a Decentralized Computer System*. PhD thesis, MIT. 237
- Reiss, F. and Hellerstein, J. (2005). Data triage: an adaptive architecture for load shedding in telegraphCQ. In *Proc. 21st Int. Conf. on Data Engineering*, pages 155–156. 498
- Rekatsinas, T., Joglekar, M., Garcia-Molina, H., Parameswaran, A. G., and Ré, C. (2017). Slimfast: Guaranteed results for data fusion and source reliability. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 1399–1414. 591, 594
- Revilak, S., O’Neil, P. E., and O’Neil, E. J. (2011). Precisely serializable snapshot isolation (PSSI). In *Proc. 27th Int. Conf. on Data Engineering*, pages 482–493. 238
- Ribeiro-Neto, B. A. and Barbosa, R. A. (1998). Query performance for tightly coupled distributed digital libraries. In *Proc. 3rd ACM Int. Conf. on Digital Libraries*, pages 182–190. 593
- Richter, S., Quiané-Ruiz, J.-A., Schuh, S., and Dittrich, J. (2013). Towards zero-overhead static and adaptive indexing in Hadoop. *VLDB J.*, 23(3):469–494. 86

- Ritter, J. (2001). Why Gnutella can't scale, no, really. <http://www.darkridge.com/jpr5/doc/gnutella.html>. Last accessed June 2019. 433
- Rivera-Vega, P., Varadarajan, R., and Navathe, S. B. (1990). Scheduling data redistribution in distributed databases. In *Proc. Int. Conf. on Data Eng.*, pages 166–173. 85
- Rjaibi, W. (2004). An introduction to multilevel secure relational database management systems. In *Proc. Conf. of the IBM Centre for Advanced Studies on Collaborative Research*, pages 232–241. 124
- Robinson, I., Webber, J., and Eifrem, E. (2015). *Graph Databases*. O'Reilly, 2 edition. 535
- Röhm, U., Böhm, K., and Schek, H.-J. (2000). OLAP query routing and physical design in a database cluster. In *Advances in Database Technology, Proc. 7th Int. Conf. on Extending Database Technology*, pages 254–268. 382
- Röhm, U., Böhm, K., and Schek, H.-J. (2001). Cache-aware query routing in a cluster of databases. In *Proc. 17th Int. Conf. on Data Engineering*, pages 641–650. 382
- Röhm, U., Böhm, K., Schek, H.-J., and Schuldt, H. (2002). FAS - A freshness-sensitive coordination middleware for a cluster of OLAP components. In *Proc. 28th Int. Conf. on Very Large Data Bases*, pages 754–765. 273
- Roitman, H. and Gal, A. (2006). Ontobuilder: Fully automatic extraction and consolidation of ontologies from web sources using sequence semantics. In *Proc. EDBT Workshops*, volume 4254 of *LNCS*, pages 573–576. 327
- Rosenkrantz, D. J. and Hunt, H. B. (1980). Processing conjunctive predicates and queries. In *Proc. 6th Int. Conf. on Very Data Bases*, pages 64–72. B-6
- Rosenkrantz, D. J., Stearns, R. E., and Lewis, P. M. (1978). System level concurrency control for distributed database systems. *ACM Trans. Database Syst.*, 3(2):178–198. C-36
- Roth, J. P., Bouricius, W. G., Carter, E. C., and Schneider, P. R. (1967). Phase ii of an architectural study for a self-repairing computer. Report SAMS0-TR-67-106, U. S. Air Force Space and Missile Division, El Segundo, Calif. Cited in [Siewiorek and Swarz, 1982]. C-38
- Roth, M. and Schwartz, P. (1997). Don't scrap it, wrap it! a wrapper architecture for legacy data sources. In *Proc. 23th Int. Conf. on Very Large Data Bases*, pages 266–275. 327
- Roth, M. T., Ozcan, F., and Haas, L. M. (1999). Cost models do matter: Providing cost information for diverse data sources in a federated system. In *Proc. 25th Int. Conf. on Very Large Data Bases*, pages 599–610. 328
- Rothermel, K. and Mohan, C. (1989). Aries/nt: A recovery method based on write-ahead logging for nested transactions. In *Proc. 15th Int. Conf. on Very Large Data Bases*, pages 337–346. 238
- Roubini, N. (2018). Testimony for the hearing of the US senate committee on banking, housing and community affairs on exploring the cryptocurrency and blockchain ecosystem. Accessible at <https://www.banking.senate.gov/imo/media/doc/Roubini%20Testimony%2010-11-18.pdf>. Last accessed February 2019. 434
- Roy, A., Mihailovic, I., and Zwaenepoel, W. (2013). X-stream: edge-centric graph processing using streaming partitions. In *Proc. 24th ACM Symp. on Operating System Principles*, pages 472–488. 500

- Ryvkina, E., Maskey, A., Adams, I., Sandler, B., Fuchs, C., Cherniack, M., and Zdonik, S. (2006). Revision processing in a stream processing engine: A high-level design. In *Proc. 22nd Int. Conf. on Data Engineering*, page 141. 498
- Sacca, D. and Wiederhold, G. (1985). Database partitioning in a cluster of processors. *ACM Trans. Database Syst.*, 10(1):29–56. 85
- Sacco, M. S. and Yao, S. B. (1982). Query optimization in distributed data base systems. In Yovits, M., editor, *Advances in Computers*, volume 21, pages 225–273. 31
- Saito, Y. and Shapiro, M. (2005). Optimistic replication. *ACM Comput. Surv.*, 37(1): 42–81. 272
- Sakr, S., Liu, A., and Fayoumi, A. G. (2013). The family of MapReduce and large-scale data processing systems. *ACM Comput. Surv.*, 46(1):11:1–11:44. 497
- Salihoglu, S. and Widom, J. (2013). GPS: a graph processing system. In *Proc. 25th Int. Conf. on Scientific and Statistical Database Management*, pages 22:1–22:12. 500
- Salihoglu, S. and Widom, J. (2014). Optimizing graph algorithms on Pregel-like systems. *Proc. VLDB Endowment*, 7(7):577–588. 500
- Salton, G. (1989). *Automatic Text Processing – The Transformation, Analysis, and Retrieval of Information by Computer*. Addison–Wesley. 593
- Sandhu, R. S., Coyne, E. J., Feinstein, H. L., and Youman, C. E. (1996). Role-based access control models. *IEEE Computer*, 29(2):38–47. 124
- Schenkel, R., Weikum, G., Weißberg, N., and Wu, X. (2000). Federated transaction management with snapshot isolation. In Saake, G., Schwarz, K., and Türker, C., editors, *Transactions and Database Dynamics*, pages 1–25. Springer. 238
- Schmachtenberg, M., Bizer, C., and Paulheim, H. (2014). Adoption of best data practices in different topical domains. In *Proc. 13th Int. Semantic Web Conf.*, pages 245–260. 593
- Schmidt, C. and Parashar, M. (2004). Enabling flexible queries with guarantees in P2P systems. *IEEE Internet Computing*, 8(3):19–26. 433
- Schreiber, F. (1977). A framework for distributed database systems. In *Proc. Int. Computing Symposium*, pages 475–482. 31
- Schuhknecht, F. M., Jindal, A., and Dittrich, J. (2013). The uncracked pieces in database cracking. *Proc. VLDB Endowment*, 7(2):97–108. 86
- Selinger, P. G. and Adiba, M. (1980). Access path selection in distributed data base management systems. In *Proc. First Int. Conf. on Data Bases*, pages 204–215. 177, B-21
- Selinger, P. G., Astrahan, M. M., Chamberlin, D. D., Lorie, R. A., and Price, T. G. (1979). Access path selection in a relational database management system. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 23–34. 176, B-2, B-20, B-27
- Sequeda, J. F., Arenas, M., and Miranker, D. P. (2014). OBDA: query rewriting or materialization? in practice, both! In *Proc. 13th Int. Semantic Web Conf.*, pages 535–551. 593
- Severence, D. G. and Lohman, G. M. (1976). Differential files: Their application to the maintenance of large databases. *ACM Trans. Database Syst.*, 1(3):256–261. C-44
- Shah, M. A., Hellerstein, J. M., Chandrasekaran, S., and Franklin, M. J. (2003). Flux: An adaptive partitioning operator for continuous query systems. In *Proc. 19th Int. Conf. on Data Engineering*, pages 25–36. 177, 328, 499

- Shao, B., Wang, H., and Li, Y. (2013). Trinity: a distributed graph engine on a memory cloud. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 505–516. 500
- Shatdal, A. and Naughton, J. F. (1993). Using shared virtual memory for parallel join processing. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 119–128. 381
- Shatdal, A., Kant, C., and Naughton, J. F. (1994). Cache conscious algorithms for relational query processing. In *Proc. 20th Int. Conf. on Very Large Data Bases*, pages 510–521. 381
- Shekita, E. J., Young, H. C., and Tan, K. L. (1993). Multi-join optimization for symmetric multiprocessor. In *Proc. 19th Int. Conf. on Very Large Data Bases*, pages 479–492. 381
- Sheth, A., Larson, J., Cornello, A., and Navathe, S. B. (1988a). A tool for integrating conceptual schemas and user views. In *Proc. 4th Int. Conf. on Data Engineering*, pages 176–183. 124, 327
- Sheth, A., Larson, J., and Watkins, E. (1988b). Tailor, a tool for updating views. In *Advances in Database Technology, Proc. 1st Int. Conf. on Extending Database Technology*, pages 190–213. 124
- Sheth, A. P. and Kashyap, V. (1992). So far (schematically) yet so near (semantically). In *Proc. IFIP WG 2.6 Database Semantics Conf. on Interoperable Database Systems*, pages 283–312. 327
- Sheth, A. P. and Larson, J. (1990). Federated database systems for managing distributed, heterogeneous, and autonomous databases. *ACM Comput. Surv.*, 22(3):183–236. 31, 32, 326
- Shute, J., Vingralek, R., Samwel, B., Handy, B., Whipkey, C., Rollins, E., Oancea, M., Littlefield, K., Menestrina, D., Ellner, S., Cieslewicz, J., Rae, I., Stancescu, T., and Apte, H. (2013). F1: A distributed SQL database that scales. *Proc. VLDB Endowment*, 6(11):1068–1079. 536
- Sidell, J., Aoki, P. M., Sah, A., Staelin, C., Stonebraker, M., and Yu, A. (1996). Data replication in Mariposa. In *Proc. 12th Int. Conf. on Data Eng.*, pages 485–494. 273
- Sidirourgos, L., Goncalves, R., Kersten, M., Nes, N., and Manegold, S. (2008). Column-store support for RDF data management: not all swans are white. *Proc. VLDB Endowment*, 1(2):1553–1563. 593
- Silberschatz, A., Korth, H., and Sudarshan, S. (2019). *Database System Concepts*. McGraw-Hill, 7 edition. 239
- Simitsis, A., Wilkinson, K., Castellanos, M., and Dayal, U. (2009). QoX-driven ETL design: reducing the cost of ETL consulting engagements. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 953–960. 536
- Simitsis, A., Wilkinson, K., Castellanos, M., and Dayal, U. (2012). Optimizing analytic data flows for multiple execution engines. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 829–840. 536
- Simon, E. and Valduriez, P. (1984). Design and implementation of an extendible integrity subsystem. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 9–17. 124
- Simon, E. and Valduriez, P. (1986). Integrity control in distributed database systems. In *Proc. 19th Hawaii Int. Conf. on System Sciences*, pages 622–632. 124
- Simon, E. and Valduriez, P. (1987). Design and analysis of a relational integrity subsystem. Technical Report DB-015-87, Microelectronics and Computer Corporation, Austin, Tex. 124

- Singhal, M. (1989). Deadlock detection in distributed systems. *Computer*, 22(11): 37–48. 238
- Skarra, A. (1989). Concurrency control for cooperating transactions in an object-oriented database. In *Proc. ACM SIGPLAN Workshop on Object-Based Concurrent Programming*, pages 145–147. 238
- Skarra, A., Zdonik, S., and Reiss, S. (1986). An object server for an object-oriented database system. In *Proc. of the 1st Int. Workshop on Object-Oriented Database Systems*, pages 196–204. 238
- Skeen, D. (1981). Nonblocking commit protocols. In *ACM SIGMOD Int. Conf. on Management of Data*, pages 133–142. 237, 239
- Skeen, D. (1982a). A quorum-based commit protocol. In *Proc. 6th Berkeley Workshop on Distributed Data Management and Computer Networks*, pages 69–80. 237
- Skeen, D. (1982b). *Crash Recovery in a Distributed Database Management System*. Ph.D. thesis, Department of Electrical Engineering and Computer Science, University of California at Berkeley, Berkeley, Calif. 239
- Skeen, D. and Stonebraker, M. (1983). A formal model of crash recovery in a distributed system. *IEEE Trans. Softw. Eng.*, SE-9(3):219–228. 237, 239
- Skeen, D. and Wright, D. (1984). Increasing availability in partitioned networks. In *Proc. 3rd ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 290–299. 273
- Smith, J. M. and Chang, P. Y. (1975). Optimizing the performance of a relational algebra database interface. *Commun. ACM*, 18(10):568–579. B-10
- Somani, A., Choy, D., and Kleewein, J. C. (2002). Bringing together content and data management systems: Challenges and opportunities. *IBM Systems J.*, 41(4): 686–696. 329
- Sousa, A., Oliveira, R., Moura, F., and Pedone, F. (2001). Partial replication in the database state machine. In *Proc. IEEE Int. Symp. Network Computing and Applications*, pages 298–309. 382
- Srivastava, U. and Widom, J. (2004a). Flexible time management in data stream systems. In *Proc. ACM SIGACT-SIGMOD Symp. on Principles of Database Systems*, pages 263–274. 498
- Srivastava, U. and Widom, J. (2004b). Memory-limited execution of windowed stream joins. In *Proc. 30th Int. Conf. on Very Large Data Bases*, pages 324–335. 498
- Stanoi, I., Agrawal, D., and El Abbadi, A. (1998). Using broadcast primitives in replicated databases. In *Proc. 8th IEEE Int. Conf. on Distributed Computing Systems*, pages 148–155. 273
- Stearns, R. E., II, P. M. L., and Rosenkrantz, D. J. (1976). Concurrency controls for database systems. In *Proc. 17th Symp. on Foundations of Computer Science*, pages 19–32. C-14
- Stöhr, T., Märtens, H., and Rahm, E. (2000). Multi-dimensional database allocation for parallel data warehouses. In *Proc. 26th Int. Conf. on Very Large Data Bases*, pages 273–284. 382
- Stoica, I., Morris, R., Karger, D., Kaashoek, M., and Balakrishnan, H. (2001). Chord: A scalable peer-to-peer lookup service for internet applications. In *Proc. 2001 Conf. on Applications, Technologies, Architectures, and Protocols for Computer Communication*, pages 149–160. 433

- Stoica, I., Morris, R., Liben-Nowell, D., Karger, D., Kaashoek, M., Dabek, F., and Balakrishnan, H. (2003). Chord: a scalable peer-to-peer lookup protocol for internet applications. *IEEE/ACM Trans. Netw.*, 11(1):17–32. 433
- Stonebraker, M. (1975). Implementation of integrity constraints and views by query modification. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 65–78. 123, 124
- Stonebraker, M. (1981). Operating system support for database management. *Commun. ACM*, 24(7):412–418. 31, 497, C-40
- Stonebraker, M. (1986). The case for shared nothing. *Q. Bull. IEEE TC on Data Eng.*, 9(1):4–9. 381
- Stonebraker, M. and Neuhold, E. (1977). A distributed database version of INGRES. In *Proc. 2nd Berkeley Workshop on Distributed Data Management and Computer Networks*, pages 9–36. 31, 273
- Stonebraker, M., Kreps, P., Wong, W., and Held, G. (1976). The design and implementation of ingres. *ACM Trans. Database Syst.*, 1(3):198–222. B-24
- Stonebraker, M., Abadi, D. J., DeWitt, D. J., Madden, S., Paulson, E., Pavlo, A., and Rasin, A. (2010). MapReduce and parallel DBMSs: friends or foes? *Commun. ACM*, 53(1):64–71. 497
- Strauch, C. (2011). *NoSQL Databases*. Stuttgart Media University. 535
- Sullivan, M. and Heybey, A. (1998). Tribeca: A system for managing large databases of network traffic. In *Proc. USENIX 1998 Annual Technical Conf.* 499
- Swami, A. (1989). Optimization of large join queries: combining heuristics and combinatorial techniques. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 367–376. 176, 381, B-2, B-17
- Taft, R., Mansour, E., Serafini, M., Duggan, J., Elmore, A. J., Abounaga, A., Pavlo, A., and Stonebraker, M. (2014). E-Store: Fine-Grained Elastic Partitioning for Distributed Transaction Processing. *Proc. VLDB Endowment*, 8(3):245–256. 85
- Taft, R., El-Sayed, N., Serafini, M., Lu, Y., Abounaga, A., Stonebraker, M., Mayerhofer, R., and Andrade, F. (2018). P-store: An elastic database system with predictive provisioning. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 205–219. 86
- Tandem. (1987). NonStop SQL – a distributed high-performance, high-availability implementation of sql. In *Proc. Int. Workshop on High Performance Transaction Systems*, pages 60–104. 237
- Tandem. (1988). A benchmark of NonStop SQL on the debit credit transaction. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 337–341. 237
- Tanenbaum, A. S. (2003). *Computer Networks*. Prentice-Hall, 4th edition. D-2
- Tanenbaum, A. S. and van Renesse, R. (1988). Voting with ghosts. In *Proc. 8th IEEE Int. Conf. on Distributed Computing Systems*, pages 456–461. 273
- Tang, W., Zhao, X., Rafique, W., Qi, L., Dou, W., and Ni, Q. (2019). An offloading method using decentralized P2P-enabled mobile edge servers in edge computing. *Journal of Systems Architecture – Embedded Systems Design*, 94:1–13. 434
- Tatarinov, I., Ives, Z. G., Madhavan, J., Halevy, A. Y., Suciu, D., Dalvi, N. N., Dong, X., Kadiyska, Y., Miklau, G., and Mork, P. (2003). The piazza peer data management project. *ACM SIGMOD Rec.*, 32(3):47–52. 433
- Tatbul, N., Cetintemel, U., Zdonik, S., Cherniack, M., and Stonebraker, M. (2003). Load shedding in a data stream manager. In *Proc. 29th Int. Conf. on Very Large Data Bases*, pages 309–320. 498

- Thiran, P., Hainaut, J.-L., Houben, G.-J., and Benslimane, D. (2006). Wrapper-based evolution of legacy information systems. *ACM Trans. Softw. Eng. and Meth.*, 15(4):329–359. 327
- Thomas, R. H. (1979). A majority consensus approach to concurrency control for multiple copy databases. *ACM Trans. Database Syst.*, 4(2):180–209. 237, 273
- Thomasian, A. (1996). *Database Concurrency Control: Methods, Performance, and Analysis*. Kluwer Academic Publishers. 238
- Thomson, A. and Abadi, D. J. (2010). The case for determinism in database systems. *Proc. VLDB Endowment*, 3(1):70–80. 239
- Thuraisingham, B. (2001). Secure distributed database systems. *Information Security Technical Report*, 6(2). 124
- Thusoo, A., Sarma, J. S., Jain, N., Shao, Z., Chakka, P., Anthony, S., Liu, H., Wyckoff, P., and Murthy, R. (2009). Hive: a warehousing solution over a map-reduce framework. *Proc. VLDB Endowment*, 2(2):1626–1629. 497
- Tian, F. and DeWitt, D. J. (2003). Tuple routing strategies for distributed eddies. In *Proc. 29th Int. Conf. on Very Large Data Bases*, pages 333–344. 177, 328
- Tian, Y., Balmin, A., Corsten, S. A., Tatikonda, S., and McPherson, J. (2013). From “think like a vertex” to “think like a graph”. *Proc. VLDB Endowment*, 7(3):193–204. 500
- Tian, Y., Özcan, F., Zou, T., Goncalves, R., and Pirahesh, H. (2016). Building a hybrid warehouse: Efficient joins between data stored in HDFS and enterprise warehouse. *ACM Trans. Database Syst.*, 41(4):21:1–21:38. 536
- Tomasic, A., Raschid, L., and Valduriez, P. (1996). Scaling heterogeneous databases and the design of disco. In *Proc. 16th IEEE Int. Conf. on Distributed Computing Systems*, pages 449–457. 328
- Tomasic, A., Amouroux, R., Bonnet, P., Kapitskaia, O., Naacke, H., and Raschid, L. (1997). The distributed information search component (disco) and the world-wide web – prototype demonstration. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 546–548. 327
- Tomasic, A., Raschid, L., and Valduriez, P. (1998). Scaling access to distributed heterogeneous data sources with Disco. In *IEEE Trans. Knowl. and Data Eng.* in press. 328
- Toshniwal, A., Taneja, S., Shukla, A., Ramasamy, K., Patel, J. M., Kulkarni, S., Jackson, J., Gade, K., Fu, M., Donham, J., Bhagat, N., Mittal, S., and Ryaboy, D. (2014). Storm@twitter. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 147–156. 499
- Traiger, I. L., Gray, J., Galtieri, C. A., and Lindsay, B. G. (1982). Transactions and recovery in distributed database systems. *ACM Trans. Database Syst.*, 7(3):323–342. 239
- Triantafillou, P. and Taylor, D. J. (1995). The location-based paradigm for replication: Achieving efficiency and availability in distributed systems. *IEEE Trans. Softw. Eng.*, 21(1):1–18. 273
- Tu, S., Zheng, W., Kohler, E., Liskov, B., and Madden, S. (2013). Speedy transactions in multicore in-memory databases. In *Proc. 24th ACM Symp. on Operating System Principles*, pages 18–32. 239
- Tucker, P., Maier, D., Sheard, T., and Faragas, L. (2003). Exploiting punctuation semantics in continuous data streams. *IEEE Trans. Knowl. and Data Eng.*, 15(3):555–568. 498

- Ugander, J. and Backstrom, L. (2013). Balanced label propagation for partitioning massive graphs. In *Proc. 6th ACM Int. Conf. Web Search and Data Mining*, pages 507–516. 499
- Ullman, J. (1997). Information integration using logical views. In *Proc. 6th Int. Conf. on Database Theory*, pages 19–40. 328
- Ullman, J. D. (1982). *Principles of Database Systems*. Computer Science Press, 2nd edition. 177, B-6, B-10, B-12
- Ullman, J. D. (1988). *Principles of Database and Knowledge-Base Systems*, volume 1. Computer Science Press. C-2
- Ulusoy, Ö. (2007). Research issues in peer-to-peer data management. In *Proc. 22nd Int. Symp. on Computer and Information Science*, pages 1–8. 433
- Umbrich, J., Hose, K., Karnstedt, M., Harth, A., and Polleres, A. (2011). Comparing data summaries for processing live queries over linked data. *World Wide Web J.*, 14(5-6):495–544. 594
- Urhan, T. and Franklin, M. (2000). XJoin: A reactively-scheduled pipelined join operator. *Q. Bull. IEEE TC on Data Eng.*, 23:27. 498
- Urhan, T., Franklin, M. J., and Amsaleg, L. (1998). Cost based query scrambling for initial delays. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 130–141. 177, 328
- Valduriez, P. (1982). Semi-join algorithms for distributed database machines. In Schneider, J.-J., editor, *Distributed Data Bases*. pages 23–37. 177
- Valduriez, P. (1993). Parallel database systems: Open problems and new issues. *Distrib. Parall. Databases*, 1:137–16. 381
- Valduriez, P. and Gardarin, G. (1984). Join and semi-join algorithms for a multi processor database machine. *ACM Trans. Database Syst.*, 9(1):133–161. 381
- Valduriez, P. and Pacitti, E. (2004). Data management in large-scale P2P systems. In *Proc. 6th Int. Conf. High Performance Comp. for Computational Sci.*, pages 104–118. 433
- Valiant, L. G. (1990). A bridging model for parallel computation. *Commun. ACM*, 33(8):103–111. 500
- van Hee, K. (2002). *Workflow Management*. M.I.T. Press. 238
- Van Renesse, R. and Altinbuken, D. (2015). Paxos made moderately complex. *ACM Comput. Surv.*, 47(3):42:1–42:36. 239
- Varadarajan, R., Rivera-Vega, P., and Navathe, S. B. (1989). Data redistribution scheduling in fully connected networks. In *Proc. 27th Annual Allerton Conf. on Communication, Control, and Computing*. 85
- Velegarakis, Y., Miller, R. J., and Popa, L. (2004). Preserving mapping consistency under schema changes. *VLDB J.*, 13(3):274–293. 327, 329
- Verhofstadt, J. S. (1978). Recovery techniques for database systems. *ACM Comput. Surv.*, 10(2):168–195. 31, 239, C-44
- Verma, S., Leslie, L. M., Shin, Y., and Gupta, I. (2017). An experimental comparison of partitioning strategies in distributed graph processing. *Proc. VLDB Endowment*, 10(5):493–504. 499
- Vermeer, M. (1997). *Semantic Interoperability for Legacy Databases*. Ph.D. thesis, Department of Computer Science, University of Twente, Enschede, Netherlands. 327
- Viglas, S., Naughton, J., and Burger, J. (2003). Maximizing the output rate of multi-join queries over streaming information sources. In *Proc. 29th Int. Conf. on Very Large Data Bases*, pages 285–296. 498

- Voulgaris, S., Jelasity, M., and van Steen, M. (2003). A robust and scalable peer-to-peer gossiping protocol. In *Agents and Peer-to-Peer Computing, Second Int. Workshop, (AP2PC)*, pages 47–58. 433
- Vu, Q. H., Lupu, M., and Ooi, B. C. (2009). *Peer-to-Peer Computing: Principles and Applications*. Springer. 433
- Wah, B. W. and Lien, Y. N. (1985). Design of distributed databases on local computer systems. *IEEE Trans. Softw. Eng.*, SE-11(7):609–619. 177
- Walton, C., Dale, A., and Jenevin, R. (1991). A taxonomy and performance model of data skew effects in parallel joins. In *Proc. 17th Int. Conf. on Very Large Data Bases*, pages 537–548. 381, 382
- Wang, G., Xie, W., Demers, A. J., and Gehrke, J. (2013). Asynchronous large-scale graph processing made easy. In *Proc. 6th Biennial Conf. on Innovative Data Systems Research*. 500
- Wang, H., Zaniolo, C., and Luo, R. (2003). Atlas: A small but complete SQL extension for data mining and data streams. In *Proc. 29th Int. Conf. on Very Large Data Bases*, pages 1113–1116. 498
- Wang, L., Xiao, Y., Shao, B., and Wang, H. (2014). How to partition a billion-node graph. In *Proc. 30th Int. Conf. on Data Engineering*, pages 568–579. 499
- Wang, W., Li, J., Zhang, D., and Guo, L. (2004). Processing sliding window join aggregate in continuous queries over data streams. In *Proc. 8th East European Conf. Advances in Databases and Information Systems*, pages 348–363. 498
- Weikum, G. and Vossen, G. (2001). *Transactional Information Systems: Theory, Algorithms, and the Practice of Concurrency Control*. Morgan Kaufmann. 237
- Weil, S. A., Brandt, S. A., Miller, E. L., Long, D. D. E., and Maltzahn, C. (2006). Ceph: A scalable, high-performance distributed file system. In *Proc. 7th USENIX Symp. on Operating System Design and Implementation*, pages 307–320. 497
- Weiss, C., Karras, P., and Bernstein, A. (2008). Hexastore: sextuple indexing for semantic web data management. *Proc. VLDB Endowment*, 1(1):1008–1019. 593
- Wiederhold, G. (1992). Mediators in the architecture of future information systems. *Computer*, 25(3):38–49. 25, 31, 328
- Wiesmann, M., Schiper, A., Pedone, F., Kemme, B., and Alonso, G. (2000). Database replication techniques: A three parameter classification. In *Proc. 28th Symp. on Reliable Distributed Systems*, pages 206–215. 272
- Wilkinson, K. (2006). Jena property table implementation. Technical Report HPL-2006-140, HP Laboratories Palo Alto. 593
- Wilms, P. F. and Lindsay, B. G. (1981). A database authorization mechanism supporting individual and group authorization. Research Report RJ 3137, IBM Almaden Research Laboratory, San Jose, Calif. 124
- Wilschut, A. and Apers, P. (1991). Dataflow query execution in a parallel main-memory environment. In *Proc. 1st Int. Conf. on Parallel and Distributed Information Systems*, pages 68–77. 381, 498
- Wilson, B. and Navathe, S. B. (1986). An analytical framework for the redesign of distributed databases. In *Proc. 6th Advanced Database Symposium*, pages 77–83. 85
- Wolfson, O. (1987). The overhead of locking (and commit) protocols in distributed databases. *ACM Trans. Database Syst.*, 12(3):453–471. 273
- Wong, E. (1977). Retrieving dispersed data from SDD-1. In *Proc. 2nd Berkeley Workshop on Distributed Data Management and Computer Networks*, pages 217–235. 31, 176

- Wong, E. and Youssefi, K. (1976). Decomposition: A strategy for query processing. *ACM Trans. Database Syst.*, 1(3):223–241. 177, B-24
- Wright, D. D. (1983). Managing distributed databases in partitioned networks. Technical Report TR83-572, Department of Computer Science, Cornell University, Ithaca, N.Y. 273
- Wu, E., Diao, Y., and Rizvi, S. (2006). High-performance complex event processing over streams. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 407–418. 498
- Wu, K.-L., Yu, P. S., and Pu, C. (1997). Divergence control algorithms for epsilon serializability. *IEEE Trans. Knowl. and Data Eng.*, 9(2):262–274. 238, 272
- Wu, S., Yu, G., Yu, Y., Ou, Z., Yang, X., and Gu, Y. (2005). A deadline-sensitive approach for real-time processing of sliding windows. In *Proc. 6th Int. Conf. on Web-Age Information Management*, pages 566–577. 498
- Xing, Y., Hwang, J.-H., Çetintemel, U., and Zdonik, S. (2006). Providing resiliency to load variations in distributed stream processing. In *Proc. 32nd Int. Conf. on Very Large Data Bases*, pages 775–786. 499
- Yan, D., Cheng, J., Lu, Y., and Ng, W. (2014). Blogel: A block-centric framework for distributed computation on real-world graphs. *Proc. VLDB Endowment*, 7(14):1981–1992. 500
- Yan, D., Bu, Y., Tian, Y., and Deshpande, A. (2017). Big graph analytics platforms. *Foundations and Trends in Databases*, 7(1-2):1–195. 499
- Yan, L. L. (1997). Towards efficient and scalable mediation: The AURORA approach. In *Proc. IBM CASCON Conference*, pages 15–29. 329
- Yan, L.-L., Özsü, M. T., and Liu, L. (1997). Accessing heterogeneous data through homogenization and integration mediators. In *Proc. Int. Conf. on Cooperative Inf. Syst.*, pages 130–139. 329
- Yan, L. L., Miller, R. J., Haas, L. M., and Fagin, R. (2001). Data-driven understanding and refinement of schema mappings. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 485–496. 327
- Yang, B. and Garcia-Molina, H. (2002). Improving search in peer-to-peer networks. In *Proc. 22nd IEEE Int. Conf. on Distributed Computing Systems*, pages 5–14. 433
- Yang, X., Lee, M.-L., and Ling, T. W. (2003). Resolving structural conflicts in the integration of XML schemas: A semantic approach. In *Proc. 22nd Int. Conf. on Conceptual Modeling*, pages 520–533. 327
- Yao, S. B., Waddle, V., and Housel, B. (1982). View modeling and integration using the functional data model. *IEEE Trans. Softw. Eng.*, SE-8(6):544–554. 327
- Yu, C. and Meng, W. (1998). *Principles of Query Processing for Advanced Database Applications*. Morgan Kaufmann. 328
- Zaharia, M. (2016). *An Architecture for Fast and General Data Processing on Large Clusters*. ACM Books. 498
- Zaharia, M., Chowdhury, M., Franklin, M. J., Shenker, S., and Stoica, I. (2010). Spark: Cluster computing with working sets. In *Proc. 2nd USENIX Workshop on Hot Topics in Cloud Computing*, pages 10–10. 498
- Zaharia, M., Das, T., and Timothy Hunter, H. L., Shenker, S., and Stoica, I. (2013). Discretized streams: Fault-tolerant streaming computation at scale. In *Proc. 24th ACM Symp. on Operating System Principles*, pages 423–438. 498, 499
- Zhao, B., Huang, L., Stribling, J., Rhea, S., Joseph, A. D., and Kubiatowicz, J. (2004). Tapestry: A resilient global-scale overlay for service deployment. *IEEE J. Selected Areas in Comm.*, 22(1):41–53. 433

- Zhu, M. and Risch, T. (2011). Querying combined cloud-based and relational databases. In *Proc. 2011 Int. Conf. on Cloud and Service Comp.*, pages 330–335. 536
- Zhu, Q. (1995). *Estimating Local Cost Parameters for Global Query Optimization in a Multidatabase System*. Ph.D. thesis, Department of Computer Science, University of Waterloo, Waterloo, Canada. 328
- Zhu, Q. and Larson, P.-Å. (1994). A query sampling method of estimating local cost parameters in a multidatabase system. In *Proc. 10th Int. Conf. on Data Engineering*, pages 144–153. 328
- Zhu, Q. and Larson, P. A. (1996a). Global query processing and optimization in the CORDS multidatabase system. In *Proc. Int. Conf. on Parallel and Distributed Computing Systems*, pages 640–647. 328
- Zhu, Q. and Larson, P. A. (1996b). Developing regression cost models for multidatabase systems. In *Proc. 4th Int. Conf. on Parallel and Distributed Information Systems*, pages 220–231. 328
- Zhu, Q. and Larson, P. A. (1998). Solving local cost estimation problem for global query optimization in multidatabase systems. *Distrib. Paralle. Databases*, 6(4): 373–420. 328
- Zhu, Q., Sun, Y., and Motheramgari, S. (2000). Developing cost models with qualitative variables for dynamic multidatabase environments. In *Proc. 16th Int. Conf. on Data Engineering*, pages 413–424. 328
- Zhu, Q., Motheramgari, S., and Sun, Y. (2003). Cost estimation for queries experiencing multiple contention states in dynamic multidatabase environments. *Knowledge and Information Systems*, 5(1):26–49. 328
- Zhu, Y., Rundensteiner, E., and Heineman, G. (2004). Dynamic plan migration for continuous queries over data streams. In *Proc. ACM SIGMOD Int. Conf. on Management of Data*, pages 431–442. 499
- Zhu, Y., Zhang, H., Qin, L., and Cheng, H. (2017). Efficient MapReduce algorithms for triangle listing in billion-scale graphs. *Distrib. Paralle. Databases*, 35(2): 149–176. 499
- Ziane, M., Zaït, M., and Borla-Salamet, P. (1993). Parallel query processing with zigzag trees. *VLDB J.*, 2(3):277–301. 381
- Zilio, D. C. (1998). *Physical Database Design Decision Algorithms and Concurrent Reorganization for Parallel Database Systems*. PhD thesis, University of Toronto. 85
- Zou, L. and Özsu, M. T. (2017). Graph-based RDF data management. *Data Science and Engineering*, 2(1):56–70. 593
- Zou, L., Mo, J., Chen, L., Özsu, M. T., and Zhao, D. (2011). gStore: answering SPARQL queries via subgraph matching. *Proc. VLDB Endowment*, 4(8):482–493. 593
- Zou, L., Özsu, M. T., Chen, L., Shen, X., Huang, R., and Zhao, D. (2014). gStore: A graph-based SPARQL query engine. *VLDB J.*, 23(4):565–590. 593