Designing Data-Intensive Applications

บทที่ 1

---

[1] Richard T. Kouzes, Gordon A. Anderson, Stephen T. Elbert, Ian Gorton, andDeborah K. Gracio. “The Changing Paradigm of Data-Intensive Computing.” *IEEE Computer*, volume 42, issue 1, pages 26–34, January 2009. *doi:10.1109/MC.2009.26*

[2] Martin Kleppmann, Adam Wiggins, Peter van Hardenberg, and Mark McGra-naghan. “Local-First Software: You Own Your Data, in Spite of the Cloud.” At *2019 ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Onward!), October 2019. *doi:10.1145/3359591.3359737*

[3] Joe Reis and Matt Housley. *Fundamentals of Data Engineering*. O’Reilly Media, 2022. ISBN: 9781098108304

[4] Rui Pedro Machado and Helder Russa. *Analytics Engineering with SQL and dbt*. O’Reilly Media, 2023. ISBN: 9781098142384

[5] Edgar F. Codd, S. B. Codd, and C. T. Salley. “Providing OLAP to User-Analysts: An IT Mandate.” E. F. Codd Associates, 1993. Archived at *perma.cc/RKX8-2GEE*

[6] Chinmay Soman and Neha Pawar. “Comparing Three Real-Time OLAP Databases: Apache Pinot, Apache Druid, and ClickHouse.” *startree.ai*, April 2023. Archived at *perma.cc/8BZP-VWPA*

[7] Surajit Chaudhuri and Umeshwar Dayal. “An Overview of Data Warehousing and OLAP Technology.” *ACM SIGMOD Record*, volume 26, issue 1, pages 65–74, March 1997. *doi:10.1145/248603.248616*

[8] Fatma Özcan, Yuanyuan Tian, and Pinar Tözün. “Hybrid Transactional/Analytical Processing: A Survey.” At *ACM International Conference on Management of Data* (SIGMOD), May 2017. *doi:10.1145/3035918.3054784*

[9] Adam Prout, Szu-Po Wang, Joseph Victor, Zhou Sun, Yongzhu Li, Jack Chen,Evan Bergeron, Eric Hanson, Robert Walzer, Rodrigo Gomes, and Nikita Shamgunov. “Cloud-Native Transactions and Analytics in SingleStore.” At *International Conference on Management of Data* (SIGMOD), June 2022. *doi:10.1145/3514221.3526055*

[10] [10] Chao Zhang, Guoliang Li, Jintao Zhang, Xinning Zhang, and Jianhua Feng. “HTAP Databases: A Survey.” *IEEE Transactions on Knowledge and Data Engineering*, volume 36, issue 11, pages 6410–6429, April 2024. *doi:10.1109/TKDE.2024.3389693*

[11] Michael Stonebraker and Uğur Çetintemel. “‘One Size Fits All’: An Idea Whose Time Has Come and Gone.” At *21st International Conference on Data Engineering* (ICDE), April 2005. *doi:10.1109/ICDE.2005.1*

[12] Jeffrey Cohen, Brian Dolan, Mark Dunlap, Joseph M. Hellerstein, andCaleb Welton. “MAD Skills: New Analysis Practices for Big Data.” *Proceedings of the VLDB Endowment*, volume 2, issue 2, pages 1481–1492, August 2009. *doi:10.14778/1687553.1687576*

[13] Dan Olteanu. “The Relational Data Borg Is Learning.” *Proceedings of the VLDB Endowment*, volume 13, issue 12, pages 3502–3515, August 2020. *doi:10.14778/3415478.3415572*

[14] Matt Bornstein, Martin Casado, and Jennifer Li. “Emerging Architectures for Modern Data Infrastructure: 2020.” *future.a16z.com*, October 2020. Archived at *perma.cc/LF8W-KDCC*

[15] Rihan Hai, Christos Koutras, Christoph Quix, and Matthias Jarke. “Data Lakes: A Survey of Functions and Systems.” *IEEE Transactions on Knowledge and Data Engineering* (TKDE), volume 35, issue 12, pages 12571–12590, December 2023. *doi:10.1109/TKDE.2023.3270101*

[16] Martin Fowler. “Data Lake.” *martinfowler.com*, February 2015. Archived at *perma.cc/4WKN-CZUK*

[17] Bobby Johnson and Joseph Adler. “The Sushi Principle: Raw Data Is Better.” At *Strata+Hadoop World*, February 2015.

[18] DataKitchen, Inc. “The DataOps Manifesto.” *dataopsmanifesto.org*, 2017. Archived at *perma.cc/3F5N-FUQ4*

[19] Tejas Manohar. “What Is Reverse ETL: A Definition & Why It’s Taking Off.” *hightouch.io*, November 2021. Archived at *perma.cc/A7TN-GLYJ*

[20] Camille Fournier. “Why Is It So Hard to Decide to Buy?” *skamille.medium.com*, July 2021. Archived at *perma.cc/6VSG-HQ5X*

[21] David Heinemeier Hansson. “Why We’re Leaving the Cloud.” *world.hey.com*, October 2022. Archived at *perma.cc/82E6-UJ65*

[22] Nima Badizadegan. “Use One Big Server.” *specbranch.com*, August 2022. Archived at *perma.cc/M8NB-95UK*

[23] Steve Yegge. “Dear Google Cloud: Your Deprecation Policy Is Killing You.” *steve-yegge.medium.com*, August 2020. Archived at *perma.cc/KQP9-SPGU*

[24] Alexandre Verbitski, Anurag Gupta, Debanjan Saha, Murali Brahmadesam,Kamal Gupta, Raman Mittal, Sailesh Krishnamurthy, Sandor Maurice, Tengiz Kharatishvili, and Xiaofeng Bao. “Amazon Aurora: Design Considerations for High Throughput Cloud-Native Relational Databases.” At *ACM International Conference on Management of Data* (SIGMOD), May 2017. *doi:10.1145/3035918.3056101*

[25] Panagiotis Antonopoulos, Alex Budovski, Cristian Diaconu, Alejandro Hernandez Saenz, Jack Hu, Hanuma Kodavalla, Donald Kossmann, Sandeep Lingam, Umar Farooq Minhas, Naveen Prakash, Vijendra Purohit, Hugh Qu, Chaitanya Sreenivas Ravella, Krystyna Reisteter, Sheetal Shrotri, Dixin Tang, and Vikram Wakade. “Socrates: The New SQL Server in the Cloud.” At *ACM International Conference on Management of Data* (SIGMOD), June 2019. *doi:10.1145/3299869.3314047*

[26] Midhul Vuppalapati, Justin Miron, Rachit Agarwal, Dan Truong, Ashish Moti-vala, and Thierry Cruanes. “Building an Elastic Query Engine on Disaggregated Storage.” At *17th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), February 2020.

[27] Nick Van Wiggeren. “The Real Failure Rate of EBS.” *planetscale.com*, March 2025. Archived at *perma.cc/43CR-SAH5*

[28] Colin Breck. “Predicting the Future of Distributed Systems.” *blog.colinbreck.com*, August 2024. Archived at *perma.cc/K5FC-4XX2*

[29] Gwen Shapira. “Compute-Storage Separation Explained.” *thenile.dev*, January 2023. Archived at perma.cc/QCV3-XJNZ

[30] Ravi Murthy and Gurmeet Goindi. “AlloyDB for PostgreSQL Under the Hood: Intelligent, Database-Aware Storage.” *cloud.google.com*, May 2022. Archived at archive.org

[31] Jack Vanlightly. “The Architecture of Serverless Data Systems.” *jackvanlightly.com*, November 2023. Archived at *perma.cc/UDV4-TNJ5*

[32] Eric Jonas, Johann Schleier-Smith, Vikram Sreekanti, Chia-Che Tsai, AnuragKhandelwal, Qifan Pu, Vaishaal Shankar, Joao Carreira, Karl Krauth, Neeraja Yadwadkar, Joseph E. Gonzalez, Raluca Ada Popa, Ion Stoica, and David A. Patterson. “Cloud Programming Simplified: A Berkeley View on Serverless Computing.” *arXiv:1902.03383*, February 2019.

[33] Betsy Beyer, Jennifer Petoff, Chris Jones, and Niall Richard Murphy. *Site Reliability Engineering: How Google Runs Production Systems*. O’Reilly Media, 2016. ISBN: 9781491929124

[34] Thomas Limoncelli. “The Time I Stole $10,000 from Bell Labs.” *ACM Queue*, volume 18, issue 5, November 2020. *doi:10.1145/3434571.3434773*

[35] Charity Majors. “The Future of Ops Jobs.” *acloudguru.com*, August 2020. Archived at *perma.cc/GRU2-CZG3*

[36] Boris Cherkasky. “(Over)Pay as You Go for Your Datastore.” *medium.com*, September 2021. Archived at *perma.cc/Q8TV-2AM2*

[37] Shlomi Kushchi. “Serverless Doesn’t Mean DevOpsLess or NoOps.” *thenewstack.io*, February 2023. Archived at *perma.cc/3NJR-AYYU*

[38] Erik Bernhardsson. “Storm in the Stratosphere: How the Cloud Will Be Reshuffled.” *erikbern.com*, November 2021. Archived at *perma.cc/SYB2-99P3*

[39] Benn Stancil. “The Data OS.” *benn.substack.com*, September 2021. Archived at *perma.cc/WQ43-FHS6*

[40] Maria Korolov. “Data Residency Laws Pushing Companies Toward Residency as a Service.” *csoonline.com*, January 2022. Archived at *perma.cc/CHE4-XZZ2*

[41] Severin Borenstein. “Can Data Centers Flex Their Power Demand?” *energyathaas.wordpress.com*, April 2025. Archived at *perma.cc/MUD3-A6FF*

[42] Bilge Acun, Benjamin Lee, Fiodar Kazhamiaka, Aditya Sundarrajan, KiwanMaeng, Manoj Chakkaravarthy, David Brooks, and Carole-Jean Wu. “Carbon Dependencies in Datacenter Design and Management.” *ACM SIGENERGY Energy Informatics Review*, volume 3, issue 3, pages 21–26, October 2023. *doi:10.1145/3630614.3630619*

[43] Kousik Nath. “These Are the Numbers Every Computer Engineer Should Know.” *freecodecamp.org*, September 2019. Archived at *perma.cc/RW73-36RL*

[44] Joseph M. Hellerstein, Jose Faleiro, Joseph E. Gonzalez, Johann Schleier-Smith,Vikram Sreekanti, Alexey Tumanov, and Chenggang Wu. “Serverless Computing: One Step Forward, Two Steps Back.” *arXiv:1812.03651*, December 2018.

[45] Frank McSherry, Michael Isard, and Derek G. Murray. “Scalability! But at What COST?” At *15th USENIX Workshop on Hot Topics in Operating Systems* (HotOS), May 2015.

[46] Cindy Sridharan. *Distributed Systems Observability: A Guide to Building Robust Systems*. Report, O’Reilly Media, 2018. Archived at *perma.cc/M6JL-XKCM*

[47] Charity Majors. “Observability—A 3-Year Retrospective.” *thenewstack.io*, August 2019. Archived at *perma.cc/CG62-TJWL*

[48] Benjamin H. Sigelman, Luiz André Barroso, Mike Burrows, Pat Stephenson,Manoj Plakal, Donald Beaver, Saul Jaspan, and Chandan Shanbhag. “Dapper, a Large-Scale Distributed Systems Tracing Infrastructure.” Google Technical Report dapper-2010-1, April 2010. Archived at *perma.cc/K7KU-2TMH*

[49] Rodrigo Laigner, Yongluan Zhou, Marcos Antonio Vaz Salles, Yijian Liu, andMarcos Kalinowski. “Data Management in Microservices: State of the Practice, Challenges, and Research Directions.” *Proceedings of the VLDB Endowment*, volume 14, issue 13, pages 3348–3361, September 2021. *doi:10.14778/3484224.3484232*

[50] Jordan Tigani. “Big Data Is Dead.” *motherduck.com*, February 2023. Archived at *perma.cc/HT4Q-K77U*

[51] Sam Newman. *Building Microservices*, 2nd edition. O’Reilly Media, 2021. ISBN: 9781492034025

[52] Chris Richardson. “Microservices: Decomposing Applications for Deployability and Scalability.” *infoq.com*, May 2014. Archived at *perma.cc/CKN4-YEQ2*

[53] Mohammad Shahrad, Rodrigo Fonseca, Íñigo Goiri, Gohar Chaudhry, PaulBatum, Jason Cooke, Eduardo Laureano, Colby Tresness, Mark Russinovich, and Ricardo Bianchini. “Serverless in the Wild: Characterizing and Optimizing the Serverless Workload at a Large Cloud Provider.” At *USENIX Annual Technical Conference* (ATC), July 2020.

[54] Luiz André Barroso, Urs Hölzle, and Parthasarathy Ranganathan. *The Datacenter as a Computer: Designing Warehouse-Scale Machines*, 3rd edition. Springer Nature, 2019. ISBN: 9783031017612

[55] David Fiala, Frank Mueller, Christian Engelmann, Rolf Riesen, Kurt Ferreira,and Ron Brightwell. “Detection and Correction of Silent Data Corruption for Large-Scale High-Performance Computing.” At *International Conference for High Performance Computing, Networking, Storage and Analysis* (SC), November 2012.
doi:10.1109/SC.2012.49

[56] Anna Kornfeld Simpson, Adriana Szekeres, Jacob Nelson, and Irene Zhang.“Securing RDMA for High-Performance Datacenter Storage Systems.” At *12th USENIX Workshop on Hot Topics in Cloud Computing* (HotCloud), July 2020.

[57] Arjun Singh, Joon Ong, Amit Agarwal, Glen Anderson, Ashby Armistead,Roy Bannon, Seb Boving, Gaurav Desai, Bob Felderman, Paulie Germano, Anand Kanagala, Jeff Provost, Jason Simmons, Eiichi Tanda, Jim Wanderer, Urs Hölzle, Stephen Stuart, and Amin Vahdat. “Jupiter Rising: A Decade of Clos Topologies and Centralized Control in Google’s Datacenter Network.” At *Annual Conference of the ACM Special Interest Group on Data Communication* (SIGCOMM), August 2015. *doi:10.1145/2785956.2787508*

[58] Glenn K. Lockwood. “Hadoop’s Uncomfortable Fit in HPC.” *glennklockwood.blogspot.co.uk*, May 2014. Archived at *perma.cc/S8XX-Y67B*

[59] Cathy O’Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishing, 2016. ISBN: 9780553418811

[60] Supreeth Shastri, Vinay Banakar, Melissa Wasserman, Arun Kumar, and VijayChidambaram. “Understanding and Benchmarking the Impact of GDPR on Database Systems.” *Proceedings of the VLDB Endowment*, volume 13, issue 7, pages 1064–1077, March 2020. *doi:10.14778/3384345.3384354*

[61] Martin Fowler. “Datensparsamkeit.” *martinfowler.com*, December 2013. Archived at *perma.cc/R9QX-CME6*

[62] “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016 (General Data Protection Regulation).” *Official Journal of the European Union* L 119/1, May 2016.

บทที่ 2

[1] Mike Cvet. “How We Learned to Stop Worrying and Love Fan-in at Twitter.” At *QCon San Francisco*, December 2016.

[2] Raffi Krikorian. “Timelines at Scale.” At *QCon San Francisco*, November 2012. Archived at *perma.cc/V9G5-KLYK*

[3] Twitter. “Twitter’s Recommendation Algorithm.” *blog.x.com*, March 2023. Archived at *perma.cc/L5GT-229T*

[4] Raffi Krikorian. “New Tweets per Second Record, and How!” *blog.x.com*, August 2013. Archived at *perma.cc/6JZN-XJYN*

[5] Jaz Volpert. “When Imperfect Systems Are Good, Actually: Bluesky’s Lossy Timelines.” *jazco.dev*, February 2025. Archived at *perma.cc/2PVE-L2MX*

[6] Samuel Axon. “3% of Twitter’s Servers Dedicated to Justin Bieber.” *mashable.com*, September 2010. Archived at *perma.cc/F35N-CGVX*

[7] Nathan Bronson, Abutalib Aghayev, Aleksey Charapko, and Timothy Zhu. “Metastable Failures in Distributed Systems.” At *Workshop on Hot Topics in Operating Systems* (HotOS), May 2021. *doi:10.1145/3458336.3465286*

[8] Marc Brooker. “Metastability and Distributed Systems.” *brooker.co.za*, May 2021. Archived at *perma.cc/7FGJ-7XRK*

[9] Lexiang Huang, Matthew Magnusson, Abishek Bangalore Muralikrishna, SalmanEstyak, Rebecca Isaacs, Abutalib Aghayev, Timothy Zhu, and Aleksey Charapko. "Metastable Failures in the Wild." At *16th USENIX Symposium on Operating Systems Design and Implementation* (OSDI), July 2022.

[10] Marc Brooker. "Exponential Backoff and Jitter." *aws.amazon.com*, March 2015. Archived at *perma.cc/R6MS-AZKH*

[11] Marc Brooker. "What Is Backoff For?" *brooker.co.za*, August 2022. Archived at *perma.cc/PW9N-55Q5*

[12] Michael T. Nygard. *Release It!*, 2nd edition. Pragmatic Bookshelf, 2018. ISBN: 9781680502398

[13] Frank Chen. "Slowing Down to Speed Up—Circuit Breakers for Slack's CI/CD." *slack.engineering*, August 2022. Archived at *perma.cc/5FGS-ZPH3*

[14] Marc Brooker. "Fixing Retries with Token Buckets and Circuit Breakers." *brooker.co.za*, February 2022. Archived at *perma.cc/MD6N-GW26*

[15] David Yanacek. "Using Load Shedding to Avoid Overload." Amazon Builders' Library, *aws.amazon.com*. Archived at *perma.cc/9SAW-68MP*

[16] Matthew Sackman. "Pushing Back." *wellquite.org*, May 2016. Archived at *perma.cc/3KCZ-RUFY*

[17] Dmitry Kopytkov and Patrick Lee. "Meet Bandaid, the Dropbox Service Proxy." *dropbox.tech*, March 2018. Archived at *perma.cc/KUU6-YG4S*

[18] Haryadi S. Gunawi, Riza O. Suminto, Russell Sears, Casey Golliher, SwaminathanSundararaman, Xing Lin, Tim Emami, Weiguang Sheng, Nematollah Bidokhti, Caitie McCaffrey, Gary Grider, Parks M. Fields, Kevin Harms, Robert B. Ross, Andree Jacobson, Robert Ricci, Kirk Webb, Peter Alvaro, H. Birali Runesha, Mingzhe Hao, and Huaicheng Li. "Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems." At *16th USENIX Conference on File and Storage Technologies*, February 2018.

[19] Marc Brooker. "Is the Mean Really Useless?" *brooker.co.za*, December 2017. Archived at *perma.cc/U5AE-CVEM*

[20] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakula-pati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. "Dynamo: Amazon's Highly Available Key-Value Store." At *21st ACM Symposium on Operating Systems Principles* (SOSP), October 2007. *doi:10.1145/1294261.1294281*

[21] Kathryn Whitenton. "The Need for Speed, 23 Years Later." *nngroup.com*, May 2020. Archived at *perma.cc/C4ER-LZYA*

[22] Greg Linden. "Marissa Mayer at Web 2.0." *glinden.blogspot.com*, November 2005. Archived at *perma.cc/V7EA-3VXB*

[23] Jake Brutlag. "Speed Matters for Google Web Search." *services.google.com*, June 2009. Archived at *perma.cc/BK7R-X7M2*

[24] Eric Schurman and Jake Brutlag. "Performance Related Changes and Their User Impact." Talk at *Velocity 2009*.

[25] Akamai Technologies, Inc. "The State of Online Retail Performance." *akamai.com*, April 2017. Archived at *perma.cc/UEK2-HYCS*

[26] Xiao Bai, Ioannis Arapakis, B. Barla Cambazoglu, and Ana Freire. "Understanding and Leveraging the Impact of Response Latency on User Behaviour in Web Search." *ACM Transactions on Information Systems*, volume 36, issue 2, article 21, April 2018. *doi:10.1145/3106372*

[27] Jeffrey Dean and Luiz André Barroso. "The Tail at Scale." *Communications of the ACM*, volume 56, issue 2, pages 74–80, February 2013. *doi:10.1145/2408776.2408794*

[28] Alex Hidalgo. *Implementing Service Level Objectives: A Practical Guide to SLIs, SLOs, and Error Budgets*. O'Reilly Media, 2020. ISBN: 9781492076813

[29] Jeffrey C. Mogul and John Wilkes. “Nines Are Not Enough: Meaningful Metrics for Clouds.” At *17th Workshop on Hot Topics in Operating Systems* (HotOS), May 2019. *doi:10.1145/3317550.3321432*

[30] Tamás Hauer, Philipp Hoffmann, John Lunney, Dan Ardelean, and Amer Diwan.“Meaningful Availability.” At *17th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), February 2020.

[31] Gil Tene. “HdrHistogram: A High Dynamic Range Histogram.” *hdrhistogram.github.io/HdrHistogram*

[32] Ted Dunning. “The *t*-digest: Efficient Estimates of Distributions.” *Software Impacts*, volume 7, article 100049, February 2021. *doi:10.1016/j.simpa.2020.100049*

[33] David Kohn. “How Percentile Approximation Works (and Why It’s More Useful than Averages).” *timescale.com*, September 2021. Archived at *perma.cc/3PDP-NR8B*

[34] Heinrich Hartmann and Theo Schlossnagle. “Circllhist—A Log-Linear Histogram Data Structure for IT Infrastructure Monitoring.” *arXiv:2001.06561*, January 2020.

[35] Charles Masson, Jee E. Rim, and Homin K. Lee. “DDSketch: A Fast and Fully-Mergeable Quantile Sketch with Relative-Error Guarantees.” *Proceedings of the VLDB Endowment*, volume 12, issue 12, pages 2195–2205, August 2019. *doi:10.14778/3352063.3352135*

[36] Baron Schwartz. “Why Percentiles Don’t Work the Way You Think.” *solarwinds.com*, November 2016. Archived at *perma.cc/469T-6UGB*

[37] Walter L. Heimerdinger and Charles B. Weinstock. “A Conceptual Framework for System Fault Tolerance.” Technical Report CMU/SEI-92-TR-033, Software Engineering Institute, Carnegie Mellon University, October 1992. Archived at *perma.cc/GD2V-DMJW*

[38] Felix C. Gärtner. “Fundamentals of Fault-Tolerant Distributed Computing in Asynchronous Environments.” *ACM Computing Surveys*, volume 31, issue 1, pages 1–26, March 1999. *doi:10.1145/311531.311532*

[39] Algirdas Avižienis, Jean-Claude Laprie, Brian Randell, and Carl Landwehr.“Basic Concepts and Taxonomy of Dependable and Secure Computing.” *IEEE Transactions on Dependable and Secure Computing*, volume 1, issue 1, pages 11–33, January 2004. *doi:10.1109/TDSC.2004.2*

[40] Ding Yuan, Yu Luo, Xin Zhuang, Guilherme Renna Rodrigues, Xu Zhao, YongleZhang, Pranay U. Jain, and Michael Stumm. “Simple Testing Can Prevent Most Critical Failures: An Analysis of Production Failures in Distributed Data-Intensive Systems.” At *11th USENIX Symposium on Operating Systems Design and Implementation* (OSDI), October 2014.

[41] Casey Rosenthal and Nora Jones. *Chaos Engineering*. O’Reilly Media, 2020. ISBN: 9781492043867

[42] Eduardo Pinheiro, Wolf-Dietrich Weber, and Luiz Andre Barroso. “Failure Trends in a Large Disk Drive Population.” At *5th USENIX Conference on File and Storage Technologies* (FAST), February 2007.

[43] Bianca Schroeder and Garth A. Gibson. “Disk Failures in the Real World: What Does an Mttf of 1,000,000 Hours Mean to You?” At *5th USENIX Conference on File and Storage Technologies* (FAST), February 2007.

[44] Andy Klein. “Backblaze Drive Stats for Q2 2021.” *backblaze.com*, August 2021. Archived at *perma.cc/2943-UD5E*

[45] Iyswarya Narayanan, Di Wang, Myeongjae Jeon, Bikash Sharma, Laura Caul-field, Anand Sivasubramaniam, Ben Cutler, Jie Liu, Badriddine Khessib, and Kushagra Vaid. “SSD Failures in Datacenters: What? When? And Why?” At *9th ACM International on Systems and Storage Conference* (SYSTOR), June 2016. *doi:10.1145/2928275.2928278*

[46] Alibaba Cloud Storage Team. “Storage System Design Analysis: Factors Affecting NVMe SSD Performance (1).” *alibabacloud.com*, January 2019. Archived at *archive.org*

[47] Bianca Schroeder, Raghav Lagisetty, and Arif Merchant. “Flash Reliability in Production: The Expected and the Unexpected.” At *14th USENIX Conference on File and Storage Technologies* (FAST), February 2016.

[48] Jacob Alter, Ji Xue, Alma Dimnaku, and Evgenia Smirni. “SSD Failures in the Field: Symptoms, Causes, and Prediction Models.” At *International Conference for High Performance Computing, Networking, Storage and Analysis* (SC), November 2019. *doi:10.1145/3295500.3356172*

[49] Daniel Ford, François Labelle, Florentina I. Popovici, Murray Stokely, Van-AnhTruong, Luiz Barroso, Carrie Grimes, and Sean Quinlan. “Availability in Globally Distributed Storage Systems.” At *9th USENIX Symposium on Operating Systems Design and Implementation* (OSDI), October 2010.

[50] Kashi Venkatesh Vishwanath and Nachiappan Nagappan. “Characterizing Cloud Computing Hardware Reliability.” At *1st ACM Symposium on Cloud Computing* (SoCC), June 2010. *doi:10.1145/1807128.1807161*

[51] Peter H. Hochschild, Paul Turner, Jeffrey C. Mogul, Rama Govindaraju, Par-thasarathy Ranganathan, David E. Culler, and Amin Vahdat. “Cores That Don’t Count.” At *Workshop on Hot Topics in Operating Systems* (HotOS), June 2021. *doi:10.1145/3458336.3465297*

[52] Harish Dattatraya Dixit, Sneha Pendharkar, Matt Beadon, Chris Mason, TejasviChakravarthy, Bharath Muthiah, and Sriram Sankar. Silent Data Corruptions at Scale. *arXiv:2102.11245*, February 2021.

[53] Diogo Behrens, Marco Serafini, Sergei Arnautov, Flavio P. Junqueira, and Chris-tof Fetzer. “Scalable Error Isolation for Distributed Systems.” At *12th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), May 2015.

[54] Bianca Schroeder, Eduardo Pinheiro, and Wolf-Dietrich Weber. “DRAM Errors in the Wild: A Large-Scale Field Study.” At *11th International Joint Conference on Measurement and Modeling of Computer Systems* (SIGMETRICS), June 2009. *doi:10.1145/1555349.1555372*

[55] Yoongu Kim, Ross Daly, Jeremie Kim, Chris Fallin, Ji Hye Lee, DonghyukLee, Chris Wilkerson, Konrad Lai, and Onur Mutlu. “Flipping Bits in Memory Without Accessing Them: An Experimental Study of DRAM Disturbance Errors.” At *41st Annual International Symposium on Computer Architecture* (ISCA), June 2014. *doi:10.5555/2665671.2665726*

[56] Tim Bray. “Worst Case.” *tbray.org*, October 2021. Archived at *perma.cc/4QQM-RTHN*

[57] Sangeetha Abdu Jyothi. “Solar Superstorms: Planning for an Internet Apocalypse.” At *ACM SIGCOMM Conference*, August 2021. *doi:10.1145/3452296.3472916*

[58] Adrian Cockcroft. “Failure Modes and Continuous Resilience.” *adrianco.medium.com*, November 2019. Archived at *perma.cc/7SYS-BVJP*

[59] Shujie Han, Patrick P. C. Lee, Fan Xu, Yi Liu, Cheng He, and Jiongzhou Liu. “An In-Depth Study of Correlated Failures in Production SSD-Based Data Centers.” At *19th USENIX Conference on File and Storage Technologies* (FAST), February 2021.

[60] Edmund B. Nightingale, John R. Douceur, and Vince Orgovan. “Cycles, Cells and Platters: An Empirical Analysis of Hardware Failures on a Million Consumer PCs.” At *6th European Conference on Computer Systems* (EuroSys), April 2011. *doi:10.1145/1966445.1966477*

[61] Haryadi S. Gunawi, Mingzhe Hao, Tanakorn Leesatapornwongsa, Tiratat Patanaanake, Thanh Do, Jeffry Adityatama, Kurnia J. Eliazar, Agung Laksono, Jeffrey F. Lukman, Vincentius Martin, and Anang D. Satria. “What Bugs Live in the Cloud? A Study of 3000+ Issues in Cloud Systems.” At *5th ACM Symposium on Cloud Computing* (SoCC), November 2014. *doi:10.1145/2670979.2670986*

[62] Jay Kreps. “Getting Real About Distributed System Reliability.” *blog.empathybox.com*, March 2012. Archived at *perma.cc/9B5Q-AEBW*

[63] Nelson Minar. “Leap Second Crashes Half the Internet.” *somebits.com*, July 2012. Archived at *perma.cc/2WB8-D6EU*

[64] Hewlett Packard Enterprise. “Support Alerts—Customer Bulletin a00092491en_us.” *support.hpe.com*, November 2019. Archived at *perma.cc/ S5F6-7ZAC*

[65] Lorin Hochstein. “Awesome Limits.” *github.com*, November 2020. Archived at *perma.cc/3R5M-E5Q4*

[66] Caitie McCaffrey. “Clients Are Jerks: AKA How Halo 4 DoSed the Services at Launch & How We Survived.” *caitiem.com*, June 2015. Archived at *perma.cc/MXX4-W373*

[67] Lilia Tang, Chaitanya Bhandari, Yongle Zhang, Anna Karanika, Shuyang Ji,Indranil Gupta, and Tianyin Xu. “Fail Through the Cracks: Cross-System Interaction Failures in Modern Cloud Systems.” At *18th European Conference on Computer Systems* (EuroSys), May 2023. *doi:10.1145/3552326.3587448*

[68] Mike Ulrich. Addressing Cascading Failures. In Betsy Beyer, Jennifer Petoff, Chris Jones, and Niall Richard Murphy (ed). *Site Reliability Engineering: How Google Runs Production Systems*. O’Reilly Media, 2016. ISBN: 9781491929124

[69] Harri Faßbender. “Cascading Failures in Large-Scale Distributed Systems.” *blog.mi.hdm-stuttgart.de*, March 2022. Archived at *perma.cc/K7VY-YJRX*

[70] Richard I. Cook. “How Complex Systems Fail.” Cognitive Technologies Laboratory, April 2000. Archived at *perma.cc/RDS6-2YVA*

[71] David D. Woods. “STELLA: Report from the SNAFUcatchers Workshop on Coping with Complexity.” *snafucatchers.github.io*, March 2017. Archived at *archive.org*

[72] David Oppenheimer, Archana Ganapathi, and David A. Patterson. “Why Do Internet Services Fail, and What Can Be Done About It?” At *4th USENIX Symposium on Internet Technologies and Systems* (USITS), March 2003.

[73] Sidney Dekker. *The Field Guide to Understanding “Human Error”*, 3rd edition. CRC Press, 2017. ISBN: 9781472439055

[74] Sidney Dekker. *Drift into Failure: From Hunting Broken Components to Understanding Complex Systems*. CRC Press, 2011. ISBN: 9781315257396

[75] John Allspaw. “Blameless PostMortems and a Just Culture.” *etsy.com*, May 2012. Archived at *perma.cc/YMJ7-NTAP*

[76] Itzy Sabo. “Uptime Guarantees—A Pragmatic Perspective.” *world.hey.com*, March 2023. Archived at *perma.cc/F7TU-78JB*

[77] Michael Jurewitz. “The Human Impact of Bugs.” *jury.me*, March 2013. Archived at *perma.cc/5KQ4-VDYL*

[78] Mark Halper. “How Software Bugs Led to ‘One of the Greatest Miscarriages of Justice’ in British History.” *Communications of the ACM*, volume 68, issue 3, pages 12–14, January 2025. *doi:10.1145/3703779*

[79] Nicholas Bohm, James Christie, Peter Bernard Ladkin, Bev Littlewood, PaulMarshall, Stephen Mason, Martin Newby, Steven J. Murdoch, Harold Thimbleby, and Martyn Thomas. “The Legal Rule That Computers Are Presumed to be Operating Correctly—Unforeseen and Unjust Consequences.” Briefing note, *benthamsgaze.org*, June 2022. Archived at *perma.cc/WQ6X-TMW4*

[80] Dan McKinley. “Choose Boring Technology.” *mcfunley.com*, March 2015. Archived at *perma.cc/7QW7-J4YP*

[81] Andy Warfield. “Building and Operating a Pretty Big Storage System Called S3.” *allthingsdistributed.com*, July 2023. Archived at *perma.cc/7LPK-TP7V*

[82] Marc Brooker. “Surprising Scalability of Multitenancy.” *brooker.co.za*, March 2023. Archived at *perma.cc/ZZD9-VV8T*

[83] Ben Stopford. “Shared Nothing vs. Shared Disk Architectures: An Independent View.” *benstopford.com*, November 2009. Archived at *perma.cc/7BXH-EDUR*

[84] Michael Stonebraker. “The Case for Shared Nothing.” *IEEE Database Engineering Bulletin*, volume 9, issue 1, pages 4–9, March 1986. perma.cc/P9YL-C4PS

[85] Panagiotis Antonopoulos, Alex Budovski, Cristian Diaconu, Alejandro Hernan-dez Saenz, Jack Hu, Hanuma Kodavalla, Donald Kossmann, Sandeep Lingam, Umar Farooq Minhas, Naveen Prakash, Vijendra

Purohit, Hugh Qu, Chaitanya Sreenivas Ravella, Krystyna Reisteter, Sheetal Shrotri, Dixin Tang, and Vikram Wakade. “Socrates: The New SQL Server in the Cloud.” At *ACM International Conference on Management of Data* (SIGMOD), June 2019. *doi:10.1145/3299869.3314047*

[86] Sam Newman. *Building Microservices*, 2nd edition. O’Reilly Media, 2021. ISBN: 9781492034025

[87] Nathan Ensmenger. “When Good Software Goes Bad: The Surprising Durability of an Ephemeral Technology.” At *The Maintainers Conference*, April 2016. Archived at *perma.cc/ZXT4-HGZB*

[88] Robert L. Glass. *Facts and Fallacies of Software Engineering*. Addison-Wesley Professional, 2002. ISBN: 9780321117427

[89] Marianne Bellotti. *Kill It with Fire*. No Starch Press, 2021. ISBN: 9781718501188

[90] Lisanne Bainbridge. “Ironies of Automation.” *Automatica*, volume 19, issue 6, pages 775–779, November 1983. *doi:10.1016/0005-1098(83)90046-8*

[91] James Hamilton. “On Designing and Deploying Internet-Scale Services.” At *21st Large Installation System Administration Conference* (LISA), November 2007.

[92] Dotan Horovits. “Open Source for Better Observability.” *horovits.medium.com*, October 2021. Archived at *perma.cc/R2HD-U2ZT*

[93] Brian Foote and Joseph Yoder. “Big Ball of Mud.” At *4th Conference on Pattern Languages of Programs* (PLoP), September 1997. Archived at *perma.cc/4GUP-2PBV*

[94] Marc Brooker. “What Is a Simple System?” *brooker.co.za*, May 2022. Archived at *perma.cc/U72T-BFVE*

[95] Frederick P. Brooks. “No Silver Bullet—Essence and Accident in Software Engineering.” In *The Mythical Man-Month*, Anniversary edition, Addison-Wesley, 1995. ISBN: 9780201835953

[96] Dan Luu. “Against Essential and Accidental Complexity.” *danluu.com*, December 2020. Archived at *perma.cc/H5ES-69KC*

[97] Erich Gamma, Richard Helm, Ralph Johnson, and John Vlissides. *Design Patterns: Elements of Reusable Object-Oriented Software*. Addison-Wesley Professional, 1994. ISBN: 9780201633610

[98] Eric Evans. *Domain-Driven Design: Tackling Complexity in the Heart of Software*. Addison-Wesley Professional, 2003. ISBN: 9780321125217

[99] Hongyu Pei Breivold, Ivica Crnkovic, and Peter J. Eriksson. “Analyzing Software Evolvability.” At *32nd Annual IEEE International Computer Software and Applications Conference* (COMPSAC), July 2008. *doi:10.1109/COMPSAC.2008.50*

[100] Enrico Zaninotto. “From X Programming to the X Organisation.” At *XP Conference*, May 2002. Archived at *perma.cc/R9AR-QCKZ*

บทที่ 3

[1] Jamie Brandon. “Unexplanations: Query Optimization Works Because SQL Is Declarative.” *scattered-thoughts.net*, February 2024. Archived at *perma.cc/P6W2-WMFZ*

[2] Neel Krishnaswami. “What Declarative Languages Are.” *semantic-domain.blogspot.com*, July 2013. Archived at *perma.cc/R4LP-T2RV*

[3] Joseph M. Hellerstein. “The Declarative Imperative: Experiences and Conjectures in Distributed Logic.” Tech report UCB/EECS-2010-90, Electrical Engineering and Computer Sciences, University of California at Berkeley, June 2010. Archived at *perma.cc/K56R-VVQM*

[4] Edgar F. Codd. “A Relational Model of Data for Large Shared Data Banks.” *Communications of the ACM*, volume 13, issue 6, pages 377–387, June 1970. *doi:10.1145/362384.362685*

[5] Michael Stonebraker and Joseph M. Hellerstein. “What Goes Around Comes Around.” In *Readings in Database Systems*, 4th edition, MIT Press, 2005, pages 2–41. ISBN: 9780262693141

[6] Markus Winand. “Modern SQL: Beyond Relational.” *modern-sql.com*, 2015. Archived at *perma.cc/D63V-WAPN*

[7] Martin Fowler. “Orm Hate.” *martinfowler.com*, May 2012. Archived at *perma.cc/VCM8-PKNG*

[8] Vlad Mihalcea. “N+1 Query Problem with JPA and Hibernate.” *vladmihalcea.com*, January 2023. Archived at *perma.cc/79EV-TZKB*

[9] Jens Schauder. “This Is the Beginning of the End of the N+1 Problem: Introducing Single Query Loading.” *spring.io*, August 2023. Archived at *perma.cc/6V96-R333*

[10] Jamie Brandon. “SQL Needed Structure.” *scattered-thoughts.net*, September 2025. Archived at *perma.cc/9EVK-HLVR*

[11] William Zola. “6 Rules of Thumb for MongoDB Schema Design.” *mongodb.com*, June 2014. Archived at *perma.cc/T2BZ-PPJB*

[12] Sidney Andrews and Christopher McClister. “Data Modeling in Azure Cosmos DB.” *learn.microsoft.com*, February 2023. Archived at *archive.org*

[13] Raffi Krikorian. “Timelines at Scale.” At *QCon San Francisco*, November 2012. Archived at *perma.cc/V9G5-KLYK*

[14] Ralph Kimball and Margy Ross. *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, 3rd edition. John Wiley & Sons, 2013. ISBN: 9781118530801

[15] Michael Kaminsky. “Data Warehouse Modeling: Star Schema vs. OBT.” *fivetran.com*, August 2022. Archived at *perma.cc/2PZK-BFFP*

[16] Joe Nelson. “User-defined Order in SQL.” *begriffs.com*, March 2018. Archived at *perma.cc/GS3W-F7AD*

[17] Evan Wallace. “Realtime Editing of Ordered Sequences.” *figma.com*, March 2017. Archived at *perma.cc/K6ER-CQZW*

[18] David Greenspan. “Implementing Fractional Indexing.” *observablehq.com*, October 2020. Archived at *perma.cc/5N4R-MREN*

[19] Martin Fowler. “Schemaless Data Structures.” *martinfowler.com*, January 2013.

[20] Amr Awadallah. “Schema-on-Read vs. Schema-on-Write.” At *Berkeley EECS RAD Lab Retreat*, May 2009. Archived at *perma.cc/DTB2-JCFR*

[21] Martin Odersky. “The Trouble with Types.” At *Strange Loop*, September 2013. Archived at *perma.cc/85QE-PVEP*

[22] Conrad Irwin. “MongoDB—Confessions of a PostgreSQL Lover.” At *HTML5DevConf*, October 2013. Archived at *perma.cc/C2J6-3AL5*

[23] “Percona Toolkit Documentation: pt-online-schema-change.” *docs.percona.com*, 2023. Archived at *perma.cc/9K8R-E5UH*

[24] Shlomi Noach. “gh-ost: GitHub’s Online Schema Migration Tool for MySQL.” *github.blog*, August 2016. Archived at *perma.cc/7XAG-XB72*

[25] Shayon Mukherjee. “pg-osc: Zero Downtime Schema Changes in PostgreSQL.” *shayon.dev*, February 2022. Archived at *perma.cc/35WN-7WMY*

[26] Carlos Pérez-Aradros Herce. “Introducing pgroll: Zero-Downtime, Reversible, Schema Migrations for Postgres.” *xata.io*, October 2023. Archived at *archive.org*

[27] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, Christopher Frost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean

Quinlan, Rajesh Rao, Lindsay Rolig, Dale Woodford, Yasushi Saito, Christopher Taylor, Michal Szymaniak, and Ruth Wang. “Spanner: Google’s Globally-Distributed Database.” At *10th USENIX Symposium on Operating System Design and Implementation* (OSDI), October 2012.

[28] Donald K. Burleson. “Reduce I/O with Oracle Cluster Tables.” *dba-oracle.com*. Archived at *perma.cc/7LBJ-9X2C*

[29] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wal-lach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. “Bigtable: A Distributed Storage System for Structured Data.” At *7th USENIX Symposium on Operating System Design and Implementation* (OSDI), November 2006.

[30] Priscilla Walmsley. *XQuery*, 2nd edition. O’Reilly Media, 2015. ISBN: 9781491915080

[31] Paul C. Bryan, Kris Zyp, and Mark Nottingham. “JavaScript Object Notation (JSON) Pointer.” RFC 6901, IETF, April 2013.

[32] Stefan Gössner, Glyn Normington, and Carsten Bormann. “JSONPath: Query Expressions for JSON.” RFC 9535, IETF, February 2024.

[33] Michael Stonebraker and Andrew Pavlo. “What Goes Around Comes Around… And Around….” *ACM SIGMOD Record*, volume 53, issue 2, pages 21–37, July 2024. *doi:10.1145/3685980.3685984*

[34] Lawrence Page, Sergey Brin, Rajeev Motwani, and Terry Winograd. “The PageRank Citation Ranking: Bringing Order to the Web.” Technical Report 1999-66, Stanford University InfoLab, November 1999. Archived at *perma.cc/UML9-UZHW*

[35] Nathan Bronson, Zach Amsden, George Cabrera, Prasad Chakka, Peter Dimov, Hui Ding, Jack Ferris, Anthony Giardullo, Sachin Kulkarni, Harry Li, Mark Marchukov, Dmitri Petrov, Lovro Puzar, Yee Jiun Song, and Venkat Venkataramani. “TAO: Facebook’s Distributed Data Store for the Social Graph.” At *USENIX Annual Technical Conference* (ATC), June 2013.

[36] Natasha Noy, Yuqing Gao, Anshu Jain, Anant Narayanan, Alan Patterson,and Jamie Taylor. “Industry-Scale Knowledge Graphs: Lessons and Challenges.” *Communications of the ACM*, volume 62, issue 8, pages 36–43, August 2019. *doi:10.1145/3331166*

[37] Xiyang Feng, Guodong Jin, Ziyi Chen, Chang Liu, and Semih Salihoğlu. “KÙZU Graph Database Management System.” At *13th Annual Conference on Innovative Data Systems Research* (CIDR 2023), January 2023. Archived at *perma.cc/PS6J-ZBZU*

[38] Maciej Besta, Emanuel Peter, Robert Gerstenberger, Marc Fischer, Michał Podstawski, Claude Barthels, Gustavo Alonso, Torsten Hoefler. “Demystifying Graph Databases: Analysis and Taxonomy of Data Organization, System Designs, and Graph Queries.” *arXiv:1910.09017*, October 2019.

[39] “Apache TinkerPop. TinkerPop 3.6.3 Documentation.” *tinkerpop.apache.org*, May 2023. Archived at *perma.cc/KM7W-7PAT*

[40] Nadime Francis, Alastair Green, Paolo Guagliardo, Leonid Libkin, TobiasLindaaker, Victor Marsault, Stefan Plantikow, Mats Rydberg, Petra Selmer, and Andrés Taylor. “Cypher: An Evolving Query Language for Property Graphs.” At *International Conference on Management of Data* (SIGMOD), May 2018. *doi:10.1145/3183713.3190657*

[41] Emil Eifrem. Twitter correspondence, January 2014. Archived at *perma.cc/WM4S-BW64*

[42] Francesco Tisiot. “Explore the New SEARCH and CYCLE Features in PostgreSQL® 14.” *aiven.io*, December 2021. Archived at *perma.cc/J6BT-83UZ*

[43] Gaurav Goel. “Understanding Hierarchies in Oracle.” *towardsdatascience.com*, May 2020. Archived at *perma.cc/5ZLR-Q7EW*

[44] Alin Deutsch, Yu Xu, and Mingxi Wu. “Seamless Syntactic and Semantic Integration of Query Primitives over Relational and Graph Data in GSQL.” *tigergraph.com*, November 2018. Archived at *perma.cc/JG7J-Y35X*

[45] Oskar van Rest, Sungpack Hong, Jinha Kim, Xuming Meng, and HassanChafi. “PGQL: A Property Graph Query Language.” At *4th International Workshop on Graph Data Management Experiences and Systems* (GRADES), June 2016. *doi:10.1145/2960414.2960421*

[46] Philip Rathle and Brad Bebee. “GQL: The ISO Standard for Graphs Has Arrived.” *aws.amazon.com*, April 2024. Archived at *perma.cc/5TEU-N2Y8*

[47] Alin Deutsch, Nadime Francis, Alastair Green, Keith Hare, Bei Li, LeonidLibkin, Tobias Lindaaker, Victor Marsault, Wim Martens, Jan Michels, Filip Murlak, Stefan Plantikow, Petra Selmer, Oskar van Rest, Hannes Voigt, Domagoj Vrgoč, Mingxi Wu, and Fred Zemke. “Graph Pattern Matching in GQL and SQL/ PGQ.” At *International Conference on Management of Data* (SIGMOD), June 2022. *doi:10.1145/3514221.3526057*

[48] Alastair Green. “SQL...And Now GQL.” *opencypher.org*, September 2019. Archived at *perma.cc/AFB2-3SY7*

[49] Amazon Web Services. “Neptune Graph Data Model.” Amazon Neptune User Guide, *docs.aws.amazon.com*. Archived at *perma.cc/CX3T-EZU9*

[50] Cognitect. “Datomic Data Model.” Datomic Cloud Documentation, *docs.datomic.com*. Archived at *perma.cc/LGM9-LEUT*

[51] David Beckett and Tim Berners-Lee. “Turtle—Terse RDF Triple Language.” W3C Team Submission, March 2011.

[52] Sinclair Target. “Whatever Happened to the Semantic Web?” *twobithistory.org*, May 2018. Archived at *perma.cc/M8GL-9KHS*

[53] Gavin Mendel-Gleason. “The Semantic Web Is Dead—Long Live the Semantic Web!” *terminusdb.com*, August 2022. Archived at *perma.cc/G2MZ-DSS3*

[54] Manu Sporny. “JSON-LD and Why I Hate the Semantic Web.” *manu.sporny.org*, January 2014. Archived at *perma.cc/7PT4-PJKF*

[55] University of Michigan Library. “Biomedical Ontologies and Controlled Vocabularies.” *guides.lib.umich.edu/ontology*. Archived at *perma.cc/Q5GA-F2N8*

[56] Facebook. “The Open Graph Protocol.” *ogp.me*. Archived at *perma.cc/C49A-GUSY*

[57] Matt Haughey. “Everything You Ever Wanted to Know About Unfurling but Were Afraid to Ask /or/ How to Make Your Site Previews Look Amazing in Slack.” *medium.com*, November 2015. Archived at *perma.cc/C7S8-4PZN*

[58] W3C RDF Working Group. “Resource Description Framework (RDF).” *w3.org*, February 2004.

[59] Steve Harris, Andy Seaborne, and Eric Prud’hommeaux. “SPARQL 1.1 Query Language.” W3C Recommendation, March 2013.

[60] Todd J. Green, Shan Shan Huang, Boon Thau Loo, and Wenchao Zhou. “Datalog and Recursive Query Processing.” *Foundations and Trends in Databases*, volume 5, issue 2, pages 105–195, November 2013. *doi:10.1561/1900000017*

[61] Stefano Ceri, Georg Gottlob, and Letizia Tanca. “What You Always Wanted to Know About Datalog (And Never Dared to Ask).” *IEEE Transactions on Knowledge and Data Engineering*, volume 1, issue 1, pages 146–166, March 1989. *doi:10.1109/69.43410*

[62] Serge Abiteboul, Richard Hull, and Victor Vianu. *Foundations of Databases*. Addison-Wesley, 1995. ISBN: 9780201537710. Available online at *webdam.inria.fr/ Alice*.

[63] Scott Meyer, Andrew Carter, and Andrew Rodriguez. “LIquid: The Soul of a New Graph Database, Part 2.” *engineering.linkedin.com*, September 2020. Archived at *perma.cc/K9M4-PD6Q*

[64] Matt Bessey. “Why, After 6 Years, I’m over GraphQL.” *bessey.dev*, May 2024. Archived at *perma.cc/2PAU-JYRA*

[65] Dominic Betts, Julián Domínguez, Grigori Melnik, Fernando Simonazzi, andMani Subramanian. *Exploring CQRS and Event Sourcing*. Microsoft Patterns & Practices, 2012. ISBN: 9781621140164. Archived at *perma.cc/7A39-3NM8*

[66] Greg Young. “CQRS and Event Sourcing.” At *Code on the Beach*, August 2014.

[67] Greg Young. “CQRS Documents.” *cqrs.wordpress.com*, November 2010. Archived at *perma.cc/X5R6-R47F*

[68] Brent Robinson. “Crypto Shredding: How It Can Solve Modern Data Retention Challenges.” *medium.com*, January 2019. Archived at *perma.cc/4LFK-S6XE*

[69] Devin Petersohn, Stephen Macke, Doris Xin, William Ma, Doris Lee, Xiangxi Mo, Joseph E. Gonzalez, Joseph M. Hellerstein, Anthony D. Joseph, and Aditya Parameswaran. “Towards Scalable Dataframe Systems.” *Proceedings of the VLDB Endowment*, volume 13, issue 11, pages 2033–2046, July 2020. *doi:10.14778/3407790.3407807*

[70] Stavros Papadopoulos, Kushal Datta, Samuel Madden, and Timothy Mattson. “The TileDB Array Data Storage Manager.” *Proceedings of the VLDB Endowment*, volume 10, issue 4, pages 349–360, November 2016. *doi:10.14778/3025111.3025117*

[71] Florin Rusu. “Multidimensional Array Data Management.” *Foundations and Trends in Databases*, volume 12, issues 2–3, pages 69–220, February 2023. *doi:10.1561/1900000069*

[72] Ed Targett. “Bloomberg, Man Group Team Up to Develop Open Source ‘ArcticDB’ Database.” *thestack.technology*, March 2023. Archived at *perma.cc/M5YD-QQYV*

[73] Dennis A. Benson, Ilene Karsch-Mizrachi, David J. Lipman, James Ostell, andDavid L. Wheeler. GenBank. *Nucleic Acids Research*, volume 36, issue suppl_1, pages D25–D30, January 2008. *doi:10.1093/nar/gkm929*

บทที่ 4

[1] Nikolay Samokhvalov. “How Partial, Covering, and Multicolumn Indexes May Slow Down UPDATEs in PostgreSQL.” *postgres.ai*, October 2021. Archived at *perma.cc/PBK3-F4G9*

[2] Goetz Graefe. “Modern B-Tree Techniques.” *Foundations and Trends in Databases*, volume 3, issue 4, pages 203–402, August 2011. *doi:10.1561/1900000028*

[3] Evan Jones. “Why Databases Use Ordered Indexes but Programming Uses Hash Tables.” *evanjones.ca*, December 2019. Archived at *perma.cc/NJX8-3ZZD*

[4] Branimir Lambov. “CEP-25: Trie-Indexed SSTable Format.” *cwiki.apache.org*, November 2022. Archived at *perma.cc/HD7W-PW8U* (linked Google Doc archived at *perma.cc/UL6C-AAAE*)

[5] Thomas H. Cormen, Charles E. Leiserson, Ronald L. Rivest, and Clifford Stein.*Introduction to Algorithms*, 3rd edition. MIT Press, 2009. ISBN: 9780262533058

[6] Branimir Lambov. “Trie Memtables in Cassandra.” *Proceedings of the VLDB Endowment*, volume 15, issue 12, pages 3359–3371, August 2022. *doi:10.14778/3554821.3554828*

[7] Dhruba Borthakur. “The History of RocksDB.” *rocksdb.blogspot.com*, November 2013. Archived at *perma.cc/Z7C5-JPSP*

[8] Matteo Bertozzi. “Apache HBase I/O—HFile.” *blog.cloudera.com*, June 2012. Archived at *perma.cc/U9XH-L2KL*

[9] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wal-lach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. “Bigtable: A Distributed Storage System for Structured Data.” At *7th USENIX Symposium on Operating System Design and Implementation* (OSDI), November 2006.

[10] Patrick O’Neil, Edward Cheng, Dieter Gawlick, and Elizabeth O’Neil. “The Log-Structured Merge-Tree (LSM-Tree).” *Acta Informatica*, volume 33, issue 4, pages 351–385, June 1996. *doi:10.1007/s002360050048*

[11] Mendel Rosenblum and John K. Ousterhout. “The Design and Implementation of a Log-Structured File System.” *ACM Transactions on Computer Systems*, volume 10, issue 1, pages 26–52, February 1992. *doi:10.1145/146941.146943*

[12] Michael Armbrust, Tathagata Das, Liwen Sun, Burak Yavuz, Shixiong Zhu,Mukul Murthy, Joseph Torres, Herman van Hovell, Adrian Ionescu, Alicja Łuszczak, Michał Świtakowski, Michał Szafrański, Xiao Li, Takuya Ueshin, Mostafa Mokhtar, Peter Boncz, Ali Ghodsi, Sameer Paranjpye, Pieter Senster, Reynold Xin, and Matei Zaharia. “Delta Lake: High-Performance ACID Table Storage over Cloud Object Stores.” *Proceedings of the VLDB Endowment*, volume 13, issue 12, pages 3411–3424, August 2020. *doi:10.14778/3415478.3415560*

[13] Burton H. Bloom. “Space/Time Trade-offs in Hash Coding with Allowable Errors.” *Communications of the ACM*, volume 13, issue 7, pages 422–426, July 1970. *doi:10.1145/362686.362692*

[14] Adam Kirsch and Michael Mitzenmacher. “Less Hashing, Same Performance: Building a Better Bloom Filter.” *Random Structures & Algorithms*, volume 33, issue 2, pages 187–218, September 2008. *doi:10.1002/rsa.20208*

[15] Thomas Hurst. “Bloom Filter Calculator.” *hur.st*, September 2023. Archived at *perma.cc/L3AV-6VC2*

[16] Chen Luo and Michael J. Carey. “LSM-Based Storage Techniques: a Survey.” *The VLDB Journal*, volume 29, pages 393–418, July 2019. *doi:10.1007/s00778-019-00555-y*

[17] Subhadeep Sarkar and Manos Athanassoulis. “Dissecting, Designing, and Optimizing LSM-Based Data Stores.” Tutorial at *ACM International Conference on Management of Data* (SIGMOD), June 2022. Slides archived at

[18] Mark Callaghan. “Name That Compaction Algorithm.” *smalldatum.blogspot.com*, August 2018. Archived at *perma.cc/CN4M-82DY*

[19] Prashanth Rao. “Embedded Databases (1): The Harmony of DuckDB, KùzuDB and LanceDB.” thedataquarry.com, August 2023. Archived at perma.cc/PA28-2R35

[20] Hacker News discussion. “Bluesky Migrates to Single-Tenant SQLite.” news.ycombinator.com, October 2023. Archived at perma.cc/69LM-5P6X

[21] Rudolf Bayer and Edward M. McCreight. “Organization and Maintenance of Large Ordered Indices.” Boeing Scientific Research Laboratories, Mathematical and Information Sciences Laboratory, report no. 20, July 1970. *doi:10.1145/1734663.1734671*

[22] Douglas Comer. “The Ubiquitous B-Tree.” *ACM Computing Surveys*, volume 11, issue 2, pages 121–137, June 1979. *doi:10.1145/356770.356776*

[23] Alex Miller. “Torn Write Detection and Protection.” *transactional.blog*, April 2025. Archived at *perma.cc/G7EB-33EW*

[24] C. Mohan and Frank Levine. “ARIES/IM: An Efficient and High Concurrency Index Management Method Using Write-Ahead Logging.” At *ACM International Conference on Management of Data* (SIGMOD), June 1992. *doi:10.1145/130283.130338*

[25] Hironobu Suzuki. “The Internals of PostgreSQL.” *interdb.jp*, 2017. Archived at *archive.org*

[26] Howard Chu. “LDAP at Lightning Speed.” At *Build Stuff ’14*, November 2014. Archived at *perma.cc/GB6Z-P8YH*

[27] Manos Athanassoulis, Michael S. Kester, Lukas M. Maas, Radu Stoica, StratosIdreos, Anastasia Ailamaki, and Mark Callaghan. “Designing Access Methods: The RUM Conjecture.” At *19th International Conference on Extending Database Technology* (EDBT), March 2016. *doi:10.5441/002/edbt.2016.42*

[28] Ben Stopford. “Log Structured Merge Trees.” *benstopford.com*, February 2015. Archived at *perma.cc/E5BV-KUJ6*

[29] Mark Callaghan. “The Advantages of an LSM vs. a B-Tree.” *smalldatum.blogspot.co.uk*, January 2016. Archived at *perma.cc/3TYZ-EFUD*

[30] Oana Balmau, Florin Dinu, Willy Zwaenepoel, Karan Gupta, RavishankarChandhiramoorthi, and Diego Didona. “SILK: Preventing Latency Spikes in LogStructured Merge Key-Value Stores.” At *USENIX Annual Technical Conference*, July 2019.

[31] Igor Canadi, Siying Dong, Mark Callaghan, et al. “RocksDB Tuning Guide.” *github.com*, 2023. Archived at *perma.cc/UNY4-MK6C*

[32] Gabriel Haas and Viktor Leis. “What Modern NVMe Storage Can Do, and How to Exploit It: High-Performance I/O for High-Performance Storage Engines.” *Proceedings of the VLDB Endowment*, volume 16, issue 9, pages 2090–2102. May 2023. *doi:10.14778/3598581.3598584*

[33] Emmanuel Goossaert. “Coding for SSDs.” *codecapsule.com*, February 2014.

[34] Jack Vanlightly. “Is Sequential IO Dead in the Era of the NVMe Drive?” *jackvanlightly.com*, May 2023. Archived at *perma.cc/7TMZ-TAPU*

[35] Alibaba Cloud Storage Team. “Storage System Design Analysis: Factors Affecting NVMe SSD Performance (2).” *alibabacloud.com*, January 2019. Archived at *archive.org*

[36] Xiao-Yu Hu and Robert Haas. “The Fundamental Limit of Flash Random Write Performance: Understanding, Analysis and Performance Modelling.” *dominoweb.draco.res.ibm.com*, March 2010. Archived at *perma.cc/8JUL-4ZDS*

[37] Lanyue Lu, Thanumalayan Sankaranarayana Pillai, Andrea C. Arpaci-Dusseau,and Remzi H. Arpaci-Dusseau. “WiscKey: Separating Keys from Values in SSDConscious Storage.” At *4th USENIX Conference on File and Storage Technologies* (FAST), February 2016.

[38] Peter Zaitsev. “Innodb Double Write.” *percona.com*, August 2006. Archived at *perma.cc/NT4S-DK7T*

[39] Tomas Vondra. “On the Impact of Full-Page Writes.” *2ndquadrant.com*, November 2016. Archived at *perma.cc/7N6B-CVL3*

[40] Mark Callaghan. “Read, Write & Space Amplification—B-Tree vs. LSM.” *smalldatum.blogspot.com*, November 2015. Archived at *perma.cc/S487-WK5P*

[41] Mark Callaghan. “Choosing Between Efficiency and Performance with RocksDB.” At *Code Mesh*, November 2016

[42] Subhadeep Sarkar, Tarikul Islam Papon, Dimitris Staratzis, Zichen Zhu, andManos Athanassoulis. “Enabling Timely and Persistent Deletion in LSM-Engines.” *ACM Transactions on Database Systems*, volume 48, issue 3, article no. 8, August 2023. *doi:10.1145/3599724*

[43] Lukas Fittl. “Postgres vs. SQL Server: B-Tree Index Differences & the Benefit of Deduplication.” *pganalyze.com*, April 2025. Archived at *perma.cc/XY6T-LTPX*

[44] Drew Silcock. “How Postgres Stores Data on Disk—This One’s a Page Turner.” *drew.silcock.dev*, August 2024. Archived at *perma.cc/8K7K-7VJ2*

[45] Joe Webb. “Using Covering Indexes to Improve Query Performance.” *simpletalk.com*, September 2008. Archived at *perma.cc/6MEZ-R5VR*

[46] Michael Stonebraker, Samuel Madden, Daniel J. Abadi, Stavros Harizopoulos,Nabil Hachem, and Pat Helland. “The End of an Architectural Era (It’s Time for a Complete Rewrite).” At *33rd International Conference on Very Large Data Bases* (VLDB), September 2007.

[47] “VoltDB Technical Overview White Paper.” VoltDB, 2017. Archived at *perma.cc/ B9SF-SK5G*

[48] Stephen M. Rumble, Ankita Kejriwal, and John K. Ousterhout. “Log-Structured Memory for DRAM-Based Storage.” At *12th USENIX Conference on File and Storage Technologies* (FAST), February 2014.

[49] Stavros Harizopoulos, Daniel J. Abadi, Samuel Madden, and Michael Stone-braker. “OLTP Through the Looking Glass, and What We Found There.” At *ACM International Conference on Management of Data* (SIGMOD), June 2008.
doi:10.1145/1376616.1376713

[50] Per-Åke Larson, Cipri Clinciu, Campbell Fraser, Eric N. Hanson, Mostafa Mokh-tar, Michal Nowakiewicz, Vassilis Papadimos, Susan L. Price, Srikumar Rangarajan, Remus Rusanu, and Mayukh Saubhasik. “Enhancements to SQL Server Column Stores.” At *ACM International Conference on Management of Data* (SIGMOD), June 2013. *doi:10.1145/2463676.2463708*

[51] Franz Färber, Norman May, Wolfgang Lehner, Philipp Große, Ingo Müller,Hannes Rauhe, and Jonathan Dees. “The SAP HANA Database—An Architecture Overview.” *IEEE Data Engineering Bulletin*, volume 35, issue 1, pages 28–33, March 2012. Archived at *perma.cc/H2WC-YQZY*

[52] Michael Stonebraker. “The Traditional RDBMS Wisdom Is (Almost Certainly) All Wrong.” Presentation at *EPFL*, May 2013.

[53] Adam Prout, Szu-Po Wang, Joseph Victor, Zhou Sun, Yongzhu Li, JackChen, Evan Bergeron, Eric Hanson, Robert Walzer, Rodrigo Gomes, and Nikita Shamgunov. “Cloud-Native Transactions and Analytics in SingleStore.” At *ACM International Conference on Management of Data* (SIGMOD), June 2022. *doi:10.1145/3514221.3526055*

[54] Tino Tereshko and Jordan Tigani. “BigQuery Under the Hood.” *cloud.google.com*, January 2016. Archived at *perma.cc/WP2Y-FUCF*

[55] Wes McKinney. “The Road to Composable Data Systems: Thoughts on the Last 15 Years and the Future.” *wesmckinney.com*, September 2023. Archived at *perma.cc/ 6L2M-GTJX*

[56] Michael Stonebraker, Daniel J. Abadi, Adam Batkin, Xuedong Chen, MitchCherniack, Miguel Ferreira, Edmond Lau, Amerson Lin, Sam Madden, Elizabeth O’Neil, Pat O’Neil, Alex Rasin, Nga Tran, and Stan Zdonik. “C-Store: A ColumnOriented DBMS.” At *31st International Conference on Very Large Data Bases* (VLDB), September 2005.

[57] Julien Le Dem. “Dremel Made Simple with Parquet.” *blog.x.com*, September 2013. Archived at *archive.org*

[58] Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivaku-mar, Matt Tolton, and Theo Vassilakis. “Dremel: Interactive Analysis of Web-Scale Datasets.” At *36th International Conference on Very Large Data Bases* (VLDB), September 2010. *doi:10.14778/1920841.1920886*

[59] Joe Kearney. “Understanding Record Shredding: Storing Nested Data in Columns.” *joekearney.co.uk*, December 2016. Archived at *perma.cc/ZD5N-AX5D*

[60] Jamie Brandon. “A Shallow Survey of OLAP and HTAP Query Engines.” scattered-thoughts.net, September 2023. Archived at perma.cc/L3KH-J4JF

[61] Benoit Dageville, Thierry Cruanes, Marcin Zukowski, Vadim Antonov, ArtinAvanes, Jon Bock, Jonathan Claybaugh, Daniel Engovatov, Martin Hentschel, Jiansheng Huang, Allison W. Lee, Ashish Motivala, Abdul Q. Munir, Steven Pelley, Peter Povinec, Greg Rahn, Spyridon Triantafyllis, and Philipp Unterbrunner. “The Snowflake Elastic Data Warehouse.” At *ACM International Conference on Management of Data* (SIGMOD), June 2016. *doi:10.1145/2882903.2903741*

[62] Mark Raasveldt and Hannes Mühleisen. “Data Management for Data Science Towards Embedded Analytics.” At *10th Conference on Innovative Data Systems Research* (CIDR), January 2020. Archived at *perma.cc/65G2-NYDT*

[63] Jean-François Im, Kishore Gopalakrishna, Subbu Subramaniam, Mayank Shriv-astava, Adwait Tumbde, Xiaotian Jiang, Jennifer Dai, Seunghyun Lee, Neha Pawar, Jialiang Li, and Ravi Aringunram. “Pinot: Realtime OLAP for 530 Million Users.” At *ACM International Conference on Management of Data* (SIGMOD), May 2018. *doi:10.1145/3183713.3190661*

[64] Fangjin Yang, Eric Tschetter, Xavier Léauté, Nelson Ray, Gian Merlino, and DeepGanguli. “Druid: A Real-Time Analytical Data Store.” At *ACM International Conference on Management of Data* (SIGMOD), June 2014. *doi:10.1145/2588555.2595631*

[65] Chunwei Liu, Anna Pavlenko, Matteo Interlandi, and Brandon Haynes. “Deep Dive into Common Open Formats for Analytical DBMSs.” *Proceedings of the VLDB Endowment*, volume 16, issue 11, pages 3044–3056, July 2023. *doi:10.14778/3611479.3611507*

[66] Xinyu Zeng, Yulong Hui, Jiahong Shen, Andrew Pavlo, Wes McKinney,and Huanchen Zhang. “An Empirical Evaluation of Columnar Storage Formats.” *Proceedings of the VLDB Endowment*, volume 17, issue 2, pages 148–161. *doi:10.14778/3626292.3626298*

[67] Weston Pace. “Lance v2: A Columnar Container Format for Modern Data.” blog.lancedb.com, April 2024. Archived at perma.cc/ZK3Q-S9VJ

[68] Yoav Helfman. “Nimble, A New Columnar File Format.” At *VeloxCon*, April 2024.

[69] Wes McKinney. “Apache Arrow: High-Performance Columnar Data Framework.” At *CMU Database Group—Vaccination Database Tech Talks*, December 2021.

[70] Wes McKinney. *Python for Data Analysis*, 3rd edition. O’Reilly Media, 2022. ISBN: 9781098104023

[71] Paul Dix. “The Design of InfluxDB IOx: An In-Memory Columnar Database Written in Rust with Apache Arrow.” At *CMU Database Group—Vaccination Database Tech Talks*, May 2021.

[72] Carlota Soto and Mike Freedman. “Building Columnar Compression for Large PostgreSQL Databases.” *timescale.com*, March 2024. Archived at *perma.cc/7KTF-V3EH*

[73] Daniel J. Abadi, Peter Boncz, Stavros Harizopoulos, Stratos Idreos, and SamuelMadden. “The Design and Implementation of Modern Column-Oriented Database Systems.” *Foundations and Trends in Databases*, volume 5, issue 3, pages 197–280, December 2013. *doi:10.1561/1900000024*

[74] Daniel Lemire, Gregory Ssi-Yan-Kai, and Owen Kaser. “Consistently Faster and Smaller Compressed Bitmaps with Roaring.” *Software: Practice and Experience*, volume 46, issue 11, pages 1547–1569, November 2016. *doi:10.1002/spe.2402*

[75] Jaz Volpert. “An Entire Social Network in 1.6GB (GraphD Part 2).” *jazco.dev*, April 2024. Archived at *perma.cc/L27Z-QVMG*

[76] Andrew Lamb, Matt Fuller, Ramakrishna Varadarajan, Nga Tran, Ben Vandiver,Lyric Doshi, and Chuck Bear. “The Vertica Analytic Database: C-Store 7 Years Later.” *Proceedings of the VLDB Endowment*, volume 5, issue 12, pages 1790–1801, August 2012. *doi:10.14778/2367502.2367518*

[77] Timo Kersten, Viktor Leis, Alfons Kemper, Thomas Neumann, AndrewPavlo, and Peter Boncz. “Everything You Always Wanted to Know About Compiled and Vectorized Queries But Were Afraid to Ask.” *Proceedings of the VLDB Endowment*, volume 11, issue 13, pages 2209–2222, September 2018. *doi:10.14778/3275366.3284966*

[78] Forrest Smith. “Memory Bandwidth Napkin Math.” *forrestthewoods.com*, February 2020. Archived at *perma.cc/Y8U4-PS7N*

[79] Peter Boncz, Marcin Zukowski, and Niels Nes. “MonetDB/X100: HyperPipelining Query Execution.” At *2nd Biennial Conference on Innovative Data Systems Research* (CIDR), January 2005. Archived at *perma.cc/R4KF-QKHF*

[80] Jingren Zhou and Kenneth A. Ross. “Implementing Database Operations Using SIMD Instructions.” At *ACM International Conference on Management of Data* (SIGMOD), June 2002. *doi:10.1145/564691.564709*

[81] Kevin Bartley. “OLTP Queries: Transfer Expensive Workloads to Materialize.” *materialize.com*, August 2024. Archived at *perma.cc/4TYM-TYD8*

[82] Jim Gray, Surajit Chaudhuri, Adam Bosworth, Andrew Layman, Don Reichart,Murali Venkatrao, Frank Pellow, and Hamid Pirahesh. “Data Cube: A Relational Aggregation Operator Generalizing Group-By, Cross-Tab, and Sub-Totals.” *Data Mining and Knowledge Discovery*, volume 1, issue 1, pages 29–53, March 2007. *doi:10.1023/A:1009726021843*

[83] Frank Ramsak, Volker Markl, Robert Fenk, Martin Zirkel, Klaus Elhardt, andRudolf Bayer. “Integrating the UB-Tree into a Database System Kernel.” At *26th International Conference on Very Large Data Bases* (VLDB), September 2000.

[84] Octavian Procopiuc, Pankaj K. Agarwal, Lars Arge, and Jeffrey Scott Vitter. “Bkd-Tree: A Dynamic Scalable kd-Tree.” At *8th International Symposium on Spatial and Temporal Databases* (SSTD), July 2003. *doi:10.1007/978-3-540-45072-6_4*

[85] Joseph M. Hellerstein, Jeffrey F. Naughton, and Avi Pfeffer. “Generalized Search Trees for Database Systems.” At *21st International Conference on Very Large Data Bases* (VLDB), September 1995.

[86] Isaac Brodsky. “H3: Uber’s Hexagonal Hierarchical Spatial Index.” *eng.uber.com*, June 2018. Archived at *archive.org*

[87] Robert Escriva, Bernard Wong, and Emin Gün Sirer. “HyperDex: A Distributed, Searchable Key-Value Store.” At *ACM SIGCOMM Conference*, August 2012. *doi:10.1145/2377677.2377681*

[88] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze.*Introduction to Information Retrieval*. Cambridge University Press, 2008. ISBN: 9780521865715. Available online at *nlp.stanford.edu/IR-book*.

[89] Jianguo Wang, Chunbin Lin, Yannis Papakonstantinou, and Steven Swanson. “An Experimental Study of Bitmap Compression vs. Inverted List Compression.” At *ACM International Conference on Management of Data* (SIGMOD), May 2017. *doi:10.1145/3035918.3064007*

[90] Adrien Grand. “What Is in a Lucene Index?” At *Lucene/Solr Revolution*, November 2013. Archived at *perma.cc/Z7QN-GBYY*

[91] Michael McCandless. “Visualizing Lucene’s Segment Merges.” *blog.mikemccandless.com*, February 2011. Archived at *perma.cc/3ZV8-72W6*

[92] Lukas Fittl. “Understanding Postgres GIN Indexes: The Good and the Bad.” *pganalyze.com*, December 2021. Archived at *perma.cc/V3MW-26H6*

[93] Jimmy Angelakos. “The State of (Full) Text Search in PostgreSQL 12.” At *FOSDEM*, February 2020. Archived at *perma.cc/J6US-3WZS*

[94] Alexander Korotkov. “Index Support for Regular Expression Search.” At *PGConf.EU Prague*, October 2012. Archived at *perma.cc/5RFZ-ZKDQ*

[95] Michael McCandless. “Lucene’s FuzzyQuery Is 100 Times Faster in 4.0.” *blog.mikemccandless.com*, March 2011. Archived at *perma.cc/E2WC-GHTW*

[96] Steffen Heinz, Justin Zobel, and Hugh E. Williams. “Burst Tries: A Fast, Efficient Data Structure for String Keys.” *ACM Transactions on Information Systems*, volume 20, issue 2, pages 192–223, April 2002. *doi:10.1145/506309.506312*

[97] Klaus U. Schulz and Stoyan Mihov. “Fast String Correction with Levenshtein Automata.” *International Journal on Document Analysis and Recognition*, volume 5, issue 1, pages 67–85, November 2002. *doi:10.1007/s10032-002-0082-8*

[98] Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. “Efficient Estimation of Word Representations in Vector Space.” *arXiv:1301.3781*, September 2013

[99] Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. “BERT: Pre-Training of Deep Bidirectional Transformers for Language Understanding.” At *Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, June 2019. *doi:10.18653/v1/N19-1423*

[100] Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. “Improving Language Understanding by Generative Pre-Training.” *openai.com*, June 2018. Archived at *perma.cc/5N3C-DJ4C*

[101] Matthijs Douze, Maria Lomeli, and Lucas Hosseini. “Faiss Indexes.” *github.com*, August 2024. Archived at *perma.cc/2EWG-FPBS*

[102] Varik Matevosyan. “Understanding pgvector’s HNSW Index Storage in Postgres.” *lantern.dev*, August 2024. Archived at *perma.cc/B2YB-JB59*

[103] Dmitry Baranchuk, Artem Babenko, and Yury Malkov. “Revisiting the Inverted Indices for Billion-Scale Approximate Nearest Neighbors.” At *European Conference on Computer Vision* (ECCV), September 2018. *doi:10.1007/978-3-030-01258-8_13*

[104] Yury A. Malkov and Dmitry A. Yashunin. “Efficient and Robust Approximate Nearest Neighbor Search Using Hierarchical Navigable Small World Graphs.” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, volume 42, issue 4, pages 824–836, April 2020. *doi:10.1109/TPAMI.2018.2889473*

บทที่ 5

[1] “CWE-502: Deserialization of Untrusted Data.” Common Weakness Enumeration, *cwe.mitre.org*, July 2006. Archived at *perma.cc/26EU-UK9Y*

[2] Steve Breen. “What Do WebLogic, WebSphere, JBoss, Jenkins, OpenNMS, and Your Application Have in Common? This Vulnerability.” *foxglovesecurity.com*, November 2015. Archived at *perma.cc/9U97-UVVD*

[3] Patrick McKenzie. “What the Rails Security Issue Means for Your Startup.” *kalzumeus.com*, January 2013. Archived at *perma.cc/2MBJ-7PZ6*

[4] Brian Goetz. “Towards Better Serialization.” *openjdk.org*, June 2019. Archived at *perma.cc/UK6U-GQDE*

[5] Eishay Smith. “jvm-serializers Wiki.” *github.com*, October 2023. Archived at *perma.cc/PJP7-WCNG*

[6] “XML Is a Poor Copy of S-Expressions.” *wiki.c2.com*, May 2013. Archived at *perma.cc/7FAN-YBKL*

[7] Julia Evans. “Examples of Floating Point Problems.” *jvns.ca*, January 2023. Archived at *perma.cc/M57L-QKKW*

[8] Matt Harris. “Snowflake: An Update and Some Very Important Information.” Email to *Twitter Development Talk* mailing list, October 2010. Archived at *perma.cc/8UBV-MZ3D*

[9] Yakov Shafranovich. “RFC 4180: Common Format and MIME Type for CommaSeparated Values (CSV) Files.” IETF, October 2005.

[10] Andy Coates. “Evolving JSON Schemas—Part I.” *creekservice.org*, January 2024. Archived at *perma.cc/MZW3-UA54*

[11] Andy Coates. “Evolving JSON Schemas—Part II.” *creekservice.org*, January 2024. Archived at *perma.cc/GT5H-WKZ5*

[12] Pierre Genevès, Nabil Layaïda, and Vincent Quint. “Ensuring Query Compatibility with Evolving XML Schemas.” INRIA Technical Report 6711, November 2008. Archived at *arxiv.org*

[13] Tim Bray. “Bits on the Wire.” *tbray.org*, November 2019. Archived at *perma.cc/ 3BT3-BQU3*

[14] Mark Slee, Aditya Agarwal, and Marc Kwiatkowski. “Thrift: Scalable CrossLanguage Services Implementation.” Facebook Technical Report, April 2007. Archived at *perma.cc/22BS-TUFB*

[15] Martin Kleppmann. “Schema Evolution in Avro, Protocol Buffers and Thrift.” *martin.kleppmann.com*, December 2012. Archived at *perma.cc/E4R2-9RJT*

[16] Doug Cutting et al. “[PROPOSAL] New Subproject: Avro.” Email thread on *hadoop-general* mailing list, *lists.apache.org*, April 2009. Archived at *perma.cc/4A79-BMEB*

[17] Apache Software Foundation. “Apache Avro 1.12.0 Specification.” *avro.apache.org*, August 2024. Archived at *perma.cc/C36P-5EBQ*

[18] Apache Software Foundation. “Avro Schemas as LL(1) CFG Definitions.” *avro.apache.org*, August 2024. Archived at *perma.cc/JB44-EM9Q*

[19] Tony Hoare. “Null References: The Billion Dollar Mistake.” At *QCon London*, March 2009.

[20] Confluent, Inc. “Schema Registry Overview.” *docs.confluent.io*, 2024. Archived at *perma.cc/92C3-A9JA*

[21] Aditya Auradkar and Tom Quiggle. “Introducing Espresso—LinkedIn’s Hot New Distributed Document Store.” *engineering.linkedin.com*, January 2015. Archived at *perma.cc/FX4P-VW9T*

[22] Jay Kreps. “Putting Apache Kafka to Use: A Practical Guide to Building a Stream Data Platform (Part 2).” *confluent.io*, February 2015. Archived at *perma.cc/8UA4-ZS5S*

[23] Gwen Shapira. “The Problem of Managing Schemas.” *oreilly.com*, November 2014. Archived at *perma.cc/BY8Q-RYV3*

[24] John Larmouth. *ASN.1 Complete*. Morgan Kaufmann, 1999. ISBN: 9780122334351. Archived at *perma.cc/GB7Y-XSXQ*

[25] Burton S. Kaliski Jr. “A Layman’s Guide to a Subset of ASN.1, BER, and DER.” Technical Note, RSA Data Security, Inc., November 1993. Archived at *perma.cc/2LMN-W9U8*

[26] Jacob Hoffman-Andrews. “A Warm Welcome to ASN.1 and DER.” *letsencrypt.org*, April 2020. Archived at *perma.cc/CYT2-GPQ8*

[27] Lev Walkin. “Question: Extensibility and Dropping Fields.” *lionet.info*, September 2010. Archived at *perma.cc/VX8E-NLH3*

[28] Jacqueline Xu. “Online Migrations at Scale.” *stripe.com*, February 2017. Archived at *perma.cc/X59W-DK7Y*

[29] Geoffrey Litt, Peter van Hardenberg, and Orion Henry. “Project Cambria: Translate Your Data with Lenses.” Technical Report, October 2020. Archived at *perma.cc/WA4V-VKDB*

[30] Pat Helland. “Data on the Outside Versus Data on the Inside.” At *2nd Biennial Conference on Innovative Data Systems Research* (CIDR), January 2005. Archived at *perma.cc/GH56-WYZS*

[31] Roy Thomas Fielding. “Architectural Styles and the Design of Network-Based Software Architectures.” PhD thesis, University of California, Irvine, 2000. Archived at *perma.cc/LWY9-7BPE*

[32] Roy Thomas Fielding. “REST APIs Must Be Hypertext-Driven.” *roy.gbiv.com*, October 2008. Archived at *perma.cc/M2ZW-8ATG*

[33] “OpenAPI Specification Version 3.1.0.” *swagger.io*, February 2021. Archived at *perma.cc/3S6S-K5M4*

[34] Michi Henning. “The Rise and Fall of CORBA.” *Communications of the ACM*, volume 51, issue 8, pages 52–57, August 2008. *doi:10.1145/1378704.1378718*

[35] Pete Lacey. “The S Stands for Simple.” *harmful.cat-v.org*, November 2006. Archived at *perma.cc/4PMK-Z9X7*

[36] Stefan Tilkov. “Interview: Pete Lacey Criticizes Web Services.” *infoq.com*, December 2006. Archived at *perma.cc/JWF4-XY3P*

[37] Tim Bray. “The Loyal WS-Opposition.” *tbray.org*, September 2004. Archived at *perma.cc/J5Q8-69Q2*

[38] Andrew D. Birrell and Bruce Jay Nelson. “Implementing Remote Procedure Calls.” *ACM Transactions on Computer Systems* (TOCS), volume 2, issue 1, pages 39–59, February 1984. *doi:10.1145/2080.357392*

[39] Jim Waldo, Geoff Wyant, Ann Wollrath, and Sam Kendall. “A Note on Distributed Computing.” Sun Microsystems Laboratories, Inc., Technical Report TR-94-29, November 1994. Archived at *perma.cc/8LRZ-BSZR*

[40] Steve Vinoski. “Convenience over Correctness.” *IEEE Internet Computing*, volume 12, issue 4, pages 89–92, July 2008. *doi:10.1109/MIC.2008.75*

[41] Brandur Leach. “Designing Robust and Predictable APIs with Idempotency.” *stripe.com*, February 2017. Archived at *perma.cc/JD22-XZQT*

[42] Sam Rose. “Load Balancing.” *samwho.dev*, April 2023. Archived at *perma.cc/ Q7BA-9AE2*

[43] Troy Hunt. “Your API Versioning Is Wrong, Which Is Why I Decided to Do It 3 Different Wrong Ways.” *troyhunt.com*, February 2014. Archived at *perma.cc/9DSW-DGR5*

[44] Brandur Leach. “APIs As Infrastructure: Future-Proofing Stripe with Versioning.” *stripe.com*, August 2017. Archived at *perma.cc/L63K-USFW*

[45] AOASIS Web Services Business Process Execution Language (WSBPEL) Techni-cal Committee. “Web Services Business Process Execution Language Version 2.0.” *docs.oasis-open.org*, April 2007.

[46] “Temporal. Temporal Service.” *docs.temporal.io*, 2024. Archived at *perma.cc/ 32P3-CJ9V*

[47] Stephan Ewen. “Why We Built Restate.” *restate.dev*, August 2023. Archived at *perma.cc/BJJ2-X75K*

[48] Keith Tenzer and Joshua Smith. “Understanding Idempotency in Distributed Systems.” *temporal.io*, February 2024. Archived at *perma.cc/TY4U-EH3W*

[49] “Temporal. Temporal Workflow.” *docs.temporal.io*, 2024. Archived at *perma.cc/ B5C5-Y396*

[50] Jack Kleeman. “Solving Durable Execution’s Immutability Problem.” *restate.dev*, February 2024. Archived at *perma.cc/G55L-EYH5*

[51] Srinath Perera. “Exploring Event-Driven Architecture: A Beginner’s Guide for Cloud Native Developers.” *wso2.com*, August 2023. Archived at *archive.org*

[52] Philip A. Bernstein, Sergey Bykov, Alan Geller, Gabriel Kliot, and Jorgen Thelin.“Orleans: Distributed Virtual Actors for Programmability and Scalability.” Microsoft Research Technical Report MSR-TR-2014-41, March 2014. Archived at *perma.cc/PD3U-WDMF*

บทที่ 6

[1] B. G. Lindsay, P. G. Selinger, C. Galtieri, J. N. Gray, R. A. Lorie, T. G. Price, F. Putzolu, I. L. Traiger, and B. W. Wade. “Notes on Distributed Databases.” IBM Research, Research Report RJ2571(33471), July 1979. Archived at perma.cc/EPZ3-MHDD

[2] Kenny Gryp. “MySQL Terminology Updates.” dev.mysql.com, July 2020. Archived at perma.cc/S62G-6RJ2

[3] Oracle Corporation. “Oracle (Active) Data Guard 19c: Real-Time Data Protection and Availability.” White Paper, oracle.com, March 2019. Archived at perma.cc/P5ST-RPKE

[4] Microsoft. “What Is an Always On Availability Group?” *learn.microsoft.com*, September 2024. Archived at *perma.cc/ABH6-3MXF*

[5] Mostafa Elhemali, Niall Gallagher, Nicholas Gordon, Joseph Idziorek, RichardKrog, Colin Lazier, Erben Mo, Akhilesh Mritunjai, Somu Perianayagam, Tim Rath, Swami Sivasubramanian, James Christopher Sorenson III, Sroaj Sosothikul, Doug Terry, and Akshat Vig. “Amazon DynamoDB: A Scalable, Predictably Performant, and Fully Managed NoSQL Database Service.” At *USENIX Annual Technical Conference* (ATC), July 2022.

[6] Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis,Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, Paul Bardea, Amruta Ranade, Ben Darnell, Bram Gruneir, Justin Jaffray, Lucy Zhang, and Peter Mattis. “CockroachDB: The Resilient Geo-Distributed SQL Database.” At *ACM SIGMOD International Conference on Management of Data* (SIGMOD), June 2020. *doi:10.1145/3318464.3386134*

[7] Dongxu Huang, Qi Liu, Qiu Cui, Zhuhe Fang, Xiaoyu Ma, Fei Xu, Li Shen,Liu Tang, Yuxing Zhou, Menglong Huang, Wan Wei, Cong Liu, Jian Zhang, Jianjun Li, Xuelian Wu, Lingyu Song, Ruoxi Sun, Shuaipeng Yu, Lei Zhao, Nicholas Cameron, Liquan Pei, and Xin Tang. “TiDB: A Raft-Based HTAP Database.” *Proceedings of the VLDB Endowment*, volume 13, issue 12, pages 3072–3084, August 2020. *doi:10.14778/3415478.3415535*

[8] Mallory Knodel and Niels ten Oever. “Terminology, Power, and Inclusive Language in Internet-Drafts and RFCs.” *IETF Internet-Draft*, August 2023. Archived at *perma.cc/5ZY9-725E*

[9] Buck Hodges. “Postmortem: VSTS 4 September 2018.” *devblogs.microsoft.com*, September 2018. Archived at *perma.cc/ZF5R-DYZS*

[10] Gunnar Morling. “Leader Election with S3 Conditional Writes.” *www.morling.dev*, August 2024. Archived at *perma.cc/7V2N-J78Y*

[11] Vignesh Chandramohan, Rohan Desai, and Chris Riccomini. “SlateDB Manifest Design.” *github.com*, May 2024. Archived at *perma.cc/8EUY-P32Z*

[12] Stas Kelvich. “Why Does Neon Use Paxos Instead of Raft, and What’s the Difference?” *neon.tech*, August 2022. Archived at *perma.cc/SEZ4-2GXU*

[13] Dimitri Fontaine. “An Introduction to the pg_auto_failover Project.” *tapoueh.org*, November 2021. Archived at *perma.cc/3WH5-6BAF*

[14] Jesse Newland. “GitHub Availability This Week.” *github.blog*, September 2012. Archived at *perma.cc/3YRF-FTFJ*

[15] Mark Imbriaco. “Downtime Last Saturday.” *github.blog*, December 2012. Archived at *perma.cc/M7X5-E8SQ*

[16] John Hugg. “‘All In’ with Determinism for Performance and Testing in Distributed Systems.” At *Strange Loop*, September 2015.

[17] Hironobu Suzuki. “The Internals of PostgreSQL.” *interdb.jp*, 2017. Archived at *archive.org*

[18] Amit Kapila. “WAL Internals of PostgreSQL.” At *PostgreSQL Conference* (PGCon), May 2012. Archived at *perma.cc/6225-3SUX*

[19] Amit Kapila. “Evolution of Logical Replication.” *amitkapila16.blogspot.com*, September 2023. Archived at *perma.cc/F9VX-JLER*

[20] Aru Petchimuthu. “Upgrade Your Amazon RDS for PostgreSQL or Amazon Aurora PostgreSQL Database, Part 2: Using the pglogical Extension.” *aws.amazon.com*, August 2021. Archived at *perma.cc/RXT8-FS2T*

[21] Yogeshwer Sharma, Philippe Ajoux, Petchean Ang, David Callies, AbhishekChoudhary, Laurent Demailly, Thomas Fersch, Liat Atsmon Guz, Andrzej Kotulski, Sachin Kulkarni, Sanjeev Kumar, Harry Li, Jun Li, Evgeniy Makeev, Kowshik Prakasam, Robbert van Renesse, Sabyasachi Roy, Pratyush Seth, Yee Jiun Song, Benjamin Wester, Kaushik Veeraraghavan, and Peter Xie. “Wormhole: Reliable Pub-Sub to Support Geo-Replicated Internet Services.” At *12th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), May 2015.

[22] Douglas B. Terry. “Replicated Data Consistency Explained Through Baseball.” Microsoft Research, Technical Report MSR-TR-2011-137, October 2011. Archived at *perma.cc/F4KZ-AR38*

[23] Douglas B. Terry, Alan J. Demers, Karin Petersen, Mike J. Spreitzer, Marvin M.Theher, and Brent B. Welch. “Session Guarantees for Weakly Consistent Replicated Data.” At *3rd International Conference on Parallel and Distributed Information Systems* (PDIS), September 1994. *doi:10.1109/PDIS.1994.331722*

[24] Werner Vogels. “Eventually Consistent.” *ACM Queue*, volume 6, issue 6, pages 14–19, October 2008. *doi:10.1145/1466443.1466448*

[25] Simon Willison. Reply to: “My thoughts about Fly.io (so far) and other newish technology I’m getting into”. *news.ycombinator.com*, May 2022.

[26] Nithin Tharakan. “Scaling Bitbucket’s Database.” *atlassian.com*, October 2020. Archived at *perma.cc/JAB7-9FGX*

[27] Terry Pratchett. *Reaper Man: A Discworld Novel*. Victor Gollancz, 1991. ISBN: 9780575049796

[28] Peter Bailis, Alan Fekete, Michael J. Franklin, Ali Ghodsi, Joseph M. Heller-stein, and Ion Stoica. “Coordination Avoidance in Database Systems.” *Proceedings of the VLDB Endowment*, volume 8, issue 3, pages 185–196, November 2014. *doi:10.14778/2735508.2735509*

[29] Yaser Raja and Peter Celentano. “PostgreSQL Bi-Directional Replication Using pglogical.” *aws.amazon.com*, January 2022. Archived at *perma.cc/BUQ2-5QWN*

[30] Robert Hodges. “If You *Must* Deploy Multi-Master Replication, Read This First.” *scale-out-blog.blogspot.com*, April 2012. Archived at *perma.cc/C2JN-F6Y8* [31] Lars Hofhansl. “HBASE-7709: Infinite Loop Possible in Master/Master Replication.” *issues.apache.org*, January 2013. Archived at *perma.cc/24G2-8NLC*

[32] John Day-Richter. “What’s Different About the New Google Docs: Making Collaboration Fast.” *drive.googleblog.com*, September 2010. Archived at *perma.cc/5TL8TSJ2*

[33] Evan Wallace. “How Figma’s Multiplayer Technology Works.” *figma.com*, October 2019. Archived at *perma.cc/L49H-LY4D*

[34] Tuomas Artman. “Scaling the Linear Sync Engine.” *linear.app*, June 2023.

[35] Amr Saafan. “Why Sync Engines Might Be the Future of Web Applications.” *nilebits.com*, September 2024. Archived at *perma.cc/5N73-5M3V*

[36] Isaac Hagoel. “Are Sync Engines the Future of Web Applications?” *dev.to*, July 2024. Archived at *perma.cc/R9HF-BKKL*

[37] Sujay Jayakar. “A Map of Sync.” *stack.convex.dev*, October 2024. Archived at *perma.cc/82R3-H42A*

[38] Alex Feyerke. “Designing Offline-First Web Apps.” *alistapart.com*, December 2013. Archived at *perma.cc/WH7R-S2DS*

[39] Martin Kleppmann, Adam Wiggins, Peter van Hardenberg, and Mark McGrana-ghan. “Local-First Software: You Own Your Data, in Spite of the Cloud.” At *ACM SIGPLAN International Symposium on New Ideas, New Paradigms, and Reflections on Programming and Software* (Onward!), October 2019. *doi:10.1145/3359591.3359737*

[40] Martin Kleppmann. “The Past, Present, and Future of Local-First.” At *Local-First Conference*, May 2024.

[41] Conrad Hofmeyr. “API Calling Is to Sync Engines as jQuery Is to React.” *powersync.com*, November 2024. Archived at *perma.cc/2FP9-7WJJ*

[42] Peter van Hardenberg and Martin Kleppmann. “PushPin: Towards ProductionQuality Peer-to-Peer Collaboration.” At *7th Workshop on Principles and Practice of Consistency for Distributed Data* (PaPoC), April 2020. *doi:10.1145/3380787.3393683*

[43] Leonard Kawell, Jr., Steven Beckhardt, Timothy Halvorsen, Raymond Ozzie, and Irene Greif. “Replicated Document Management in a Group Communication System.” At *ACM Conference on Computer-Supported Cooperative Work* (CSCW), September 1988. *doi:10.1145/62266.1024798*

[44] Ricky Pusch. “Explaining How Fighting Games Use Delay-Based and Rollback Netcode.” *words.infil.net* and *arstechnica.com*, October 2019. Archived at *perma.cc/DE7W-RDJ8*

[45] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakula-pati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. “Dynamo: Amazon’s Highly Available Key-Value Store.” At *21st ACM Symposium on Operating Systems Principles* (SOSP), October 2007. *doi:10.1145/1323293.1294281*

[46] Marc Shapiro, Nuno Preguiça, Carlos Baquero, and Marek Zawirski.“Conflict-Free Replicated Data Types.” At *13th International Symposium on Stabilization, Safety, and Security of Distributed Systems* (SSS), October 2011. *doi:10.1007/978-3-642-24550-3_29*

[47] Chengzheng Sun and Clarence Ellis. “Operational Transformation in RealTime Group Editors: Issues, Algorithms, and Achievements.” At *ACM Conference on Computer Supported Cooperative Work* (CSCW), November 1998. *doi:10.1145/289444.289469*

[48] Joseph Gentle and Martin Kleppmann. “Collaborative Text Editing with Egwalker: Better, Faster, Smaller.” At *20th European Conference on Computer Systems* (EuroSys), March 2025. *doi:10.1145/3689031.3696076*

[49] Dharma Shukla. “Azure Cosmos DB: Pushing the Frontier of Globally Distributed Databases.” *azure.microsoft.com*, September 2018. Archived at *perma.cc/UT3B-HH6R*

[50] David K. Gifford. “Weighted Voting for Replicated Data.” At *7th ACM Symposium on Operating Systems Principles* (SOSP), December 1979. *doi:10.1145/800215.806583*

[51] Marc Brooker. “Dynamo, DynamoDB, and Aurora DSQL.” *brooker.co.za*, August 2025. Archived at *perma.cc/XG3C-ALDQ*

[52] Heidi Howard, Dahlia Malkhi, and Alexander Spiegelman. “Flexible Paxos: Quorum Intersection Revisited.” At *20th International Conference on Principles of Distributed Systems* (OPODIS), December 2016. *doi:10.4230/LIPIcs.OPODIS.2016.25*

[53] Joseph Blomstedt. “Bringing Consistency to Riak.” At *RICON West*, October 2012. Archived at *archive.org*

[54] Peter Bailis, Shivaram Venkataraman, Michael J. Franklin, Joseph M. Hellerstein, and Ion Stoica. “Quantifying Eventual Consistency with PBS.” *The VLDB Journal*, volume 23, issue 2, pages 279–302, April 2014. *doi:10.1007/s00778-013-0330-1*

[55] Colin Breck. “Shared-Nothing Architectures for Server Replication and Synchronization.” *blog.colinbreck.com*, December 2019. Archived at *perma.cc/48P3-J6CJ*

[56] Jeffrey Dean and Luiz André Barroso. “The Tail at Scale.” *Communications of the ACM*, volume 56, issue 2, pages 74–80, February 2013. *doi:10.1145/2408776.2408794*

[57] Peng Huang, Chuanxiong Guo, Lidong Zhou, Jacob R. Lorch, Yingnong Dang, Murali Chintalapati, and Randolph Yao. “Gray Failure: The Achilles’ Heel of CloudScale Systems.” At *16th Workshop on Hot Topics in Operating Systems* (HotOS), May 2017. *doi:10.1145/3102980.3103005*

[58] Leslie Lamport. “Time, Clocks, and the Ordering of Events in a Distributed System.” *Communications of the ACM*, volume 21, issue 7, pages 558–565, July 1978. *doi:10.1145/359545.359563*

[59] D. Stott Parker Jr., Gerald J. Popek, Gerard Rudisin, Allen Stoughton, BruceJ. Walker, Evelyn Walton, Johanna M. Chow, David Edwards, Stephen Kiser, and Charles Kline. “Detection of Mutual Inconsistency in Distributed Systems.” *IEEE Transactions on Software Engineering*, volume SE-9, issue 3, pages 240–247, May 1983. *doi:10.1109/TSE.1983.236733*

[60] Nuno Preguiça, Carlos Baquero, Paulo Sérgio Almeida, Victor Fonte, andRicardo Gonçalves. “Dotted Version Vectors: Logical Clocks for Optimistic Replication.” *arXiv:1011.5808*, November 2010.

[61] Giridhar Manepalli. “Clocks and Causality—Ordering Events in Distributed Systems.” *exhypothesi.com*, November 2022. Archived at *perma.cc/8REU-KVLQ*

[62] Sean Cribbs. “A Brief History of Time in Riak.” At *RICON*, October 2014. Archived at *perma.cc/7U9P-6JFX*

[63] Russell Brown. “Vector Clocks Revisited Part 2: Dotted Version Vectors.” *riak.com*, November 2015. Archived at *perma.cc/96QP-W98R*

[64] Carlos Baquero. “Version Vectors Are Not Vector Clocks.” *haslab.wordpress.com*, July 2011. Archived at *perma.cc/7PNU-4AMG*

[65] Reinhard Schwarz and Friedemann Mattern. “Detecting Causal Relationships in Distributed Computations: In Search of the Holy Grail.” *Distributed Computing*, volume 7, issue 3, pages 149–174, March 1994. *doi:10.1007/BF02277859*

บทที่ 7

---

[1] Claire Giordano. “Understanding Partitioning and Sharding in Postgres and Citus.” *citusdata.com*, August 2023. Archived at *perma.cc/8BTK-8959*

[2] Brandur Leach. “Partitioning in Postgres, 2022 Edition.” *brandur.org*, October 2022. Archived at *perma.cc/Z5LE-6AKX*

[3] Raph Koster. “Database ‘Sharding’ Came from UO?” *raphkoster.com*, January 2009. Archived at *perma.cc/4N9U-5KYF*

[4] Garrett Fidalgo. “Herding Elephants: Lessons Learned from Sharding Postgres at Notion.” *notion.com*, October 2021. Archived at *perma.cc/5J5V-W2VX*

[5] Ulrich Drepper. “What Every Programmer Should Know About Memory.” *akkadia.org*, November 2007. Archived at *perma.cc/NU6Q-DRXZ*

[6] Jingyu Zhou, Meng Xu, Alexander Shraer, Bala Namasivayam, Alex Miller, EvanTschannen, Steve Atherton, Andrew J. Beamon, Rusty Sears, John Leach, Dave Rosenthal, Xin Dong, Will Wilson, Ben Collins, David Scherer, Alec Grieser, Young Liu, Alvin Moore, Bhaskar Muppana, Xiaoge Su, and Vishesh Yadav. “FoundationDB: A Distributed Unbundled Transactional Key Value Store.” At *ACM International Conference on Management of Data* (SIGMOD), June 2021. *doi:10.1145/3448016.3457559*

[7] Marco Slot. “Citus 12: Schema-Based Sharding for PostgreSQL.” *citusdata.com*, July 2023. Archived at *perma.cc/R874-EC9W*

[8] Robisson Oliveira. “Reducing the Scope of Impact with Cell-Based Architecture.” AWS Well-Architected White Paper, Amazon Web Services, September 2023. Archived at *perma.cc/4KWW-47NR*

[9] Gwen Shapira. “Things DBs Don’t Do—But Should.” *thenile.dev*, February 2023. Archived at *perma.cc/C3J4-JSFW*

[10] Malte Schwarzkopf, Eddie Kohler, M. Frans Kaashoek, and Robert Morris. “Position: GDPR Compliance by Construction.” At *Towards Polystores That Manage Multiple Databases, Privacy, Security and/or Policy Issues for Heterogenous Data* (Poly), August 2019. *doi:10.1007/978-3-030-33752-0_3*

[11] Gwen Shapira. “Introducing pg_karnak: Transactional Schema Migration Across Tenant Databases.” *thenile.dev*, November 2024. Archived at *perma.cc/R5RD-8HR9*

[12] Arka Ganguli, Guido Iaquinti, Maggie Zhou, and Rafael Chacón. “Scaling Datastores at Slack with Vitess.” *slack.engineering*, December 2020. Archived at *perma.cc/UW8F-ALJK*

[13] Ikai Lan. “App Engine Datastore Tip: Monotonically Increasing Values Are Bad.” *ikaisays.com*, January 2011. Archived at *perma.cc/BPX8-RPJB*

[14] Enis Soztutar. “Apache HBase Region Splitting and Merging.” *cloudera.com*, February 2013. Archived at *perma.cc/S9HS-2X2C*

[15] Eric Evans. “Rethinking Topology in Cassandra.” At *Cassandra Summit*, June 2013. Archived at *perma.cc/2DKM-F438*

[16] Martin Kleppmann. “Java’s hashCode Is Not Safe for Distributed Systems.” *martin.kleppmann.com*, June 2012. Archived at *perma.cc/LK5U-VZSN*

[17] Mostafa Elhemali, Niall Gallagher, Nicholas Gordon, Joseph Idziorek, RichardKrog, Colin Lazier, Erben Mo, Akhilesh Mritunjai, Somu Perianayagam, Tim Rath, Swami Sivasubramanian, James Christopher Sorenson III, Sroaj Sosothikul, Doug Terry, and Akshat Vig. “Amazon DynamoDB: A Scalable, Predictably Performant, and Fully Managed NoSQL Database Service.” At *USENIX Annual Technical Conference* (ATC), July 2022.

[18] David Karger, Eric Lehman, Tom Leighton, Rina Panigrahy, Matthew Levine,and Daniel Lewin. “Consistent Hashing and Random Trees: Distributed Caching Protocols for Relieving Hot Spots on the World Wide Web.” At *29th Annual ACM Symposium on Theory of Computing* (STOC), May 1997. *doi:10.1145/258533.258660*

[19] Damian Gryski. “Consistent Hashing: Algorithmic Tradeoffs.” *dgryski.medium.com*, April 2018. Archived at *perma.cc/B2WF-TYQ8*

[20] David G. Thaler and Chinya V. Ravishankar. “Using Name-Based Mappings to Increase Hit Rates.” *IEEE/ACM Transactions on Networking*, volume 6, issue 1, pages 1–14, February 1998. *doi:10.1109/90.663936*

[21] John Lamping and Eric Veach. “A Fast, Minimal Memory, Consistent Hash Algorithm.” *arXiv:1406.2294*, June 2014.

[22] Samuel Axon. “3% of Twitter’s Servers Dedicated to Justin Bieber.” *mashable.com*, September 2010. Archived at *perma.cc/F35N-CGVX*

[23] Gerald Guo and Thawan Kooburat. “Scaling Services with Shard Manager.” engineering.fb.com, August 2020. Archived at perma.cc/EFS3-XQYT

[24] Sangmin Lee, Zhenhua Guo, Omer Sunercan, Jun Ying, Thawan Kooburat,Suryadeep Biswal, Jun Chen, Kun Huang, Yatpang Cheung, Yiding Zhou, Kaushik Veeraraghavan, Biren Damani, Pol Mauri Ruiz, Vikas Mehta, and Chunqiang Tang. “Shard Manager: A Generic Shard Management Framework for Geo-Distributed Applications.” At *28th ACM SIGOPS Symposium on Operating Systems Principles* (SOSP), October 2021. *doi:10.1145/3477132.3483546*

[25] Scott Lystig Fritchie. “A Critique of Resizable Hash Tables: Riak Core & Random Slicing.” *infoq.com*, August 2018. Archived at *perma.cc/RPX7-7BLN*

[26] Andy Warfield. “Building and Operating a Pretty Big Storage System Called S3.” *allthingsdistributed.com*, July 2023. Archived at *perma.cc/6S7P-GLM4*

[27] Rich Houlihan. “DynamoDB Adaptive Capacity: Smooth Performance for Chaotic Workloads (DAT327).” At *AWS re:Invent*, November 2017.

[28] Kostja Osipov. “ScyllaDB’s Safe Topology and Schema Changes on Raft.” scylladb.com, June 2024. Archived at perma.cc/4S82-M277

[29] Christopher D. Manning, Prabhakar Raghavan, and Hinrich Schütze.*Introduction to Information Retrieval*. Cambridge University Press, 2008. ISBN: 9780521865715. Available online at *nlp.stanford.edu/IR-book.*

[30] Michael Busch, Krishna Gade, Brian Larson, Patrick Lok, Samuel Luckenbill,and Jimmy Lin. “Earlybird: Real-Time Search at Twitter.” At *28th IEEE International Conference on Data Engineering* (ICDE), April 2012. *doi:10.1109/ICDE.2012.149*

[31] Nadav Har’El. “Indexing in Cassandra 3.” *github.com*, April 2017. Archived at *perma.cc/3ENV-8T9P*

[32] Zachary Tong. “Customizing Your Document Routing.” *elastic.co*, June 2013. Archived at *perma.cc/97VM-MREN*

[33] Andrew Pavlo. “H-Store Documentation: Frequently Asked Questions.” *hstore.cs.brown.edu*, October 2013. Archived at *perma.cc/X3ZA-DW6Z*

บทที่ 8

[1] Steven J. Murdoch. “What Went Wrong with Horizon: Learning from the Post Office Trial.” benthamsgaze.org, July 2021. Archived at perma.cc/CNM4-553F

[2] Donald D. Chamberlin, Morton M. Astrahan, Michael W. Blasgen, James N. Gray, W. Frank King, Bruce G. Lindsay, Raymond Lorie, James W. Mehl, Thomas G. Price, Franco Putzolu, Patricia Griffiths Selinger, Mario Schkolnick, Donald R. Slutz, Irving L. Traiger, Bradford W. Wade, and Robert A. Yost. “A History and Evaluation of System R.” Communications of the ACM, volume 24, issue 10, pages 632–646, October 1981. doi:10.1145/358769.358784

[3] Jim N. Gray, Raymond A. Lorie, Gianfranco R. Putzolu, and Irving L. Traiger.“Granularity of Locks and Degrees of Consistency in a Shared Data Base.” In *Modelling in Data Base Management Systems: Proceedings of the IFIP Working Conference on Modelling in Data Base Management Systems*, edited by G. M. Nijssen, pages 364–394, Elsevier/North Holland Publishing, 1976. Also in *Readings in Database*

Systems, 4th edition, edited by Joseph M. Hellerstein and Michael Stonebraker, MIT Press, 2005. ISBN: 9780262693141

[4] Kapali P. Eswaran, Jim N. Gray, Raymond A. Lorie, and Irving L. Traiger."The Notions of Consistency and Predicate Locks in a Database System." *Communications of the ACM*, volume 19, issue 11, pages 624–633, November 1976. *doi:10.1145/360363.360369*

[5] Rebecca Taft, Irfan Sharif, Andrei Matei, Nathan VanBenschoten, Jordan Lewis,Tobias Grieger, Kai Niemi, Andy Woods, Anne Birzin, Raphael Poss, Paul Bardea, Amruta Ranade, Ben Darnell, Bram Gruneir, Justin Jaffray, Lucy Zhang, and Peter Mattis. "CockroachDB: The Resilient Geo-Distributed SQL Database." At *ACM SIGMOD International Conference on Management of Data* (SIGMOD), June 2020. *doi:10.1145/3318464.3386134*

[6] Dongxu Huang, Qi Liu, Qiu Cui, Zhuhe Fang, Xiaoyu Ma, Fei Xu, Li Shen,Liu Tang, Yuxing Zhou, Menglong Huang, Wan Wei, Cong Liu, Jian Zhang, Jianjun Li, Xuelian Wu, Lingyu Song, Ruoxi Sun, Shuaipeng Yu, Lei Zhao, Nicholas Cameron, Liquan Pei, and Xin Tang. "TiDB: A Raft-Based HTAP Database." *Proceedings of the VLDB Endowment*, volume 13, issue 12, pages 3072–3084, August 2020. *doi:10.14778/3415478.3415535*

[7] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, ChristopherFrost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Dale Woodford, Yasushi Saito, Christopher Taylor, Michal Szymaniak, and Ruth Wang. "Spanner: Google's Globally-Distributed Database." At *10th USENIX Symposium on Operating System Design and Implementation* (OSDI), October 2012.

[8] Jingyu Zhou, Meng Xu, Alexander Shraer, Bala Namasivayam, Alex Miller, EvanTschannen, Steve Atherton, Andrew J. Beamon, Rusty Sears, John Leach, Dave Rosenthal, Xin Dong, Will Wilson, Ben Collins, David Scherer, Alec Grieser, Young Liu, Alvin Moore, Bhaskar Muppana, Xiaoge Su, and Vishesh Yadav. "FoundationDB: A Distributed Unbundled Transactional Key Value Store." At *ACM International Conference on Management of Data* (SIGMOD), June 2021. *doi:10.1145/3448016.3457559*

[9] Theo Härder and Andreas Reuter. "Principles of Transaction-Oriented Database Recovery." *ACM Computing Surveys*, volume 15, issue 4, pages 287–317, December 1983. *doi:10.1145/289.291*

[10] Peter Bailis, Alan Fekete, Ali Ghodsi, Joseph M. Hellerstein, and Ion Stoica."HAT, not CAP: Towards Highly Available Transactions." At *14th USENIX Workshop on Hot Topics in Operating Systems* (HotOS), May 2013.

[11] Armando Fox, Steven D. Gribble, Yatin Chawathe, Eric A. Brewer, and Paul Gauthier. "Cluster-Based Scalable Network Services." At *16th ACM Symposium on Operating Systems Principles* (SOSP), October 1997. *doi:10.1145/268998.266662*

[12] Tony Andrews. "Enforcing Complex Constraints in Oracle." *tonyandrews.blogspot.co.uk*, October 2004. Archived at *archive.org*

[13] Philip A. Bernstein, Vassos Hadzilacos, and Nathan Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987. ISBN: 9780201107159. Available online at *microsoft.com*.

[14] Alan Fekete, Dimitrios Liarokapis, Elizabeth O'Neil, Patrick O'Neil, and DennisShasha. "Making Snapshot Isolation Serializable." *ACM Transactions on Database Systems*, volume 30, issue 2, pages 492–528, June 2005. *doi:10.1145/1071610.1071615*

[15] Mai Zheng, Joseph Tucek, Feng Qin, and Mark Lillibridge. "Understanding the Robustness of SSDs Under Power Fault." At *11th USENIX Conference on File and Storage Technologies* (FAST), February 2013.

[16] Laurie Denness. "SSDs: A Gift and a Curse." *laur.ie*, June 2015. Archived at *perma.cc/6GLP-BX3T*

[17] Adam Surak. "When Solid State Drives Are Not That Solid." *blog.algolia.com*, June 2015. Archived at *perma.cc/CBR9-QZEE*

[18] Hewlett Packard Enterprise. “Bulletin: (Revision) HPE SAS Solid State Drives— Critical Firmware Upgrade Required for Certain HPE SAS Solid State Drive Models to Prevent Drive Failure at 32,768 Hours of Operation.” *support.hpe.com*, November 2019. Archived at *perma.cc/CZR4-AQBS*

[19] Craig Ringer et al. “PostgreSQL’s Handling of fsync() Errors Is Unsafe and Risks Data Loss at Least on XFS.” Email thread on *pgsql-hackers* mailing list, *postgresql.org*, March 2018. Archived at *perma.cc/5RKU-57FL*

[20] Anthony Rebello, Yuvraj Patel, Ramnatthan Alagappan, Andrea C. ArpaciDusseau, and Remzi H. Arpaci-Dusseau. “Can Applications Recover from fsync Failures?” At *USENIX Annual Technical Conference* (ATC), July 2020.

[21] Thanumalayan Sankaranarayana Pillai, Vijay Chidambaram, Ramnatthan Ala-gappan, Samer Al-Kiswany, Andrea C. Arpaci-Dusseau, and Remzi H. ArpaciDusseau. “Crash Consistency: Rethinking the Fundamental Abstractions of the File System.” *ACM Queue*, volume 13, issue 7, pages 20–28, July 2015. *doi:10.1145/2800695.2801719*

[22] Thanumalayan Sankaranarayana Pillai, Vijay Chidambaram, Ramnatthan Ala-gappan, Samer Al-Kiswany, Andrea C. Arpaci-Dusseau, and Remzi H. ArpaciDusseau. “All File Systems Are Not Created Equal: On the Complexity of Crafting Crash-Consistent Applications.” At *11th USENIX Symposium on Operating Systems Design and Implementation* (OSDI), October 2014.

[23] Chris Siebenmann. “Unix’s File Durability Problem.” *utcc.utoronto.ca*, April 2016. Archived at perma.cc/VSS8-5MC4

[24] Aishwarya Ganesan, Ramnatthan Alagappan, Andrea C. Arpaci-Dusseau, andRemzi H. Arpaci-Dusseau. “Redundancy Does Not Imply Fault Tolerance: Analysis of Distributed Storage Reactions to Single Errors and Corruptions.” At *15th USENIX Conference on File and Storage Technologies* (FAST), February 2017.

[25] Lakshmi N. Bairavasundaram, Garth R. Goodson, Bianca Schroeder, Andrea C.Arpaci-Dusseau, and Remzi H. Arpaci-Dusseau. “An Analysis of Data Corruption in the Storage Stack.” At *6th USENIX Conference on File and Storage Technologies* (FAST), February 2008.

[26] Richard van der Hoff. “How We Discovered, and Recovered from, Postgres Corruption on the matrix.org Homeserver.” *matrix.org*, July 2025. Archived at *perma.cc/CDF5-NRBK*

[27] Bianca Schroeder, Raghav Lagisetty, and Arif Merchant. “Flash Reliability in Production: The Expected and the Unexpected.” At *14th USENIX Conference on File and Storage Technologies* (FAST), February 2016.

[28] Don Allison. “SSD Storage—Ignorance of Technology Is No Excuse.” *blog.korelogic.com*, March 2015. Archived at *perma.cc/9QN4-9SNJ*

[29] Gordon Mah Ung. “Debunked: Your SSD Won’T Lose Data If Left Unplugged After All.” *pcworld.com*, May 2015. Archived at *perma.cc/S46H-JUDU*

[30] Martin Kleppmann. “Hermitage: Testing the ‘I’ in ACID.” *martin.kleppmann.com*, November 2014. Archived at *perma.cc/KP2Y-AQGK*

[31] Vlad Mihalcea. “The Race Condition That Led to Flexcoin Bankruptcy.” *vladmihalcea.com*, February 2025. Archived at *perma.cc/RRK5-TFAU*

[32] Todd Warszawski and Peter Bailis. “ACIDRain: Concurrency-Related Attacks on Database-Backed Web Applications.” At *ACM International Conference on Management of Data* (SIGMOD), May 2017. *doi:10.1145/3035918.3064037*

[33] Tristan D’Agosta. “BTC Stolen from Poloniex.” *bitcointalk.org*, March 2014. Archived at *perma.cc/YHA6-4C5D*

[34] bitcointhief2. “How I Stole Roughly 100 BTC from an Exchange and How I Could Have Stolen More!” *reddit.com*, February 2014. Archived at *archive.org*

[35] Sudhir Jorwekar, Alan Fekete, Krithi Ramamritham, and S. Sudarshan. “Automating the Detection of Snapshot Isolation Anomalies.” At *33rd International Conference on Very Large Data Bases* (VLDB), September 2007.

[36] Michael Melanson. “Transactions: The Limits of Isolation.” *michaelmelanson.net*, November 2014. Archived at *perma.cc/RG5R-KMYZ*

[37] Edward Kim. “How ACH Works: A Developer Perspective—Part 1.” *engineering.gusto.com*, April 2014. Archived at *perma.cc/7B2H-PU94*

[38] Hal Berenson, Philip A. Bernstein, Jim N. Gray, Jim Melton, ElizabethO’Neil, and Patrick O’Neil. “A Critique of ANSI SQL Isolation Levels.” At *ACM International Conference on Management of Data* (SIGMOD), May 1995. *doi:10.1145/568271.223785*

[39] Atul Adya. “Weak Consistency: A Generalized Theory and Optimistic Implementations for Distributed Transactions.” PhD thesis, Massachusetts Institute of Technology, March 1999. Archived at *perma.cc/E97M-HW5Q*

[40] Peter Bailis, Aaron Davidson, Alan Fekete, Ali Ghodsi, Joseph M. Hellerstein,and Ion Stoica. “Highly Available Transactions: Virtues and Limitations.” *Proceedings of the VLDB Endowment*, volume 7, issue 3, pages 181–192, November 2013. *doi:10.14778/2732232.2732237*.

[41] Natacha Crooks, Youer Pu, Lorenzo Alvisi, and Allen Clement. “Seeing Is Believing: A Client-Centric Specification of Database Isolation.” At *ACM Symposium on Principles of Distributed Computing* (PODC), July 2017. *doi:10.1145/3087801.3087802*

[42] Bruce Momjian. “MVCC Unmasked.” *momjian.us*, July 2014. Archived at *perma.cc/KQ47-9GYB*

[43] Peter Alvaro and Kyle Kingsbury. “MySQL 8.0.34.” *jepsen.io*, December 2023. Archived at *perma.cc/HGE2-Z878*

[44] Egor Rogov. *PostgreSQL 14 Internals*. Postgres Professional, April 2023. Archived at *perma.cc/FRK2-D7WB*

[45] Hironobu Suzuki. “The Internals of PostgreSQL.” *interdb.jp*, 2017.

[46] Rohan Reddy Alleti. “Internals of MVCC in Postgres: Hidden Costs of Updates vs Inserts.” *medium.com*, March 2025. Archived at *perma.cc/3ACX-DFXT*

[47] Andy Pavlo and Bohan Zhang. “The Part of PostgreSQL We Hate the Most.” cs.cmu.edu, April 2023. Archived at perma.cc/XSP6-3JBN

[48] Yingjun Wu, Joy Arulraj, Jiexi Lin, Ran Xian, and Andrew Pavlo. “An Empirical Evaluation of In-Memory Multi-Version Concurrency Control.” *Proceedings of the VLDB Endowment*, volume 10, issue 7, pages 781–792, March 2017. *doi:10.14778/3067421.3067427*

[49] Nikita Prokopov. “Unofficial Guide to Datomic Internals.” *tonsky.me*, May 2014. Archived at *perma.cc/ULM2-T2FW*

[50] Daniil Svetlov. “A Practical Guide to Taming Postgres Isolation Anomalies.” *dansvetlov.me*, March 2025. Archived at *perma.cc/L7LE-TDLS*

[51] Nate Wiger. “An Atomic Rant.” *nateware.com*, February 2010. Archived at *perma.cc/5ZYB-PE44*

[52] James Coglan. “Reading and Writing, Part 3: Web Applications.” *blog.jcoglan.com*, October 2020. Archived at *perma.cc/A7EK-PJVS*

[53] Peter Bailis, Alan Fekete, Michael J. Franklin, Ali Ghodsi, Joseph M. Hellerstein,and Ion Stoica. “Feral Concurrency Control: An Empirical Investigation of Modern Application Integrity.” At *ACM International Conference on Management of Data* (SIGMOD), June 2015. *doi:10.1145/2723372.2737784*

[54] Jaana Dogan. “Things I Wished More Developers Knew About Databases.” rakyll.medium.com, April 2020. Archived at perma.cc/6EFK-P2TD

[55] Michael J. Cahill, Uwe Röhm, and Alan Fekete. “Serializable Isolation for Snapshot Databases.” At *ACM International Conference on Management of Data* (SIGMOD), June 2008. *doi:10.1145/1376616.1376690*

[56] Dan R. K. Ports and Kevin Grittner. “Serializable Snapshot Isolation in PostgreSQL.” *Proceedings of the VLDB Endowment*, volume 5, issue 12, pages 1850–1861, August 2012. *doi:10.14778/2367502.2367523*

[57] Douglas B. Terry, Marvin M. Theimer, Karin Petersen, Alan J. Demers, Mike J.Spreitzer and Carl H. Hauser. “Managing Update Conflicts in Bayou, a Weakly Connected Replicated Storage System.” At *15th ACM Symposium on Operating Systems Principles* (SOSP), December 1995. *doi:10.1145/224056.224070*

[58] Hans-Jürgen Schönig. “Constraints over Multiple Rows in PostgreSQL.” *cybertecpostgresql.com*, June 2021. Archived at *perma.cc/2TGH-XUPZ*

[59] Michael Stonebraker, Samuel Madden, Daniel J. Abadi, Stavros Harizopoulos,Nabil Hachem, and Pat Helland. “The End of an Architectural Era (It’s Time for a Complete Rewrite).” At *33rd International Conference on Very Large Data Bases* (VLDB), September 2007.

[60] John Hugg. “H-Store/VoltDB Architecture vs. CEP Systems and Newer Streaming Architectures.” At *Data @Scale Boston*, November 2014.

[61] Robert Kallman, Hideaki Kimura, Jonathan Natkins, Andrew Pavlo, AlexanderRasin, Stanley Zdonik, Evan P. C. Jones, Samuel Madden, Michael Stonebraker, Yang Zhang, John Hugg, and Daniel J. Abadi. “H-Store: A High-Performance, Distributed Main Memory Transaction Processing System.” *Proceedings of the VLDB Endowment*, volume 1, issue 2, pages 1496–1499, August 2008. *doi:10.14778/1454159.1454211*

[62] Rich Hickey. “The Architecture of Datomic.” *infoq.com*, November 2012. Archived at *perma.cc/5YWU-8XJK*

[63] John Hugg. “Debunking Myths About the VoltDB In-Memory Database.” *dzone.com*, May 2014. Archived at *perma.cc/2Z9N-HPKF*

[64] Xinjing Zhou, Viktor Leis, Xiangyao Yu, and Michael Stonebraker. “OLTP Through the Looking Glass 16 Years Later: Communication Is the New Bottleneck.” At *15th Annual Conference on Innovative Data Systems Research* (CIDR), January 2025. Archived at *perma.cc/Q33D-K9YE*

[65] Xinjing Zhou, Xiangyao Yu, Goetz Graefe, and Michael Stonebraker. “Lotus: Scalable Multi-Partition Transactions On Single-Threaded Partitioned Databases.” *Proceedings of the VLDB Endowment* (PVLDB), volume 15, issue 11, pages 2939– 2952, July 2022. *doi:10.14778/3551793.3551843*

[66] Joseph M. Hellerstein, Michael Stonebraker, and James Hamilton. “Architecture of a Database System.” *Foundations and Trends in Databases*, volume 1, issue 2, pages 141–259, November 2007. *doi:10.1561/1900000002*

[67] Michael J. Cahill. “Serializable Isolation for Snapshot Databases.” PhD thesis, University of Sydney, July 2009. Archived at *perma.cc/727J-NTMP*

[68] Cristian Diaconu, Craig Freedman, Erik Ismert, Per-Åke Larson, Pravin Mit-tal, Ryan Stonecipher, Nitin Verma, and Mike Zwilling. “Hekaton: SQL Server’s Memory-Optimized OLTP Engine.” At *ACM SIGMOD International Conference on Management of Data* (SIGMOD), June 2013. *doi:10.1145/2463676.2463710*

[69] Thomas Neumann, Tobias Mühlbauer, and Alfons Kemper. “Fast Serializable Multi-Version Concurrency Control for Main-Memory Database Systems.” At *ACM SIGMOD International Conference on Management of Data* (SIGMOD), May 2015. *doi:10.1145/2723372.2749436*

[70] D. Z. Badal. “Correctness of Concurrency Control and Implications in Distributed Databases.” At *3rd International IEEE Computer Software and Applications Conference* (COMPSAC), November 1979. *doi:10.1109/CMPSAC.1979.762563*

[71] Rakesh Agrawal, Michael J. Carey, and Miron Livny. “Concurrency Control Performance Modeling: Alternatives and Implications.” *ACM Transactions on Database Systems* (TODS), volume 12, issue 4, pages 609–654, December 1987. *doi:10.1145/32204.32220*

[72] Marc Brooker. “Snapshot Isolation vs. Serializability.” *brooker.co.za*, December 2024. Archived at *perma.cc/5TRC-CR5G*

[73] B. G. Lindsay, P. G. Selinger, C. Galtieri, J. N. Gray, R. A. Lorie, T. G. Price, F. Put-zolu, I. L. Traiger, and B. W. Wade. “Notes on Distributed Databases.” IBM Research, Research Report RJ2571(33471), July 1979. Archived at *perma.cc/EPZ3-MHDD*

[74] C. Mohan, Bruce G. Lindsay, and Ron Obermarck. “Transaction Management in the R* Distributed Database Management System.” *ACM Transactions on Database Systems*, volume 11, issue 4, pages 378–396, December 1986. *doi:10.1145/7239.7266*

[75] X/Open Company Ltd. “Distributed Transaction Processing: The XA Specification.” Technical Standard XO/CAE/91/300, December 1991. ISBN: 9781872630243, archived at *perma.cc/Z96H-29JB*

[76] Ivan Silva Neto and Francisco Reverbel. “Lessons Learned from Implementing WS-Coordination and WS-AtomicTransaction.” At *7th IEEE/ACIS International Conference on Computer and Information Science* (ICIS), May 2008. *doi:10.1109/ ICIS.2008.75*

[77] James E. Johnson, David E. Langworthy, Leslie Lamport, and Friedrich H.Vogt. “Formal Specification of a Web Services Protocol.” At *1st International Workshop on Web Services and Formal Methods* (WS-FM), February 2004. *doi:10.1016/ j.entcs.2004.02.022*

[78] Jim Gray. “The Transaction Concept: Virtues and Limitations.” At *7th International Conference on Very Large Data Bases* (VLDB), September 1981.

[79] Dale Skeen. “Nonblocking Commit Protocols.” At *ACM International Conference on Management of Data* (SIGMOD), April 1981. *doi:10.1145/582318.582339*

[80] Gregor Hohpe. “Your Coffee Shop Doesn’t Use Two-Phase Commit.” *IEEE Software*, volume 22, issue 2, pages 64–66, March 2005. *doi:10.1109/MS.2005.52*

[81] Pat Helland. “Life Beyond Distributed Transactions: An Apostate’s Opinion.” At *3rd Biennial Conference on Innovative Data Systems Research* (CIDR), January 2007. Archived at *perma.cc/FC4F-AHGH*

[82] Jonathan Oliver. “My Beef with MSDTC and Two-Phase Commits.” *blog.jonathanoliver.com*, April 2011. Archived at *perma.cc/K8HF-Z4EN*

[83] Oren Eini (Ahende Rahien). “The Fallacy of Distributed Transactions.” ayende.com, July 2014. Archived at perma.cc/VB87-2JEF

[84] Clemens Vasters. “Transactions in Windows Azure (with Service Bus)—An Email Discussion.” *learn.microsoft.com*, July 2012. Archived at *perma.cc/4EZ9-5SKW*

[85] Ajmer Dhariwal. “Orphaned MSDTC Transactions (-2 spids).” *eraofdata.com*, December 2008. Archived at *perma.cc/YG6F-U34C*

[86] Paul Randal. “Real World Story of DBCC PAGE Saving the Day.” *sqlskills.com*, June 2013. Archived at *perma.cc/2MJN-A5QH*

[87] Guozhang Wang, Lei Chen, Ayusman Dikshit, Jason Gustafson, Boyang Chen,Matthias J. Sax, John Roesler, Sophie Blee-Goldman, Bruno Cadonna, Apurva Mehta, Varun Madan, and Jun Rao. “Consistency and Completeness: Rethinking Distributed Stream Processing in Apache Kafka.” At *ACM International Conference on Management of Data* (SIGMOD), June 2021. *doi:10.1145/3448016.3457556*

บทที่ 9

[1] Mark Cavage. “There’s Just No Getting Around It: You’re Building a Distributed System.” *ACM Queue*, volume 11, issue 4, pages 80–89, April 2013. *doi:10.1145/2466486.2482856*

[2] Jay Kreps. “Getting Real About Distributed System Reliability.” *blog.empathybox.com*, March 2012. Archived at *perma.cc/9B5Q-AEBW*

[3] Coda Hale. “You Can’t Sacrifice Partition Tolerance.” *codahale.com*, October 2010. Archived at *perma.cc/6GJU-X4G5*

[4] Jeff Hodges. “Notes on Distributed Systems for Young Bloods.” *somethingsimilar.com*, January 2013. Archived at *perma.cc/B636-62CE*

[5] Van Jacobson. “Congestion Avoidance and Control.” At *ACM Symposium on Communications Architectures and Protocols* (SIGCOMM), August 1988. *doi:10.1145/52324.52356*

[6] Bert Hubert. “The Ultimate SO_LINGER Page, or: Why Is My TCP Not Reliable.” *blog.netherlabs.nl*, January 2009. Archived at *perma.cc/6HDX-L2RR*

[7] Jerome H. Saltzer, David P. Reed, and David D. Clark. “End-To-End Arguments in System Design.” *ACM Transactions on Computer Systems*, volume 2, issue 4, pages 277–288, November 1984. *doi:10.1145/357401.357402*

[8] Peter Bailis and Kyle Kingsbury. “The Network Is Reliable.” *ACM Queue*, volume 12, issue 7, pages 48–55, July 2014. *doi:10.1145/2639988.2639988*

[9] Joshua B. Leners, Trinabh Gupta, Marcos K. Aguilera, and Michael Wal-fish. “Taming Uncertainty in Distributed Systems with Help from the Network.” At *10th European Conference on Computer Systems* (EuroSys), April 2015. *doi:10.1145/2741948.2741976*

[10] Phillipa Gill, Navendu Jain, and Nachiappan Nagappan. “Understanding Network Failures in Data Centers: Measurement, Analysis, and Implications.” At *ACM SIGCOMM Conference*, August 2011. *doi:10.1145/2018436.2018477*

[11] Urs Hölzle. “But recently a farmer had started grazing a herd of cows nearby. And whenever they stepped on the fiber link, they bent it enough to cause a blip.” *x.com*, May 2020. Archived at *perma.cc/WX8X-ZZA5*

[12] CBC News. “Hundreds Lose Internet Service in Northern B.C. After Beaver Chews Through Cable.” *cbc.ca*, April 2021. Archived at *perma.cc/UW8C-H2MY*

[13] Will Oremus. “The Global Internet Is Being Attacked by Sharks, Google Confirms.” *slate.com*, August 2014. Archived at *perma.cc/P6F3-C6YG*

[14] Jess Auerbach Jahajeeah. “Down to the Wire: The Ship Fixing Our Internet.” *continent.substack.com*, November 2023. Archived at *perma.cc/DP7B-EQ7S*

[15] Santosh Janardhan. “More Details About the October 4 Outage.” *engineering.fb.com*, October 2021. Archived at *perma.cc/WW89-VSXH*

[16] Tom Parfitt. “Georgian Woman Cuts off Web Access to Whole of Armenia.” *theguardian.com*, April 2011. Archived at *perma.cc/KMC3-N3NZ*

[17] Antonio Voce, Tural Ahmedzade and Ashley Kirk. “‘Shadow Fleets’ and Subaquatic Sabotage: Are Europe’s Undersea Internet Cables Under Attack?” *theguardian.com*, March 2025. Archived at *perma.cc/HA7S-ZDBV*

[18] Shengyun Liu, Paolo Viotti, Christian Cachin, Vivien Quéma, and MarkoVukolić. “XFT: Practical Fault Tolerance Beyond Crashes.” At *12th USENIX Symposium on Operating Systems Design and Implementation* (OSDI), November 2016.

[19] Mark Imbriaco. “Downtime Last Saturday.” *github.blog*, December 2012. Archived at *perma.cc/M7X5-E8SQ*

[20] Tom Lianza and Chris Snook. “A Byzantine Failure in the Real World.” *blog.cloudflare.com*, November 2020. Archived at *perma.cc/83EZ-ALCY*

[21] Mohammed Alfatafta, Basil Alkhatib, Ahmed Alquraan, and Samer Al-Kiswany.“Toward a Generic Fault Tolerance Technique for Partial Network Partitioning.” At *14th USENIX Symposium on Operating Systems Design and Implementation* (OSDI), November 2020.

[22] Marc A. Donges. “Re: bnx2 Cards Intermittantly Going Offline.” Message to Linux *netdev* mailing list, *spinics.net*, September 2012. Archived at *perma.cc/TXP6-H8R3*

[23] Troy Toman. “Inside a CODE RED: Network Edition.” *signalvnoise.com*, September 2020. Archived at *perma.cc/BET6-FY25*

[24] Kyle Kingsbury. “Jepsen: Elasticsearch.” *aphyr.com*, June 2014. Archived at *perma.cc/JK47-S89J*

[25] Salvatore Sanfilippo. “A Few Arguments About Redis Sentinel Properties and Fail Scenarios.” *antirez.com*, October 2014. Archived at *perma.cc/8XEU-CLM8*

[26] Nicolas Liochon. “CAP: If All You Have Is a Timeout, Everything Looks Like a Partition.” *blog.thislongrun.com*, May 2015. Archived at *perma.cc/FS57-V2PZ*

[27] Matthew P. Grosvenor, Malte Schwarzkopf, Ionel Gog, Robert N. M. Watson, Andrew W. Moore, Steven Hand, and Jon Crowcroft. “Queues Don’t Matter When You Can JUMP Them!” At *12th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), May 2015.

[28] Theo Julienne. “Debugging Network Stalls on Kubernetes.” *github.blog*, November 2019. Archived at *perma.cc/K9M8-XVGL*

[29] Guohui Wang and T. S. Eugene Ng. “The Impact of Virtualization on Network Performance of Amazon EC2 Data Center.” At *29th IEEE International Conference on Computer Communications* (INFOCOM), March 2010. *doi:10.1109/ INFCOM.2010.5461931*

[30] Brandon Philips. “etcd: Distributed Locking and Service Discovery.” At *Strange Loop*, September 2014.

[31] Steve Newman. “A Systematic Look at EC2 I/O.” *blog.scalyr.com*, October 2012. Archived at *perma.cc/FL4R-H2VE*

[32] Naohiro Hayashibara, Xavier Défago, Rami Yared, and Takuya Katayama.“The ϕ Accrual Failure Detector.” Japan Advanced Institute of Science and Technology, School of Information Science, Technical Report IS-RR-2004-010, May 2004. Archived at *perma.cc/NSM2-TRYA*

[33] Jeffrey Wang. “Phi Accrual Failure Detector.” *ternarysearch.blogspot.co.uk*, August 2013. Archived at *perma.cc/L452-AMLV*

[34] Srinivasan Keshav. *An Engineering Approach to Computer Networking: ATM Networks, the Internet, and the Telephone Network*. Addison-Wesley Professional, 1997. ISBN: 9780201634426

[35] Othmar Kyas. *ATM Networks*. International Thomson Publishing, 1995. ISBN: 9781850321286

[36] Jialin Li, Naveen Kr. Sharma, Dan R. K. Ports, and Steven D. Gribble. “Tales of the Tail: Hardware, OS, and Application-level Sources of Tail Latency.” At *ACM Symposium on Cloud Computing* (SOCC), November 2014. *doi:10.1145/2670979.2670988*

[37] Mellanox Technologies. “InfiniBand FAQ, Rev 1.3.” *network.nvidia.com*, December 2014. Archived at *perma.cc/LQJ4-QZVK*

[38] Jose Renato Santos, Yoshio Turner, and G. (John) Janakiraman. “End-toEnd Congestion Control for InfiniBand.” At *22nd Annual Joint Conference of the IEEE Computer and Communications Societies* (INFOCOM), April 2003. Also published by HP Laboratories Palo Alto, Tech Report HPL-2002-359. *doi:10.1109/ INFCOM.2003.1208949*

[39] Ulrich Windl, David Dalton, Marc Martinec, and Dale R. Worley. “The NTP FAQ and HOWTO.” *ntp.org*, November 2006. Archived at *archive.org*

[40] John Graham-Cumming. “How and Why the Leap Second Affected Cloudflare DNS.” *blog.cloudflare.com*, January 2017. Archived at *archive.org*

[41] David Holmes. “Inside the Hotspot VM: Clocks, Timers and Scheduling Events —Part I— Windows.” *blogs.oracle.com*, October 2006. Archived at *archive.org*

[42] Joran Dirk Greef. “Three Clocks Are Better than One.” *tigerbeetle.com*, August 2021. Archived at *perma.cc/5RXG-EU6B*

[43] Oliver Yang. “Pitfalls of TSC Usage.” *oliveryang.net*, September 2015. Archived at *perma.cc/Z2QY-5FRA*

[44] Steve Loughran. “Time on Multi-Core, Multi-Socket Servers.” *steveloughran.blogspot.co.uk*, September 2015. Archived at *perma.cc/7M4S-D4U6*

[45] James C. Corbett, Jeffrey Dean, Michael Epstein, Andrew Fikes, ChristopherFrost, JJ Furman, Sanjay Ghemawat, Andrey Gubarev, Christopher Heiser, Peter Hochschild, Wilson Hsieh, Sebastian Kanthak, Eugene Kogan, Hongyi Li, Alexander Lloyd, Sergey Melnik, David Mwaura, David Nagle, Sean Quinlan, Rajesh Rao, Lindsay Rolig, Dale Woodford, Yasushi Saito, Christopher Taylor, Michal Szymaniak, and Ruth Wang. “Spanner: Google’s Globally-Distributed Database.” At *10th USENIX Symposium on Operating System Design and Implementation* (OSDI), October 2012.

[46] M. Caporaloni and R. Ambrosini. “How Closely Can a Personal Computer Clock Track the UTC Timescale Via the Internet?” *European Journal of Physics*, volume 23, issue 4, pages L17–L21, June 2012. *doi:10.1088/0143-0807/23/4/103*

[47] Nelson Minar. “A Survey of the NTP Network.” *alumni.media.mit.edu*, December 1999. Archived at *perma.cc/EV76-7ZV3*

[48] Viliam Holub. “Synchronizing Clocks in a Cassandra Cluster Pt. 1—The Problem.” *blog.rapid7.com*, March 2014. Archived at *perma.cc/N3RV-5LNL*

[49] Poul-Henning Kamp. “The One-Second War (What Time Will You Die?)” *ACM Queue*, volume 9, issue 4, pages 44–48, April 2011. *doi:10.1145/1966989.1967009*

[50] Nelson Minar. “Leap Second Crashes Half the Internet.” *somebits.com*, July 2012. Archived at *perma.cc/2WB8-D6EU*

[51] Christopher Pascoe. “Time, Technology and Leaping Seconds.” *googleblog.blogspot.co.uk*, September 2011. Archived at *perma.cc/U2JL-7E74*

[52] Mingxue Zhao and Jeff Barr. “Look Before You Leap—The Coming Leap Second and AWS.” *aws.amazon.com*, May 2015. Archived at *perma.cc/KPE9-XMFM*

[53] Darryl Veitch and Kanthaiah Vijayalayan. “Network Timing and the 2015 Leap Second.” At *17th International Conference on Passive and Active Measurement* (PAM), April 2016. *doi:10.1007/978-3-319-30505-9_29*

[54] VMware, Inc. “Timekeeping in VMware Virtual Machines.” *vmware.com*, October 2008. Archived at *perma.cc/HM5R-T5NF*

[55] Victor Yodaiken. “Clock Synchronization in Finance and Beyond.” *yodaiken.com*, November 2017. Archived at *perma.cc/9XZD-8ZZN*

[56] Mustafa Emre Acer, Emily Stark, Adrienne Porter Felt, Sascha Fahl, RadhikaBhargava, Bhanu Dev, Matt Braithwaite, Ryan Sleevi, and Parisa Tabriz. “Where the Wild Warnings Are: Root Causes of Chrome HTTPS Certificate Errors.” At *ACM SIGSAC Conference on Computer and Communications Security* (CCS), October 2017. *doi:10.1145/3133956.3134007*

[57] European Securities and Markets Authority. “MiFID II / MiFIR: Regulatory Technical and Implementing Standards—Annex I.” *esma.europa.eu*, Report ESMA/ 2015/1464, September 2015. Archived at *perma.cc/ZLX9-FGQ3*

[58] Luke Bigum. “Solving MiFID II Clock Synchronisation with Minimum Spend (Part 1).” *catach.blogspot.com*, November 2015. Archived at *perma.cc/4J5W-FNM4*

[59] Oleg Obleukhov and Ahmad Byagowi. “How Precision Time Protocol Is Being Deployed at Meta.” *engineering.fb.com*, November 2022. Archived at *perma.cc/29G6-UJNW*

[60] John Wiseman. “GPSJAM: Daily Maps of GPS Interference.” *gpsjam.org*

[61] Josh Levinson, Julien Ridoux, and Chris Munns. “It’s About Time: MicrosecondAccurate Clocks on Amazon EC2 Instances.” *aws.amazon.com*, November 2023. Archived at *perma.cc/56M6-5VMZ*

[62] Kyle Kingsbury. “Jepsen: Cassandra.” *aphyr.com*, September 2013. Archived at *perma.cc/4MBR-J96V*

[63] John Daily. “Clocks Are Bad, or, Welcome to the Wonderful World of Distributed Systems.” *riak.com*, November 2013. Archived at *perma.cc/4XB5-UCXY*

[64] Marc Brooker. “It’s About Time!” *brooker.co.za*, November 2023. Archived at *perma.cc/N6YK-DRPA*

[65] Kyle Kingsbury. “The Trouble with Timestamps.” *aphyr.com*, October 2013. Archived at *perma.cc/W3AM-5VAV*

[66] Leslie Lamport. “Time, Clocks, and the Ordering of Events in a Distributed System.” *Communications of the ACM*, volume 21, issue 7, pages 558–565, July 1978. *doi:10.1145/359545.359563*

[67] Justin Sheehy. “There Is No Now: Problems with Simultaneity in Distributed Systems.” *ACM Queue*, volume 13, issue 3, pages 36–41, March 2015. *doi:10.1145/2733108*

[68] Murat Demirbas. “Spanner: Google’s Globally-Distributed Database.” *muratbuffalo.blogspot.co.uk*, July 2013. Archived at *perma.cc/6VWR-C9WB*

[69] Dahlia Malkhi and Jean-Philippe Martin. “Spanner’s Concurrency Control.” *ACM SIGACT News*, volume 44, issue 3, pages 73–77, September 2013. *doi:10.1145/2527748.2527767*

[70] Franck Pachot. “Achieving Precise Clock Synchronization on AWS.” *yugabyte.com*, December 2024. Archived at *perma.cc/UYM6-RNBS*

[71] Spencer Kimball. “Living Without Atomic Clocks: Where CockroachDB and Spanner Diverge.” *cockroachlabs.com*, January 2022. Archived at *perma.cc/AWZ7-RXFT*

[72] Murat Demirbas. “Use of Time in Distributed Databases (Part 4): Synchronized Clocks in Production Databases.” *muratbuffalo.blogspot.com*, January 2025. Archived at *perma.cc/9WNX-Q9U3*

[73] Cary G. Gray and David R. Cheriton. “Leases: An Efficient Fault-Tolerant Mechanism for Distributed File Cache Consistency.” At *12th ACM Symposium on Operating Systems Principles* (SOSP), December 1989. *doi:10.1145/74850.74870*

[74] Daniel Sturman, Scott Delap, Max Ross, et al. “Roblox Return to Service.” *corp.roblox.com*, January 2022. Archived at *perma.cc/8ALT-WAS4*

[75] Todd Lipcon. “Avoiding Full GCs with MemStore-Local Allocation Buffers.” *slideshare.net*, February 2011. Archived at *perma.cc/CH62-2EWJ*

[76] Christopher Clark, Keir Fraser, Steven Hand, Jacob Gorm Hansen, Eric Jul,Christian Limpach, Ian Pratt, and Andrew Warfield. “Live Migration of Virtual Machines.” At *2nd USENIX Symposium on Symposium on Networked Systems Design & Implementation* (NSDI), May 2005.

[77] Mike Shaver. “fsyncers and Curveballs.” *shaver.off.net*, May 2008. Archived at *archive.org*

[78] Zhenyun Zhuang and Cuong Tran. “Eliminating Large JVM GC Pauses Caused by Background IO Traffic.” *engineering.linkedin.com*, February 2016. Archived at *perma.cc/ML2M-X9XT*

[79] Martin Thompson. “Java Garbage Collection Distilled.” *mechanicalsympathy.blogspot.co.uk*, July 2013. Archived at *perma.cc/DJT3-NQLQ*

[80] David Terei and Amit Levy. “Blade: A Data Center Garbage Collector.” *arXiv:1504.02578*, April 2015.

[81] Martin Maas, Tim Harris, Krste Asanović, and John Kubiatowicz. “Trash Day: Coordinating Garbage Collection in Distributed Systems.” At *15th USENIX Workshop on Hot Topics in Operating Systems* (HotOS), May 2015.

[82] Martin Fowler. “The LMAX Architecture.” *martinfowler.com*, July 2011. Archived at *perma.cc/5AV4-N6RJ*

[83] Joseph Y. Halpern and Yoram Moses. “Knowledge and Common Knowledge in a Distributed Environment.” *Journal of the ACM* (JACM), volume 37, issue 3, pages 549–587, July 1990. *doi:10.1145/79147.79161*

[84] Chuzhe Tang, Zhaoguo Wang, Xiaodong Zhang, Qianmian Yu, Binyu Zang,Haibing Guan, and Haibo Chen. “Ad Hoc Transactions in Web Applications: The Good, the Bad, and the Ugly.” At *ACM International Conference on Management of Data* (SIGMOD), June 2022. *doi:10.1145/3514221.3526120*

[85] Flavio P. Junqueira and Benjamin Reed. *ZooKeeper: Distributed Process Coordination*. O’Reilly Media, 2013. ISBN: 9781449361303

[86] Enis Söztutar. “HBase and HDFS: Understanding Filesystem Usage in HBase.” At *HBaseCon*, June 2013. Archived at *perma.cc/4DXR-9P88*

[87] SUSE LLC. “SUSE Linux Enterprise High Availability 15 SP6 Administration Guide, Section 12: Fencing and STONITH.” *documentation.suse.com*, March 2025. Archived at *perma.cc/8LAR-EL9D*

[88] Mike Burrows. “The Chubby Lock Service for Loosely-Coupled Distributed Systems.” At *7th USENIX Symposium on Operating System Design and Implementation* (OSDI), November 2006.

[89] Kyle Kingsbury. “etcd 3.4.3.” *jepsen.io*, January 2020. Archived at *perma.cc/2P3Y-MPWU*

[90] Ensar Basri Kahveci. “Distributed Locks Are Dead; Long Live Distributed Locks!” *hazelcast.com*, April 2019. Archived at *perma.cc/7FS5-LDXE*

[91] Martin Kleppmann. “How to Do Distributed Locking.” *martin.kleppmann.com*, February 2016. Archived at *perma.cc/Y24W-YQ5L*

[92] Salvatore Sanfilippo. “Is Redlock Safe?” *antirez.com*, February 2016. Archived at *perma.cc/B6GA-9Q6A*

[93] Gunnar Morling. “Leader Election with S3 Conditional Writes.” *morling.dev*, August 2024. Archived at *perma.cc/7V2N-J78Y*

[94] Leslie Lamport, Robert Shostak, and Marshall Pease. “The Byzantine Generals Problem.” *ACM Transactions on Programming Languages and Systems* (TOPLAS), volume 4, issue 3, pages 382–401, July 1982. *doi:10.1145/357172.357176*

[95] Jim N. Gray. “Notes on Data Base Operating Systems.” In *Operating Systems: An Advanced Course*, Lecture Notes in Computer Science, volume 60, edited by R. Bayer, R. M. Graham, and G. Seegmüller, pages 393–481, Springer-Verlag, 1978. ISBN: 9783540087557. Archived at *perma.cc/7S9M-2LZU*

[96] Brian Palmer. “How Complicated Was the Byzantine Empire?” *slate.com*, October 2011. Archived at *perma.cc/AN7X-FL3N*

[97] Leslie Lamport. “My Writings.” *lamport.azurewebsites.net*, December 2014. Archived at *perma.cc/5NNM-SQGR*

[98] John Rushby. “Bus Architectures for Safety-Critical Embedded Systems.” At *1st International Workshop on Embedded Software* (EMSOFT), October 2001. *doi:10.1007/3-540-45449-7_22*

[99] Jake Edge. “ELC: SpaceX Lessons Learned.” *lwn.net*, March 2013. Archived at *perma.cc/AYX8-QP5X*

[100] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, Patrick McCorry, Sarah Meiklejohn, and George Danezis. “SoK: Consensus in the Age of Blockchains.” At *1st ACM Conference on Advances in Financial Technologies* (AFT), October 2019. *doi:10.1145/3318041.3355458*

[101] Ezra Feilden, Adi Oltean, and Philip Johnston. “Why We Should Train AI in Space.” White Paper, *starcloud.com*, September 2024. Archived at *perma.cc/ 7Y3S-8UB6*

[102] James Mickens. “The Saddest Moment.” *USENIX ;login*, May 2013. Archived at *perma.cc/T7BZ-XCFR*

[103] Martin Kleppmann and Heidi Howard. “Byzantine Eventual Consistency and the Fundamental Limits of Peer-to-Peer Databases.” *arXiv:2012.00472*, December 2020.

[104] Martin Kleppmann. “Making CRDTs Byzantine Fault Tolerant.” At *9th Workshop on Principles and Practice of Consistency for Distributed Data* (PaPoC), April 2022. *doi:10.1145/3517209.3524042*

[105] Evan Gilman. “The Discovery of Apache ZooKeeper’s Poison Packet.” *pagerduty.com*, May 2015. Archived at *perma.cc/RV6L-Y5CQ*

[106] Jonathan Stone and Craig Partridge. “When the CRC and TCP Checksum Disagree.” At *ACM Conference on Applications, Technologies, Architectures, and Protocols for Computer Communication* (SIGCOMM), August 2000. *doi:10.1145/347059.347561*

[107] Evan Jones. “How Both TCP and Ethernet Checksums Fail.” *evanjones.ca*, October 2015. Archived at *perma.cc/9T5V-B8X5*

[108] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. “Consensus in the Presence of Partial Synchrony.” *Journal of the ACM*, volume 35, issue 2, pages 288–323, April 1988. *doi:10.1145/42282.42283*

[109] Richard D. Schlichting and Fred B. Schneider. “Fail-Stop Processors: An Approach to Designing Fault-Tolerant Computing Systems.” *ACM Transactions on Computer Systems* (TOCS), volume 1, issue 3, pages 222–238, August 1983. *doi:10.1145/357369.357371*

[110] Thanh Do, Mingzhe Hao, Tanakorn Leesatapornwongsa, Tiratat Patana-anake,and Haryadi S. Gunawi. “Limplock: Understanding the Impact of Limpware on Scaleout Cloud Systems.” At *4th ACM Symposium on Cloud Computing* (SoCC), October 2013. *doi:10.1145/2523616.2523627*

[111] Josh Snyder and Joseph Lynch. “Garbage Collecting Unhealthy JVMs, a Proactive Approach.” *netflixtechblog.medium.com*, November 2019. Archived at *perma.cc/ 8BTA-N3YB*

[112] Haryadi S. Gunawi, Riza O. Suminto, Russell Sears, Casey Golliher, Swamina-than Sundararaman, Xing Lin, Tim Emami, Weiguang Sheng, Nematollah Bidokhti, Caitie McCaffrey, Gary Grider, Parks M. Fields, Kevin Harms, Robert B. Ross, Andree Jacobson, Robert Ricci, Kirk Webb, Peter Alvaro, H. Birali Runesha, Mingzhe Hao, and Huaicheng Li. “Fail-Slow at Scale: Evidence of Hardware Performance Faults in Large Production Systems.” At *16th USENIX Conference on File and Storage Technologies*, February 2018.

[113] Peng Huang, Chuanxiong Guo, Lidong Zhou, Jacob R. Lorch, Yingnong Dang,Murali Chintalapati, and Randolph Yao. “Gray Failure: The Achilles’ Heel of CloudScale Systems.” At *16th Workshop on Hot Topics in Operating Systems* (HotOS), May 2017. *doi:10.1145/3102980.3103005*

[114] Chang Lou, Peng Huang, and Scott Smith. “Understanding, Detecting and Localizing Partial Failures in Large System Software.” At *17th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), February 2020.

[115] Peter Bailis and Ali Ghodsi. “Eventual Consistency Today: Limitations, Extensions, and Beyond.” *ACM Queue*, volume 11, issue 3, pages 55–63, March 2013. *doi:10.1145/2460276.2462076*

[116] Bowen Alpern and Fred B. Schneider. “Defining Liveness.” *Information Processing Letters*, volume 21, issue 4, pages 181–185, October 1985. *doi:10.1016/0020-0190(85)90056-0*

[117] Flavio P. Junqueira. “Dude, Where’s My Metadata?” *fpj.me*, May 2015. Archived at *perma.cc/D2EU-Y9S5*

[118] Scott Sanders. “January 28th Incident Report.” *github.com*, February 2016. Archived at *perma.cc/5GZR-88TV*

[119] Jay Kreps. “A Few Notes on Kafka and Jepsen.” *blog.empathybox.com*, September 2013. Archived at *perma.cc/XJ5C-F583*

[120] Marc Brooker and Ankush Desai. “Systems Correctness Practices at AWS.” *ACM Queue*, volume 22, issue 6, pages 79–96, November/December 2024. *doi:10.1145/3712057*

[121] Andrey Satarin. “Testing Distributed Systems: Curated list of Resources on Testing Distributed Systems.” *asatarin.github.io*. Archived at *perma.cc/U5V8-XP24*

[122] Phil Eaton and Joran Dirk Greef. “We Put a Distributed Database in the Browser—And Made a Game of It!” *tigerbeetle.com*, June 2023. Archived at *perma.cc/L7M7-X4HD*

[123] Apple, Inc. and the FoundationDB project authors. “FoundationDB—Simulation and Testing.” *apple.github.io*. Archived at *perma.cc/4C4L-AUH3*

[124] Jack Vanlightly. “Verifying Kafka Transactions—Diary Entry 2—Writing an Initial TLA+ Spec.” *jack-vanlightly.com*, December 2024. Archived at *perma.cc/NSQ8-MQ5N*

[125] Siddon Tang. “From Chaos to Order—Tools and Techniques for Testing TiDB, A Distributed NewSQL Database.” *pingcap.com*, April 2018. Archived at *perma.cc/ 5EJB-R29F*

[126] Nathan VanBenschoten. “Parallel Commits: An Atomic Commit Protocol for Globally Distributed Transactions.” *cockroachlabs.com*, November 2019. Archived at *perma.cc/5FZ7-QK6J*

[127] Jack Vanlightly. “Paper: VR Revisited—State Transfer (Part 3).” *jackvanlightly.com*, December 2022. Archived at *perma.cc/KNK3-K6WS*

[128] Hillel Wayne. “What If the Spec Doesn’t Match the Code?” *buttondown.com*, March 2024. Archived at *perma.cc/8HEZ-KHER*

[129] Lingzhi Ouyang, Xudong Sun, Ruize Tang, Yu Huang, Madhav Jivrajani, Xiaox-ing Ma, Tianyin Xu. “Multi-Grained Specifications for Distributed System Model Checking and Verification.” At *20th European Conference on Computer Systems* (EuroSys), March 2025. *doi:10.1145/3689031.3696069*

[130] Yury Izrailevsky and Ariel Tseitlin. “The Netflix Simian Army.” *netflixtechblog.com*, July, 2011. Archived at *perma.cc/M3NY-FJW6*

[131] Kyle Kingsbury. “Jepsen: On the Perils of Network Partitions.” *aphyr.com*, May, 2013. Archived at *perma.cc/W98G-6HQP*

[132] Kyle Kingsbury. Analyses. *jepsen.io*, 2024. Archived at *perma.cc/8LDN-D2T8* [133] Rupak Majumdar and Filip Niksic. “Why Is Random Testing Effective for Partition Tolerance Bugs?” *Proceedings of the ACM on Programming Languages* (PACMPL), volume 2, issue POPL, article no. 46, December 2017. *doi:10.1145/3158134*

[134] FoundationDB project authors. “Simulation and Testing.” *apple.github.io*. Archived at *perma.cc/NQ3L-PM4C*

[135] Alex Kladov. “Simulation Testing for Liveness.” *tigerbeetle.com*, July 2023. Archived at *perma.cc/RKD4-HGCR*

[136] Alfonso Subiotto Marqués. “(Mostly) Deterministic Simulation Testing in Go.” *polarsignals.com*, May 2024. Archived at *perma.cc/ULD6-TSA4*

บทที่ 10

[1] Maurice P. Herlihy and Jeannette M. Wing. “Linearizability: A Correctness Condition for Concurrent Objects.” *ACM Transactions on Programming Languages and Systems* (TOPLAS), volume 12, issue 3, pages 463–492, July 1990. *doi:10.1145/78969.78972*

[2] Leslie Lamport. “On Interprocess Communication.” *Distributed Computing*, volume 1, issue 2, pages 77–101, June 1986. *doi:10.1007/BF01786228*

[3] David K. Gifford. “Information Storage in a Decentralized Computer System.” Xerox Palo Alto Research Centers, CSL-81-8, June 1981. Archived at *perma.cc/ 2XXP-3JPB*

[4] Martin Kleppmann. “Please Stop Calling Databases CP or AP.” *martin.kleppmann.com*, May 2015. Archived at *perma.cc/MJ5G-75GL*

[5] Kyle Kingsbury. “Jepsen: MongoDB Stale Reads.” *aphyr.com*, April 2015. Archived at *perma.cc/DXB4-J4JC*

[6] Kyle Kingsbury. “Computational Techniques in Knossos.” *aphyr.com*, May 2014. Archived at *perma.cc/2X5M-EHTU*

[7] Kyle Kingsbury and Peter Alvaro. “Elle: Inferring Isolation Anomalies from Experimental Observations.” *Proceedings of the VLDB Endowment*, volume 14, issue 3, pages 268–280, November 2020. *doi:10.14778/3430915.3430918*

[8] Paolo Viotti and Marko Vukolić. “Consistency in Non-Transactional Distributed Storage Systems.” *ACM Computing Surveys* (CSUR), volume 49, issue 1, article no. 19, June 2016. *doi:10.1145/2926965*

[9] Peter Bailis. “Linearizability Versus Serializability.” *bailis.org*, September 2014. Archived at *perma.cc/386B-KAC3*

[10] Daniel Abadi. “Correctness Anomalies Under Serializable Isolation.” *dbmsmusings.blogspot.com*, June 2019. Archived at *perma.cc/JGS7-BZFY*

[11] Peter Bailis, Aaron Davidson, Alan Fekete, Ali Ghodsi, Joseph M. Hellerstein,and Ion Stoica. “Highly Available Transactions: Virtues and Limitations.” *Proceedings of the VLDB Endowment*, volume 7, issue 3, pages 181–192, November 2013. *doi:10.14778/2732232.2732237*, extended version published as *arXiv:1302.0309*

[12] Philip A. Bernstein, Vassos Hadzilacos, and Nathan Goodman. *Concurrency Control and Recovery in Database Systems*. Addison-Wesley, 1987. ISBN: 9780201107159. Available online at *microsoft.com*.

[13] Andrei Matei. “CockroachDB’s Consistency Model.” *cockroachlabs.com*, February 2021. Archived at *perma.cc/MR38-883B*

[14] Murat Demirbas. “Strict-Serializability, but at What Cost, for What Purpose?” *muratbuffalo.blogspot.com*, August 2022. Archived at *perma.cc/T8AY-N3U9*

[15] Doug Judd. “Spanner Under the Hood: Understanding Strict Serializability and External Consistency.” *cloud.google.com*, April 2023. Archived at *perma.cc/KJ9F-BJ5T*

[16] FoundationDB project authors. “Developer Guide.” *apple.github.io*. Archived at *perma.cc/F53L-TM9P*

[17] Ben Darnell. “How to Talk About Consistency and Isolation in Distributed DBs.” *cockroachlabs.com*, February 2022. Archived at *perma.cc/53SV-JBGK*

[18] Daniel Abadi. “An Explanation of the Difference Between Isolation Levels vs. Consistency Levels.” *dbmsmusings.blogspot.com*, August 2019. Archived at *perma.cc/ QSF2-CD4P*

[19] Mike Burrows. “The Chubby Lock Service for Loosely-Coupled Distributed Systems.” At *7th USENIX Symposium on Operating System Design and Implementation* (OSDI), November 2006.

[20] Flavio P. Junqueira and Benjamin Reed. *ZooKeeper: Distributed Process Coordination*. O’Reilly Media, 2013. ISBN: 9781449361303

[21] Murali Vallath. *Oracle 10g RAC Grid, Services & Clustering*. Elsevier Digital Press, 2006. ISBN: 9781555583217

[22] Peter Bailis, Alan Fekete, Michael J. Franklin, Ali Ghodsi, Joseph M. Heller-stein, and Ion Stoica. “Coordination Avoidance in Database Systems.” *Proceedings of the VLDB Endowment*, volume 8, issue 3, pages 185–196, November 2014. *doi:10.14778/2735508.2735509*, extended version published as *arXiv:1402.2237*

[23] Kyle Kingsbury. “Jepsen: etcd and Consul.” *aphyr.com*, June 2014. Archived at *perma.cc/XL7U-378K*

[24] Flavio P. Junqueira, Benjamin C. Reed, and Marco Serafini. “Zab: HighPerformance Broadcast for Primary-Backup Systems.” At *41st IEEE International Conference on Dependable Systems and Networks* (DSN), June 2011. *doi:10.1109/ DSN.2011.5958223*

[25] Diego Ongaro and John K. Ousterhout. “In Search of an Understandable Consensus Algorithm.” At *USENIX Annual Technical Conference* (ATC), June 2014.

[26] Hagit Attiya, Amotz Bar-Noy, and Danny Dolev. “Sharing Memory Robustly in Message-Passing Systems.” *Journal of the ACM*, volume 42, issue 1, pages 124–142, January 1995. *doi:10.1145/200836.200869*

[27] Nancy Lynch and Alex Shvartsman. “Robust Emulation of Shared Memory Using Dynamic Quorum-Acknowledged Broadcasts.” At *27th Annual International Symposium on Fault-Tolerant Computing* (FTCS), June 1997. *doi:10.1109/ FTCS.1997.614100*

[28] Christian Cachin, Rachid Guerraoui, and Luís Rodrigues. *Introduction to Reliable and Secure Distributed Programming*, 2nd edition. Springer, 2011. ISBN: 9783642152597, *doi:10.1007/978-3-642-15260-3*

[29] Niklas Ekström, Mikhail Panchenko, and Jonathan Ellis. "Possible Issue with Read Repair?" Email thread on *cassandra-dev* mailing list, October 2012. Archived at *perma.cc/49GF-QMWA*

[30] Maurice P. Herlihy. "Wait-Free Synchronization." *ACM Transactions on Programming Languages and Systems* (TOPLAS), volume 13, issue 1, pages 124–149, January 1991. *doi:10.1145/114005.102808*

[31] Armando Fox and Eric A. Brewer. "Harvest, Yield, and Scalable Tolerant Systems." At *7th Workshop on Hot Topics in Operating Systems* (HotOS), March 1999. *doi:10.1109/HOTOS.1999.798396*

[32] Seth Gilbert and Nancy Lynch. "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-Tolerant Web Services." *ACM SIGACT News*, volume 33, issue 2, pages 51–59, June 2002. *doi:10.1145/564585.564601*

[33] Seth Gilbert and Nancy Lynch. "Perspectives on the CAP Theorem." *IEEE Computer Magazine*, volume 45, issue 2, pages 30–36, February 2012. *doi:10.1109/ MC.2011.389*

[34] Eric A. Brewer. "CAP Twelve Years Later: How the 'Rules' Have Changed." *IEEE Computer Magazine*, volume 45, issue 2, pages 23–29, February 2012. *doi:10.1109/ MC.2012.37*

[35] Susan B. Davidson, Hector Garcia-Molina, and Dale Skeen. "Consistency in Partitioned Networks." *ACM Computing Surveys*, volume 17, issue 3, pages 341–370, September 1985. *doi:10.1145/5505.5508*

[36] Paul R. Johnson and Robert H. Thomas. "RFC 677: The Maintenance of Duplicate Databases." Network Working Group, January 1975.

[37] Michael J. Fischer and Alan Michael. "Sacrificing Serializability to Attain High Availability of Data in an Unreliable Network." At *1st ACM Symposium on Principles of Database Systems* (PODS), March 1982. *doi:10.1145/588111.588124*

[38] Eric A. Brewer. "NoSQL: Past, Present, Future." At *QCon San Francisco*, November 2012.

[39] Eric Brewer. "Spanner, TrueTime & The CAP Theorem." *research.google.com*, February 2017. Archived at *perma.cc/59UW-RH7N*

[40] Daniel J. Abadi. "Consistency Tradeoffs in Modern Distributed Database System Design." *IEEE Computer Magazine*, volume 45, issue 2, pages 37–42, February 2012. *doi:10.1109/MC.2012.33*

[41] Nancy A. Lynch. "A Hundred Impossibility Proofs for Distributed Computing." At *8th ACM Symposium on Principles of Distributed Computing* (PODC), August 1989. *doi:10.1145/72981.72982*

[42] Prince Mahajan, Lorenzo Alvisi, and Mike Dahlin. "Consistency, Availability, and Convergence." University of Texas at Austin, Department of Computer Science, Tech Report UTCS TR-11-22, May 2011. Archived at *perma.cc/SAV8-9JAJ*

[43] Hagit Attiya, Faith Ellen, and Adam Morrison. "Limitations of Highly-Available Eventually-Consistent Data Stores." At *ACM Symposium on Principles of Distributed Computing* (PODC), July 2015. *doi:10.1145/2767386.2767419*

[44] Adrian Cockcroft. "Migrating to Microservices." At *QCon London*, March 2014.

[45] Martin Kleppmann. "A Critique of the CAP Theorem." *arXiv:1509.05393*, September 2015.

[46] Daniel Abadi. "Problems with CAP, and Yahoo's Little Known NoSQL System." *dbmsmusings.blogspot.com*, April 2010. Archived at *perma.cc/4NTZ-CLM9*

[47] Daniel Abadi. "Hazelcast and the Mythical PA/EC System." *dbmsmusings.blogspot.com*, October 2017. Archived at *perma.cc/J5XM-U5C2*

[48] Peter Sewell, Susmit Sarkar, Scott Owens, Francesco Zappa Nardelli, and MagnusO. Myreen. "x86-TSO: A Rigorous and Usable Programmer's Model for x86 Multiprocessors." *Communications of the ACM*, volume 53, issue 7, pages 89–97, July 2010. *doi:10.1145/1785414.1785443*

[49] Martin Thompson. “Memory Barriers/Fences.” *mechanical-sympathy.blogspot.co.uk*, July 2011. Archived at *perma.cc/7NXM-GC5U*

[50] Ulrich Drepper. “What Every Programmer Should Know About Memory.” *akkadia.org*, November 2007. Archived at *perma.cc/NU6Q-DRXZ*

[51] Hagit Attiya and Jennifer L. Welch. “Sequential Consistency Versus Linearizability.” *ACM Transactions on Computer Systems* (TOCS), volume 12, issue 2, pages 91– 122, May 1994. *doi:10.1145/176575.176576*

[52] Kyzer R. Davis, Brad G. Peabody, and Paul J. Leach. “Universally Unique IDentifiers (UUIDs).” RFC 9562, IETF, May 2024.

[53] Ryan King. “Announcing Snowflake.” *blog.x.com*, June 2010. Archived at *archive.org*

[54] Alizain Feerasta. “Universally Unique Lexicographically Sortable Identifier.” *github.com*, 2016. Archived at *perma.cc/NV2Y-ZP8U*

[55] Rob Conery. “A Better ID Generator for PostgreSQL.” *bigmachine.io*, May 2014. Archived at *perma.cc/K7QV-3KFC*

[56] Leslie Lamport. “Time, Clocks, and the Ordering of Events in a Distributed System.” *Communications of the ACM*, volume 21, issue 7, pages 558–565, July 1978. *doi:10.1145/359545.359563*

[57] Sandeep S. Kulkarni, Murat Demirbas, Deepak Madeppa, Bharadwaj Avva, andMarcelo Leone. “Logical Physical Clocks.” *18th International Conference on Principles of Distributed Systems* (OPODIS), December 2014. *doi:10.1007/978-3-319-14472-6_2*

[58] Manuel Bravo, Nuno Diegues, Jingna Zeng, Paolo Romano, and Luís Rodrigues. “On the Use of Clocks to Enforce Consistency in the Cloud.” *IEEE Data Engineering Bulletin*, volume 38, issue 1, pages 18–31, March 2015. Archived at *perma.cc/ 68ZU-45SH*

[59] Daniel Peng and Frank Dabek. “Large-Scale Incremental Processing Using Distributed Transactions and Notifications.” At *9th USENIX Conference on Operating Systems Design and Implementation* (OSDI), October 2010.

[60] Tushar Deepak Chandra, Robert Griesemer, and Joshua Redstone. “Paxos Made Live—An Engineering Perspective.” At *26th ACM Symposium on Principles of Distributed Computing* (PODC), June 2007. *doi:10.1145/1281100.1281103*

[61] Will Portnoy. “Lessons Learned from Implementing Paxos.” *blog.willportnoy.com*, June 2012. Archived at *perma.cc/QHD9-FDD2*

[62] Brian M. Oki and Barbara H. Liskov. “Viewstamped Replication: A New Primary Copy Method to Support Highly-Available Distributed Systems.” At *7th ACM Symposium on Principles of Distributed Computing* (PODC), August 1988. *doi:10.1145/62546.62549*

[63] Barbara H. Liskov and James Cowling. “Viewstamped Replication Revisited.” Massachusetts Institute of Technology, Tech Report MIT-CSAIL-TR-2012-021, July 2012. Archived at *perma.cc/56SJ-WENQ*

[64] Leslie Lamport. “The Part-Time Parliament.” *ACM Transactions on Computer Systems*, volume 16, issue 2, pages 133–169, May 1998. *doi:10.1145/279227.279229*

[65] Leslie Lamport. “Paxos Made Simple.” *ACM SIGACT News*, volume 32, issue 4, pages 51–58, December 2001. Archived at *perma.cc/82HP-MNKE*

[66] Robbert van Renesse and Deniz Altinbuken. “Paxos Made Moderately Complex.” *ACM Computing Surveys* (CSUR), volume 47, issue 3, article no. 42, February 2015. *doi:10.1145/2673577*

[67] Diego Ongaro. “Consensus: Bridging Theory and Practice.” PhD thesis, Stanford University, August 2014. Archived at *perma.cc/5VTZ-2ADH*

[68] Heidi Howard, Malte Schwarzkopf, Anil Madhavapeddy, and Jon Crowcroft.“Raft Refloated: Do We Have Consensus?” *ACM SIGOPS Operating Systems Review*, volume 49, issue 1, pages 12–21, January 2015. *doi:10.1145/2723872.2723876*

[69] André Medeiros. “ZooKeeper’s Atomic Broadcast Protocol: Theory and Practice.” Aalto University School of Science, March 2012. Archived at *perma.cc/FVL4-JMVA*

[70] Robbert van Renesse, Nicolas Schiper, and Fred B. Schneider. “Vive la Différence: Paxos vs. Viewstamped Replication vs. Zab.” *IEEE Transactions on Dependable and Secure Computing*, volume 12, issue 4, pages 472–484, September 2014. *doi:10.1109/TDSC.2014.2355848*

[71] Heidi Howard and Richard Mortier. “Paxos vs Raft: Have We Reached Consensus on Distributed Consensus?” At *7th Workshop on Principles and Practice of Consistency for Distributed Data* (PaPoC), April 2020. *doi:10.1145/3380787.3393681*

[72] Miguel Castro and Barbara H. Liskov. “Practical Byzantine Fault Tolerance and Proactive Recovery.” *ACM Transactions on Computer Systems*, volume 20, issue 4, pages 396–461, November 2002. *doi:10.1145/571637.571640*

[73] Shehar Bano, Alberto Sonnino, Mustafa Al-Bassam, Sarah Azouvi, PatrickMcCorry, Sarah Meiklejohn, and George Danezis. “SoK: Consensus in the Age of Blockchains.” At *1st ACM Conference on Advances in Financial Technologies* (AFT), October 2019. *doi:10.1145/3318041.3355458*

[74] Michael J. Fischer, Nancy Lynch, and Michael S. Paterson. “Impossibility of Distributed Consensus with One Faulty Process.” *Journal of the ACM*, volume 32, issue 2, pages 374–382, April 1985. *doi:10.1145/3149.214121*

[75] Tushar Deepak Chandra and Sam Toueg. “Unreliable Failure Detectors for Reliable Distributed Systems.” *Journal of the ACM*, volume 43, issue 2, pages 225–267, March 1996. *doi:10.1145/226643.226647*

[76] Michael Ben-Or. “Another Advantage of Free Choice: Completely Asynchronous Agreement Protocols.” At *2nd ACM Symposium on Principles of Distributed Computing* (PODC), August 1983. *doi:10.1145/800221.806707*

[77] Cynthia Dwork, Nancy Lynch, and Larry Stockmeyer. “Consensus in the Presence of Partial Synchrony.” *Journal of the ACM*, volume 35, issue 2, pages 288–323, April 1988. *doi:10.1145/42282.42283*

[78] Xavier Défago, André Schiper, and Péter Urbán. “Total Order Broadcast and Multicast Algorithms: Taxonomy and Survey.” *ACM Computing Surveys*, volume 36, issue 4, pages 372–421, December 2004. *doi:10.1145/1041680.1041682*

[79] Hagit Attiya and Jennifer Welch. *Distributed Computing: Fundamentals, Simulations and Advanced Topics*, 2nd edition. John Wiley & Sons, 2004. ISBN: 9780471453246, *doi:10.1002/0471478210*

[80] Rachid Guerraoui. “Revisiting the Relationship Between Non-Blocking Atomic Commitment and Consensus.” At *9th International Workshop on Distributed Algorithms* (WDAG), September 1995. *doi:10.1007/BFb0022140*

[81] Jim N. Gray and Leslie Lamport. “Consensus on Transaction Commit.” *ACM Transactions on Database Systems* (TODS), volume 31, issue 1, pages 133–160, March 2006. *doi:10.1145/1132863.1132867*

[82] Fred B. Schneider. “Implementing Fault-Tolerant Services Using the State Machine Approach: A Tutorial.” *ACM Computing Surveys*, volume 22, issue 4, pages 299–319, December 1990. *doi:10.1145/98163.98167*

[83] Alexander Thomson, Thaddeus Diamond, Shu-Chun Weng, Kun Ren, PhilipShao, and Daniel J. Abadi. “Calvin: Fast Distributed Transactions for Partitioned Database Systems.” At *ACM International Conference on Management of Data* (SIGMOD), May 2012. *doi:10.1145/2213836.2213838*

[84] Mahesh Balakrishnan, Dahlia Malkhi, Ted Wobber, Ming Wu, Vijayan Prabha-karan, Michael Wei, John D. Davis, Sriram Rao, Tao Zou, and Aviad Zuck. “Tango: Distributed Data Structures over a Shared Log.” At *24th ACM Symposium on Operating Systems Principles* (SOSP), November 2013. *doi:10.1145/2517349.2522732*

[85] Mahesh Balakrishnan, Dahlia Malkhi, Vijayan Prabhakaran, Ted Wobber, Michael Wei, and John D. Davis. “CORFU: A Shared Log Design for Flash Clusters.” At *9th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), April 2012.

[86] Iulian Moraru, David G. Andersen, and Michael Kaminsky. “There Is More Consensus in Egalitarian Parliaments.” At *24th ACM Symposium on Operating Systems Principles* (SOSP), November 2013. *doi:10.1145/2517349.2517350*

[87] Vasilis Gavrielatos, Antonios Katsarakis, and Vijay Nagarajan. “Odyssey: the Impact of Modern Hardware on Strongly-Consistent Replication Protocols.” At *16th European Conference on Computer Systems* (EuroSys), April 2021. *doi:10.1145/3447786.3456240*

[88] Heidi Howard, Dahlia Malkhi, and Alexander Spiegelman. “Flexible Paxos: Quorum Intersection Revisited.” At *20th International Conference on Principles of Distributed Systems* (OPODIS), December 2016. *doi:10.4230/LIPIcs.OPODIS.2016.25*

[89] Martin Kleppmann. “Distributed Systems.” Lecture Notes. *University of Cambridge*, October 2024. Archived at *perma.cc/SS3Q-FNS5*

[90] Kyle Kingsbury. “Jepsen: Elasticsearch 1.5.0.” *aphyr.com*, April 2015. Archived at *perma.cc/37MZ-JT7H*

[91] Heidi Howard and Jon Crowcroft. “Coracle: Evaluating Consensus at the Internet Edge.” At *Annual Conference of the ACM Special Interest Group on Data Communication* (SIGCOMM), August 2015. *doi:10.1145/2829988.2790010*

[92] Tom Lianza and Chris Snook. “A Byzantine failure in the Real World.” *blog.cloudflare.com*, November 2020. Archived at *perma.cc/83EZ-ALCY*

[93] Ivan Kelly. “BookKeeper Tutorial.” *github.com*, October 2014. Archived at *perma.cc/37Y6-VZWU*

[94] Jack Vanlightly. “Apache BookKeeper Insights Part 1—External Consensus and Dynamic Membership.” *medium.com*, November 2021. Archived at *perma.cc/ 3MDB-8GFB*

บทที่ 11

[1] Nathan Marz. “How to Beat the CAP Theorem.” *nathanmarz.com*, October 2011. Archived at *perma.cc/4BS9-R9A4*

[2] Molly Bartlett Dishman and Martin Fowler. “Agile Architecture.” At *O’Reilly Software Architecture Conference*, March 2015.

[3] Jeffrey Dean and Sanjay Ghemawat. “MapReduce: Simplified Data Processing on Large Clusters.” At *6th USENIX Symposium on Operating System Design and Implementation* (OSDI), December 2004.

[4] Shivnath Babu and Herodotos Herodotou. “Massively Parallel Databases and MapReduce Systems.” *Foundations and Trends in Databases*, volume 5, issue 1, pages 1–104, November 2013. *doi:10.1561/1900000036*

[5] David J. DeWitt and Michael Stonebraker. “MapReduce: A Major Step Backwards.” Originally published at *databasecolumn.vertica.com*, January 2008. Archived at *perma.cc/U8PA-K48V*

[6] Henry Robinson. “The Elephant Was a Trojan Horse: On the Death of MapReduce at Google.” *the-paper-trail.org*, June 2014. Archived at *perma.cc/9FEM-X787*

[7] Urs Hölzle. “R.I.P. MapReduce. After having served us well since 2003, today we removed the remaining internal codebase for good.” *x.com*, September 2019.
Archived at *perma.cc/B34T-LLY7*

[8] Adam Drake. “Command-Line Tools Can Be 235x Faster than Your Hadoop Cluster.” *aadrake.com*, January 2014. Archived at *perma.cc/87SP-ZMCY*

[9] "`sort`: Sort Text Files.” GNU Coreutils 9.7 Documentation, Free Software Foundation, Inc., 2025. Archived at *perma.cc/68KN-E8TL*

[10] Michael Ovsiannikov, Silvius Rus, Damian Reeves, Paul Sutter, Sriram Rao, andJim Kelly. “The Quantcast File System.” *Proceedings of the VLDB Endowment*, volume 6, issue 11, pages 1092–1101, August 2013. *doi:10.14778/2536222.2536234*

[11] Andrew Wang, Zhe Zhang, Kai Zheng, Uma Maheswara G., and VinayakumarB. “Introduction to HDFS Erasure Coding in Apache Hadoop.” *blog.cloudera.com*, September 2015. Archived at *archive.org*

[12] Andy Warfield. “Building and Operating a Pretty Big Storage System Called S3.” *allthingsdistributed.com*, July 2023. Archived at *perma.cc/7LPK-TP7V*

[13] Vinod Kumar Vavilapalli, Arun C. Murthy, Chris Douglas, Sharad Agarwal,Mahadev Konar, Robert Evans, Thomas Graves, Jason Lowe, Hitesh Shah, Siddharth Seth, Bikas Saha, Carlo Curino, Owen O’Malley, Sanjay Radia, Benjamin Reed, and Eric Baldeschwieler. “Apache Hadoop YARN: Yet Another Resource Negotiator.” At *4th Annual Symposium on Cloud Computing* (SoCC), October 2013. *doi:10.1145/2523616.2523633*

[14] Richard M. Karp. “Reducibility Among Combinatorial Problems.” *Complexity of Computer Computations. The IBM Research Symposia Series*. Springer, 1972. *doi:10.1007/978-1-4684-2001-2_9*

[15] J. D. Ullman. “NP-Complete Scheduling Problems.” *Journal of Computer and System Sciences*, volume 10, issue 3, pages 384–393, June 1975. *doi:10.1016/ S0022-0000(75)80008-0*

[16] Gilad David Maayan. “The Complete Guide to Spot Instances on AWS, Azure and GCP.” *datacenterdynamics.com*, March 2021. Archived at *archive.org*

[17] Abhishek Verma, Luis Pedrosa, Madhukar Korupolu, David Oppenheimer,Eric Tune, and John Wilkes. “Large-Scale Cluster Management at Google with Borg.” At *10th European Conference on Computer Systems* (EuroSys), April 2015. *doi:10.1145/2741948.2741964*

[18] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma,Murphy McCauley, Michael J. Franklin, Scott Shenker, and Ion Stoica. “Resilient Distributed Datasets: A Fault-Tolerant Abstraction for In-Memory Cluster Computing.” At *9th USENIX Symposium on Networked Systems Design and Implementation* (NSDI), April 2012.

[19] Paris Carbone, Stephan Ewen, Seif Haridi, Asterios Katsifodimos, Volker Markl,and Kostas Tzoumas. “Apache Flink: Stream and Batch Processing in a Single Engine.” *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering*, volume 38, issue 4, pages 28–38, December 2015. Archived at *perma.cc/G3N3-BKX5*

[20] Mark Grover, Ted Malaska, Jonathan Seidman, and Gwen Shapira. *Hadoop Application Architectures*. O’Reilly Media, 2015. ISBN: 9781491900048

[21] Jules S. Damji, Brooke Wenig, Tathagata Das, and Denny Lee. *Learning Spark*, 2nd edition. O’Reilly Media, 2020. ISBN: 9781492050049

[22] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly.“Dryad: Distributed Data-Parallel Programs from Sequential Building Blocks.” At *2nd European Conference on Computer Systems* (EuroSys), March 2007. *doi:10.1145/1272996.1273005*

[23] Daniel Warneke and Odej Kao. “Nephele: Efficient Parallel Data Processing in the Cloud.” At *2nd Workshop on Many-Task Computing on Grids and Supercomputers* (MTAGS), November 2009. *doi:10.1145/1646468.1646476*

[24] Hossein Ahmadi. “In-Memory Query Execution in Google BigQuery.” *cloud.google.com*, August 2016. Archived at *perma.cc/DGG2-FL9W*

[25] Tom White. *Hadoop: The Definitive Guide*, 4th edition. O’Reilly Media, 2015. ISBN: 9781491901632

[26] Fabian Hüske. “Peeking into Apache Flink’s Engine Room.” *flink.apache.org*, March 2015. Archived at *perma.cc/44BW-ALJX*

[27] Mostafa Mokhtar. “Hive 0.14 Cost Based Optimizer (CBO) Technical Overview.” *hortonworks.com*, March 2015. Archived at *archive.org*

[28] Michael Armbrust, Reynold S. Xin, Cheng Lian, Yin Huai, Davies Liu,Joseph K. Bradley, Xiangrui Meng, Tomer Kaftan, Michael J. Franklin, Ali Ghodsi, and Matei Zaharia. “Spark SQL: Relational Data

Processing in Spark.” At *ACM International Conference on Management of Data* (SIGMOD), June 2015. *doi:10.1145/2723372.2742797*

[29] Kaya Kupferschmidt. “Spark vs. Pandas, Part 2—Spark.” *towardsdatascience.com*, October 2020. Archived at *perma.cc/5BRK-G4N5*

[30] Ammar Chalifah. “Tracking Payments at Scale.” *bolt.eu.com*, June 2025. Archived at *perma.cc/Q4KX-8K3J*

[31] Nafi Ahmet Turgut, Hamza Akyıldız, Hasan Burak Yel, Mehmet İkbal Özmen,Mutlu Polatcan, Pinar Baki, and Esra Kayabali. “Demand Forecasting at Getir Built with Amazon Forecast.” *aws.amazon.com*, May 2023. Archived at *perma.cc/H3H6GNL7*

[32] Jason (Siyu) Zhu. “Enhancing Homepage Feed Relevance by Harnessing the Power of Large Corpus Sparse ID Embeddings.” *linkedin.com*, August 2023. Archived at *archive.org*

[33] Avery Ching, Sital Kedia, and Shuojie Wang. “Apache Spark @Scale: A 60 TB+ Production Use Case.” *engineering.fb.com*, August 2016. Archived at *perma.cc/F7R5-YFAV*

[34] Edward Kim. “How ACH Works: A Developer Perspective—Part 1.” *engineering.gusto.com*, April 2014. Archived at *perma.cc/F67P-VBLK*

[35] Zhamak Dehghani. “How to Move Beyond a Monolithic Data Lake to a Distributed Data Mesh.” *martinfowler.com*, May 2019. Archived at *perma.cc/LN2L-L4VC*

[36] Chris Riccomini. “What the Heck Is a Data Mesh?!” *cnr.sh*, June 2021. Archived at *perma.cc/NEJ2-BAX3*

[37] Chad Sanderson, Mark Freeman, and B. E. Schmidt. *Data Contracts*. O’Reilly Media, 2025. ISBN: 9781098157623

[38] Daniel Abadi. “Data Fabric vs. Data Mesh: What’s the Difference?” *starburst.io*, November 2021. Archived at *perma.cc/RSK3-HXDK*

[39] Michael Armbrust, Ali Ghodsi, Reynold Xin, and Matei Zaharia. “Lakehouse: A New Generation of Open Platforms That Unify Data Warehousing and Advanced Analytics.” At *11th Annual Conference on Innovative Data Systems Research* (CIDR), January 2021. Archived at *perma.cc/7C6D-T9NR*

[40] Leslie G. Valiant. “A Bridging Model for Parallel Computation.” *Communications of the ACM*, volume 33, issue 8, pages 103–111, August 1990. *doi:10.1145/79173.79181*

[41] Stephan Ewen, Kostas Tzoumas, Moritz Kaufmann, and Volker Markl. “Spinning Fast Iterative Data Flows.” *Proceedings of the VLDB Endowment*, volume 5, issue 11, pages 1268–1279, July 2012. *doi:10.14778/2350229.2350245*

[42] Grzegorz Malewicz, Matthew H. Austern, Aart J. C. Bik, James C. Dehnert,Ilan Horn, Naty Leiser, and Grzegorz Czajkowski. “Pregel: A System for Large-Scale Graph Processing.” At *ACM International Conference on Management of Data* (SIGMOD), June 2010. *doi:10.1145/1807167.1807184*

[43] Richard MacManus. “OpenAI Chats About Scaling LLMs at Anyscale’s Ray Summit.” *thenewstack.io*, September 2023. Archived at *perma.cc/YJD6-KUXU* [44] Jay Kreps. “Why Local State Is a Fundamental Primitive in Stream Processing.” *oreilly.com*, July 2014. Archived at *perma.cc/P8HU-R5LA*

[45] Félix GV. “Open Sourcing Venice—LinkedIn’s Derived Data Platform.” *linkedin.com*, September 2022. Archived at *archive.org*

บทที่ 12

[1] Tyler Akidau, Robert Bradshaw, Craig Chambers, Slava Chernyak, Rafael J.Fernández-Moctezuma, Reuven Lax, Sam McVeety, Daniel Mills, Frances Perry, Eric Schmidt, and Sam Whittle. “The Dataflow Model: A Practical Approach to Balancing Correctness, Latency, and Cost in Massive-Scale, Unbounded, Out-of-Order Data Processing.” *Proceedings of the VLDB Endowment*, volume 8, issue 12, pages 1792– 1803, August 2015. *doi:10.14778/2824032.2824076*

[2] Harold Abelson, Gerald Jay Sussman, and Julie Sussman. *Structure and Interpretation of Computer Programs*, 2nd edition. MIT Press, 1996. ISBN: 9780262510875. Archived at *archive.org*

[3] Patrick Th. Eugster, Pascal A. Felber, Rachid Guerraoui, and Anne-Marie Kermar-rec. "The Many Faces of Publish/Subscribe." *ACM Computing Surveys*, volume 35, issue 2, pages 114–131, June 2003. *doi:10.1145/857076.857078*

[4] Don Carney, Uğur Çetintemel, Mitch Cherniack, Christian Convey, SangdonLee, Greg Seidman, Michael Stonebraker, Nesime Tatbul, and Stan Zdonik. "Monitoring Streams—A New Class of Data Management Applications." At *28th International Conference on Very Large Data Bases* (VLDB), August 2002. *doi:10.1016/ B978-155860869-6/50027-5*

[5] Matthew Sackman. "Pushing Back." *wellquite.org*, May 2016. Archived at *perma.cc/3KCZ-RUFY*

[6] Thomas Figg (tef). "How (Not) to Write a Pipeline." *cohost.org*, June 2023. Archived at *perma.cc/A3V8-NYCM*

[7] Vicent Martí. "Brubeck, a statsd-Compatible Metrics Aggregator." *github.blog*, June 2015. Archived at *perma.cc/TP3Q-DJYM*

[8] Seth Lowenberger. "MoldUDP64 Protocol Specification V 1.00." *nasdaqtrader.com*, July 2009. Archived at *perma.cc/7CRQ-QBD7*

[9] Ian Malpass. "Measure Anything, Measure Everything." *codeascraft.com*, February 2011. Archived at *archive.org*

[10] Dieter Plaetinck. "25 Graphite, Grafana and statsd Gotchas." *grafana.com*, March 2016. Archived at *perma.cc/3NP3-67U7*

[11] Jeff Lindsay. "Web Hooks to Revolutionize the Web." *progrium.com*, May 2007. Archived at *perma.cc/BF9U-XNX4*

[12] Jim N. Gray. "Queues Are Databases." Microsoft Research Technical Report MSR-TR-95-56, December 1995. Archived at *arxiv.org*

[13] Mark Hapner, Rich Burridge, Rahul Sharma, Joseph Fialli, Kate Stout, and NigelDeakin. "JSR-343 Java Message Service (JMS) 2.0 Specification." *jms-spec.java.net*, March 2013. Archived at *perma.cc/E4YG-46TA*

[14] Sanjay Aiyagari, Matthew Arrott, Mark Atwell, Jason Brome, Alan Conway,Robert Godfrey, Robert Greig, Pieter Hintjens, John O'Hara, Matthias Radestock, Alexis Richardson, Martin Ritchie, Shahrokh Sadjadi, Rafael Schloming, Steven Shaw, Martin Sustrik, Carl Trieloff, Kim van der Riet, and Steve Vinoski. "AMQP: Advanced Message Queuing Protocol Specification." Version 0-9-1, November 2008. Archived at *perma.cc/6YJJ-GM9X*

[15] "Architectural Overview of Pub/Sub." *cloud.google.com*, 2025. Archived at *perma.cc/VWF5-ABP4*

[16] Aris Tzoumas. "Lessons from Scaling PostgreSQL Queues to 100k Events Per Second." *rudderstack.com*, July 2025. Archived at *perma.cc/QD8C-VA4Y*

[17] Robin Moffatt. "Kafka Connect Deep Dive—Error Handling and Dead Letter Queues." *confluent.io*, March 2019. Archived at *perma.cc/KQ5A-AB28*

[18] Dunith Danushka. " Message Reprocessing: How to Implement the Dead Letter Queue." *redpanda.com*. Archived at *perma.cc/R7UB-WEWF*

[19] Damien Gasparina, Loic Greffier, and Sebastien Viale. "KIP-1034: Dead Letter Queue in Kafka Streams." *cwiki.apache.org*, April 2024. Archived at *perma.cc/3VXV-QXAN*

[20] Jay Kreps, Neha Narkhede, and Jun Rao. "Kafka: A Distributed Messaging System for Log Processing." At *6th International Workshop on Networking Meets Databases* (NetDB), June 2011. Archived at *perma.cc/CSW7-TCQ5*

[21] Jay Kreps. "Benchmarking Apache Kafka: 2 Million Writes Per Second (On Three Cheap Machines)." *engineering.linkedin.com*, April 2014. Archived at archive.org

[22] Kartik Paramasivam. “How We’re Improving and Advancing Kafka at LinkedIn.” *engineering.linkedin.com*, September 2015. Archived at *perma.cc/3S3V-JCYJ* [23] Philippe Dobbelaere and Kyumars Sheykh Esmaili. “Kafka Versus RabbitMQ: A Comparative Study of Two Industry Reference Publish/Subscribe Implementations.” At *11th ACM International Conference on Distributed and Event-Based Systems* (DEBS), June 2017. *doi:10.1145/3093742.3093908*

[24] Kate Holterhoff. “Why Message Queues Endure: A History.” *redmonk.com*, December 2024. Archived at *perma.cc/6DX8-XK4W*

[25] Andrew Schofield. “KIP-932: Queues for Kafka.” *cwiki.apache.org*, May 2023. Archived at *perma.cc/LBE4-BEMK*

[26] Jack Vanlightly. “The Advantages of Queues on Logs.” *jack-vanlightly.com*, October 2023. Archived at *perma.cc/WJ7V-287K*

[27] Jay Kreps. “The Log: What Every Software Engineer Should Know About Real-Time Data’s Unifying Abstraction.” *engineering.linkedin.com*, December 2013. Archived at *perma.cc/2JHR-FR64*

[28] Andy Hattemer. “Change Data Capture Is Having a Moment. Why?” *materialize.com*, September 2021. Archived at *perma.cc/AL37-P53C*

[29] Prem Santosh Udaya Shankar. “Streaming MySQL Tables in Real-Time to Kafka.” *engineeringblog.yelp.com*, August 2016. Archived at *perma.cc/5ZR3-2GVV*

[30] Andreas Andreakis, Ioannis Papapanagiotou. “DBLog: A Watermark Based Change-Data-Capture Framework.” Archived at *arXiv:2010.12597*, October 2020.

[31] Jiri Pechanec. “Percolator.” *debezium.io*, October 2021. Archived at *perma.cc/EQ8E-W6KQ*

[32] Debezium maintainers. “Debezium Connector for Cassandra.” *debezium.io*. Archived at *perma.cc/WR6K-EKMD*

[33] Neha Narkhede. “Announcing Kafka Connect: Building Large-Scale LowLatency Data Pipelines.” *confluent.io*, February 2016. Archived at *perma.cc/8WXJL6GF*

[34] Chris Riccomini. “Kafka Change Data Capture Breaks Database Encapsulation.” *cnr.sh*, November 2018. Archived at *perma.cc/P572-9MKF*

[35] Gunnar Morling. “‘Change Data Capture Breaks Encapsulation’. Does It, Though?” *decodable.co*, November 2023. Archived at *perma.cc/YX2P-WNWR* [36] Gunnar Morling. “Revisiting the Outbox Pattern.” *decodable.co*, October 2024. Archived at *perma.cc/M5ZL-RPS9*

[37] Ashish Gupta and Inderpal Singh Mumick. “Maintenance of Materialized Views: Problems, Techniques, and Applications.” *IEEE Data Engineering Bulletin*, volume 18, issue 2, pages 3–18, June 1995. Archived at *archive.org*

[38] Mihai Budiu, Tej Chajed, Frank McSherry, Leonid Ryzhyk, and Val Tan-nen. “DBSP: Incremental Computation on Streams and Its Applications to Databases.” *SIGMOD Record*, volume 53, issue 1, pages 87–95, March 2024. *doi:10.1145/3665252.3665271*

[39] Jim Gray and Andreas Reuter. *Transaction Processing: Concepts and Techniques*. Morgan Kaufmann, 1992. ISBN: 9781558601901

[40] Martin Kleppmann. “Accounting for Computer Scientists.” *martin.kleppmann.com*, March 2011. Archived at *perma.cc/9EGX-P38N*

[41] Pat Helland. “Immutability Changes Everything.” At *7th Biennial Conference on Innovative Data Systems Research* (CIDR), January 2015. Archived at *perma.cc/ 33WX-3669*

[42] Martin Kleppmann. *Making Sense of Stream Processing*. Report, O’Reilly Media, May 2016. Archived at *perma.cc/RAY4-JDVX*

[43] Kartik Paramasivam. “Stream Processing Hard Problems—Part 1: Killing Lambda.” *engineering.linkedin.com*, June 2016. Archived at *archive.org*

[44] Stéphane Derosiaux. “CQRS: What? Why? How?” *sderosiaux.medium.com*, September 2019. Archived at *perma.cc/FZ3U-HVJ4*

[45] Baron Schwartz. “Immutability, MVCC, and Garbage Collection.” *xaprb.com*, December 2013. Archived at *archive.org*

[46] Daniel Eloff, Slava Akhmechet, Jay Kreps, et al. “Re: Turning the Database Inside-out with Apache Samza.” Hacker News discussion, *news.ycombinator.com*, March 2015. Archived at *perma.cc/ML9E-JC83*

[47] Cognitect, Inc. “Datomic Documentation: Excision.” *docs.datomic.com*. Archived at *perma.cc/J5QQ-SH32*

[48] “Fossil Documentation: Deleting Content from Fossil.” *fossil-scm.org*, 2025. Archived at *perma.cc/DS23-GTNG*

[49] Jay Kreps. “The irony of distributed systems is that data loss is really easy but deleting data is surprisingly hard.” *x.com*, March 2015. Archived at *perma.cc/7RRZ-V7B7*

[50] Brent Robinson. “Crypto Shredding: How It Can Solve Modern Data Retention Challenges.” *medium.com*, January 2019. Archived at *perma.cc/4LFK-S6XE*

[51] Matthew D. Green and Ian Miers. “Forward Secure Asynchronous Messaging from Puncturable Encryption.” At *IEEE Symposium on Security and Privacy*, May 2015. *doi:10.1109/SP.2015.26*

[52] David C. Luckham. “What’s the Difference Between ESP and CEP?” *complexevents.com*, June 2019. Archived at *perma.cc/E7PZ-FDEF*

[53] Arvind Arasu, Shivnath Babu, and Jennifer Widom. “The CQL Continuous Query Language: Semantic Foundations and Query Execution.” *The VLDB Journal*, volume 15, issue 2, pages 121–142, June 2006. *doi:10.1007/s00778-004-0147-z*

[54] Julian Hyde. “Data in Flight: How Streaming SQL Technology Can Help Solve the Web 2.0 Data Crunch.” *ACM Queue*, volume 7, issue 11, December 2009. *doi:10.1145/1661785.1667562*

[55] Philippe Flajolet, Éric Fusy, Olivier Gandouet, and Frédéric Meunier. "HyperLogLog: The Analysis of a Near-Optimal Cardinality Estimation Algorithm.” At *Conference on Analysis of Algorithms* (AofA), June 2007. *doi:10.46298/dmtcs.3545*

[56] Jay Kreps. “Questioning the Lambda Architecture.” *oreilly.com*, July 2014. Archived at *perma.cc/2WY5-HC8Y*

[57] Ian Reppel. “An Overview of Apache Streaming Technologies.” *ianreppel.org*, March 2016. Archived at *perma.cc/BB3E-QJLW*

[58] Jay Kreps. “Why Local State Is a Fundamental Primitive in Stream Processing.” *oreilly.com*, July 2014. Archived at *perma.cc/P8HU-R5LA*

[59] RisingWave Labs. “Deep Dive into the RisingWave Stream Processing Engine —Part 2: Computational Model.” *risingwave.com*, November 2023. Archived at *perma.cc/LM74-XDEL*

[60] Frank McSherry, Derek G. Murray, Rebecca Isaacs, and Michael Isard. “Differential Dataflow.” At *6th Biennial Conference on Innovative Data Systems Research* (CIDR), January 2013. Archived at *perma.cc/T83W-ZBR2*

[61] Andy Hattemer. “Incremental Computation in the Database.” *materialize.com*, March 2020. Archived at *perma.cc/AL94-YVRN*

[62] Shay Banon. “Percolator.” *elastic.co*, February 2011. Archived at *perma.cc/LS5R-4FQX*

[63] Alan Woodward and Martin Kleppmann. “Real-Time Full-Text Search with Luwak and Samza.” *martin.kleppmann.com*, April 2015. Archived at *perma.cc/2U92Q7R4*

[64] Tyler Akidau. “The World Beyond Batch: Streaming 102.” *oreilly.com*, January 2016. Archived at *perma.cc/4XF9-8M2K*

[65] Stephan Ewen. “Streaming Analytics with Apache Flink.” At *Kafka Summit*, April 2016. Archived at *perma.cc/QBQ4-F9MR*

[66] Tyler Akidau, Alex Balikov, Kaya Bekiroğlu, Slava Chernyak, Josh Haber-man, Reuven Lax, Sam McVeety, Daniel Mills, Paul Nordstrom, and Sam Whittle. “MillWheel: Fault-Tolerant Stream Processing

at Internet Scale." *Proceedings of the VLDB Endowment*, volume 6, issue 11, pages 1033–1044, August 2013. *doi:10.14778/2536222.2536229*

[67] Alex Dean. "Improving Snowplow's Understanding of Time." *snowplow.io*, September 2015. Archived at *perma.cc/6CT9-Z3Q2*

[68] "Azure Stream Analytics: Windowing Functions." Microsoft Azure Reference, *learn.microsoft.com*, July 2025. Archived at *archive.org*

[69] Rajagopal Ananthanarayanan, Venkatesh Basker, Sumit Das, Ashish Gupta, Hai-feng Jiang, Tianhao Qiu, Alexey Reznichenko, Deomid Ryabkov, Manpreet Singh, and Shivakumar Venkataraman. "Photon: Fault-Tolerant and Scalable Joining of Continuous Data Streams." At *ACM International Conference on Management of Data* (SIGMOD), June 2013. *doi:10.1145/2463676.2465272*

[70] Ben Kirwin. "Doing the Impossible: Exactly-Once Messaging Patterns in Kafka." *ben.kirw.in*, November 2014. Archived at *perma.cc/A5QL-QRX7*

[71] Pat Helland. "Data on the Outside Versus Data on the Inside." At *2nd Biennial Conference on Innovative Data Systems Research* (CIDR), January 2005. Archived at *perma.cc/K9AH-LQPS*

[72] Ralph Kimball and Margy Ross. *The Data Warehouse Toolkit: The Definitive Guide to Dimensional Modeling*, 3rd edition. John Wiley & Sons, 2013. ISBN: 9781118530801

[73] Viktor Klang. "I'm coining the phrase 'effectively-once' for message processing with at-least-once + idempotent operations." *x.com*, October 2016. Archived at *perma.cc/7DT9-TDG2*

[74] Matei Zaharia, Tathagata Das, Haoyuan Li, Scott Shenker, and Ion Stoica."Discretized Streams: An Efficient and Fault-Tolerant Model for Stream Processing on Large Clusters." At *4th USENIX Conference in Hot Topics in Cloud Computing* (HotCloud), June 2012.

[75] Kostas Tzoumas, Stephan Ewen, and Robert Metzger. "High-Throughput, LowLatency, and Exactly-Once Stream Processing with Apache Flink." *ververica.com*, August 2015. Archived at *archive.org*

[76] Paris Carbone, Gyula Fóra, Stephan Ewen, Seif Haridi, and Kostas Tzoumas."Lightweight Asynchronous Snapshots for Distributed Dataflows." *arXiv:1506.08603*, June 2015.

[77] Ryan Betts and John Hugg. *Fast Data: Smart and at Scale*. Report, O'Reilly Media, October 2015. Archived at *perma.cc/VQ6S-XQQY*

[78] Neha Narkhede and Guozhang Wang. "Exactly-Once Semantics Are Possible: Here's How Kafka Does It." *confluent.io*, June 2019. Archived at *perma.cc/Q2AU-Q2ED*

[79] Jason Gustafson, Flavio Junqueira, Apurva Mehta, Sriram Subramanian, andGuozhang Wang. "KIP-98—Exactly Once Delivery and Transactional Messaging." *cwiki.apache.org*, November 2016. Archived at *perma.cc/95PT-RCTG*

[80] Pat Helland. "Idempotence Is Not a Medical Condition." *Communications of the ACM*, volume 55, issue 5, pages 56–65, May 2012. *doi:10.1145/2160718.2160734*

[81] Jay Kreps. "Re: Trying to Achieve Deterministic Behavior on Recovery/Rewind." Email to *samza-dev* mailing list, September 2014. Archived at *perma.cc/7DPD-GJNL*

[82] E. N. (Mootaz) Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. "A Survey of Rollback-Recovery Protocols in Message-Passing Systems." *ACM Computing Surveys*, volume 34, issue 3, pages 375–408, September 2002. *doi:10.1145/568522.568525*

[83] Adam Warski. "Kafka Streams—How Does It Fit the Stream Processing Landscape?" *softwaremill.com*, June 2016. Archived at *perma.cc/WQ5Q-H2J2*

[84] Stephan Ewen, Fabian Hueske, and Xiaowei Jiang. "Batch as a Special Case of Streaming and Alibaba's contribution of Blink." *flink.apache.org*, February 2019.
Archived at *perma.cc/A529-SKA9*

บทที่ 13

[1] Rachid Belaid. “Postgres Full-Text Search Is Good Enough!” *rachbelaid.com*, July 2015. Archived at *perma.cc/ZVP9-YDCB*

[2] Philippe Ajoux, Nathan Bronson, Sanjeev Kumar, Wyatt Lloyd, and Kaushik Veer-araghavan. “Challenges to Adopting Stronger Consistency at Scale.” At *15th USENIX Workshop on Hot Topics in Operating Systems* (HotOS), May 2015.

[3] Pat Helland and Dave Campbell. “Building on Quicksand.” At *4th Biennial Conference on Innovative Data Systems Research* (CIDR), January 2009. Archived at *arxiv.org*

[4] Jessica Kerr. “Provenance and Causality in Distributed Systems.” *jessitron.com*, September 2016. Archived at *perma.cc/DTD2-F8ZM*

[5] Jay Kreps. “The Log: What Every Software Engineer Should Know About Real-Time Data’s Unifying Abstraction.” *engineering.linkedin.com*, December 2013. Archived at *perma.cc/2JHR-FR64*

[6] Pat Helland. “Life Beyond Distributed Transactions: An Apostate’s Opinion.” At *3rd Biennial Conference on Innovative Data Systems Research* (CIDR), January 2007. Archived at *perma.cc/2GZG-UZ65*

[7] Lionel A. Smith. “The Broad Gauge Story.” *Journal of the Monmouthshire Railway Society*, Summer 1985. Archived at *perma.cc/DDK9-JA6X*

[8] Jacqueline Xu. “Online Migrations at Scale.” *stripe.com*, February 2017. Archived at *perma.cc/ZQY2-EAU2*

[9] Flavio Santos and Robert Stephenson. “Changing the Wheels on a Moving Bus—Spotify’s Event Delivery Migration.” *engineering.atspotify.com*, October 2021. Archived at *perma.cc/5C4V-G8EV*

[10] Molly Bartlett Dishman and Martin Fowler. “Agile Architecture.” At *O’Reilly Software Architecture Conference*, March 2015.

[11] Nathan Marz and James Warren. *Big Data: Principles and Best Practices of Scalable Real-Time Data Systems*. Manning, 2015. ISBN: 9781617290343

[12] Jay Kreps. “Questioning the Lambda Architecture.” *oreilly.com*, July 2014. Archived at *perma.cc/PGH6-XUCH*

[13] Raul Castro Fernandez, Peter Pietzuch, Jay Kreps, Neha Narkhede, Jun Rao, JoelKoshy, Dong Lin, Chris Riccomini, and Guozhang Wang. “Liquid: Unifying Nearline and Offline Big Data Integration.” At *7th Biennial Conference on Innovative Data Systems Research* (CIDR), January 2015. Archived at *perma.cc/QMA9-8PKL*

[14] Dennis M. Ritchie and Ken Thompson. “The UNIX Time-Sharing System.” *Communications of the ACM*, volume 17, issue 7, pages 365–375, July 1974. *doi:10.1145/361011.361061*

[15] Wes McKinney. “The Road to Composable Data Systems: Thoughts on the Last 15 Years and the Future.” *wesmckinney.com*, September 2023. Archived at *perma.cc/ J9SJ-886N*

[16] Eric A. Brewer and Joseph M. Hellerstein. “CS262a: Advanced Topics in Computer Systems.” Lecture notes, University of California, Berkeley, *cs.berkeley.edu*, August 2011. Archived at *perma.cc/TE79-LGWU*

[17] Michael Stonebraker. “The Case for Polystores.” *wp.sigmod.org*, July 2015. Archived at *perma.cc/G7J2-KR45*

[18] Jennie Duggan, Aaron J. Elmore, Michael Stonebraker, Magda Balazinska, BillHowe, Jeremy Kepner, Sam Madden, David Maier, Tim Mattson, and Stan Zdonik. “The BigDAWG Polystore System.” *ACM SIGMOD Record*, volume 44, issue 2, pages 11–16, June 2015. *doi:10.1145/2814710.2814713*

[19] David B. Lomet, Alan Fekete, Gerhard Weikum, and Mike Zwilling. “Unbundling Transaction Services in the Cloud.” At *4th Biennial Conference on Innovative Data Systems Research* (CIDR), January 2009. Archived at *arxiv.org*

[20] Martin Kleppmann and Jay Kreps. “Kafka, Samza and the Unix Philosophy of Distributed Data.” *IEEE Data Engineering Bulletin*, volume 38, issue 4, pages 4–14, December 2015. Archived at *perma.cc/BJM5-TJ4Z*

[21] John Hugg. “Winning Now and in the Future: Where Volt Active Data Shines.” *voltactivedata.com*, March 2016. Archived at *perma.cc/44MP-3MWM*

[22] Felienne Hermans. “Spreadsheets Are Code.” At *Code Mesh*, November 2015.

[23] Dan Bricklin and Bob Frankston. “VisiCalc: Information from Its Creators.” *danbricklin.com*. Archived at *archive.org*

[24] D. Sculley, Gary Holt, Daniel Golovin, Eugene Davydov, Todd Phillips, DietmarEbner, Vinay Chaudhary, and Michael Young. “Machine Learning: The High-Interest Credit Card of Technical Debt.” At *NIPS Workshop on Software Engineering for Machine Learning* (SE4ML), December 2014. Archived at *perma.cc/M3MD-U7WL*

[25] Peter Bailis, Alan Fekete, Michael J. Franklin, Ali Ghodsi, Joseph M. Hellerstein, and Ion Stoica. “Feral Concurrency Control: An Empirical Investigation of Modern Application Integrity.” At *ACM International Conference on Management of Data* (SIGMOD), June 2015. *doi:10.1145/2723372.2737784*

[26] Guy Steele. “Re: Need for Macros (Was Re: Icon).” Email to *ll1-discuss* mailing list, *people.csail.mit.edu*, December 2001. Archived at *perma.cc/K9X8-CJ65*

[27] Ben Stopford. “Microservices in a Streaming World.” At *QCon London*, March 2016.

[28] Adam Bellemare. *Building Event-Driven Microservices*, 2nd edition. O’Reilly Media, 2025. ISBN: 9798341622180

[29] Christian Posta. “Why Microservices Should Be Event Driven: Autonomy vs Authority.” *blog.christianposta.com*, May 2016. Archived at *perma.cc/E6N9-3X92* [30] Alex Feyerke. “Designing Offline-First Web Apps.” *alistapart.com*, December 2013. Archived at *perma.cc/WH7R-S2DS*

[31] Martin Kleppmann. “Turning the Database Inside-out with Apache Samza.” At *Strange Loop*, September 2014. Archived at *perma.cc/U6E8-A9MT*

[32] Sebastian Burckhardt, Daan Leijen, Jonathan Protzenko, and Manuel Fähndrich.“Global Sequence Protocol: A Robust Abstraction for Replicated Shared State.” At *29th European Conference on Object-Oriented Programming* (ECOOP), July 2015. *doi:10.4230/LIPIcs.ECOOP.2015.568*

[33] Evan Czaplicki and Stephen Chong. “Asynchronous Functional Reactive Programming for GUIs.” At *34th ACM SIGPLAN Conference on Programming Language Design and Implementation* (PLDI), June 2013. *doi:10.1145/2491956.2462161*

[34] Eno Thereska, Damian Guy, Michael Noll, and Neha Narkhede. “Unifying Stream Processing and Interactive Queries in Apache Kafka.” *confluent.io*, October 2016. Archived at *perma.cc/W8JG-EAZF*

[35] Frank McSherry. “Dataflow as Database.” *github.com*, July 2016. Archived at *perma.cc/384D-DUFH*

[36] Peter Alvaro. “I See What You Mean.” At *Strange Loop*, September 2015.

[37] Nathan Marz. “Trident: A High-Level Abstraction for Realtime Computation.” *blog.x.com*, August 2012. Archived at *archive.org*

[38] Edi Bice. “Low Latency Web Scale Fraud Prevention with Apache Samza, Kafka and Friends.” At *Merchant Risk Council MRC Vegas Conference*, March 2016. Archived at *perma.cc/T3H5-QN3R*

[39] Charity Majors. “The Accidental DBA.” *charity.wtf*, October 2016. Archived at *perma.cc/6ANP-ARB6*

[40] Arthur J. Bernstein, Philip M. Lewis, and Shiyong Lu. “Semantic Conditions for Correctness at Different Isolation Levels.” At *16th International Conference on Data Engineering* (ICDE), February 2000. *doi:10.1109/ICDE.2000.839387*

[41] Sudhir Jorwekar, Alan Fekete, Krithi Ramamritham, and S. Sudarshan. “Automating the Detection of Snapshot Isolation Anomalies.” At *33rd International Conference on Very Large Data Bases* (VLDB), September 2007.

[42] Kyle Kingsbury. “Distributed Systems Safety Research.” *jepsen.io*.

[43] Michael Jouravlev. “Redirect After Post.” *theserverside.com*, August 2004. Archived at *archive.org*

[44] Jerome H. Saltzer, David P. Reed, and David D. Clark. “End-to-End Arguments in System Design.” *ACM Transactions on Computer Systems*, volume 2, issue 4, pages 277–288, November 1984. *doi:10.1145/357401.357402*

[45] Peter Bailis, Alan Fekete, Michael J. Franklin, Ali Ghodsi, Joseph M. Heller-stein, and Ion Stoica. “Coordination Avoidance in Database Systems.” *Proceedings of the VLDB Endowment*, volume 8, issue 3, pages 185–196, November 2014. *doi:10.14778/2735508.2735509*, extended version published as *arXiv:1402.2237*

[46] Alex Yarmula. “Strong Consistency in Manhattan.” *blog.x.com*, March 2016. Archived at *archive.org*

[47] Martin Kleppmann, Alastair R. Beresford, and Boerge Svingen. “Online Event Processing: Achieving Consistency Where Distributed Transactions Have Failed.” *Communications of the ACM*, volume 62, issue 5, pages 43–49, May 2019. *doi:10.1145/3312527*

[48] Jim Gray. “The Transaction Concept: Virtues and Limitations.” At *7th International Conference on Very Large Data Bases* (VLDB), September 1981. Archived at *perma.cc/8VPT-N5H6*

[49] Hector Garcia-Molina and Kenneth Salem. “Sagas.” At *ACM International Conference on Management of Data* (SIGMOD), May 1987. *doi:10.1145/38713.38742*

[50] Annamalai Gurusami and Daniel Price. “Bug #73170: Duplicates in Unique Secondary Index Because of Fix of Bug#68021.” *bugs.mysql.com*, July 2014. Archived at *perma.cc/P6BV-W7JJ*

[51] Gary Fredericks. “Postgres Serializability Bug.” *github.com*, September 2015. Archived at *perma.cc/N8UP-2822*

[52] Xiao Chen. “HDFS DataNode Scanners and Disk Checker Explained.” *blog.cloudera.com*, December 2016. Archived at *perma.cc/6S36-X98L*

[53] Daniel Persson. “How Does Ceph Scrubbing Work?” *youtube.com*, March 2022.

[54] Jay Kreps. “Getting Real About Distributed System Reliability.” *blog.empathybox.com*, March 2012. Archived at *perma.cc/9B5Q-AEBW*

[55] Martin Fowler. “The LMAX Architecture.” *martinfowler.com*, July 2011. Archived at *perma.cc/5AV4-N6RJ*

[56] Sam Stokes. “Move Fast with Confidence.” *five-eights.com*, July 2016. Archived at *perma.cc/J8C6-DHXB*

[57] Ralph C. Merkle. “A Digital Signature Based on a Conventional Encryption Function.” At *CRYPTO ’87*, August 1987. *doi:10.1007/3-540-48184-2_32*

[58] Ben Laurie. “Certificate Transparency.” *ACM Queue*, volume 12, issue 8, pages 10–19, August 2014. *doi:10.1145/2668152.2668154*

[59] Mark D. Ryan. “Enhanced Certificate Transparency and End-to-End Encrypted Mail.” At *Network and Distributed System Security Symposium* (NDSS), February 2014. *doi:10.14722/ndss.2014.23379*

บทที่ 14

[1] David Schmudde. “What If Data Is a Bad Idea?” *schmud.de*, August 2024. Archived at *perma.cc/ZXU5-XMCT*

[2] Association for Computing Machinery. “ACM Code of Ethics and Professional Conduct.” *acm.org*, 2018. Archived at *perma.cc/SEA8-CMB8*

[3] Igor Perisic. “Making Hard Choices: The Quest for Ethics in Machine Learning.” *linkedin.com*, November 2016. Archived at *perma.cc/DGF8-KNT7*

[4] John Naughton. “Algorithm Writers Need a Code of Conduct.” *theguardian.com*, December 2015. Archived at *perma.cc/TBG2-3NG6*

[5] Deborah G. Johnson and Mario Verdicchio. “Ethical AI Is Not About AI.” *Communications of the ACM*, volume 66, issue 2, pages 32–34, January 2023. *doi:10.1145/3576932*

[6] Ben Green. “‘Good’ Isn’t Good Enough.” At *NeurIPS Joint Workshop on AI for Social Good*, December 2019. Archived at *perma.cc/H4LN-7VY3*

[7] Marc Steen. “Ethics as a Participatory and Iterative Process.” *Communications of the ACM*, volume 66, issue 5, pages 27–29, April 2023. *doi:10.1145/3550069*

[8] Logan Kugler. “What Happens When Big Data Blunders?” *Communications of the ACM*, volume 59, issue 6, pages 15–16, June 2016. *doi:10.1145/2911975*

[9] Miri Zilka. “Algorithms and the Criminal Justice System: Promises and Challenges in Deployment and Research.” At *University of Cambridge Security Seminar Series*, March 2023. Archived at *archive.org*

[10] Bill Davidow. “Welcome to Algorithmic Prison.” *theatlantic.com*, February 2014. Archived at *archive.org*

[11] Don Peck. “They’re Watching You at Work.” *theatlantic.com*, December 2013. Archived at *perma.cc/YR9T-6M38*

[12] Leigh Alexander. “Is an Algorithm Any Less Racist Than a Human?” *theguardian.com*, August 2016. Archived at *perma.cc/XP93-DSVX*

[13] Jesse Emspak. “How a Machine Learns Prejudice.” *scientificamerican.com*, December 2016. perma.cc/R3L5-55E6

[14] Rohit Chopra, Kristen Clarke, Charlotte A. Burrows, and Lina M. Khan. “Joint Statement on Enforcement Efforts Against Discrimination and Bias in Automated Systems.” *ftc.gov*, April 2023. Archived at *perma.cc/YY4Y-RCCA*

[15] Maciej Cegłowski. “The Moral Economy of Tech.” *idlewords.com*, June 2016. Archived at *perma.cc/L8XV-BKTD*

[16] Greg Nichols. “Artificial Intelligence in Healthcare Is Racist.” *zdnet.com*, November 2020. Archived at *perma.cc/3MKW-YKRS*

[17] Cathy O’Neil. *Weapons of Math Destruction: How Big Data Increases Inequality and Threatens Democracy*. Crown Publishing, 2016. ISBN: 9780553418811

[18] Julia Angwin. “Make Algorithms Accountable.” *nytimes.com*, August 2016. Archived at *archive.org*

[19] Bryce Goodman and Seth Flaxman. “European Union Regulations on Algorithmic Decision-Making and a ‘Right to Explanation.’” At *ICML Workshop on Human Interpretability in Machine Learning*, June 2016. Archived at *arxiv.org*

[20] United States Senate Committee on Commerce, Science, and Transportation,Office of Oversight and Investigations, Majority Staff. “A Review of the Data Broker Industry: Collection, Use, and Sale of Consumer Data for Marketing Purposes.” Staff Report, *commerce.senate.gov*, December 2013. Archived at *perma.cc/32NV-YWLQ*

[21] Stephanie Assad, Robert Clark, Daniel Ershov, and Lei Xu. “Algorithmic Pricing and Competition: Empirical Evidence from the German Retail Gasoline Market.” *Journal of Political Economy*, volume 132, issue 3, pages 723–771, March 2024. *doi:10.1086/726906*

[22] Donella H. Meadows and Diana Wright. *Thinking in Systems: A Primer*. Chelsea Green Publishing, 2008. ISBN: 9781603580557

[23] Daniel J. Bernstein. “Listening to a ‘big data’/‘data science’ talk. Mentally translating ‘data’ to "‘surveillance’: ‘…everything starts with surveillance…’” *x.com*, May 2015. Archived at *perma.cc/EY3D-WBBJ*

[24] Marc Andreessen. “Why Software Is Eating the World.” *a16z.com*, August 2011. Archived at *perma.cc/3DCC-W3G6*

[25] J. M. Porup. “‘Internet of Things’ Security Is Hilariously Broken and Getting Worse.” *arstechnica.com*, January 2016. Archived at *archive.org*

[26] Bruce Schneier. *Data and Goliath: The Hidden Battles to Collect Your Data and Control Your World*. W. W. Norton, 2015. ISBN: 9780393352177

[27] The Grugq. “Nothing to Hide.” *grugq.tumblr.com*, April 2016. Archived at *perma.cc/BL95-8W5M*

[28] Federal Trade Commission. “FTC Takes Action Against General Motors for Sharing Drivers’ Precise Location and Driving Behavior Data Without Consent.” *ftc.gov*, January 2025. Archived at *perma.cc/3XGV-3HRD*

[29] Tony Beltramelli. “Deep-Spying: Spying Using Smartwatch and Deep Learning.” Masters thesis, IT University of Copenhagen, December 2015. Archived at *arxiv.org*

[30] Shoshana Zuboff. “Big Other: Surveillance Capitalism and the Prospects of an Information Civilization.” *Journal of Information Technology*, volume 30, issue 1, pages 75–89, April 2015. *doi:10.1057/jit.2015.5*

[31] Michiel Rhoen. “Beyond Consent: Improving Data Protection Through Consumer Protection Law.” *Internet Policy Review*, volume 5, issue 1, March 2016. *doi:10.14763/2016.1.404*

[32] “Regulation (EU) 2016/679 of the European Parliament and of the Council of 27 April 2016.” *Official Journal of the European Union*, L 119/1, May 2016.

[33] UK Information Commissioner’s Office. “What Is the ‘Legitimate Interests’ Basis?” *ico.org.uk*. Archived at *perma.cc/W8XR-F7ML*

[34] Tristan Harris. “How a Handful of Tech Companies Control Billions of Minds Every Day.” At *TED2017*, April 2017. Archived at *archive.org*

[35] Carina C. Zona. “Consequences of an Insightful Algorithm.” At *GOTO Berlin*, November 2016.

[36] Imanol Arrieta Ibarra, Leonard Goff, Diego Jiménez Hernández, Jaron Lanier,and E. Glen Weyl. “Should We Treat Data as Labor? Moving Beyond ‘Free.’” *American Economic Association Papers Proceedings*, volume 108, pages 38–42, May 2018. *doi:10.1257/pandp.20181003*

[37] Bruce Schneier. “Data Is a Toxic Asset, So Why Not Throw It Out?” *schneier.com*, March 2016. Archived at *perma.cc/4GZH-WR3D*

[38] Cory Scott. “Data is not toxic—which implies no benefit—but rather hazardous material, where we must balance need vs. want.” *x.com*, March 2016. Archived at *perma.cc/CLV7-JF2E*

[39] Mark Pesce. “Data Is The New Uranium—Incredibly Powerful And Amazingly Dangerous.” *theregister.com*, November 2024. Archived at *perma.cc/NV8B-GYGV*

[40] Bruce Schneier. “Mission Creep: When Everything Is Terrorism.” *schneier.com*, July 2013. Archived at *perma.cc/QB2C-5RCE*

[41] Lena Ulbricht and Maximilian von Grafenstein. “Big Data: Big Power Shifts?” *Internet Policy Review*, volume 5, issue 1, March 2016. *doi:10.14763/2016.1.406*

[42] Ellen P. Goodman and Julia Powles. “Facebook and Google: Most Powerful and Secretive Empires We’ve Ever Known.” *theguardian.com*, September 2016. Archived at *perma.cc/8UJA-43G6*

[43] Judy Estrin and Sam Gill. “The World Is Choking on Digital Pollution.” *washingtonmonthly.com*, January 2019. Archived at *perma.cc/3VHF-C6UC*

[44] A. Michael Froomkin. “Regulating Mass Surveillance as Privacy Pollution: Learning from Environmental Impact Statements.” *University of Illinois Law Review*, volume 2015, issue 5, August 2015. Archived at *perma.cc/24ZL-VK2T*

[45] Pengyuan Wang, Li Jiang, and Jian Yang. “The Early Impact of GDPR Compliance on Display Advertising: The Case of an Ad Publisher.” *Journal of Marketing Research*, volume 61, issue 1, April 2023. *doi:10.1177/00222437231171848*

[46] Johnny Ryan. “Don’t Be Fooled by Meta’s Fine for Data Breaches.” *The Economist*, May 2023. Archived at *perma.cc/VCR6-55HR*

[47] Jessica Leber. “Your Data Footprint Is Affecting Your Life in Ways You Can’t Even Imagine.” *fastcompany.com*, March 2016. Archived at *archive.org*

[48] Maciej Cegłowski. “Haunted by Data.” *idlewords.com*, October 2015. Archived at *archive.org*

[49] Sam Thielman. “You Are Not What You Read: Librarians Purge User Data to Protect Privacy.” *theguardian.com*, January 2016. Archived at *archive.org*

[50] Jez Humble. “It’s a cliché that people get into tech to ‘change the world.’ So then, you have to actually consider what the impact of your work is on the world. The idea that you can or should exclude societal and political discussions in tech is idiotic. It means you’re not doing your job.” *x.com*, April 2021. Archived at *perma.cc/3NYS-MHLC*