

- [10] Surajit Chaudhuri, Ravi Krishnamurthy, Spyros Potamianos, and Kyuseok Shim. 1995. Optimizing Queries with Materialized Views. In *Proceedings of the Eleventh International Conference on Data Engineering (ICDE '95)*. IEEE Computer Society, Washington, DC, USA, 190–200.
- [11] E. F. Codd. 1970. A Relational Model of Data for Large Shared Data Banks. *Commun. ACM* 13, 6 (June 1970), 377–387.
- [12] Alex Şuhan. Fast and Flexible Query Analysis at MapD with Apache Calcite. (Feb 2017). Retrieved November 20, 2017 from <https://www.mapd.com/blog/2017/02/08/fast-and-flexible-query-analysis-at-mapd-with-apache-calcite-2/>
- [13] Drill. Apache Drill. (Nov. 2017). Retrieved November 20, 2017 from <http://drill.apache.org/>
- [14] Druid. Druid. (Nov. 2017). Retrieved November 20, 2017 from <http://druid.io/>
- [15] Elastic. Elasticsearch. (Nov. 2017). Retrieved November 20, 2017 from <https://www.elastic.co>
- [16] Flink. Apache Flink. <https://flink.apache.org>. (Nov. 2017).
- [17] Yupeng Fu, Kian Win Ong, Yannis Papakonstantinou, and Michalis Petropoulos. 2011. The SQL-based all-declarative FORWARD web application development framework. In *CIDR*.
- [18] Jonathan Goldstein and Per-Åke Larson. 2001. Optimizing Queries Using Materialized Views: A Practical, Scalable Solution. *SIGMOD Rec.* 30, 2 (May 2001), 331–342.
- [19] Goetz Graefe. 1995. The Cascades Framework for Query Optimization. *IEEE Data Eng. Bull.* (1995).
- [20] Goetz Graefe and William J. McKenna. 1993. The Volcano Optimizer Generator: Extensibility and Efficient Search. In *Proceedings of the Ninth International Conference on Data Engineering*. IEEE Computer Society, Washington, DC, USA, 209–218.
- [21] Daniel Halperin, Victor Teixeira de Almeida, Lee Lee Choo, Shumo Chu, Paraschos Koutris, Dominik Moritz, Jennifer Ortiz, Vaspol Ruamviboonsuk, Jingjing Wang, Andrew Whitaker, Shengliang Xu, Magdalena Balazinska, Bill Howe, and Dan Suciu. 2014. Demonstration of the Myria Big Data Management Service. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD '14)*. ACM, New York, NY, USA, 881–884.
- [22] Venky Harinarayan, Anand Rajaraman, and Jeffrey D. Ullman. 1996. Implementing Data Cubes Efficiently. *SIGMOD Rec.* 25, 2 (June 1996), 205–216.
- [23] HBase. Apache HBase. (Nov. 2017). Retrieved November 20, 2017 from <http://hbase.apache.org/>
- [24] Hive. Apache Hive. (Nov. 2017). Retrieved November 20, 2017 from <http://hive.apache.org/>
- [25] Yin Huai, Ashutosh Chauhan, Alan Gates, Gunther Hagleitner, Eric N. Hanson, Owen O'Malley, Jitendra Pandey, Yuan Yuan, Rubao Lee, and Xiaodong Zhang. 2014. Major Technical Advancements in Apache Hive. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD '14)*. ACM, New York, NY, USA, 1235–1246.
- [26] Julian Hyde. 2010. Data in Flight. *Commun. ACM* 53, 1 (Jan. 2010), 48–52.
- [27] Janino. Janino: A super-small, super-fast Java compiler. (Nov. 2017). Retrieved November 20, 2017 from <http://www.janino.net/>
- [28] Kylin. Apache Kylin. (Nov. 2017). Retrieved November 20, 2017 from <http://kylin.apache.org/>
- [29] Avinash Lakshman and Prashant Malik. 2010. Cassandra: A Decentralized Structured Storage System. *SIGOPS Oper. Syst. Rev.* 44, 2 (April 2010), 35–40.
- [30] Lingual. Lingual. (Nov. 2017). Retrieved November 20, 2017 from <http://www.cascading.org/projects/lingual/>
- [31] Lucene. Apache Lucene. (Nov. 2017). Retrieved November 20, 2017 from <https://lucene.apache.org/>
- [32] MapD. MapD. (Nov. 2017). Retrieved November 20, 2017 from <https://www.mapd.com>
- [33] Erik Meijer, Brian Beckman, and Gavin Bierman. 2006. LINQ: Reconciling Object, Relations and XML in the .NET Framework. In *Proceedings of the 2006 ACM SIGMOD International Conference on Management of Data (SIGMOD '06)*. ACM, New York, NY, USA, 706–706.
- [34] Sergey Melnik, Andrey Gubarev, Jing Jing Long, Geoffrey Romer, Shiva Shivakumar, Matt Tolton, and Theo Vassilakis. 2010. Dremel: Interactive Analysis of Web-Scale Datasets. *PVLDB* 3, 1 (2010), 330–339. <http://www.comp.nus.edu.sg/~vlb2010/proceedings/files/papers/R29.pdf>
- [35] Marcelo RN Mendes, Pedro Bizarro, and Paulo Marques. 2009. A performance study of event processing systems. In *Technology Conference on Performance Evaluation and Benchmarking*. Springer, 221–236.
- [36] Mongo. MongoDB. (Nov. 2017). Retrieved November 28, 2017 from <https://www.mongodb.com/>
- [37] Christopher Olston, Benjamin Reed, Utkarsh Srivastava, Ravi Kumar, and Andrew Tomkins. 2008. Pig Latin: A Not-so-foreign Language for Data Processing. In *Proceedings of the 2008 ACM SIGMOD International Conference on Management of Data (SIGMOD '08)*. ACM, New York, NY, USA, 1099–1110.
- [38] Kian Win Ong, Yannis Papakonstantinou, and Romain Vernoux. 2014. The SQL++ query language: Configurable, unifying and semi-structured. *arXiv preprint arXiv:1405.3631* (2014).
- [39] Open Geospatial Consortium. OpenGIS Implementation Specification for Geographic information - Simple feature access - Part 2: SQL option. http://portal.opengeospatial.org/files/?artifact_id=25355. (2010).
- [40] Phoenix. Apache Phoenix. (Nov. 2017). Retrieved November 20, 2017 from <http://phoenix.apache.org/>
- [41] Pig. Apache Pig. (Nov. 2017). Retrieved November 20, 2017 from <http://pig.apache.org/>
- [42] Qubole Quark. Qubole Quark. (Nov. 2017). Retrieved November 20, 2017 from <https://github.com/qubole/quark>
- [43] Bikas Saha, Hitesh Shah, Siddharth Seth, Gopal Vijayaraghavan, Arun C. Murthy, and Carlo Curino. 2015. Apache Tez: A Unifying Framework for Modeling and Building Data Processing Applications. In *Proceedings of the 2015 ACM SIGMOD International Conference on Management of Data (SIGMOD '15)*. ACM, New York, NY, USA, 1357–1369.
- [44] Samza. Apache Samza. (Nov. 2017). Retrieved November 20, 2017 from <http://samza.apache.org/>
- [45] Mohamed A. Soliman, Lyublena Antova, Venkatesh Raghavan, Amr El-Helw, Zhongxian Gu, Entong Shen, George C. Caragea, Carlos Garcia-Alvarado, Foyzur Rahman, Michalis Petropoulos, Florian Waas, Sivaramakrishnan Narayanan, Konstantinos Krikellias, and Rhonda Baldwin. 2014. Orca: A Modular Query Optimizer Architecture for Big Data. In *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data (SIGMOD '14)*. ACM, New York, NY, USA, 337–348.
- [46] Solr. Apache Solr. (Nov. 2017). Retrieved November 20, 2017 from <http://lucene.apache.org/solr/>
- [47] Spark. Apache Spark. (Nov. 2017). Retrieved November 20, 2017 from <http://spark.apache.org/>
- [48] Splunk. Splunk. (Nov. 2017). Retrieved November 20, 2017 from <https://www.splunk.com/>
- [49] Michael Stonebraker and Ugur Çetintemel. 2005. "One size fits all": an idea whose time has come and gone. In *21st International Conference on Data Engineering (ICDE'05)*. IEEE Computer Society, Washington, DC, USA, 2–11.
- [50] Storm. Apache Storm. (Nov. 2017). Retrieved November 20, 2017 from <http://storm.apache.org/>
- [51] Tez. Apache Tez. (Nov. 2017). Retrieved November 20, 2017 from <http://tez.apache.org/>
- [52] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Suresh Anthony, Hao Liu, Pete Wyckoff, and Raghotham Murthy. 2009. Hive: a warehousing solution over a map-reduce framework. *VLDB* (2009), 1626–1629.
- [53] Immanuel Trummer and Christoph Koch. 2017. Multi-objective parametric query optimization. *The VLDB Journal* 26, 1 (2017), 107–124.
- [54] Ashwin Kumar Vajjantri, Kunwar Deep Singh Toor, and Edmon Begoli. 2017. An Apache Calcite-based Polystore Variation for Federated Querying of Heterogeneous Healthcare Sources. In *2nd Workshop on Methods to Manage Heterogeneous Big Data and Polystore Databases*. IEEE Computer Society, Washington, DC, USA.
- [55] Katherine Yu, Vijay Gadepally, and Michael Stonebraker. 2017. Database engine integration and performance analysis of the BigDAWG polystore system. In *2017 IEEE High Performance Extreme Computing Conference (HPEC)*. IEEE Computer Society, Washington, DC, USA, 1–7.
- [56] Matei Zaharia, Mosharaf Chowdhury, Michael J. Franklin, Scott Shenker, and Ion Stoica. 2010. Spark: Cluster Computing with Working Sets. In *HotCloud*.
- [57] Jingren Zhou, Per-Åke Larson, and Ronnie Chaiken. 2010. Incorporating partitioning and parallel plans into the SCOPE optimizer. In *2010 IEEE 26th International Conference on Data Engineering (ICDE 2010)*. IEEE Computer Society, Washington, DC, USA, 1060–1071.