Bibliography

[1] Daniel J. Abadi. Data management in the cloud: Limitations and opportunities. *IEEE Data Eng. Bull.*, 32(1):3–12, 2009. III.1

[2] Foto N. Afrati and Jeffrey D. Ullman. Optimizing joins in a map-reduce environment. In *Proceedings of the 13th International Conference on Extending Database Technology*, EDBT '10, pages 99–110, New York, NY, USA, 2010. ACM. V.1

[3] Sameer Agarwal, Dhruba Borthakur, and Ion Stoica. Snapshots in hadoop distributed file system. Technical report, EECS Department, University of California, Berkeley, Nov 2010. II.3(b)

[4] Divyakant Agrawal, Amr El Abbadi, Shyam Antony, and Sudipto Das. Data management challenges in cloud computing infrastructures. In Kikuchi et al. [73], pages 1–10. III.1

[5] Divyakant Agrawal, Sudipto Das, and Amr El Abbadi. Big data and cloud computing: current state and future opportunities. In *Proceedings of the 14th International Conference on Extending Database Technology*, EDBT/ICDT '11, pages 530–533, New York, NY, USA, 2011. ACM. II.2(b)

[6] Mona Ahuja, Cheng Che Chen, Ravi Gottapu, Jörg Hallmann, Waqar Hasan, Richard Johnson, Maciek Kozyrczak, Ramesh Pabbati, Neeta Pandit, Sreenivasulu Pokuri, and Krishna Uppala. Peta-scale data warehousing at yahoo! In *SIGMOD '09: Proceedings of the 35th SIGMOD international conference on Management of data*, pages 855–862, New York, NY, USA, 2009. ACM. II.2(b)

[7] Michael Armbrust, Armando Fox, Rean Griffith, Anthony D. Joseph, Randy Katz, Andy Konwinski, Gunho Lee, David Patterson, Ariel Rabkin, Ion Stoica, and Matei Zaharia. Above the clouds: A berkeley view of cloud computing. Technical report, UC Berkeley, February 2009. II.1

[8] Jay Kreps at Linkedin. Project voldemort. a distributed database. `http://project-voldemort.com/`, June 2010. II.2(a)

[9] Cliff Moon at Powerset. Dynomite. `http://project-voldemort.com/`, June 2010. II.2(a)

[10] Francis R. Bach and Michael I. Jordan. Kernel independent component analysis. *Journal of Machine Learning Research*, 3:1–48, 2002. II.6(c)

[11] Jost Berthold, Mischa Dieterle, and Rita Loogen. Implementing parallel google map-reduce in eden. In *Proceedings of the 15th International Euro-Par Conference on Parallel Processing*, Euro-Par '09, pages 990–1002, Berlin, Heidelberg, 2009. Springer-Verlag. II.3(a)

[12] Ken Birman, Gregory Chockler, and Robbert van Renesse. Toward a cloud computing research agenda. *SIGACT News*, 40(2):68–80, 2009. II.1

[13] Spyros Blanas, Jignesh M. Patel, Vuk Ercegovac, Jun Rao, Eugene J. Shekita, and Yuanyuan Tian. A comparison of join algorithms for log processing in mapreduce. In *Proceedings of the 2010 international conference on Management of data*, SIGMOD '10, pages 975–986, New York, NY, USA, 2010. ACM. V.1

[14] M. W. Bright, A. R. Hurson, and S. Pakzad. Automated resolution of semantic heterogeneity in multidatabases. *ACM Trans. Database Syst.*, 19:212–253, June 1994. III.3, IV.3(b), V.1

[15] Randal E. Bryant. Data-intensive Supercomputing: The Case for DISC. Technical report, Carnegie Mellon, School of Computer Science, 2007. III.3

[16] Mike Burrows. The chubby lock service for loosely-coupled distributed systems. In *Proceedings of the 7th symposium on Operating systems design and implementation*, OSDI '06, pages 335–350, Berkeley, CA, USA, 2006. USENIX Association.

[17] Nokia Research Center. Disco. `http://discoproject.org/`. II.3(a)

[18] Ronnie Chaiken, Bob Jenkins, Per-Ake Larson, Bill Ramsey, Darren Shakib, Simon Weaver, and Jingren Zhou. Scope: easy and efficient parallel processing of massive data sets. *Proc. VLDB Endow.*, 1:1265–1276, August 2008. II.4

[19] Jairam Chandar. Join algorithms using mapreduce. Master's thesis, School of Informatics, University of Edinburgh, United Kingdom, Jan 2010. V.1

[20] Fay Chang, Jeffrey Dean, Sanjay Ghemawat, Wilson C. Hsieh, Deborah A. Wallach, Mike Burrows, Tushar Chandra, Andrew Fikes, and Robert E. Gruber. Bigtable: a distributed storage system for structured data. In *Proceedings of the 7th USENIX Symposium on Operating Systems Design and Implementation - Volume 7*, OSDI '06, pages 15–15, Berkeley, CA, USA, 2006. USENIX Association. II.2(a)

[21] Brian F. Cooper, Raghu Ramakrishnan, Utkarsh Srivastava, Adam Silberstein, Philip Bohannon, Hans-Arno Jacobsen, Nick Puz, Daniel Weaver, and Ramana Yerneni. Pnuts: Yahoo!'s hosted data serving platform. *Proc. VLDB Endow.*, 1:1277–1288, August 2008.

[22] Transaction Processing Performance Council. TPC-DS: A new decision support workload. `http://www.tpc.org/tpcds/tpcds.asp`. IV.1, IV.1

[23] Transaction Processing Performance Council. Transaction processing performance council. `http://www.tpc.org/`. IV.1

[24] Jeffrey Dean and Sanjay Ghemawat. Mapreduce: simplified data processing on large clusters. In *Proceedings of the 6th conference on Symposium on Opearting Systems Design & Implementation - Volume 6*, pages 10–10, Berkeley, CA, USA, 2004. USENIX Association. I, II.3(a)

[25] Giuseppe DeCandia, Deniz Hastorun, Madan Jampani, Gunavardhan Kakulapati, Avinash Lakshman, Alex Pilchin, Swaminathan Sivasubramanian, Peter Vosshall, and Werner Vogels. Dynamo: amazon's highly available key-value store. In *Proceedings of twenty-first ACM SIGOPS symposium on Operating systems principles*, SOSP '07, pages 205–220, New York, NY, USA, 2007. ACM. II.2(a)

[26] Adam Ji Dou, Vana Kalogeraki, Dimitrios Gunopulos, Taneli Mielikäinen, and Ville H. Tuulos. Misco: a mapreduce framework for mobile systems. In *PETRA*, 2010. II.3(a)

[27] Allen B. Downey and Dror G. Feitelson. The elusive goal of workload characterization. *SIGMETRICS Perform. Eval. Rev.*, 26:14–29, March 1999. IV.4, V.2

[28] Dror G. Feitelson, Larry Rudolph, and Uwe Schwiegelshohn, editors. *Job Scheduling Strategies for Parallel Processing, 10th International*

Workshop, JSSPP 2004, New York, NY, USA, June 13, 2004, Revised Selected Papers, volume 3277 of *Lecture Notes in Computer Science*. Springer, 2005. 78

[29] FileMap. Filemap. `http://mfisk.github.com/filemap/`. II.3(a)

[30] The Apache Software Foundation. Apache hadoop. `http://hadoop.apache.org/`. I, II.3(a), III.3

[31] The Apache Software Foundation. Apache pig. `http://pig.apache.org/`. III.1, III.2, III.3, IV.3, V.1, V.4

[32] The Apache Software Foundation. The apache software foundation. `http://www.apache.org/`. V.4

[33] The Apache Software Foundation. Hadoop distributed file system. `http://wiki.apache.org/pig/PigMix/`. II.3(b), III.1

[34] The Apache Software Foundation. Hive. `http://hive.apache.org/`. III.1, IV.3, V.1, V.4

[35] The Apache Software Foundation. Nutch. `http://nutch.apache.org/`. II.3(a)

[36] The Apache Software Foundation. Pigmix. `http://wiki.apache.org/pig/PigMix/`. V.3

[37] Archana Sulochana Ganapathi. *Predicting and Optimizing System Utilization and Performance via Statistical Machine Learning*. PhD thesis, EECS Department, University of California, Berkeley, Dec 2009. III.1, V.1, V.2, V.3, V.4

[38] Archana Sulochana Ganapathi, Yanpei Chen, Armando Fox, Randy H. Katz, and David A. Patterson. Statistics-driven workload modeling for the cloud. Technical Report UCB/EECS-2009-160, EECS Department, University of California, Berkeley, Nov 2009. III.1, V.1, V.2

[39] Alan Gates. Programming Pig. `http://ofps.oreilly.com/titles/9781449302641/index.html`. II.5(a), II.5(b), IV.5(a)

[40] Lise Getoor, Benjamin Taskar, and Daphne Koller. Selectivity estimation using probabilistic models. In *Proceedings of the 2001 ACM SIGMOD international conference on Management of data*, SIGMOD '01, pages 461–472, New York, NY, USA, 2001. ACM. V.4

[41] Sanjay Ghemawat, Howard Gobioff, and Shun-Tak Leung. The google file system. In *Proceedings of the nineteenth ACM symposium on Operating systems principles*, SOSP '03, pages 29–43, New York, NY, USA, 2003. ACM. II.3(b)

[42] Computing industry set for a shocking change. retrieved may 5, 2010., 2009. II.1

[43] Robert Griesemer. Parallelism by design: data analysis with sawzall. In *Proceedings of the 6th annual IEEE/ACM international symposium on Code generation and optimization*, CGO '08, pages 3–3, New York, NY, USA, 2008. ACM. II.4, III.1

[44] Peter J. Haas, Jeffrey F. Naughton, and Arun N. Swami. On the relative cost of sampling for join selectivity estimation. In *Proceedings of the thirteenth ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, PODS '94, pages 14–24, New York, NY, USA, 1994. ACM. IV.3(b)

[45] Apache Hadoop. How many maps and reducers. `http://wiki.apache.org/hadoop/HowManyMapsAndReduces/`. IV.2

[46] Banchong Harangsri, John Shepherd, and Anne Ngu. Selectivity estimation for joins using systematic sampling. In *Proceedings of the 8th International Workshop on Database and Expert Systems Applications*, DEXA '97, pages 384–, Washington, DC, USA, 1997. IEEE Computer Society. IV.3(b)

[47] Bingsheng He, Wenbin Fang, Qiong Luo, Naga K. Govindaraju, and Tuyong Wang. Mars: a mapreduce framework on graphics processors. In *PACT*, pages 260–269, 2008. II.3(a)

[48] Apache Hive. Add ability to compute statistics on hive tables. `https://issues.apache.org/jira/browse/HIVE-33`. V.1

[49] Apache Hive. Automatically gathering stats when reading a table/partition. `https://issues.apache.org/jira/browse/HIVE-1648`. IV.3(b)

[50] Apache Hive. Column level statistics for the hive system. `https://issues.apache.org/jira/browse/HIVE-1362`. IV.3(b), V.1

[51] Apache Hive. Cost based query optimization for joins in hive. `https://issues.apache.org/jira/browse/HIVE-1938`. V.1

[52] Apache Hive. Partition level statistics for the hive system. `https://issues.apache.org/jira/browse/HIVE-1361`. IV.3(b)

[53] H. Hotelling. Analysis of a complex of statistical variables into principal components. *J. Educ. Psych.*, 24, 1933. II.6(b)

[54] Harold Hotelling. Relations Between Two Sets of Variates. *Biometrika*, 28(3/4):321–377, 1936. II.6(c)

[55] Wen-Chi Hou, Gultekin Ozsoyoglu, and Erdogan Dogdu. Error-constrained count query evaluation in relational databases. In *Proceedings of the 1991 ACM SIGMOD international conference on Management of data*, SIGMOD '91, pages 278–287, New York, NY, USA, 1991. ACM. IV.3(b)

[56] Wen-Chi Hou, Gultekin Ozsoyoglu, and Baldeo K. Taneja. Statistical estimators for relational algebra expressions. In *Proceedings of the seventh ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, PODS '88, pages 276–287, New York, NY, USA, 1988. ACM. IV.3(b)

[57] Apache Howl. Howl, a shared table management service for hadoop. `http://wiki.apache.org/pig/Howl`. IV.3(b), V.4

[58] Mei Hui, Dawei Jiang, Guoliang Li, and Yuan Zhou. Supporting database applications as a service. In *ICDE '09: Proceedings of the 2009 IEEE International Conference on Data Engineering*, pages 832–843, Washington, DC, USA, 2009. IEEE Computer Society. II.2(b)

[59] Mei Hui, Dawei Jiang, Guoliang Li, and Yuan Zhou. Supporting database applications as a service. In *Proceedings of the 2009 IEEE International Conference on Data Engineering*, pages 832–843, Washington, DC, USA, 2009. IEEE Computer Society.

[60] International Data Corporation (IDC). `http://www.idc.com/`, June 2010.

[61] Amazon Inc. Amazon. `http://www.amazon.com/`. III.1

[62] Facebook Inc. Facebook. `http://www.facebook.com/`. III.1

[63] Google Inc. Google. `http://www.google.com/`. III.1

[64] Google Inc. Google and ibm announce university initiative to address internet-scale computing challenges. `http://www.google.com/intl/en/press/pressrel/20071008_ibm_univ.html/`.

[65] Yahoo! Inc. M45 supercomputing cluster. `http://research.yahoo.com/node/1884`. V.1, V.2

[66] Yahoo Inc. Yahoo! `http://www.yahoo.com/`. III.1

[67] Zvents Inc. Hypertable. `http://hypertable.org/`, June 2010. II.2(a)

[68] Yannis E. Ioannidis and Stavros Christodoulakis. On the propagation of errors in the size of join results. In *Proceedings of the 1991 ACM SIGMOD international conference on Management of data*, SIGMOD '91, pages 268–277, New York, NY, USA, 1991. ACM. IV.3(b)

[69] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. *SIGOPS Oper. Syst. Rev.*, 41:59–72, March 2007. II.4

[70] Michael Isard, Mihai Budiu, Yuan Yu, Andrew Birrell, and Dennis Fetterly. Dryad: distributed data-parallel programs from sequential building blocks. In *Proceedings of the 2nd ACM SIGOPS/EuroSys European Conference on Computer Systems 2007*, EuroSys '07, pages 59–72, New York, NY, USA, 2007. ACM.

[71] Michael Isard and Yuan Yu. Distributed data-parallel computing using a high-level programming language. In *Proceedings of the 35th SIGMOD international conference on Management of data*, SIGMOD '09, pages 987–994, New York, NY, USA, 2009. ACM.

[72] Soila Kavulya, Jiaqi Tan, Rajeev Gandhi, and Priya Narasimhan. An analysis of traces from a production mapreduce cluster. In *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, CCGRID '10, pages 94–103, Washington, DC, USA, 2010. IEEE Computer Society. V.1, V.2, V.4

[73] Shinji Kikuchi, Shelly Sachdeva, and Subhash Bhalla, editors. *Databases in Networked Information Systems, 6th International Workshop, DNIS 2010, Aizu-Wakamatsu, Japan, March 29-31, 2010. Proceedings*, volume 5999 of *Lecture Notes in Computer Science*. Springer, 2010. 4

[74] Ioannis Konstantinou, Evangelos Angelou, Dimitrios Tsoumakos, and Nectarios Koziris. Distributed indexing of web scale datasets for the

cloud. In *MDAC '10: Proceedings of the 2010 Workshop on Massive Data Analytics on the Cloud*, pages 1–6, New York, NY, USA, 2010. ACM. II.2(b)

[75] Ioannis Konstantinou, Evangelos Angelou, Dimitrios Tsoumakos, and Nectarios Koziris. Distributed indexing of web scale datasets for the cloud. In *Proceedings of the 2010 Workshop on Massive Data Analytics on the Cloud*, MDAC '10, pages 1:1–1:6, New York, NY, USA, 2010. ACM.

[76] Avinash Lakshman and Prashant. Malik. Apache cassandra. `http://cassandra.apache.org/`, June 2010. II.2(a), III.1, IV.3

[77] Avinash Lakshman and Prashant Malik. Cassandra: a decentralized structured storage system. *SIGOPS Oper. Syst. Rev.*, 44:35–40, April 2010.

[78] Hui Li, David L. Groep, and Lex Wolters. Workload characteristics of a multi-cluster supercomputer. In Feitelson et al. [28], pages 176–193. IV.4, V.2

[79] Richard J. Lipton and Jeffrey F. Naughton. Query size estimation by adaptive sampling. In *Selected papers of the 9th annual ACM SIGACT-SIGMOD-SIGART symposium on Principles of database systems*, pages 18–25, Orlando, FL, USA, 1995. Academic Press, Inc. III.3, IV.3(b), V.1

[80] Richard J. Lipton, Jeffrey F. Naughton, and Donovan A. Schneider. Practical selectivity estimation through adaptive sampling. In *Proceedings of the 1990 ACM SIGMOD international conference on Management of data*, SIGMOD '90, pages 1–11, New York, NY, USA, 1990. ACM. IV.3(b), IV.3(b)

[81] J. B. MacQueen. Some methods for classification and analysis of multivariate observations. In L. M. Le Cam and J. Neyman, editors, *Proc. of the fifth Berkeley Symposium on Mathematical Statistics and Probability*, volume 1, pages 281–297. University of California Press, 1967. II.6(a)

[82] J. B. Macqueen. Some methods of classification and analysis of multivariate observations. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probability*, pages 281–297, 1967.

[83] Lijun Mei, W.K. Chan, and T.H. Tse. A tale of clouds: Paradigm comparisons and some thoughts on research issues. *Asia-Pacific Conference on Services Computing. 2006 IEEE*, 0:464–469, 2008. II.1

[84] P. Mell and T. Grance. The NIST Definition of Cloud Computing, 2009. I, II.1

[85] Kristi Morton, Magdalena Balazinska, and Dan Grossman. Paratimer: a progress indicator for mapreduce dags. In *Proceedings of the 2010 international conference on Management of data*, SIGMOD '10, pages 507–518, New York, NY, USA, 2010. ACM. III.1, V.1

[86] Kristi Morton, Abram Friesen, Magdalena Balazinska, and Dan Grossman. Estimating the progress of mapreduce pipelines. In Feifei Li, Mirella M. Moro, Shahram Ghandeharizadeh, Jayant R. Haritsa, Gerhard Weikum, Michael J. Carey, Fabio Casati, Edward Y. Chang, Ioana Manolescu, Sharad Mehrotra, Umeshwar Dayal, and Vassilis J. Tsotras, editors, *ICDE*, pages 681–684. IEEE, 2010. V.1

[87] MySpace. Qizmt. `http://qizmt.myspace.com/`. II.3(a)

[88] Raghunath Othayoth Nambiar and Meikel Poess. The making of tpc-ds. In *Proceedings of the 32nd international conference on Very large data bases*, VLDB '06, pages 1049–1058. VLDB Endowment, 2006. IV.1, IV.1

[89] Christopher Olston, Benjamin Reed, Adam Silberstein, and Utkarsh Srivastava. Automatic optimization of parallel dataflow programs. In *USENIX 2008 Annual Technical Conference on Annual Technical Conference*, pages 267–273, Berkeley, CA, USA, 2008. USENIX Association. III.3

[90] Christopher Olston, Benjamin Reed, Utkarsh Srivastava, Ravi Kumar, and Andrew Tomkins. 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, pages 1099–1110, New York, NY, USA, 2008. ACM. I, II.4, III.1, III.3

[91] Konstantina Palla. A comparative analysis of join algorithms using the hadoop map/reduce framework. Master's thesis, School of Informatics, University of Edinburgh, United Kingdom, Dec 2009. V.1

[92] Andrew Pavlo, Erik Paulson, Alexander Rasin, Daniel J. Abadi, David J. DeWitt, Samuel R. Madden, and Michael Stonebraker. A comparison of

approaches to large scale data analysis. In *SIGMOD*, Providence, Rhode Island, USA, 2009. II.4

[93] Adam Pisoni. Mapreduce in ruby - skynet. `http://skynet.rubyforge.org/`. II.3(a)

[94] Raghu Ramakrishnan and Johannes Gehrke. *Database Management Systems*. McGraw-Hill, Inc., New York, NY, USA, 3 edition, 2003. III.1, IV.3, IV.3

[95] Jon Sætrom and Henning Omre. Ensemble kalman filtering with shrinkage regression techniques. *Computational Geosciences*, July 2010. IV.5(b), V.4

[96] Sebastian M. Schlatt, Timo B. Hübel, Stefan Schmidt, Sebastian Reese, and Dr. Schmidt. Holumbus. `http://holumbus.fh-wedel.de/`. II.3(a)

[97] Toby Segaran and Jeff Hammerbacher. Beautiful data — the stories behind elegant data solutions.

[98] W. Smith. Prediction services for distributed computing. In *Parallel and Distributed Processing Symposium, 2007. IPDPS 2007. IEEE International*, pages 1 –10, march 2007. III.3, V.1, V.3, V.4

[99] Warren Smith, Ian T. Foster, and Valerie E. Taylor. Predicting application run times using historical information. In *Proceedings of the Workshop on Job Scheduling Strategies for Parallel Processing*, pages 122–142, London, UK, 1998. Springer-Verlag. V.1, V.3

[100] Ion Stoica, Robert Morris, David Karger, M. Frans Kaashoek, and Hari Balakrishnan. Chord: A scalable peer-to-peer lookup service for internet applications. In *Proceedings of the 2001 conference on Applications, technologies, architectures, and protocols for computer communications*, SIGCOMM '01, pages 149–160, New York, NY, USA, 2001. ACM. II.2(a)

[101] StumbleUpon. Hbase. `http://hbase.apache.org/`, June 2010. II.2(a), III.1, V.4

[102] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Suresh Anthony, Hao Liu, Pete Wyckoff, and Raghotham Murthy. Hive: a warehousing solution over a map-reduce framework. *Proc. VLDB Endow.*, 2:1626–1629, August 2009. I, II.4, III.1, III.3

[103] University of Texas. Texas Advanced Computing Center Lonestar System. `http://www.tacc.utexas.edu/resources/hpc/#lonestar`. III.3, V.1

[104] Luis M. Vaquero, Luis Rodero-Merino, Juan Caceres, and Maik Lindner. A break in the clouds: towards a cloud definition. *SIGCOMM Comput. Commun. Rev.*, 39(1):50–55, 2009. II.1

[105] Sudharshan Vazhkudai, Jennifer M. Schopf, and Ian T. Foster. Predicting the performance of wide area data transfers. In *Proceedings of the 16th International Parallel and Distributed Processing Symposium*, IPDPS '02, pages 270–, Washington, DC, USA, 2002. IEEE Computer Society. III.3, V.1, V.3

[106] Jason Venner. *Pro Hadoop*. Apress, Berkely, CA, USA, 1st edition, 2009. (document), II.1

[107] Wikipedia. Nosql. `http://en.wikipedia.org/wiki/NoSQL/`, June 2010.

[108] Fan Yang, Jayavel Shanmugasundaram, and Ramana Yerneni. A scalable data platform for a large number of small applications. In *CIDR*. www.crdrdb.org, 2009. II.2(b)

[109] Richard M. Yoo, Anthony Romano, and Christos Kozyrakis. Phoenix rebirth: Scalable mapreduce on a large-scale shared-memory system. In *Proceedings of the 2009 IEEE International Symposium on Workload Characterization (IISWC)*, IISWC '09, pages 198–207, Washington, DC, USA, 2009. IEEE Computer Society. II.3(a)

[110] L. Youseff, M. Butrico, and D. Da Silva. Toward a unified ontology of cloud computing. In *Grid Computing Environments Workshop, 2008. GCE '08*, pages 1–10, Nov 2008. II.1

[111] Yuan Yu, Michael Isard, Dennis Fetterly, Mihai Budiu, Úlfar Erlingsson, Pradeep Kumar Gunda, and Jon Currey. Dryadlinq: a system for general-purpose distributed data-parallel computing using a high-level language. In *Proceedings of the 8th USENIX conference on Operating systems design and implementation*, OSDI'08, pages 1–14, Berkeley, CA, USA, 2008. USENIX Association. II.4, III.1

[112] Qiang Zhu. An integrated method for estimating selectivities in a multidatabase system. In *Proceedings of the 1993 conference of the Centre*

for Advanced Studies on Collaborative research: distributed computing - Volume 2, CASCON '93, pages 832–847. IBM Press, 1993. III.3, IV.3(b), IV.3(b), V.1

[113] Qiang Zhu and Per-Ake Larson. A query sampling method of estimating local cost parameters in a multidatabase system. In *Proceedings of the Tenth International Conference on Data Engineering*, pages 144–153, Washington, DC, USA, 1994. IEEE Computer Society. III.3, V.1

Appendix

A Pig-Latin example queries

(a) Fragment replicate join query

```
item = LOAD 'pigData/item.dat' using PigStorage('|') AS
(i_item_sk:int, i_item_id:chararray, i_rec_start_date:chararray,
i_rec_end_date:chararray, i_item_desc:chararray,
i_current_price:double, i_wholesale_cost:double, i_brand_id:int,
i_brand:chararray, i_class_id:int, i_class:chararray,
i_category_id:int, i_category:chararray, i_manufact_id:int,
i_manufact:chararray, i_size:chararray, i_formulation:chararray,
i_color:chararray, i_units:chararray, i_container:chararray,
i_manager_id:int, i_product_name:chararray);

store_sales = LOAD 'pigData/store_sales.dat' using PigStorage('|') AS
(ss_sold_date_sk:int, ss_sold_time_sk:int, ss_item_sk:int,
ss_customer_sk:int, ss_cdemo_sk:int, ss_hdemo_sk:int, ss_addr_sk:int,
ss_store_sk:int, ss_promo_sk:int, ss_ticket_number:int,
ss_quantity:int, ss_wholesale_cost:double, ss_list_price:double,
ss_sales_price:double, ss_ext_discount_amt:double,
ss_ext_sales_price:double, ss_ext_wholesale_cost:double,
ss_ext_list_price:double, ss_ext_tax:double, ss_coupon_amt:double,
ss_net_paid:double, ss_net_paid_inc_tax:double, ss_net_profit:double);

it_proy = FOREACH item GENERATE i_item_sk;
ss_proy = FOREACH store_sales GENERATE ss_item_sk;

query52_ItSs = JOIN ss_proy BY ss_item_sk, it_proy BY i_item_sk USING
'replicated';

STORE query52_ItSs INTO 'queryResults/query.52.store_sales.item.3.1'
using PigStorage('|');
```

(b) Merge join query

```
catalog_sales = LOAD 'pigData/catalog_sales.dat' using PigStorage('|')
AS (cs_sold_date_sk:int, cs_sold_time_sk:int, cs_ship_date_sk:int,
```

```
cs_bill_customer_sk:int, cs_bill_cdemo_sk:int, cs_bill_hdemo_sk:int,
cs_bill_addr_sk:int, cs_ship_customer_sk:int,cs_ship_cdemo_sk:int,
cs_ship_hdemo_sk:int,cs_ship_addr_sk:int,
cs_call_center_sk:int,cs_catalog_page_sk:int,
cs_ship_mode_sk:int,cs_warehouse_sk:int,
cs_item_sk:int,cs_promo_sk:int,cs_order_number:int,
cs_quantity:int,cs_wholesale_cost:double, cs_list_price:double,
cs_sales_price:double, cs_ext_discount_amt:double,
cs_ext_sales_price:double, cs_ext_wholesale_cost:double,
cs_ext_list_price:double, cs_ext_tax:double, cs_coupon_amt:double,
cs_ext_ship_cost:double, cs_net_paid:double,
cs_net_paid_inc_tax:double, cs_net_paid_inc_ship:double,
cs_net_paid_inc_ship_tax:double, cs_net_profit:double);

date_dim = LOAD 'pigData/date_dim.dat' using PigStorage('|') AS
(d_date_sk:int, d_date_id:chararray, d_date:chararray,
d_month_seq:int, d_week_seq:int, d_quarter_seq:int, d_year:int,
d_dow:int, d_moy:int, d_dom:int, d_qoy:int, d_fy_year:int,
d_fy_quarter_seq:int, d_fy_week_seq:int, d_day_name:chararray,
d_quarter_name:chararray, d_holiday:chararray, d_weekend:chararray,
d_following_holiday:chararray, d_first_dom:int, d_last_dom:int,
d_same_day_ly:int, d_same_day_lq:int, d_current_day:chararray,
d_current_week:chararray, d_current_month:chararray,
d_current_quarter:chararray, d_current_year:chararray);

dd_proy = FOREACH date_dim GENERATE d_date_sk;
cs_proy = FOREACH catalog_sales GENERATE cs_sold_date_sk;

query59_DdCa = JOIN dd_proy BY d_date_sk, cs_proy BY cs_sold_date_sk
USING 'merge';
```

(c) Hash join query

```
store_returns = LOAD 'pigData/store_returns.dat' using PigStorage('|')
AS (sr_returned_date_sk:int, sr_return_time_sk:int, sr_item_sk:int,
sr_customer_sk:int, sr_cdemo_sk:int, sr_hdemo_sk:int, sr_addr_sk:int,
sr_store_sk:int, sr_reason_sk:int, sr_ticket_number:int,
sr_return_quantity:int, sr_return_amt:double, sr_return_tax:double,
sr_return_amt_inc_tax:double, sr_fee:double,
sr_return_ship_cost:double, sr_refunded_cash:double,
```

```
sr_reversed_charge:double, sr_store_credit:double,
sr_net_loss:double);

customer_demographics = LOAD 'pigData/customer_demographics.dat' using
PigStorage('|') AS (cd_demo_sk:int, cd_gender:chararray,
cd_marital_status:chararray, cd_education_status:chararray,
cd_purchase_estimate:int, cd_credit_rating:chararray,
cd_dep_count:int, cd_dep_employed_count:int,
cd_dep_college_count:int);

proy_storeReturns = FOREACH store_returns GENERATE sr_cdemo_sk;
proy_cDemo = FOREACH customer_demographics GENERATE cd_demo_sk;

join_cDemo_storeReturns = JOIN proy_cDemo BY cd_demo_sk,
proy_storeReturns BY sr_cdemo_sk;

STORE join_cDemo_storeReturns INTO
'queryResults/query.11.customer_demographics.store_returns.1.5' using
PigStorage('|');
```