About the Authors

Prof. Dr. h.c. Hasso Plattner is a co-founder of SAP AG, where he served as the CEO until 2003 and has since been chairman of the supervisory board. SAP AG is today the leading provider of enterprise software solutions. In his role as chief software advisor, he concentrates on defining the mid- and long-term technology strategy and direction of SAP. Hasso Plattner received his diploma in communications engineering from the University of Karlsruhe. In recent years, he has been focusing on teaching and research in the field of business computing and software engineering at large. In 1998, he founded the Hasso Plattner Institute (HPI) in Potsdam, Germany. At the HPI, approximately 480 students are currently pursuing their Bachelors' and Masters' degrees in IT Systems Engineering with the help of roughly 50 professors and lecturers. The HPI currently has about 100 PhD candidates. Hasso Plattner leads one of the research groups at HPI which focuses mainly on In-Memory Data Management for Enterprise Applications and Human-Centered Software Design and Engineering (see epic.hpi.uni-potsdam.de).

Dr. Alexander Zeier graduated from the University of Wuerzburg in business management and successfully completed his studies in information technology at the TU Chemnitz. He worked for a few years as a strategic IT consultant, before gaining his Ph.D. in Supply Chain Management (SCM) at the University of Erlangen-Nuremberg. He has 20 years experience with IT/SAP Systems and started working for SAP in 2002 as product manager with overall responsibility for the SCM Software, SAP's first large In-Memory Application. Since 2006 he has been Deputy Chair Enterprise Platform and Integration Concepts of Prof. Hasso Plattner at the Hasso Plattner Institute in Potsdam, focusing on real-time In-Memory Enterprise Systems. During that time he has also been Executive Director for the European Section of the MIT Forum for Supply Chain Innovation. Since March 2012 Dr. Zeier has been working at the Massachusetts Institute of Technology (MIT) as Visiting Professor, lecturing and conducting research in the area of In-Memory Technology & Applications, and Supply Chain Innovation. He is the author of more than 150 journal articles and papers and has also published six books on IT and SAP (see zeier.mit.edu).

H. Plattner and A. Zeier, *In-Memory Data Management*,
DOI: 10.1007/978-3-642-29575-1, © Springer-Verlag Berlin Heidelberg 2012

Glossary

ACID Property of a database management system to always ensure atomicity consistency, isolation, and durability of its transactions.

Active Data Data of a business transaction that is not yet completed and is therefore always kept in main memory to ensure low latency access.

Aggregation Operation on data that creates a summarized result for example, a sum, maximum, average, and so on. Aggregation operations are common in enterprise applications.

Analytical Processing Method to enable or support business decisions by giving fast and intuitive access to large amounts of enterprise data.

Application Programming Interface (API) Aninterface for application programmers to access the functionality of a software system.

Atomicity Database concept that demands that all actions of a transaction are executed or none of them.

Attribute A characteristic of an entity describing a certain detail of it.

Availability Characteristic of a system to continuously operate according to its specification measured by the ratio between the accumulated time of correct operation and the overall interval.

Available-to-Promise (ATP) Determining whether sufficient quantities of a requested product will be available in current and planned inventory levels at a required date in order to allow decision making about accepting orders for this product.

Batch Processing Method of carrying out a larger number of operations without manual intervention.

Benchmark A set of operations run on specified data in order to evaluate the performance of a system.

H. Plattner and A. Zeier, *In-Memory Data Management*,
DOI: 10.1007/978-3-642-29575-1, © Springer-Verlag Berlin Heidelberg 2012

Blade Server in a modular design to increase the density of available computing power.

Business Intelligence Methods and processes using enterprise data for analytical and planning purposes or to create reports required by management.

Business Logic Representation of the actual business tasks of the problem domain in a software system.

Business Object Representation of a real-life entity in the data model for example, a purchasing order.

Cache A fast but rather small memory that serves as buffer for larger but slower memory.

Cache Coherence State of consistency between the versions of data stored in the local caches of a CPU cache.

Cache-Conscious Algorithm An algorithm is cache conscious if program variables that are dependent on hardware configuration parameters (for example cache size and cache-line length) need to be tuned to minimize the number of cache misses.

Cache Line Smallest unit of memory that can be transferred between main memory and the processor's cache. It is of a fixed size which depends on the respective processor type.

Cache Miss Afailed request for data from a cache because it did not contain the requested data.

Cache-Oblivious Algorithm An algorithm is cache oblivious if no program variables that are dependent on hardware configuration parameters (for example cache size and cache-line length) need to be tuned to minimize the number of cache misses.

Characteristic-Oriented Database System A database system that is tailored towards the characteristics of special application areas. Examples are text mining stream processing and data warehousing.

Cloud Computing An IT provisioning model which emphasizes the ondemand, elastic pay-per-use rendering of services or provisioning of resources over a network.

Column Store Database storage engine that stores each column (attribute) of a table sequentially in a contiguous area of memory.

Compression Encoding information in such a way that its representation consumes less space in memory.

Concurrency Control Techniques that allow the simultaneous and independent execution of transactions in a database system without creating states of unwanted incorrectness.

Consistency Database concept that demands that only correct database states are visible to the user despite the execution of transactions.

Consolidation Placing the data of several customers on one server machine database or table in a multi-tenant setup.

Cube Specialized OLAP data structure that allows multi-dimensional analysis of data.

Customer Relationship Management (CRM) Business processes and respective technology used by a company to organize its interaction with its customers.

Data Aging The changeover from active data to passive data.

Data Center Facility housing servers and associated ICT components.

Data Dictionary Metadata repository.

Data Layout The structure in which data is organized in the database that is the database's physical schema.

Data Mart A database that maintains copies of data from a specific business area for example, sales or production, for analytical processing purposes.

Data Warehouse A database that maintains copies of data from operational databases for analytical processing purposes.

Database Management System (DBMS) A set of administrative programs used to create maintain and manage a database.

Database Schema Formal description of the logical structure of a database.

Demand Planning Estimating future sales by combining several sources of information.

Design Thinking A methodology that combines an end-user focus with multidisciplinary collaboration and iterative improvement. It aims at creating desirable user-friendly, and economically viable design solutions and innovative products and services.

Desirability Design thinking term expressing the practicability of a system from a human-usability point of view.

Dictionary In the context of this book the compressed and sorted repository holding all distinct data values referenced by SanssouciDB's main store.

Dictionary Encoding Light-weight compression technique that encodes variable length values by smaller fixed-length encoded values using a mapping dictionary.

Differential Buffer A write-optimized buffer to increase write performance of the SanssouciDB column store. Sometimes also referred to as differential store or delta store.

Distributed System A system consisting of a number of autonomous computers that communicate over a computer network.

Dunning The process of scanning through open invoices and identifying overdue ones in order to take appropriate steps according to the dunning level.

Durability Database concept that demands that all changes made by a transaction become permanent after this transaction has been committed.

Enterprise Application A software system that helps an organization to run its business. A key feature of an enterprise application is its ability to integrate and process up-to-the-minute data from different business areas providing a holistic real-time view of the entire enterprise.

Enterprise Resource Planning (ERP) Enterprise software to support the resource planning processes of an entire company.

Extract-Transform-Load (ETL) Process A process that extracts data required for analytical processing from various sources then transforms it (into an appropriate format, removing duplicates, sorting, aggregating, etc.) such that it can be finally loaded into the target analytical system.

Fault Tolerance Quality of a system to maintain operation according to its specification even if failures occur.

Feasibility Design thinking term expressing the practicability of a system from a technical point of view.

Front Side Bus (FSB) Bus that connects the processor with main memory (and the rest of the computer).

Horizontal Partitioning The splitting of tables with many rows into several partitions each having fewer rows.

Hybrid Store Database that allows mixing column- and row-wise storage.

In-Memory Database Adatabase system that always keeps its primary data completely in main memory.

Index Data structure in a database used to optimize read operations.

Insert-Only New and changed tuples are always appended already existing changed and deleted tuples are then marked as invalid.

Inter-Operator Parallelism Parallel execution of independent plan operators of one or multiple query plans.

Intra-Operator Parallelism Parallel execution of a single plan operation independently of any other operation of the query plan.

Isolation Database concept demanding that any two concurrently executed transactions have the illusion that they are executed alone. The effect of such an isolated execution must not differ from executing the respective transactions one after the other.

Join Database operation that is logically the cross product of two or more tables followed by a selection.

Latency The time that a storage device needs between receiving the request for a piece of data and transmitting it.

Locking A method to achieve isolation by regulating the access to a shared resource.

Logging Process of persisting change information to non-volatile storage.

Main Memory Physical memory that can be directly accessed by the central processing unit (CPU).

Main Store Read-optimized and compressed data tables of SanssouciDB that are completely stored in main memory and on which no direct inserts are allowed.

MapReduce A programming model and software framework for developing applications that allows for parallel processing of vast amounts of data on a large number of servers.

Materialized View Result set of a complex query which is persisted in the database and updated automatically.

Memory Hierarchy The hierarchy of data storage technologies characterized by increasing response time but decreasing cost.

Merge Process Process in SanssouciDB that periodically moves data from the write-optimized differential store into the main store.

Metadata Data specifying the structure of tuples in database tables (and other objects) and relationships among them in terms of physical storage.

Mixed Workload Database workload consisting both of transactional and analytical queries.

Multi-Core Processor A microprocessor that comprises more than one core (processor) in a single integrated circuit.

Multi-Tenancy The consolidation of several customers onto the operational system of the same server machine.

Multithreading Concurrently executing several threads on the same processor core.

Network Partitioning Fault Fault that separates a network into two or more subnetworks that cannot reach each other anymore.

Node Partial structure of a business object.

Normalization Designing the structure of the tables of a database in such a way that anomalies cannot occur and data integrity is maintained.

Object Data Guide A database operator and index structure introduced to allow queries on whole business objects.

Online Analytical Processing (OLAP) see Analytical Processing.

Online Transaction Processing (OLTP) see Transactional Processing.

Operational Data Store Database used to integrate data from multiple operational sources and to then update data marts and/or data warehouses.

Padding Approach to modify memory structures so that they exhibit better memory access behavior but requiring the trade-off of having additional memory consumption.

Passive Data Data of a business transaction that is closed/completed and will not be changed anymore. For SanssouciDB it may therefore be moved to non-volatile storage.

Prefetching A technique that asynchronously loads additional cache lines from main memory into the CPU cache to hide memory latency.

Query Request sent to a DBMS in order to retrieve data manipulate data, execute an operation, or change the database structure.

Query Plan The set and order of individual database operations derived by the query optimizer of the DBMS, to answer an SQL query.

Radio Frequency Identification (RFID) Wireless technology to support fast tracking and tracing of goods. The latter are equipped with tags containing a unique identifier that can be readout by reader devices.

Real Time In the context of this book defined as, within the timeliness constraints of the speed-of-thought concept.

Real-Time Analytics Analytics that have all information at its disposal the moment they are called for (within the timeliness constraints of the speed of thought concept).

Recoverability Quality of a DBMS to allow for recovery after a failure has occurred.

Recovery Process of re-attaining a correct database state and operation according to the database's specification after a failure has occurred.

Relational Database A database that organizes its data in relations (tables) as sets of tuples (rows) having the same attributes (columns) according to the relational model.

Response Time at the Speed of Thought Response time of a system that is perceived as instantaneous by a human user because of his/her own mental processes. It normally lies between 550 ms and 750 ms.

Return on Investment (ROI) Economic measure to evaluate the effciency of an investment.

Row Store Database storage engine that stores all tuples sequentially that is each memory block may contain several tuples.

Sales Analysis Process that provides an overview of historical sales numbers.

Sales Order Processing Process with the main purpose of capturing sales orders.

SanssouciDB The in-memory database described in this book.

Scalability Desired characteristic of a system to yield an efficient increase in service capacity by adding resources.

Scale-out Capable of handling increasing workloads by adding new machines and using these multiple machines to provide the given service.

Scale-up Capable of handling increasing workloads by adding new resources to a given machine to provide the given service.

Scan Database operation evaluating a simple predicate on a column.

Scheduling Process of ordering the execution of all queries (and query plan operators) of the current workload in order to maintain a given optimality criterion.

Sequential Reading Reading a given memory block by block.

Shared Database Instance Multi-tenancy implementation scheme in which each customer has its own tables and sharing takes place on the level of the database instances.

Shared Machine Multi-tenancy implementation scheme in which each customer has its own database process and these processes are executed on the same machine that is several customers share the same server.

Shared Table Multi-tenancy implementation scheme in which sharing takes place on the level of database tables that is data from different customers is stored in one and the same table.

Shared Disk All processors share one view to the non-volatile memory but computation is handled individually and privately by each computing instance.

Shared Memory All processors share direct access to a global main memory and a number of disks.

Shared Nothing Each processor has its own memory and disk(s) and acts independently of the other processors in the system.

Single Instruction Multiple Data (SIMD) A multiprocessor instruction that applies the same instructions to many data streams.

Smart Grid An electricity network that can intelligently integrate the behavior and actions of all users connected to it—generators consumers and those that do both in order to efficiently deliver sustainable, economic and secure electricity supplies.

Software-as-a-Service (SaaS) Provisioning of applications as cloud services over the Internet.

Solid-State Drive (SSD) Data storage device that uses microchips for nonvolatile high-speed storage of data and exposes itself via standard communication protocols.

Speed-Up Measure for scalability defined as the ratio between the time consumed by a sequential system and the time consumed by a parallel system to carry out the same task.

Star Schema Simplest form of a data warehouse schema with one fact table (containing the data of interest for example, sales numbers) and several accompanying dimension tables (containing the specific references to view the data of interest, for example, state, country, month) forming a star-like structure.

Stored Procedure Procedural programs that can be written in SQL or PL/SQL and that are stored and accessible within the DBMS.

Streaming SIMD Extensions (SSE) An Intel SIMD instruction set extension for the x86 processor architecture.

Structured Data Data that is described by a data model for example, business data in a relational database.

Structured Query Language (SQL) A standardized declarative language for defining querying, and manipulating data.

Supply Chain Management (SCM) Business processes and respective technology to manage the flow of inventory and goods along a company's supply chain.

Table A set of tuples having the same attributes.

Tenant (1) A set of tables or data belonging to one customer in a multitenant setup. (2) An organization with several users querying a set of tables belonging to this organization in a multi-tenant setup.

Thread Smallest schedulable unit of execution of an operating system.

Three-tier Architecture Architecture of a software system that is separated in a presentation a business logic, and a data layer (tier).

Time Travel Query Query returning only those tuples of a table that were valid at the specified point in time.

Total Cost of Ownership (TCO) Accounting technique that tries to estimate the overall life-time costs of acquiring and operating equipment for example, software or hardware assets.

Transaction A set of actions on a database executed as a single unit according to the ACID concept.

Transactional Processing Method to process every-day business operations as ACID transactions such that the database remains in a consistent state.

Translation Lookaside Buffer (TLB) A cache that is part of a CPU's memory management unit and is employed for faster virtual address translation.

Trigger A set of actions that are executed within a database when a certain event occurs for example a specific modification takes place.

Tuple A real-world entity's representation as a set of attributes stored as element in a relation. In other words a row in a table.

Unstructured Data Data without data model or that a computer program cannot easily use (in the sense of understanding its content). Examples are word processing documents or electronic mail.

Vertical Partitioning The splitting of the attribute set of a database table and distributing it across two (or more) tables.

Viability Design thinking term expressing the practicability of a system from an economic point of view.

View Virtual table in a relational database whose content is defined by a stored query.

Virtual Machine A program mimicking an entire computer by acting like a physical machine.

Virtualization Method to introduce a layer of abstraction in order to provide a common access to a set of diverse physical and thereby virtualized resources.

References

1. Abadi, D.: Query Execution in Column-Oriented Database Systems. Ph.D. thesis, MIT (2008)
2. Abadi, D., Boncz, P., Harizopoulos, S.: Column oriented database systems. PVLDB **2**, 1664–1665 (2009)
3. Advanced Micro Devices, I.: HyperTransport Technology I/O Link, a High-Bandwidth I/O Architecture. URL whitepaper. Advanced Micro Devices Inc., Sunnyvale, CA (2001)
4. Agrawal, R., Ailamaki, A., Bernstein, P.A., Brewer, E.A., Carey, M.J., Chaudhuri, S., Doan, A., Florescu, D., Franklin, M.J., Garcia-Molina, H., Gehrke, J., Gruenwald, L., Haas, L.M., Halevy, A.Y., Hellerstein, J.M., Ioannidis, Y.E., Korth, H.F., Kossmann, D., Madden, S., Magoulas, R., Ooi, B.C., O'Reilly, T., Ramakrishnan, R., Sarawagi, S., Stonebraker, M., Szalay, A.S.,Weikum, G.: The claremont report on database research. SIGMOD Rec. **37**, 9–19 (2008). doi: http://doi.acm.org/10.1145/1462571.1462573. URL http://doi.acm.org/10.1145/.1462573
5. Ailamaki, A., DeWitt, D., Hill, M.: Data page layouts for relational databases on deep memory hierarchies. VLDB J. **11**, 198–215 (2002)
6. Alsberg, P.: Space and time savings through large data base compression and dynamic restructuring. Proc. IEEE **102**, 1114–1122 (1975)
7. Amazon: Amazon Elastic Compute Cloud. http://aws.amazon.com/ec2/ (2010). Retrieved 14 Jan 2011
8. Amdahl, G.: Validity of the single processor approach to achieving large scale computing capabilities. AFIPS Conf. Proc. **30** 483–485 (1967)
9. Aulbach, S., Grust, T., Jacobs, D., Kemper, A., Rittinger, J.: Multi-tenant databases for software as a service: schema-mapping techniques. In: SIGMOD, pp. 1195–1206 (2008)
10. Banerjee, J., Kim,W., Kim, H., Korth, H.F.: Semantics and implementation of schema evolution in object-oriented databases. In: SIGMOD, pp. 311–322 (1987)
11. Barham, P., Dragovic, B., Fraser, K., Hand, S., Harris, T.L., Ho, A., Neugebauer, R., Pratt, I.,Warfield, A.: Xen and the art of virtualization. In: SOSP, pp. 164–177 (2003)
12. Becker, S.A.: Developing Quality Complex Database Systems: Practices, Techniques, and Technologies. Idea Group Inc., Hershey (2001)
13. Behling, S., Bell, R., Farrell, P., Holthoff, H., O'Connell, F., Weir, W.: The POWER4 Processor Introduction and Tuning Guide. URL IBM Redbooks (2001)
14. Binnig, C., Faerber, F., Hildenbrand, S.: Dictionary-based order-preserving string compression for main-memory column stores. In: SIGMOD, pp. 283–296 (2009)
15. Bog, A., Krueger, J., Schafner, J.: A composite benchmark for online transaction processing and operational reporting. In: AMIGE (2008)
16. Bog, A., Plattner, H., Zeier, A.: A mixed transaction processing and operational reporting benchmark. Inf. Syst. Frontiers J. **13** 301–304 (2010)

DOI: 10.1007/978-3-642-29575-1, © Springer-Verlag Berlin Heidelberg 2012

17. Borr, A.: Robustness to Crash in a Distributed Database: A Non Shared-Memory Multi-Processor Approach (1984)
18. Botezatu, B.: The Future of Processors, Painted in Multi-Core Colors. URL http://news.softpedia.com/news/The-Future-of-Processors-Painted-in-Multi-Core-Colors-78143.shtml (2008). Retrieved 14 Jan 2011
19. Brooks, F.: The Mythical Man-Month. Addison-Wesley, Boston (1975)
20. Ceri, S., Negri, M., Pelagatti, G.: Horizontal data partitioning in database design. In: SIGMOD, pp. 128–136 (1982)
21. Chamberlin, D., Boyce, R.: Sequel: A structured English query language. In: SIGFIDET (now SIGMOD), pp. 249–264 (1974)
22. Chamoni, P., Gluchowski, P.: Analytische Informationssysteme: Business Intelligence-Technologien und Anwendungen. Springer, Heidelberg (2006)
23. Chang, F., Dean, J., Ghemawat, S., Hsieh,W.C.,Wallach, D.A., Burrows, M., Chandra, T., Fikes, A., Gruber, R.E.: Bigtable:A distributed storage system for structured data. ACM Trans. Comput. Syst. **26**, 1–4 (2008)
24. Chang, S.: Database decomposition in a hierarchical computer system. In: SIGMOD, pp. 48–53 (1975)
25. Chen, Z., Gehrke, J., Korn, F.: Query optimization in compressed database systems. In: SIGMOD, pp. 271–282 (2001)
26. Cieslewicz, J., Ross, K.: Adaptive aggregation on chip multiprocessors. In: VLDB, pp. 339–350 (2007)
27. Codd, E.: A relational model of data for large shared data banks. Commun. ACM **13**, 377–387 (1970)
28. Cooper, B. F., Ramakrishnan, R., Srivastava, U., Silberstein, A., Bohannon, P., Jacobsen, H., et al.: PNUTS: Yahoo!'s hosted data serving platform. PVLDB **1**, 1277–1288 (2008)
29. Corporation, S.P.E.: SPEC CPU2006. URL Benchmark website: http://www.spec.org/cpu2006/ (2006). Retrieved 14 Jan 2011
30. Corporation, S.P.E.: SPECviewperf 11. URL Benchmark website: http://www.spec.org/gwpg/gpc.static/vp11info.html (2006). Retrieved 14 Jan 2011
31. Corporation, S.P.E.: SPECmail2009. URL Benchmark website: http://www.spec.org/mail2009/ (2009). Retrieved 14 Jan 2011
32. Corporation, S.P.E.: SPECweb2009. URL Benchmark website: http://www.spec.org/web2009/ (2009). Retrieved 14 Jan 2011
33. Corporation, S.P.E.: SPECjEnterprise2010. URL http://www.spec.org/jEnterprise2010/ (2010). Retrieved 14 Jan 2011
34. Corporation, S.P.E.: SPECvirt. URL Benchmark website: http://www.spec.org/virt_sc2010/ (2010). Retrieved 14 Jan 2011
35. Corporation, W.D.: Specification for the Seraial AT 6 Gb/s VelociRaptor Enterprise Hard Drives. URL Product website: http://wdc.custhelp.com/cgibin/wdc.cfg/php/enduser/std_adp.php?p_faqid=5377&p_created= (2010). Retrieved 14 Jan 2011
36. Council, T.P.P.: URL http://www.tpc.org (2010). Retrieved 14 Jan 2011
37. Cross, R.L.: ITJ Foreword Q1, 2002, Intel Hyper-Threading Technology (2002)
38. Date, C.: An Introduction to Database Systems, 6th edn. Addison-Wesley, Boston (1995)
39. David, J., Schuff, D., St Louis, R.: Managing your IT Total Cost of Ownership. Commun. ACM **45**, 101–106 (2002)
40. Dean, J.: Designs, Lessons and Advice from Building Large Distributed Systems. URL http://www.slideshare.net/xlight/google-designs-lessons-and-advicefrom-building-large-distributed-systems (2009). Retrieved 14 Jan 2011
41. Dean, J.: Large-scale distributed systems at Google: current systems and future directions. In: LADIS (2009)
42. Dean, J., Ghemawat, S.: MapReduce: simplified data processing on large clusters. In: OSDI, pp. 137–150 (2004)

43. DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., Sivasubramanian, S., Vosshall, P., Vogels, W.: Dynamo: Amazon's highly available key-value store. In: SIGOPS, pp. 205–220 (2007)
44. Devices, A.M.: AMD. URL http://www.amd.com. Retrieved 14 Jan 2011
45. DeWitt, D., Gray, J.: Parallel database systems: the future of high performance database systems. Commun. ACM **35** 85–98 (1992)
46. Dikaiakos, M.D., Katsaros, D., Mehra, P., Pallis, G., Vakali, A.: Cloud computing: distributed Internet computing for IT and scientific research. IEEE Internet Comput. **13**, 10–13 (2009)
47. Eager, D., Zahorjan, J., Lozowska, E.: Speedup versus efficiency in parallel systems. IEEE Trans. Comput. **38**, 408–423 (1989)
48. Elmasri, R., Navathe, S.: Fundamentals of Database Systems, 5th edn. Addison-Wesley Longman Publishing, Boston (2006)
49. Exasol: EXASolution Highspeed Database. URL Product website, http://www.exasol.com/ (2010). Retrieved 14 Jan 2011
51. Fine, C.: Clockspeed :Winning Industry Control in the Age of Temporary Advantage. Basic Books, New York (1998)
51. Florescu, D., Kossmann, D.: Rethinking cost and performance of database systems. SIGMOD Record **38**, 43–48 (2009)
52. Flynn, M.: Very high-speed computing systems. Proc. IEEE **54**(12), 1901–1909 (1966)
53. French, C.: Teaching an OLTP database Kernel advanced data warehousing techniques. In: ICDE, pp. 194–198 (1997)
54. Fujitsu: Speicher-Performance Xeon 5500 (Nehalem EP) basierter PRIMERGY Server (2009). URL Whitepaper, version 1.0
55. Fujitsu: Speicher-Performance Xeon 7500 (Nehalem EX) basierter Systeme (2010). URL Whitepaper, Version 1.0
56. Garcia-Molina, H., Salem, K.: Main memory database systems: an overview. TKDE **4**, 509–516 (1992)
57. Garcia-Molina, H., Ullman, J., Widom, J.: Database Systems: The Complete Book. Prentice Hall Press, Upper Saddle River (2008)
58. Gates, B.: Information At Your Fingertips. URL Keynote Address, Fall/-COMDEX, Las Vegas, Nevada (1994)
59. Ghemawat, S., Gobioff, H., Leung, S.: The Google file system. In: SOSP, pp. 29–43 (2003)
60. Graefe, G.: Query evaluation techniques for large databases. ACM Comput. Surv. **25**, 73–170 (1993)
61. Gray, J.: A Transaction Model. In: Colloquium on Automata, Languages and Programming, pp. 282–298 (1980)
62. Gray, J.: The Benchmark Handbook for Database and Transaction Processing Systems. Morgan Kaufmann Publishers, San Mateo (1993)
63. Gray, J.: Tape is Dead, Disk is Tape, Flash is Disk, RAM Locality is King. URL http://www.signallake.com/innovation/Flash_is_Good.pdf (2006). Retrieved 14 Jan 2011
64. Gray, J.: Distributed Computing Economics. ACM Queue **6**, 63–68 (2008)
65. Gray, J., Reuter, A.: Transaction Processing: Concepts and Techniques. Morgan Kaufmann, San Fransisco (1993)
66. Grossman, R. L. (2009). The case for cloud computing. IT Prof. **11**, 23–27
67. Grund, M., Krueger, J., Plattner, H., Zeier, A., Madden, S., Cudre-Mauroux, P.: HYRISE—A hybrid main memory storage engine. In: VLDB (2011)
68. Grund, M., Krueger, J., Tinnefeld, C., Zeier, A.: Vertical partition for insert-only scenarios in enterprise applications. In: IEEM (2009)
69. Grund, M., Schaffner, J., Krueger, J., Brunnert, J., Zeier, A.: The effects of virtualization on main memory systems. In: DaMoN, pp. 41–46 (2010)
70. Gschwind, M.: The Cell Broadband Engine: Exploiting Multiple Levels of Parallelism in a Chip Multiprocessor (2006)

71. Habich, D., Boehm, M., Thiele, M., Schlegel, B., Fischer, U., Voigt, H., Lehner, W.: Next generation database programming and execution environment. In: VLDB (2011)
72. Haerder, T., Rahm, E.: Datenbanksysteme: Konzepte und Techniken der Implementierung, 2 Auflage. Springer, Berlin (2001)
73. Haerder, T., Reuter, A.: Principles of transaction-oriented database recovery. ACM Comput. Surv. **15**, 287–317 (1983)
74. Hagel, J., Brown, J. S., Davison, L.: The Power of Pull: How Small Moves, Smartly Made, Can Set Big Things in Motion.Basic Books, New York (2010)
75. Hamilton, J.: Cost of Power in Large-Scale Data Centers. URL http://perspectives.mvdirona.com/2008/11/28/CostOfPowerInLargeScaleDataCenters.aspx (2008). Retrieved 14 Jan 2011
76. Hankins, R.A., Patel, J.M.: Data morphing: an adaptive, cache-conscious storage technique. In: VLDB, pp. 417–428 (2003)
77. Hare, C.: PC Hardware Links. URL Web page, http://mysite.verizon.net/pchardwarelinks/main.htm (2010). Retrieved 26 July 2010
78. Hellerstein, J.: Datalog Redux: Experience and Conjecture. In: PODS, pp. 1–2 (2010)
79. Hellerstein, J.M., Stonebraker, M., Hamilton, J.: Architecture of a Database System. Now Publishers Inc., Hanover (2007)
80. Hennessy, J.L., Patterson, D.A.: Computer Architecture. Morgan Kaufmann, Boston (2006)
81. Hitachi: Hitachi Global Storage Technologies, Ultrastar 15K450. URL Product website, http://www.hitachigst.com/internal-drives/enterprise/ultrastar/ultrastar-15k450 (2010). Retrieved 14 Jan 2011
82. Hohpe, G., Woolf, B.: Enterprise Integration Patterns: Designing, Building, and Deploying Messaging Solutions. Addison-Wesley Longman Publishing Co. Inc., Boston (2003)
83. Inmon, B.: Data mart does not equal dataWarehouse. DMReview May 1998. URL DM Rev. Mag. (1998)
84. Inmon, B.: ODS Types. Information management magazine January 2000. URL Inf. Manag. Mag. (2000)
85. Inmon, B.: Operational and informational reporting. Information management magazine July 2000. URL Inf. Manag. Mag. (2000)
86. Intel: An Introduction to the Intel QuickPath Interconnect. URL http://www.intel.com/technology/quickpath/introduction.pdf (2009). Retrieved 14 Jan 2011
87. Intel: Intel R X25-E SATA Solid State Drive. URL http://download.intel.com/design/flash/nand/extreme/319984.pdf (2009). Retrieved 14 Jan 2011
88. Intel: Microprocessor Quick Reference Guide. URL http://www.intel.com/pressroom/kits/quickrefyr.htm (2010). Retrieved 14 Jan 2011
89. Intel: Moore's Law: Made Real by Intel Innovation. URL Web site: http://www.intel.com/technology/mooreslaw/ (2010). Retrieved 14 Jan 2011
90. Jacobs, D.: Enterprise software as service. ACM Queue **3**, 36–42 (2005)
91. Jacobs, D., Aulbach, S.: Ruminations on multi-tenant databases. In: BTW, pp. 514–521 (2007)
92. Kimball, R.: Surrogate Keys—Keep Control Over Record Identifiers by Generating New Keys for the Data Warehouse. URL http://www.rkimball.com/html/articles_search/articles1998/9805d05.html (1998). Retrieved 14 Jan 2011
93. Kimball, R., Caserta, J.: The DataWarehouse ETL Toolkit. Wiley, New York (2004)
94. Knolmayer, G., Mertens, P., Zeier, A., Dickersbach, J.: Supply Chain Management Based on SAP Systems: Architecture and Planning Processes (2009)
95. Knuth, D.: The Art of Computer Programming, vol. 3, 2nd edn., Sorting and Searching. Addison Wesley Longman Publishing Co. Inc., Boston (1998)
96. Kossow, R.: TCO-Wirtschaftlichkeir von IT-Systemen. URL http://www.erpmanager.de/magazin/artikel_ 1339_tco_total_cost_ownership_wirtschaftlichkeit.html (2007). Retrieved 14 Jan 2011
97. Krueger, J., Grund, M., Tinnefeld, C., Zeier, A., Plattner, H.: Optimizing write performance for read optimized databases. In: DASFAA (2010)

98. Krueger, J., Grund, M.,Wust, J., Zeier, A., Plattner, H.: Merging differential updates in in-memory column store. In: DBKDA (2011)
99. Krueger, J., Grund, M., Zeier, A., Plattner, H.: Enterprise application-specific data management. In: EDOC, pp. 131–140 (2010)
100. Krueger, J., Tinnefeld, C., Grund, M., Zeier, A., Plattner, H.: A case for online mixed workload processing. In: DBTest (2010)
101. Laming, D.: Information Theory of Choice-Reaction Times. Academic Press, New York (1968)
102. Lang, W., Patel, J., Shankar, S.: Wimpy node clusters: what about non-Wimpy workloads? In: DaMoN, pp. 47–55 (2010)
103. Lee, J., Kim, K., Cha, S.: Differential logging: a commutative and associative logging scheme for highly parallel main memory databases. In: ICDE, pp. 173–182 (2001)
104. Lee, S., Moon, B.: Design of flash-based DBMS: an in-page logging approach. In: SIGMOD, pp. 55–66 (2007)
105. Lemke, C., Sattler, K.U., Faerber, F.: Compression techniques for column-oriented BI accelerator solutions. In: BTW, pp. 468–497 (2009)
106. Lemke, C., Sattler, K.U., Faerber, F., Zeier, A.: Speeding up queries in column stores—a case for compression data. In: DaWaK, pp. 117–129 (2010)
107. Lerner, B.S., Habermann, A.N.: Beyond schema evolution to database reorganization. In: OOPSLA/ECOOP, pp. 67–76 (1990)
108. Ma, H.: Distribution Design for Complex Value Databases (2007)
109. Magura, S.: Warten war frueher. SAP Spectrum **2**, 8–9 (2010)
110. Manegold, S., Boncz, P.A., Kersten, M.L.: Generic database cost models for hierarchical memory systems. In: VLDB, pp. 191–202 (2002)
111. Mehldorn, K., Sanders, P.: Algorithms and Data Structures. Springer, Berlin (2008)
112. Mertens, P.: Integrierte Informationsverarbeitung. Gabler, Wiesbaden (2009)
113. Mertens, P., Zeier, A.: ATP—available-to-promise. Wirtschaftsinformatik **41**, 378–379 (1999)
114. Microsoft: Microsoft PowerPivot. URL Product website: http://www.powerpivot.com (2010). Retrieved 14 Jan 2011
115. Microsoft: Scalability. URL http://msdn.microsoft.com/enus/library/aa29217228v=VS.7129.aspx (2010). Retrieved 14 Jan 2011
116. Molka, D., Hackenberg, D., Schone, R., Muller, M.: Memory performance and cache coherency effects on an Intel Nehalem multiprocessor system. In: PACT, pp. 261–270 (2009)
117. Moore, G.: Cramming more components onto integrated circuits. Electron. Mag. **38**, 114–117 (1965)
118. Moura, E., Ziviani, N., Baeza-Yates, R., Navarro, G.: Fast and flexible word searching on compressed text. ACM TOIS **18**, 113–139 (2000)
119. Moutsos, K.: IMS at 40: Stronger than ever. IBM Database Mag. **4** (2008)
120. MP, S.: High Performance Computing Virtualization, Virtual SMP. URL Product website, http://www.scalemp.com/ (2010). Retrieved 14 Jan 2011
121. Naffziger, S., Warnock, J., Knapp, H.: SE2 when processors hit the power wall (or When the CPU Hits the Fan). In: ISSCC (2005)
122. Navathe, S.B., Ceri, S., Wiederhold, G., Dou, J.: Vertical partitioning algorithms for database design. ACM Trans. Database Syst. **9**, 680–710 (1984)
123. North, K.: Terabytes to Petabytes: Reflections on 1999–2009. URL http://www.drdobbs.com/blog/archives/2010/01/terabytes_to_pe.html (2010). Retrieved 14 Jan 2011
124. O'Neil, P., Winter, R., French, C., Crowley, D., McKenna, W.: Data warehousing lessons from experience. In: ICDE, p. 294 (1998)
125. O'Neil, P.E., O'Neil, E.J., Chen, X.: The Star Schema Benchmark (SSB). URL http://www.cs.umb.edu/poneil/StarSchemaB.pdf (2007). Retrieved 14 Jan 2011
126. O'Neil, P.E., O'Neil, E.J., Chen, X., Revilak, S.: The star schema benchmark and augmented fact table indexing. In: TPCTC, pp. 237–252 (2009)
127. Oracle: Oracle E-Business Suite Standard Benchmark. URL http://www.oracle.com/apps-benchmark/results-166922.html (2010). Retrieved 14 Jan 2011

128. Orlowski, A.: POWER4 Debuts in IMB Regatta: Big Blue's Big Bang Eschews SMP Numbers Game. URL The Register, http://www.theregister.co.uk/2001/10/04/power4_debuts_in_ibm_regatta/ (2001), Retrieved 14 Jan 2011
129. Ousterhout, J. K., Agrawal, P., Erickson, D., Kozyrakis, C., Leverich, J., Mazières, D., et al.: The case for RAMClouds: scalable high-performance storage entirely in DRAM. Oper. Syst. Rev. **43**, 92–105 (2009)
130. Pacioli, F.L.: Treatise on Double-Entry Bookkeeping. Translated by Pietro Crivell (1494). Institute of Book-Keepers, London (1924)
131. Panchenko, O., Karstens, J., Plattner, H., Zeier, A.: Precise and Scalable Querying of Syntactical Source Code Patterns Using Sample Code Snippets and a Database. In: Proceedings of the International Conference on Program Comprehension, pp. 41–50. IEEE Computer Society, Los Alamitos, CA, USA (2011)
132. Panchenko, O., Plattner, H., Zeier, A.: Mapping terms in application and implementation domains. In: Proceedings of the Workshop on Software Reengineering (2011)
133. Paraccel: Paraccel—Column-Oriented DBMS for Decision Support and Complex Processing. URL Product website: www.paraccel.com/ (2010). Retrieved 14 Jan 2011
134. Pedersen, T.B., Jensen, C.S.: Multidimensional database technology. Computer **34**, 40–46 (2001)
135. Plattner, H.: A common database approach for OLTP and OLAP using an in-memory column database. In: SIGMOD, pp. 1–2 (2009)
136. Plattner, H.: SanssouciDB: an in-memory database for mixed-workload processing. In: BTW (2011)
137. Poess, M., Nambiar, R.O.: Tuning servers, storage and database for energy efficient data warehouses. In: ICDE, pp. 1006–1017 (2010)
138. Poess, M., Smith, B., Kollar, L., Larson, P.: TPC-DS, taking decision support benchmarking to the next level. In: SIGMOD, pp. 582–587 (2002)
139. Ponniah, P.: Data Warehousing Fundamentals: A Comprehensive Guide for IT Professionals, vol. 1. Wiley, New York (2001)
140. Power, D.: A Brief History of Decision Support Systems. URL http://dssresources.com/history/dsshistory.html (2010). Retrieved 14 Jan 2011
141. QlikView: Business Intelligence (BI) Software Solutions—Business Intelligence Reporting Software. URL Product website: http://www.qlikview.com/ (2010). Retrieved 14 Jan 2011
142. Raden, N.: Exploring the Business Imperative of Real-Time Analytics. URL Teradata Whitepaper (2003)
143. Raden, N.: Business Intelligence 2.0: Simpler, More Accessible, Inevitable. URL Intelligent enterprise, http://intelligent-enterprise.com (2007). Retrived 14 Jan 2011
144. Rahm, E.: Mehrrechner-Datenbanksysteme—Grundlagen der Verteilten und Parallelen Datenbankverarbeitung. Addison-Wesley, Boston (1994)
145. Rao, J., Ross, K.: Making B+- trees cache conscious in main memory. SIGMOD Rec. **29**, 475–486 (2000)
146. Roddick, J.: A Survey of Schema Versioning Issues for Database Systems. Information and Software Technology **37**, 383–393 (1995). URL citeseer. ist.psu.edu/roddick95survey.html
147. Rossberg, J., Redler, R.: Pro Scalable .NET 2.0 Application Designs. Designing .NET 2.0 Enterprise Applications from Conception to Deployment. Apress, New York (2005)
148. Roth, M., Van Horn, S.: Database compression. SIGMOD Rec. **22**, 31–39 (1993)
149. Salomon, D.: Data Compression. Springer, New York (2006)
150. Sanders, P., Transier, F.: Intersection in integer inverted indices. In: ALENEX (2007)
151. SAP: SAP Business By Design. URL http://www.sap.com/germany/sme/solutions/businessmanagement/businessbydesign/index.epx. Retrieved 14 Jan 2011
152. SAP: SAP BusinessObjects Explorer, Explore Your Business at the Speed of Thought. URL Product website: http://www.sap.com/solutions/sapbusinessobjects/large/businessintelligence/search-navigation/explorer/index.epx. Retrieved 14 Jan 2011
153. SAP: SAP Netweaver BusinessWarehouse Accelerator. URL Product Website, http://www.sdn.sap.com/irj/sdn/bwa (2008). Retrieved 14 Jan 2011

154. SAP: SAP Standard Application Benchmark. URL http://www.sap.com/solutions/benchmark/index.epx (2010). Retrieved 14Jan 2011
155. Schaffner, J., Bog, A., Krueger, J., Zeier, A.: A hybrid row-column OLTP database architecture for operational reporting. In: BIRTE (Informal Proceedings) (2008)
156. Schaffner, J., Eckart, B., Jacobs, D., Schwarz, C., Plattner, H., Zeier, A.: Predicting in-memory database performance for automating cluster management tasks. In: ICDE (2011)
157. Schaffner, J., Eckart, B., Schwarz, C., Brunnert, J., Jacobs, D., Zeier, A., Plattner, H.: Simulating multi-tenant olap database clusters. In: Datenbanksysteme in Business, Technologie und Web (BTW 2011), 14. Fachtagung des GI-Fachbereichs Datenbanken und Informationssysteme (DBIS), Proceedings, Kaiserslautern, Germany (2011)
158. Schapranow, M.P., Kuehne, R., Zeier, A.: Enabling real-time charging for smart grid infrastructures using in-memory databases. In: 1st IEEE LCN Workshop on Smart Grid Networking Infrastructure (2010)
159. Schapranow, M.P., Zeier, A., Plattner, H.: A dynamic mutual RFID authentication model preventing unauthorized third party access. In: The 4th International Conference on Network and System Security (2010)
160. Schapranow, M.P., Zeier, A., Plattner, H.: A formal model for enabling RFID in pharmaceutical supply chains. In: 44th Hawaii International Conference on System Sciences (2011)
161. Scheckenbach, R., Zeier, A.: Collaborative SCM in Branchen: B2B Integrationsstrategien und Realisation. Galileo, Bonn (2002)
162. Schwarz, C., Borovskiy,V., Zeier, A.: Optimizing operation scheduling for in-memory databases. In: The 2011 International Conference on Modeling, Simulation and Visualization Methods (2011)
163. Stonebraker, M.: The case for shared-nothing. IEEE Database Eng. Bull. **9**, 4–9 (1986)
164. Stonebraker, M., Abadi, D.J., Batkin, A., Chen, X., Cherniack, M., Ferreira, M., Lau, E., Lin, A., Madden, S., O'Neil, E.J., O'Neil, P.E., Rasin, A., Tran, N., Zdonik, S.B.: C-Store: a column-oriented DBMS. In: VLDB, pp. 553–564 (2005)
165. Stonebraker, M., Madden, S., Abadi, D., Harizopoulos, S., Hachem, N., Helland, P.: The End of an Architectural Era (It's Time for a Complete Rewrite). In: VLDB, pp. 1150–1160 (2007)
166. Stonebreaker, M., Rowe, L., Hirohama, M.: The implementation of postgres. IEEE Trans. Knowl. Data Eng. **2**, 125–142 (1990)
167. Sundell, H., Tsigas, P.: Fast and lock-free concurrent priority queues for multi-thread systems. In: IPDPS, p. 84 (2003)
168. Sutter, H.: The free lunch is over: a fundamental turn toward concurrency in software. Dr. Dobb's J. **30** (2005)
169. Taniar, D., Leung, C., Rahayu, J., Goel, S.: High Performance Parallel Database Processing and Grid Databases. Wiley, New York (2008)
170. Thomsen, E.: OLAP Solutions: Building Multidimensional Information. Wiley, New York (2002)
171. Thomson, A., Abadi, D.J.: The case for determinism in database systems. PVLDB **3**, 70–80 (2010)
172. Thusoo, A., Shao, Z., Anthony, S., Borthakur, D., Jain, N., Sarma, J.S., Murthy, R., Liu, H.: Data warehousing and analytics infrastructure at facebook. In: SIGMOD Conference, pp. 1013–1020 (2010)
173. Tilera: TILEPro36 Processor. URL http://www.tilera.com/products/processors/TILEPRO36 (2008). Retrieved 14 Jan 2011
174. TPC: TPC Benchmark W (Web Commerce). Standard Specification, Version 1.8, February 19, 2002. (2002)
175. TPC: TPC Benchmark C, Standard Specification Revision 5.9. (2007)
176. TPC: TPC Benchmark E, Standard Specification Version 1.5.1 (2008)
177. TPC: TPC Benchmark H (Decision Support), Standard Specification Revision 2.7.0 (2008)
178. TPC: TPC-Energy Specification. Standard Specification Version 1.1.1 (2010)

179. Transier, F.: Algorithms and data structures for in-memory text search engines. Ph.D. thesis, University of Karlsruhe (2010)
180. Transier, F., Sanders, P.: Compressed inverted indexes for in-memory search engines. In: ALENEX (2008)
181. Vassiliadis, P., Simitsis, A., Skiadopoulos, S.: Conceptual modeling for ETL processes. In: DOLAP, pp. 14–21 (2002)
182. Vogels,W.: A head in the cloud—the power of infrastructure as a service. In: CCA (2008)
183. Walsh, K.R.: Analyzing the application ASP concept: technologies, economies, and strategies. Commun. ACM **46**, 103–107 (2003)
184. Westmann, T., Kossmann, D., Helmer, S., Moerkotte, G.: The implementation and performance of compressed databases. SIGMOD Rec. **29**, 55–67 (2000)
185. Willhalm, T., Popovici, N., Boshmaf, Y., Plattner, H., Zeier, A., Schaffner, J.: Ultra fast in-memory table scan using on-chip vector processing units. In: VLDB, pp. 385–394 (2009)
186. Winter, R.: Why Are Data Warehouses Growing So Fast? An Update on the Drivers of Data Warehouse Growth. URL http://www.b-eye-network.com/view/7188 (2008). Retrieved 14 Jan 2011
187. Wust, J., Krüger, J., Blessing, S.,Tosun, C., Zeier, A., Plattner, H.: xsellerate: Supporting sales representatives with real-time information in customer dialogs. In:IMDM2011, Proceedings zur Tagung Innovative Unternehmensanwendungen mit In-Memory Data Management, Mainz (2011)
188. Zeier, A.: Ausdifferenzierung von Supply-Chain-Management-Standardsoftware in Richtung auf Betriebstypen und Branchen unter besonderer Berueckichtigung des SAP APO. Roell (2002)
189. Zicari, R.V.: Hadoop for Business: Interview with Mike Olson, Chief Executive Officer at Cloudera. URL http://www.odbms.org/blog/2011/04/hadoop-for-businessinterview-with-mike-olson-chief-executive-o_cer-at-cloudera/. Retrieved 2 Nov 2011
190. Zukowski, M., Heman, S., Nes, N., Boncz, P.: Super-scalar RAM-CPU cache compression. In: ICDE, p. 59 (2006)

Index

H. Plattner and A. Zeier, *In-Memory Data Management*,
DOI: 10.1007/978-3-642-29575-1, © Springer-Verlag Berlin Heidelberg 2012