

TEMPLATE FOR DEFINITIONAL ENTRY

Main Memory Databases

Peter Boncz

Centrum voor Wiskunde en Informatica (CWI)
Kruislaad 413, 1098 SJ Amsterdam, The Netherlands
boncz@cwi.nl

SYNONYMS

In-memory DBMS, MMDBMS

DEFINITION

A main memory database system is a DBMS that primarily relies on main memory for computer data storage. In contrast, normal database management systems employ hard disk based persistent storage.

MAIN TEXT

The main advantage of MMDBMS over normal DBMS technology is superior performance, as I/O cost is removed as a performance cost factor. With I/O as main optimization focus eliminated, the architecture of main memory database systems typically aims at optimizing CPU cost and CPU cache usage, leading to different data layout strategies (avoiding complex tuple representations) as well as indexing structures (e.g. B-trees with lower-fan-outs with nodes of one or a few CPU cache lines)..

While built on top of volatile storage, most MMDB products do offer ACID properties, via either of the following mechanisms: (i) Transaction Logging, which records changes to the database in a journal file and facilitates automatic recovery of an in-memory database, (ii) Non-volatile RAM, usually in the form of static RAM backed up with battery power (battery RAM), or an electrically erasable programmable ROM (EEPROM). With this storage, the MMDB system can recover the data store from its last consistent state upon reboot, (iii) High Availability implementations that rely on database replication, with automatic failover to an identical standby database in the event of primary database failure.

Main memory database systems were originally popular in real-time systems (used in e.g. telecommunications) for their fast and more predictable performance, and this continues to be the case. However, with increasing RAM sizes allowing more problems to be addressed using a MMDBMS, this technology is proliferating into many other areas, such as on-line transaction systems, and recently also in decision support. Main memory database systems are also deployed as drop-in systems that intercept read-only queries on cached data from an existing disk-based DBMS, thus reducing its workload and providing fast answers to a large percentage of the workload.

Examples of Main Memory database systems are MonetDB, SolidDB, TimesTen and DataBlitz. MySQL offers a main-memory backend based on Heap tables. The MySQL Cluster product is a parallel main memory system that offers ACID properties through high availability.

CROSS REFERENCES

Main Memory

Processor Cache
Disk

RECOMMENDED READING

[1] David J. DeWitt, Randy H. Katz, Frank Olken, Leonard D. Shapiro, Michael Stonebraker, David A. Wood: Implementation Techniques for Main Memory Database Systems. SIGMOD Conference 1984: 1-8

[2] Svein-Olaf Hvasshovd, Øystein Torbjørnsen, Svein Erik Bratsberg, Per Holager: The ClustRa Telecom Database: High Availability, High Throughput, and Real-Time Response. VLDB 1995:469-477

[3] Philip Bohannon, Daniel F. Liewen, Rajeev Rastogi, Abraham Silberschatz, S. Seshadri, S. Sudarshan: The Architecture of the Dalí Main-Memory Storage Manager. Multimedia Tools Appl. (MTA) 4(2):115-151 (1997)

[4] Peter A. Boncz, Martin L. Kersten: MIL Primitives for Querying a Fragmented World. VLDB J. (VLDB) 8(2):101-119 (1999)