Some Considerations about Modern Database Machines

Manole VELICANU, Daniela LIȚAN, Aura-Mihaela MOCANU (VÎRGOLICI)
Academy of Economic Studies Bucharest, Romania
manole.velicanu@ie.ase.ro, litan.daniela@gmail.com, mocanuaura@gmail.com

Optimizing the two computing resources of any computing system - time and space - has always been one of the priority objectives of any database. A current and effective solution in this respect is the computer database. Optimizing computer applications by means of database machines has been a steady preoccupation of researchers since the late seventies. Several information technologies have revolutionized the present information framework. Out of these, those which have brought a major contribution to the optimization of the databases are: efficient handling of large volumes of data (Data Warehouse, Data Mining, OLAP – On Line Analytical Processing), the improvement of DBMS – Database Management Systems facilities through the integration of the new technologies, the dramatic increase in computing power and the efficient use of it (computer networks, massive parallel computing, Grid Computing and so on). All these information technologies, and others, have favored the resumption of the research on database machines and the obtaining in the last few years of some very good practical results, as far as the optimization of the computing resources is concerned.
Keywords: *Database Optimization, Database Machines, Data Warehouse, OLAP – On Line Analytical Processing, OLTP – On Line Transaction Processing, Parallel Processing*

1 Current state of knowledge

The current DBMS are undergoing a keen competition in order to gain the market quote and that is why, for each new version, one of the key words is optimization. As the companies have to cope with a keener and keener competition nowadays, they must be more and more flexible in order to be able to adjust very fast. In this respect, the companies must be able to take into account all the events which are taking place and to foresee the future ones [1]. Therefore, such a thorough analysis implies making the best use of a company’s database. This involves the use of a great and very complex amount of data and the results need to be obtained in the shortest time possible. This thing is possible now due to the optimization mechanisms offered by DBMS.

The rapid increase in the data volume has created the problem of overloading the databases. The technology for processing and analyzing huge sets of data has legged behind the technology of collecting and stocking databases [2]. More often than not, especially in big companies, the servers can hardly meet the demands received, in spite of the mechanisms for optimizing the data allocation and retrieval. In order to cope efficiently with such situations, the solution means *database machines*. These are some computers specialized in stocking and processing a very great amount of data. They have been conceived in a special way, hardware and software, to access the data rapidly and they get connected to a main processor through a high speed channel. We can say that such computers „know” how to work only with databases.

The aim of building up and using some database machines is that of making the best use of the computing resources in the information applications with databases, which means the maximum optimization,

It is well-known that any specialized computing resource (hardware or software) offers maximum performance because it eliminates the intermediary levels. Any further level (hardware or software components) which appears between the user and the data from the computer means further time for allotting and retrieving the data from the computer, as well as an increased probability of some errors.

The current economic framework makes the managers use a greater and greater volume of data which should be accessible in the shortest time possible for the decision-taking process that is supporting systems for efficient decision-taking are necessary. The database machines meet best these requirements as: there is specialized software for handling a great volume of data (Database Management Systems – DBMS) and there is now hardware system specialized in data stocking and processing. The software component has continually developed through the integration of new information technologies, while the hardware component has stopped developing due to the lack of an adequate technology. During the past years, such technologies have appeared and have been integrated into the database machine concept: massive paralleled processing, dealing with great database volumes (data warehouse, data mining, OLAP etc.), Grid Computing etc.

The concept of database machine, like hardware and software specialized in database dates back to the seventies, period in which the first research on this subject was carried out. Ever since, the machines conceived especially for the databases have been used fluctuating, having periods when their potential was pointed out but also periods when this concept was abandoned. The favorable periods for the database machines were those when new facilities, both hardware and software were introduced on the market: parallel processing, increasing the database stocking capacity, optimizing the computing resources. Having studied the specialized literature, *we put forward* the following table of database machines evolution (Table 1).

Table 1. The evolution of the database machines

Period	Features of DB machines
The 70s	Based on parallel processing
The 80s	More and more powerful microprocessors Small capacity stocking devices
The 90s	Great capacity stocking devices
The 2000s	Cost of I-O operations Growth of the bandwidth for catching up with the processor-memory delay
2010	Modern database machines

The first two database machine manufacturers were Britton-Lee and Teradata, the first company being subsequently taken over by Teradata. He has continued his research up to now.

According to their functionality, the database machines which have been built so far can be divided into the following *categories:* multiprocessor database machines, intelligent devices for secondary storage (specialized hardware), database filters (specialized hardware) and associative memory systems.

As far as *the multiprocessor machines* are concerned, the most widely used at this moment, they can be classified as [3]:

- shared-nothing in which every processor has its own memory and disk units (examples: DBC/1012 Teradata, Gamma, Bubba, and Idioms);
- shared disk uses multiple processors, each with its memory unit but with a shared disk system;
- shared everything uses multiple processors that have shared memory and disk drive.

Shared-nothing type systems are the most performing for the decisional support on a large scale, while the shared everything system types are the best for processing *online* transactions.

It can be inferred from the specialized literature that the use of massive parallelism or associative memory are not sufficient. A revolution in terms of massive data storage devices was necessary, but also suitable devices for the reducing the loading time of the data storage device from the database [4]. Until the hardware components had evolved in this respect, the development of the database engine stopped. This evolution took place in the '90s and 2000 when new information technologies appeared concerning the data storage and processing. In this way the modern database machines have been conceived.

2 The optimization of the computing resources by using modern database machines

As DB machines are specialized computers for the development and the operation of database applications, they have a number of specific mechanisms to optimize the computing resources - time for data access and space allocation data.

The first step towards modern database machines is represented by Voltaire company's InfiniBand technology (Oracle's partner staring with 2001) which provides network technologies for Oracle RAC starting with the version Oracle 9i. The connections through network cards with protocol for InfiniBand are widely used by the supercomputers. For example, at the end of 2007 approximately 24% of the supercomputers in Top 500 were connected through InfiniBand. InfiniBand connections have a data transfer rate of 20GB per second, which means the optimization of the data access time.

Oracle Exadata V1 as the first modern engine database was created by Oracle and HP - Hewlett-Packard companies at the end of 2008. The machine consists of a family of software and hardware products designed for an efficient storage of data that improves over 10 times the query performance (access time) of a data warehouse [5].

Oracle Exadata uses a parallel architecture to increase the bandwidth necessary for the transmission of data between the database server and storage space. The storage software component is intelligent: it downloads the query intensive processes of Oracle Database 11g servers and performs the query processing closer to data. The result is a parallel processing of data which is more rapid and with less data traffic through connections with bigger bandwidth [5].

Some of the advantages of Oracle Exadata Storage product for optimizing computing resources are: the answer to queries from the data warehouses is very fast, multiple queries can be run simultaneously, the scalability is unlimited (any Oracle Exadata Storage Servers can be connected).

Oracle and HP have launched the HP Oracle Database Machine, which includes its main component the HP Oracle Exadata Storage Server. HP Oracle Database Machine contains 8 database servers with Oracle Database 11g and Oracle Real Application Clusters - RAC with the Oracle Enterprise Linux operating system. The machine includes a network storage server of 14 HP Oracle Exadata Storage Servers, each with 112 processors. The network storage has a capacity of 168 TB and a bandwidth of 14 GB per second between the storage servers and the database servers. HP Oracle Exadata Storage Server is a combination of Oracle software and HP hardware. Each storage server is based on HP ProLiant DL180 G5 server whose main specifications are shown in [6].

Oracle Exadata Storage provides the base for the dynamic building data storage grids. The network storage also forms a massive parallel query engine.

Note. CERN has tested the HP Oracle Database Machine and found that using the offload option (option selected by the producer in the initial configuration) can quickly upload large amounts of data [7], being optimized both the time and the space.

Exadata Database Machine Version 2 was produced by Sun and Oracle and launched in 2009. This database machine is the fastest server neither in the world nor only for storing a large volume of data (data warehouse and OLAP), but also for current transactions processing (OLTP), the first of this kind in the world [8]. The machine was built using standard hardware components, plus: FlashFire technology - from Sun, Oracle Database 11g Release 2 and Oracle Exadata Storage Server Software Release 11.2, Database Machine Version 2 - from Sun and Oracle. With these technologies, the new machine is two times faster than Oracle Exadata Version 1, concerning the optimization of the storage

(space) and query (time) data.
Database Machine from Sun and Oracle goes beyond the data storage applications due to the adding of Exadata Smart Flash Cache [8] based on FlashFire technology from Sun which provides exceptional performance and scalability for current transaction processing (OLTP). The core of Sun Oracle database machine is the server Oracle Exadata Storage which is optimized software for processing the retrieval requests from Oracle Database 11g.
A piece of news concerning the optimization in Oracle Exadata V2 is the component Smart *Flash Cache* (SFC), based on Sun technology. This settles the competition of the I/O operations from the disk, the random access being performed very fast, which means an efficient processing of transactions (1 million random input-output operations per second). The SFC component involves the integration of computer networks with the technology of servers and the technology of the data storage in order to create a database of performance which has not been done before.
The massive parallel machine architecture Sun Oracle database allows the modularity, that is the dynamic expansion by adding specialized servers, database servers, network nodes, resulting in a *grid* type architecture tolerant to errors. Within such architecture any component can yield, without stopping the system (usually by the distributed databases - Oracle RAC rules). Such a high performance is obtained with minimal possible computing power
Exadata Version 2 is available in four models: full rack (8 database servers and storage servers 14), half-rack (4 servers for database and 7 servers storage), quarter rack (2 database servers and 3 data storage servers) and a base system (one database server and one server storage). The applications which are running on the Database Machine from Sun and Oracle perform up to 1 million operations I / O per second to Flash storage.
The contribution of the two companies to the new database machine is as follows [5]:

- hardware from Sun: FlashFire memory cards - allowing high-performance OLTP, Intel Xeon processors (Nehalem) - 80% faster CPU, 600 GB SAS drives to 6 gigabytes per second - 50% faster drives, DDR3 - 200% faster memory, 72 GB per server database - 125% more memory, 40 gigabits per second InfiniBand - network with 100% faster, new disk capacity 100 TB (SAS) or 336 TB (SATA) per *rack*;
- Oracle software: the first flash database in the world activated - Oracle 11g Release 2, column hybrid compression, data compression of 10-50 bigger, scanning of compressed data and faster execution of queries, index storage, reducing input and output operations, transfer of processing queries to a storage device using Smart Scans, Smart Scan component of Data Mining models in the storage servers [5].

Oracle Exadata V2 is based on a massive parallel architecture that includes fast servers and interconnections that connect storage servers and databases. The level of performance of the Oracle Exadata product V2 [5] is much higher than the previous generation level (Oracle Exadata V1): 20 times more random inputs and outputs, bandwidth input / output 5 times faster - FlashFire technology, computing capacity and network performance twice bigger.
Note. According to TPC, a non-profit organization which has as objective the testing and the intercomparable performance of different database technologies, the SunFire technology is considered the second as performance in terms of decision support system. The TPC-H test is one of performance (number of queries per hour) in terms of decision support and illustrates decision support systems that examine large volumes of data, perform complex queries and give answer to critical questions of the business [9].
The launching of an advanced product such as Oracle Exadata has not remained without echo among the competitors, especially since the machine database concept was launched long ago by the company Teradata. Teradata and Netezza are two companies which make the hardware products active on the market,

especially those for data warehouses.
Michel Bruley, Marketing Director at Teradata France, shows that the *Teradata* database machine are designed for data warehousing company, which must include simultaneously several different departments: commercial, marketing, logistics, financial, human resources so as to carry out transverse analysis. Therefore, in this case there is not a problem concerning the data volume, but the complexity of data utilization.
Netezza is a database machine produced by the company with the same name, a modern competitor of the Oracle database machine. Oracle Exadata introduces an additional step, namely it filters data after reading them from the disk, so that only relevant information to the database is transferred for processing. As filtering is performed outside the database, it is much faster. In the Netezza systems the data after having been filtered continue to be processed outside the database, then retrieved and ordered at node level (for each disc), and the result is sent to a database that combines the results from different nodes. Another advantage of Netezza systems is that there is a 1 to 1 relationship between the number of processors and the number of disks, while in the HP Oracle Database Machine there are more disks than processors.
Christian Raza, Operations Manager at Netezza France concerning the HP Oracle Database Machine product performance says that what Oracle suggests represents only a package of the traditional database Oracle 11g [10], while the company Netezza has developed a completely different and new architecture compared to what existed at that moment, providing performance and ease in utilization.
A disadvantage of this system over Oracle is connected to the bandwidth: the processing modes are connected to the database through the InfiniBand mechanism means, which is much faster than the one chosen by Netezza.
Note. Transilvania Bank is the first client of Sun Oracle Exadata Database Machine V2 product. This bank uses the first Oracle database machine within their data warehouse.
LGR IT company uses Oracle Exadata Database Machine to provide to mobile telephone operators data storage services in data warehousing, analysis and reporting concerning the recordings with details of the calls performed in the network (Call Detail Records, CDRs) whose number can be hundreds of millions or even billions per day.

3 Examples of modern database machine applications

Database machines in the geographical information systems. Together with the use of databases in domain of geographical information systems (GIS - Geographical Information Systems), appeared the necessity of storage and query of new data types, the geo-spatial ones. These can be of two types: raster and vectorials.
Raster data are photos of the Earth taken either from a satellite or from the airplane: they are files that store information in discrete cells arranged in rows and columns. Each cell or pixel from photography preserves a certain value.
The vectors data consist of points, lines, and polygons. They are suitable for storing objects shape, unlike the raster data that stores their content.
Taking into account the complexity of these types of data, in a GIS project at a national or international level, a very large amount of database will be used. In addition to the choice of a database which in the terms of the software products should provide the implementation of some algorithms for the optimization of the geo-spatial data query, it would also be efficient to choose a database machine that has specialized hardware incorporated for operating with such data.
The spatial databases use a query model in two steps to solve spatial queries or spatial junctions. Two well-known operations are carried out under the name of primary filtration and secondary filtration. The primary filtering allows the quick selection of the candidate records (registration) that will enter into the second filter - secondary, after which the exact result will be returned.
Hardware acceleration mechanisms of queries and spatial junctions, based on graphics

processing are presented in the specialized literature [11], [12]. For example, the test for the intersection of two polygons will be the following:

- the first polygon will be colored with color c1;
- the second polygon with color by the color c2;
- searching for color pixels c1 + c2;
- if existing, then the two polygons intersect themselves.

A filter is proposed to be introduced as an intermediary between the primary filtration and the secondary filtration that should reduce even more the candidate set of records which will enter into the secondary filtration by using a graphics processing mechanism.

Modern database machines as a high-performance solution for e-Voting. Together with the evolution of the information technology, the governmental institutions started to take interest in the automation (computerization) of the flows of activities that they undertake [13]. Thus, the term e-Voting appeared. It is a concept widely used in the century we live, but little implemented at a global level, as information systems. An e-Voting system is part of an e-Government type information system at local or national level.

The most important function of an e-Voting type information system is the automation of the voting process (primarily the electoral vote). Such an information system requires that the citizens of a certain locality or State vote via the Internet (without going to the polling station) the president, the party or the mayor who will lead the State or the respective locality during the next mandate.

The main *reasons* for which the e-Voting type systems have a restricted use in the world are:

- it is very difficult to provide a high level of the data security;
- the need for some performing software and hardware products which could cope successfully with a large number of queries during the voting process and then while extracting the final reports with the results of the voting process and the designation of the winners.

We consider that the above mentioned difficulties can be overcome if an information system for e-Voting had as support a database machine. Therefore, we *suggest* connecting the product Oracle Exadata Storage Server V2, to an e-Voting type information system. In order to check the complete functionality and the performances of the new obtained system, it is desirable to achieve a simulation of the voting process, first on a small sample of population, for example a medium sized city in Romania. Then the implementation of this system at a national level can be done, if the obtained results are:

- the display time of the page with the list of the candidates must be reduced, less than 10 seconds (The list of candidates will be loaded in the voting page by queries of the database);
- the display time of the reports with the final results must be reduced if possible less than 10 minutes.

By joining the information e-Voting system to the Oracle Exadata Storage Server V2 product, we get a very performing product, currently without competition and that can spur institutions to adopt it.

4 Conclusions

The huge data volume and also the high complexity of data led to the emergence of modern database machines, as a solution for the optimization of the computing resources. In the new economy based on knowledge, the role of storage, queries and data analysis is crucial for any company. When the volume of these data becomes very large, the use of a database machine, hardware and software, specialized for storing, accessing and interpreting data in real time becomes a competitive advantage. The obsolescence of these machines occurs in a relatively short time, taking into account that a more performing machine can be launched simultaneously with the hardware developments in this field.

The last two years witnessed the launch on the market of modern database machines. This was possible due to the current information context, marked by the emergence and the consolidation of new information tech-

nologies concerning the storage and the efficient processing of large volumes of data.
The applicability of the modern database machines refers both to data warehouses (analytical processing - OLAP) and to the transactional databases (transactional processing - OLTP).
Modern database machines can be used for decision support systems, and also for information executive systems where it is necessary to optimize the data storage, as well as the retrieval of data (the banking system, the governmental system, marketing, etc.).
The construction of specialized computers for the optimization of the data storage space and the data access time is an efficient solution for the companies which make use of data warehouses.

References

[1] C. Combes and C. Rivat, "A modelling environment based on data warehousing to manage and to optimize the running of international company," *Int. J. Production Economics*, Vol. 112, 2008, pp. 294–308.
[2] V. S. Jacob and X. B. Li, "Adaptive data reduction for large-scale transaction data," *European Journal of Operational Research*, Vol. 188, 2008, pp. 910–924.
[3] Software documentation services from TechScribe, Available at: http://www.techscribe.co.uk
[4] H. Boral and D. J. DeWitt, *Database Machines: An idea whose time has passed? A critique of the future of database machines, in Parallel Architectures For Database Systems*, Ed. IEEE Press, Piscataway, NJ, 1983, pp. 11-28.
[5] D. Baum, "Launching Performance," *Oracle Magazine*, Vol. XXIII, No. 1, January/ February 2009.
[6] HP ProLiant DL180 Server series, Available at: http://h20000.www2.hp.com/bizsupport/TechSupport/Home.jsp?lang=en&cc=us&prodClassId=-1&prodTypeId=15351&prodSeriesId=3436778
[7] G. Eric, „Oracle and storage IOs, explanations and experience at CERN," *17th International Conference on Computing in High Energy and Nuclear Physics*, Prague, Czech Republic, March 2009, pp. 21 – 27.
[8] J. Erickson, "Data Warehousing gets extreme," *Oracle Magazine*, Vol. XXII, No. 6, November/ December 2008.
[9] TPC – Homepage, Available at: http://www.tpc.org
[10] M. Gros, "Entrepôt de données HP Oracle: les concurrents réagissent," *LeMondeInformatique.fr*, Sept. 2008.
[11] N. Bandi, C. Sun, D. Agrawal and A. E. Abbadi, „Hardware acceleration in commercial databases: a case study of spatial operations in VLDB '04," *Proceedings of the Thirtieth international conference on Very large data bases. VLDB Endowment*, 2004, pp. 1021-1032.
[12] N. K. Govindaraju, J. Gray, R. Kumar and D. Manocha, „GPUTeraSort: High Performance Graphics Coprocessor Sorting for Large Database Management," *SIGMOD Conference*, Ed. New York, NY, USA, ACM, 2006, pp. 325-336.
[13] D. Lițan, "E-Government – fundamental background and implementations," *The Proceedings of The Ninth International Conference on Informatics in Economy*, Academy of Economic Studies, Bucharest, May 2009, pp. 449.

Manole VELICANU is a Professor at the Economic Informatics Department at the Faculty of Cybernetics, Statistics and Economic Informatics from the Academy of Economic Studies of Bucharest. He has graduated the Faculty of Economic Cybernetics in 1976, holds a PhD diploma in Economic Cybernetics from 1994 and starting with 2002 he is a PhD coordinator in the field of Economic Informatics. He is the author of 19 books in the domain of economic informatics, 73 published articles (among which 4 articles ISI in-

dexed), 61 scientific papers published in conferences proceedings (among which 5 papers ISI indexed and 7 included in international databases) and 36 scientific papers presented at conferences, but unpublished. He participated (as director or as team member) in more than 49 research projects that have been financed from national research programs. He is a member of INFOREC professional association, a CNCSIS expert evaluator and a MCT expert evaluator for the program Cercetare de Excelenta - CEEX (from 2006). From 2005 he is co-manager of the master program Databases for Business Support. His fields of interest include: Advanced Databases, Design of Information Systems, Database Management Systems, Artificial Intelligence, Programming languages, Java Platform.

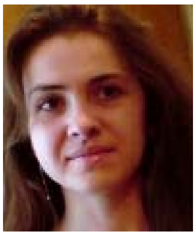

Daniela LIȚAN has graduated The Bucharest Academy of Economic Studies, Faculty of Cybernetics, Statistics and Economic Informatics in 2006. She holds a Master diploma in Databases - Support for business from 2009 and in present she is a PhD Candidate in Economic Informatics with the Doctor's Degree Thesis: The integration of new technologies in the information systems of e-Government. Her areas of interest are: Informatics and Economics.

Aura-Mihaela MOCANU (VÎRGOLICI) has graduated The Bucharest Academy of Economic Studies, Faculty of Cybernetics, Statistics and Economic Informatics in 2007. She holds a Master diploma in Databases - Support for business from 2009 and in present she is a PhD Candidate in Economic Informatics with the Doctor's Degree Thesis: Software technologies to build a Geographical Information System for a public institution. Her areas of interest are: Geographical Information Systems, Databases, Information Systems integration, Programming languages.