

## **Index and Materialized View Selection in Data Warehouses**

Kamel Aouiche<sup>1(\*)</sup> and Jérôme Darmont<sup>2</sup>

<sup>1</sup>LICEF, Université de Québec à Montréal, Canada

<sup>2</sup>ERIC, Université Lumière Lyon 2, France

{kamel.aouiche, jerome.darmont}@univ-lyon2.fr

## INTRODUCTION

Database Management Systems (DBMSs) require an administrator, whose principal tasks are data management, both at the logical and physical levels, as well as performance optimization. With the wide development of databases and data warehouses, minimizing the administration function is crucial. This function includes the selection of suitable physical structures to improve system performance.

View materialization and indexing are presumably some of the most effective optimization techniques adopted in relational implementations of data warehouses. Materialized views are physical structures that improve data access time by precomputing intermediary results. Therefore, end-user queries can be efficiently processed through data stored in views and do not need to access the original data. Indexes are also physical structures that allow direct data access. They avoid sequential scans and thereby reduce query response time. Nevertheless, these solutions require additional storage space and entail maintenance overhead. The issue is then to select an appropriate configuration of materialized views and indexes that minimizes both query response time and maintenance cost, given a limited storage space. This problem is NP-hard (Gupta & Mumick, 2005).

The aim of this article is to present an overview of the major families of state-of-the-art index and materialized view selection methods; and to discuss the issues and future trends in data warehouse performance optimization. We particularly focus on data mining-based heuristics we developed to reduce the selection problem complexity and target the most pertinent candidate indexes and materialized views.

## BACKGROUND

Today's commercial relational DBMSs provide integrated tools for automatic physical design. For a given workload, they automatically recommend configurations of indexes and materialized views (Dageville *et al.*, 2004), coupled with data partitioning (Agrawal *et al.*,

2004) or table clustering (Zilio *et al.*, 2004). However, these tools depend on the query optimizer and therefore the host DBMS, which renders their adaptation onto other systems intricate. In the remainder of this section, we detail published research about index and materialized view selection.

#### Index selection problem

The index selection problem has been studied for many years in databases (Finkelstein *et al.*, 1988; Chaudhuri *et al.*, 2004), but adaptations to data warehouses are few. In this particular context, research studies may be clustered into two families: algorithms that optimize maintenance cost (Labio *et al.*, 1997) and algorithms that optimize query response time. In both cases, optimization is realized under the storage space constraint. We particularly focus on the second family of approaches, which may be classified depending on how the set of candidate indexes and the final configuration of indexes are built.

The set of candidate indexes may be built manually by the administrator, according to his expertise of the workload (Frank *et al.*, 1992; Choenni *et al.*, 1993a; Choenni *et al.* 1993b). This is both subjective and quite hard to achieve when the number of queries is large. In opposition, candidate indexes may also be extracted automatically through a syntactic analysis of the workload (Chaudhuri & Narasayya, 1997; Valentin *et al.*, 2000; Golfarelli *et al.*, 2002).

There are several methods for building the final index configuration from the candidate indexes. Ascending methods start from an empty set of indexes (Kyu-Young, 1987; Frank *et al.*, 1992; Choenni *et al.*, 1993b; Chaudhuri & Narasayya, 1997). They increasingly select indexes minimizing workload cost until it does not decrease anymore. Descending methods start with the whole set of candidate indexes and prune indexes until workload cost increases (Kyu-Young, 1987; Choenni *et al.*, 1993a). Classical optimization algorithms have also been used to solve this problem, such as knapsack resolution (Ip *et al.*, 1983; Gündem, 1999;

Valentin *et al.*, 2000; Feldman & Reouven, 2003) and genetic algorithms (Kratika *et al.*, 2003).

#### Materialized view selection problem

The classical papers about materialized view selection in data warehouses introduce a lattice framework that models and captures ancestor/descendent dependency among aggregate views in a multidimensional context (Harinarayan *et al.*, 1996; Baralis *et al.*, 1997; Kotidis & Roussopoulos, 1999; Uchiyama *et al.*, 1999). This lattice is greedily browsed with the help of cost models to select the best views to materialize. This problem has first been addressed in one data cube, and then extended to multiple cubes (Shukla *et al.*, 2000). Another theoretical framework, called the AND-OR view graph, may also be used to capture the relationships between materialized views (Chan *et al.*, 1999; Theodoratos *et al.*, 2000; Valluri *et al.*, 2002; Gupta & Mumick, 2005). However, the majority of these solutions are theoretical and not truly scalable.

Another method decomposes data cubes into an indexed hierarchy of wavelet view elements and selects those that minimize the average processing cost of the queries defined on the data cubes (Smith *et al.*, 2004). Similarly, the Dwarf structure (Sismanis *et al.*, 2002) compresses data cubes, thereby suppressing redundancy to improve maintenance and interrogation costs. These approaches are very interesting, but they mainly focus on computing efficient data cubes by changing their physical design, which is not always convenient in practice.

Yet other approaches detect common subexpressions within workload queries that correspond to intermediary results that are suitable to materialize (Baril & Bellahsene, 2003; Goldstein & Larson, 2003). However, browsing is very costly and these methods are not truly scalable with respect to the number of queries.

Finally, the most recent approaches are workload-driven. They syntactically analyze the workload to enumerate relevant candidate views (Agrawal *et al.*, 2001). By calling the system

query optimizer, they greedily build a configuration of the most pertinent views. A real workload is indeed considered as a good starting point to predict future queries.

#### Coupling index and materialized view selection

A few research studies deal with the simultaneous selection of indexes and materialized views. Agrawal *et al.* (2001) proposed three alternative approaches. The first, MVFIRST, selects materialized views first and then indexes. The second, INDFIRST, selects indexes first and then materialized views. The third, *joint enumeration*, is claimed by the authors to be the most efficient for workload execution time optimization. It processes indexes, materialized views and indexes over these views simultaneously.

Bellatreche *et al.* (2000) studied the problem of storage space distribution among materialized views and indexes. A set of views and indexes are selected as an initial solution. Then, this solution is iteratively modified to reduce the execution cost, by redistributing storage space among indexes and materialized views.

Finally, Rizzi & Saltarelli (2003) *a priori* determine a trade-off between the storage spaces allotted to indexes and materialized views, depending on how queries are defined. Their idea is that view materialization provides the best benefit for queries involving coarse granularity aggregations, while indexing provides the best benefit with queries containing attributes with a high selectivity.

## DATA MINING-BASED INDEX AND MATERIALIZED VIEW SELECTION

### Strategy overview

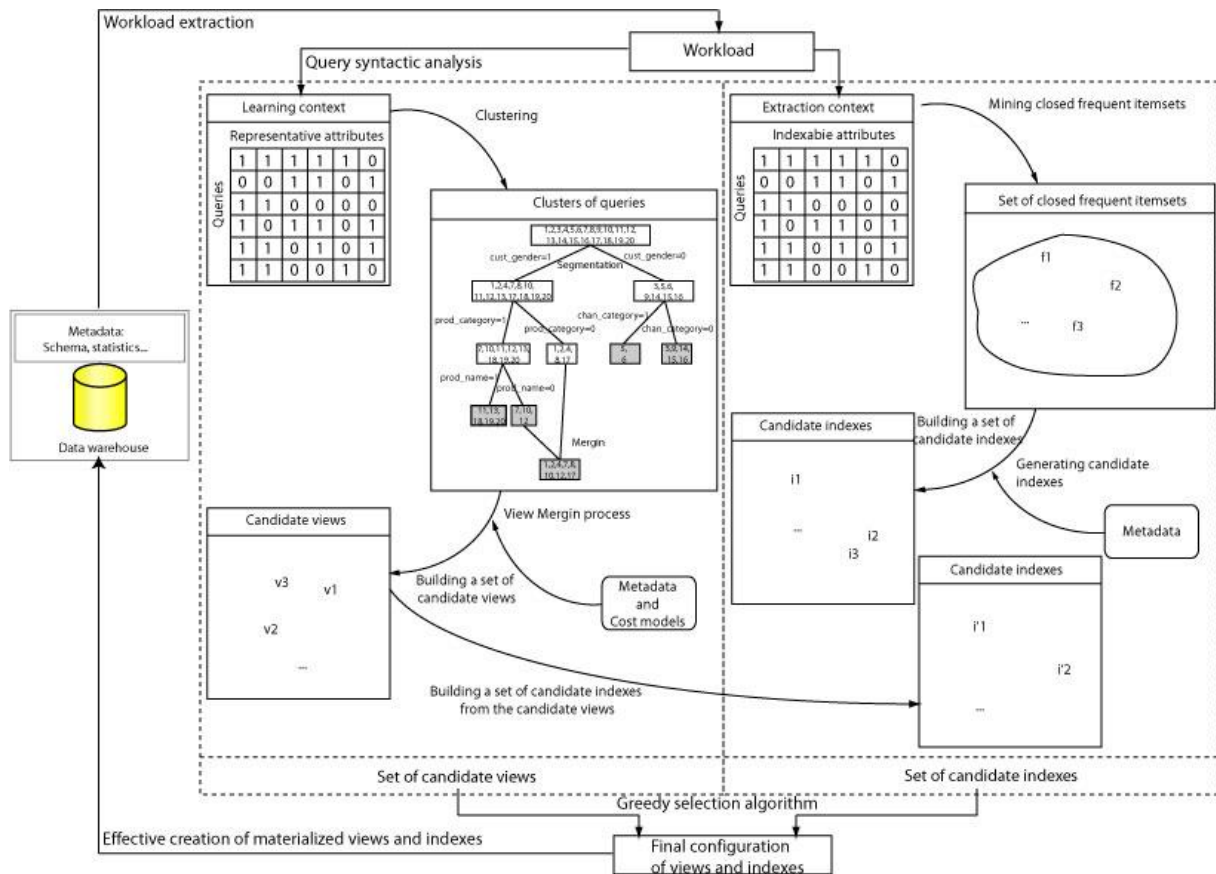
We advocate for index and materialized view selection strategies that bear the following features:

- *automatic*: the final configuration of indexes and views should be built automatically;
- *generic*: the selection strategy should not be dependent on a particular DBMS;
- *modular*: the selection strategy should be composed of independent modules;

- *scalable*: the strategy must be able to handle large workloads.

To achieve this goal, we designed a new strategy (Aouiche *et al.*, 2005; Aouiche *et al.*, 2006) that is composed of several modules: a syntactical query analyzer, a data miner, cost models, and an index and materialized view selector (Figure 1). The query analyzer syntactically processes the input workload to extract the most pertinent attributes for indexing and view materialization. It also exploits some knowledge about performance administrative tasks, formalized as if-then rules. For instance, common rules imply the selection of attributes from the “Where” and “Group by” clauses in SQL statements. The output of this process is a so-called “query-attribute” binary matrix whose lines are the analyzed queries and whose columns are the extracted attributes. The general term of this matrix is set to 1 if an extracted attribute is present in the corresponding query and to 0 otherwise.

The query-attribute matrix constitutes the extraction context for the data miner. This module builds a configuration of candidate indexes and materialized views. It may exploit any data mining technique, suiting the data structure to select (indexes or materialized views). For instance, our materialized view selection strategy exploits clustering for building sets of similar queries. The idea of exploiting clustering is motivated by the fact that several queries having a similar syntax may likely be resolved from one materialized view. Hence, workload queries are grouped into clusters that are exploited to build the set of candidate views. Furthermore, these candidate views are merged to resolve multiple queries.



**Figure 1: Materialized view and index selection system**

Our index selection strategy exploits another data mining technique, frequent itemset mining, to determine the candidate indexes. Our intuition here is that index utility is strongly correlated to the usage frequency of the corresponding attributes within a given workload, which frequent itemset mining is good at highlighting. Each itemset is analyzed to generate a set of candidate indexes, with the help of metadata (schema: primary keys, foreign keys; statistics, etc.). This process for building candidate indexes may also be applied on the candidate materialized views, since they are actually tables. Hence, we can build indexes on materialized views to maximize performance improvements.

Finally, the cost model module takes as input the data warehouse metadata, the workload and candidate indexes and materialized views. It computes the cost, in terms of access and storage cost, of each query in the presence of the candidate indexes and/or views. Since the number of candidates is generally as high as the input workload is large, it is not feasible to materialize

them all because of storage space constraints. Hence, our cost models are exploited by the index and view selector to greedily build a final configuration of indexes and materialized views. When simultaneously selecting indexes and materialized views, we exploit specific cost models that allow taking into account the interactions between indexes and materialized views and efficiently sharing storage space.

## Discussion

Thanks to its modularity, we have been able to apply our strategy in several cases. For instance, we performed B-tree index selection in a database context as well as bitmap join index selection in a data warehouse. In addition, modularity helps in gradually improving our strategy. A given module may indeed be easily replaced by another, more efficient one. For example, it is easy to replace our cost models by more accurate ones, or a data mining algorithm by a more efficient or scalable one.

In opposition to other approaches, particularly those of DBMS vendors, we aimed at remaining as generic and independent from the host DBMS as possible. Our analyzer module indeed processes standard SQL queries, for instance. Our cost models are also mathematical so that they do not depend on a query optimizer. Hence, our strategy may be instantiated within different systems.

Finally, our approach takes into account knowledge (metadata, usage statistics; knowledge extracted from the query workload, the way attributes are queried, etc.) that helps in reducing the selection problem complexity and thus targeting the most pertinent candidate indexes and materialized views. Since this approach is largely based on data mining, we benefit from the active research in this field, which now provides fast and scalable algorithms. Hence, we can process and analyze large workloads.

## FUTURE TRENDS



Our strategies are applied on a workload that is extracted from the system during a given period of time. We are thus performing static optimization. Future developments in this domain (both ours and others) should be dynamic and incremental (Kotidis & Roussopoulos, 1999). In our case, studies dealing with dynamic or incremental clustering and frequent itemset mining may be exploited to update the configuration of indexes and materialized views instead of recreating it from scratch. Entropy-based session detection could also be beneficial to determine the best moment to periodically run such a strategy.

In our work, we also only coupled the selection of indexes and materialized views, but the current trend in recent commercial systems is to exploit a mix of several optimization techniques such as buffering, physical clustering, partitioning, etc. (Agrawal *et al.*, 2004; Dageville *et al.*, 2004; Zilio *et al.*, 2004) to achieve the best performance enhancement. We also aim at integrating the selection of other optimization structures into our strategy.

Eventually, a tremendous amount of research is currently in progress to help XML-native DBMSs in becoming a credible alternative to XML-compatible, relational DBMSs. The majority of XML-native DBMSs indeed present relatively poor performances when the volume of data is very large and queries are complex. However, since XML is gaining importance for representing business data for analytics (Beyer *et al.*, 2005), it is crucial to design automatic ways of guaranteeing the best performance of XML data warehouses.

## CONCLUSION

The problem of performance optimization has been receiving significant attention since the early days of database research. However, each new class of DBMS (hierarchical, network, relational, object, XML...) or special purpose of database architecture (such as the decision-support data warehouses) invariably gives way to the reformulation or adaptation of existing techniques, and to brand new issues that require original solutions.

## REFERENCES

Agrawal, S., Chaudhuri, S., & Narasayya, V. (2001). Materialized view and index selection tool for Microsoft SQL Server 2000. In ACM SIGMOD International Conference on Management of Data (SIGMOD 2001), Santa Barbara, USA (p. 608). New York: ACM Press.

Agrawal, S., Chaudhuri, S., Koll r, L., Marathe, A.P., Narasayya, V.R., & Syamala, M. (2004). Database Tuning Advisor for Microsoft SQL Server 2005. In 30<sup>th</sup> International Conference on Very Large Data Bases (VLDB 2004), Toronto, Canada (pp. 1110-1121). San Francisco: Morgan Kaufmann.

Aouiche, K., Darmont, J., Boussa d, O., & Bentayeb, F. (2005). Automatic Selection of Bitmap Join Indexes in Data Warehouses, In 7<sup>th</sup> International Conference on Data Warehousing and Knowledge Discovery (DaWaK 2005), Copenhagen, Denmark: Vol. 3589 of LNCS (pp. 64–73). Berlin: Springer.

Aouiche, K., Jouve, P.E., & Darmont, J. (2006). Clustering-Based Materialized View Selection in Data Warehouses. In 10<sup>th</sup> East-European Conference on Advances in Databases and Information Systems (ADBIS 2006), Thessaloniki, Greece: Vol. 4152 of LNCS (pp. 81-95). Berlin: Springer.

Baralis, E., Paraboschi, S., & Teniente, E. (1997). Materialized views selection in a multidimensional database. In 23<sup>rd</sup> International Conference on Very Large Data Bases (VLDB 1997), Athens, Greece (pp. 156-165). San Francisco: Morgan Kaufmann.

Baril, X., & Bellahsene, Z. (2003). Selection of materialized views: a cost-based approach. In 15<sup>th</sup> International Conference on Advanced Information Systems Engineering (CAiSE 2003), Klagenfurt, Austria: Vol. 2681 of LNCS (pp. 665-680). Berlin: Springer.

Bellatreche, L., Karlapalem, K., & Schneider, M. (2000). On efficient storage space distribution among materialized views and indices in data warehousing environments. In 9<sup>th</sup> International Conference on Information and Knowledge Management (CIKM 2000), Mclean, USA (pp. 397-404). New York: ACM Press.

Beyer, K.S., Chamberlin, D.D., Colby, L.S., Özcan, F., Pirahesh, H., & Xu, Y. (2005). Extending XQuery for Analytics. In ACM SIGMOD International Conference on Management of Data (SIGMOD 2005), Baltimore, USA (pp. 503-514). New York: ACM Press.

Chan G. K. Y., Li Q., & L. Feng. (1999). Design and selection of materialized views in a data warehousing environment: a case study. In 2<sup>nd</sup> ACM international workshop on Data warehousing and OLAP (DOLAP 1999), Kansas City, USA (pp. 42-47).

Chaudhuri, S., & Narasayya, V.R. (1997). An efficient cost-driven index selection tool for Microsoft SQL server. In 23<sup>rd</sup> International Conference on Very Large Data Bases (VLDB 1994), Santiago de Chile, Chile (pp. 146–155). San Francisco: Morgan Kaufmann.

Chaudhuri, S., Datar, M., & Narasayya, V. (2004). Index selection for databases: A hardness study and a principled heuristic solution. *IEEE Transactions on Knowledge and Data Engineering*, 16(11), 1313-1323.

Choenni, S., Blanken, H.M., & Chang, T. (1993a). Index selection in relational databases. In 5<sup>th</sup> International Conference on Computing and Information (ICCI 1993), Ontario, Canada (pp. 491-496). Los Alamitos: IEEE Computer Society.

Choenni, S., Blanken, H.M., & Chang, T. (1993b). On the selection of secondary indices in relational databases. *Data Knowledge Engineering*, 11(3), 207-238.

Dageville, B., Das, D., Dias, K., Yagoub, K., Zaït, M., & Ziauddin, M. (2004). Automatic SQL Tuning in Oracle 10g. In 30<sup>th</sup> International Conference on Very Large Data Bases (VLDB 2004), Toronto, Canada (pp. 1098-1109). San Francisco: Morgan Kaufmann.

Feldman, Y.A., & Reouven, J. (2003). A knowledge-based approach for index selection in relational databases. *Expert System with Applications*, 25(1), 15-37.

Finkelstein, S.J., Schkolnick, M., & Tiberio, P. (1988). Physical database design for relational databases. *ACM Transactions on Database Systems*, 13(1), 91-128.

Frank, M.R., Omiecinski, E., & Navathe, S.B. (1992). Adaptive and automated index selection in RDBMS. In 3<sup>rd</sup> International Conference on Extending Database Technology, (EDBT 1992), Vienna, Austria: Vol. 580 of LNCS (pp. 277-292). Berlin: Springer.

Goldstein, J., & Larson, P.A. (2001). Optimizing queries using materialized views: a practical, scalable solution. In ACM SIGMOD International Conference on Management of Data (SIGMOD 2001), Santa Barbara, USA (pp. 331-342). New York: ACM Press.

Golfarelli, M., Rizzi, S., & Saltarelli, E. (2002). Index selection for data warehousing. In 4th International Workshop on Design and Management of Data Warehouses (DMDW 2002), Toronto, Canada: Vol. 58 of CEUR Workshop Proceedings (pp. 33-42). Aachen: CEUR-WS.org.

Gündem, T.I. (1999). Near optimal multiple choice index selection for relational databases. *Computers & Mathematics with Applications*, 37(2), 111-120.

Gupta, H., & Mumick, I.S. (2005). Selection of views to materialize in a data warehouse. *IEEE Transactions on Knowledge and Data Engineering*, 17(1), 24-43.

Harinarayan, V., Rajaraman, A., & Ullman, J.D. (1996). Implementing data cubes efficiently. In ACM SIGMOD International Conference on Management of Data (SIGMOD 1996), Montreal, Canada (pp. 205-216). New York: ACM Press.

Ip, M.Y.L., Saxton, L.V., & Raghavan, V.V. (1983). On the selection of an optimal set of indexes. *IEEE Transactions on Software Engineering*, 9(2), 135-143.

Kotidis, Y., & Roussopoulos, N. (1999). Dynamat: A dynamic view management system for data warehouses. In ACM SIGMOD International Conference on Management of Data (SIGMOD 1999), Philadelphia, USA (pp. 371–382). New York: ACM Press.

Kratika, J., Ljubic, I., & Tasic, D. (2003). A genetic algorithm for the index selection problem. In Applications of Evolutionary Computing (EvoWorkshops 2003), Essex, UK: Vol. 2611 of LNCS (pp. 281–291). Berlin: Springer.

Kyu-Young, W. (1987). Index Selection in Rational Databases. In Ghosh, S.P., Kambayashi, Y., & Tanaka, K. (Eds.), *Foundation of data organization* (pp. 497-500). New York: Plenum Publishing.

Labio, W., Quass, D., & Adelberg, B. (1997). Physical database design for data warehouses. In *13<sup>th</sup> International Conference on Data Engineering (ICDE 1997)*, Birmingham, UK (pp. 277-288). Los Alamitos: IEEE Computer Society.

Rizzi, S., & Saltarelli, E. (2003). View materialization vs. indexing: Balancing space constraints in data warehouse design. In *15<sup>th</sup> International Conference on Advanced Information Systems Engineering (CAiSE 2003)*, Klagenfurt, Austria: Vol. 2681 of LNCS (pp. 502-519). Berlin: Springer.

Shukla, A., Deshpande, P., & Naughton, J.F. (2000). Materialized view selection for multi-cube data models. In *7<sup>th</sup> International Conference on Extending Database Technology (EDBT 2000)*, Konstanz, Germany: Vol. 1777 of LNCS (pp. 269-284). Berlin: Springer

Sismanis, Y., Deligiannakis, A., Roussopoulos, N., & Kotidis, Y. (2002). Dwarf: shrinking the petacube. In *ACM SIGMOD International Conference on Management of Data (SIGMOD 2002)*, Madison, USA (pp. 464-475). New York: ACM Press.

Smith, J.R., Li, C.S., & Jhingran, A. (2004). A wavelet framework for adapting data cube views for OLAP. *IEEE Transactions on Knowledge and Data Engineering*, 16(5), 552-565.

Theodoratos D. & Bouzeghoub M. (2000). A general framework for the view selection problem for data warehouse design and evolution. In *3<sup>rd</sup> ACM International Workshop on Data Warehousing and OLAP (DOLAP 2000)*, New York, USA (pp. 1-8).

Uchiyama, H., Runapongsa, K., & Teorey, T.J. (1999). A progressive view materialization algorithm. In *2<sup>nd</sup> International Workshop on Data Warehousing and OLAP (DOLAP 1999)*, Kansas City, USA (pp. 36-41). New York: ACM Press.

Valentin, G., Zuliani, M., Zilio, D., Lohman, G., & Skelley, A. (2000). DB2 advisor: An optimizer smart enough to recommend its own indexes. In 16<sup>th</sup> International Conference on Data Engineering, (ICDE 2000), California, USA (pp. 101-110). Los Alamitos: IEEE Computer Society.

Valluri, S.R., Vadapalli, S., & Karlapalem, K. (2002). View relevance driven materialized view selection in data warehousing environment. In 13<sup>th</sup> Australasian Database Conference (ADC 2002), Melbourne, Australia (pp. 187-196).

Zilio, D.C., Rao, J., Lightstone, S., Lohman, G.M., Storm, A., Garcia-Arellano, C., & Fadden, S. (2004). DB2 Design Advisor: Integrated Automatic Physical Database Design. In 30<sup>th</sup> International Conference on Very Large Data Bases (VLDB 2004), Toronto, Canada (pp. 1087-1097). San Francisco: Morgan Kaufmann.

#### TERMS AND DEFINITIONS

**Granularity:** The aggregation level within a dimension hierarchy.

**Data cube:** Data modeled and viewed in a multidimensional space.

**Data mining:** The nontrivial extraction of implicit, previously unknown, and potentially useful information from data.

**Index:** Physical data structure that allow direct (vs. sequential) access to data.

**Materialized view:** Physical data structure that improves data access time by precomputing intermediary results.

**On-line analytical processing (OLAP):** An approach for processing decision-support, analytical queries that are dimensional in nature.

**Workload:** Set of queries that are executed over a given database or data warehouse.

**Selectivity:** The portion of accessed tuples which are effectively selected by a query.