

ETL and OLAP Systems

Krzysztof Dembczyński

Intelligent Decision Support Systems Laboratory (IDSS)
Poznań University of Technology, Poland



Software Development Technologies
Master studies, first semester
Academic year 2017/18 (winter course)

Review of the Previous Lecture

- Mining of massive datasets.
- Evolution of database systems.
- Dimensional modeling:
 - ▶ Three goals of the logical design of data warehouse: **simplicity**, **expressiveness** and **performance**.
 - ▶ The most popular conceptual schema: **star schema**.
 - ▶ Designing data warehouses is not an easy task . . .

Outline

- 1 Motivation
- 2 ETL
- 3 OLAP Systems
- 4 Summary

Outline

- 1 Motivation
- 2 ETL
- 3 OLAP Systems
- 4 Summary

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.
 - ▶ Additional data structures: sorting, indexing, summarizing, cubes.

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.
 - ▶ Additional data structures: sorting, indexing, summarizing, cubes.
 - ▶ Refreshing of data structures.

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.
 - ▶ Additional data structures: sorting, indexing, summarizing, cubes.
 - ▶ Refreshing of data structures.
- Querying multidimensional data:

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.
 - ▶ Additional data structures: sorting, indexing, summarizing, cubes.
 - ▶ Refreshing of data structures.
- Querying multidimensional data:
 - ▶ SQL extensions,

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.
 - ▶ Additional data structures: sorting, indexing, summarizing, cubes.
 - ▶ Refreshing of data structures.
- Querying multidimensional data:
 - ▶ SQL extensions,
 - ▶ Multidimensional expressions (MDX),

Motivation

- OLAP queries are usually performed in a separate system, i.e., a data warehouse.
- Transferring data to data warehouse:
 - ▶ Data warehouses combine data from multiple sources.
 - ▶ Data must be translated into a consistent format.
 - ▶ Data integration represents 80% of effort for a typical data warehouse project!
- Optimization of data warehouse:
 - ▶ Data storage: relational or multi-dimensional.
 - ▶ Additional data structures: sorting, indexing, summarizing, cubes.
 - ▶ Refreshing of data structures.
- Querying multidimensional data:
 - ▶ SQL extensions,
 - ▶ Multidimensional expressions (MDX),
 - ▶ Map-reduce-based language.

Outline

- ① Motivation
- ② ETL
- ③ OLAP Systems
- ④ Summary

ETL

- **ETL** = Extraction, Transformation, and Load
 - ▶ **Extraction** of data from source systems,
 - ▶ **Transformation** and **integration** of data into a useful format for analysis,
 - ▶ **Load** of data into the warehouse and build of additional structures.
- **Refreshment** of data warehouse is closely related to ETL process.
- The ETL process is described by metadata stored in data warehouse.
- Architecture of data warehousing:

Data sources \Rightarrow Data staging area \Rightarrow Data warehouse

ETL

BI and Reporting Tools

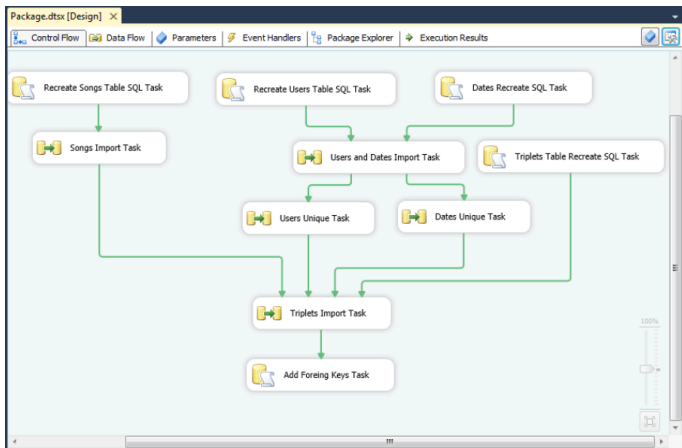


Tools for ETL

- Data extraction from heterogeneous data sources.
- Data transformation, integration, and cleansing.
- Data quality analysis and control.
- Data loading.
- High-speed data transfer.
- Data refreshment.
- Managing and analyzing metadata.
- **Examples of ETL tools:**
 - ▶ MS SQL Server Integration Services(SSIS), IBM InfoSphere DataStage, SAS ETL Studio, Oracle Warehouse Builder, Oracle Data Integrator, Business Objects Data Integrator, Pentaho Data Integration.

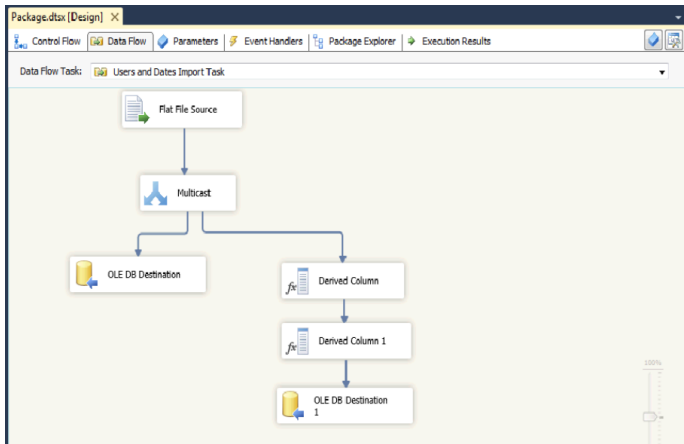
Tools for ETL

- MS SQL Server Integration Services(SSIS)



Tools for ETL

- MS SQL Server Integration Services(SSIS)



Data extraction

- Data warehouse needs extraction of data from different external data sources:

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,
 - ▶ and other documents (.txt, .doc, XML, WWW).

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,
 - ▶ and other documents (.txt, .doc, XML, WWW).
- Access to data sources can be difficult:

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,
 - ▶ and other documents (.txt, .doc, XML, WWW).
- Access to data sources can be difficult:
 - ▶ Data sources are often operational systems, providing the lowest level of data.

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,
 - ▶ and other documents (.txt, .doc, XML, WWW).
- Access to data sources can be difficult:
 - ▶ Data sources are often operational systems, providing the lowest level of data.
 - ▶ Data sources are designed for operational use, not for decision support, and the data reflect this fact.

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,
 - ▶ and other documents (.txt, .doc, XML, WWW).
- Access to data sources can be difficult:
 - ▶ Data sources are often operational systems, providing the lowest level of data.
 - ▶ Data sources are designed for operational use, not for decision support, and the data reflect this fact.
 - ▶ Multiple data sources are often from different systems, run on a wide range of hardware and much of the software is built in-house or highly customized.

Data extraction

- Data warehouse needs extraction of data from different external data sources:
 - ▶ operational databases (relational, hierarchical, network, itp.),
 - ▶ files of standard applications (Excel, COBOL applications),
 - ▶ additional databases (direct marketing databases) and data services (stock data),
 - ▶ various log files,
 - ▶ and other documents (.txt, .doc, XML, WWW).
- Access to data sources can be difficult:
 - ▶ Data sources are often operational systems, providing the lowest level of data.
 - ▶ Data sources are designed for operational use, not for decision support, and the data reflect this fact.
 - ▶ Multiple data sources are often from different systems, run on a wide range of hardware and much of the software is built in-house or highly customized.
 - ▶ Data sources can be designed using different logical structures.

Data extraction

- Identification of concepts and objects does not have to be easy.

Data extraction

- Identification of concepts and objects does not have to be easy.
- **Example:** Extract information about sales from the source system.

Data extraction

- Identification of concepts and objects does not have to be easy.
- **Example:** Extract information about sales from the source system.
 - ▶ What is meant by the term **sale**? A sale has occurred when

Data extraction

- Identification of concepts and objects does not have to be easy.
- **Example:** Extract information about sales from the source system.
 - ▶ What is meant by the term **sale**? A sale has occurred when
 - ① the order has been received by a customer,

Data extraction

- Identification of concepts and objects does not have to be easy.
- **Example:** Extract information about sales from the source system.
 - ▶ What is meant by the term **sale**? A sale has occurred when
 - 1 the order has been received by a customer,
 - 2 the order is sent to the customer,

Data extraction

- Identification of concepts and objects does not have to be easy.
- **Example:** Extract information about sales from the source system.
 - ▶ What is meant by the term **sale**? A sale has occurred when
 - ① the order has been received by a customer,
 - ② the order is sent to the customer,
 - ③ the invoice has been raised against the order.

Data extraction

- Identification of concepts and objects does not have to be easy.
- **Example:** Extract information about sales from the source system.
 - ▶ What is meant by the term **sale**? A sale has occurred when
 - ① the order has been received by a customer,
 - ② the order is sent to the customer,
 - ③ the invoice has been raised against the order.
 - ▶ It is a common problem that there is no table SALES in the operational databases; some other tables can exist like ORDER with an attribute ORDER_STATUS.

Conflicts and dirty data

- Different logical models of operational sources,

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),
- Different field lengths (address stored by using 20 or 50 chars),

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),
- Different field lengths (address stored by using 20 or 50 chars),
- Different naming conventions: homonyms and synonyms,

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),
- Different field lengths (address stored by using 20 or 50 chars),
- Different naming conventions: homonyms and synonyms,
- Inconsistent information concerning the same object,

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),
- Different field lengths (address stored by using 20 or 50 chars),
- Different naming conventions: homonyms and synonyms,
- Inconsistent information concerning the same object,
- Information concerning the same object, but indicated by different keys,

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),
- Different field lengths (address stored by using 20 or 50 chars),
- Different naming conventions: homonyms and synonyms,
- Inconsistent information concerning the same object,
- Information concerning the same object, but indicated by different keys,
- Missing values,

Conflicts and dirty data

- Different logical models of operational sources,
- Different data types (account number stored as **String** or **Numeric**),
- Different data domains (gender: **M**, **F**, **male**, **female**, 1, 0),
- Different date formats (dd-mm-yyyy or mm-dd-yyyy),
- Different field lengths (address stored by using 20 or 50 chars),
- Different naming conventions: homonyms and synonyms,
- Inconsistent information concerning the same object,
- Information concerning the same object, but indicated by different keys,
- Missing values,
- ...

Deduplication and householding

- **Deduplication** ensures that one accurate record exists for each business entity represented in a database,

Deduplication and householding

- **Deduplication** ensures that one accurate record exists for each business entity represented in a database,
- **Householding** is the technique of grouping individual customers by the household or organization of which they are a member; this technique has some interesting marketing implications, and can also support cost-saving measures of direct advertising.

Deduplication and householding

- **Deduplication** ensures that one accurate record exists for each business entity represented in a database,
- **Householding** is the technique of grouping individual customers by the household or organization of which they are a member; this technique has some interesting marketing implications, and can also support cost-saving measures of direct advertising.
- **Example:**

Deduplication and householding

- **Deduplication** ensures that one accurate record exists for each business entity represented in a database,
- **Householding** is the technique of grouping individual customers by the household or organization of which they are a member; this technique has some interesting marketing implications, and can also support cost-saving measures of direct advertising.
- **Example:**
 - ▶ Consider the following rows in a database:

Tim Jones	123	Main Street	Marlboro	MA	12234
T. Jones	123	Main St.	Marlborough	MA	12234
Timothy Jones	321	Maine Street	Marlborog	AM	12234
Jones, Timothy	123	Maine Ave	Marlborough	MA	13324

Deduplication and householding

- **Deduplication** ensures that one accurate record exists for each business entity represented in a database,
- **Householding** is the technique of grouping individual customers by the household or organization of which they are a member; this technique has some interesting marketing implications, and can also support cost-saving measures of direct advertising.
- **Example:**
 - ▶ Consider the following rows in a database:

Tim Jones	123	Main Street	Marlboro	MA	12234
T. Jones	123	Main St.	Marlborough	MA	12234
Timothy Jones	321	Maine Street	Marlborog	AM	12234
Jones, Timothy	123	Maine Ave	Marlborough	MA	13324

- ▶ The sales for around \$500 are counted for each tuple.

Deduplication and householding

- **Deduplication** ensures that one accurate record exists for each business entity represented in a database,
- **Householding** is the technique of grouping individual customers by the household or organization of which they are a member; this technique has some interesting marketing implications, and can also support cost-saving measures of direct advertising.
- **Example:**
 - ▶ Consider the following rows in a database:

Tim Jones	123	Main Street	Marlboro	MA	12234
T. Jones	123	Main St.	Marlborough	MA	12234
Timothy Jones	321	Maine Street	Marlborog	AM	12234
Jones, Timothy	123	Maine Ave	Marlborough	MA	13324

- ▶ The sales for around \$500 are counted for each tuple.
- ▶ Is it the same person?

Load of data

- After extracting, cleaning and transforming, data must be loaded into the warehouse.

Load of data

- After extracting, cleaning and transforming, data must be loaded into the warehouse.
- Loading the warehouse includes some other processing tasks: checking integrity constraints, sorting, summarizing, creating indexes, etc.

Load of data

- After extracting, cleaning and transforming, data must be loaded into the warehouse.
- Loading the warehouse includes some other processing tasks: checking integrity constraints, sorting, summarizing, creating indexes, etc.
- Batch (bulk) load utilities are used for loading.

Load of data

- After extracting, cleaning and transforming, data must be loaded into the warehouse.
- Loading the warehouse includes some other processing tasks: checking integrity constraints, sorting, summarizing, creating indexes, etc.
- Batch (bulk) load utilities are used for loading.
- A load utility must allow the administrator to monitor status, to cancel, suspend, and resume a load, and to restart after failure with no loss of data integrity.

Data warehouse refreshment

- Refreshing a warehouse means propagating updates on source data to the data stored in the warehouse.

Data warehouse refreshment

- Refreshing a warehouse means propagating updates on source data to the data stored in the warehouse.
- Follows the same structure as ETL process.

Data warehouse refreshment

- Refreshing a warehouse means propagating updates on source data to the data stored in the warehouse.
- Follows the same structure as ETL process.
- Several constraints: accessibility of data sources, size of data, size of data warehouse, frequency of data refreshing, degradation of performance of operational systems.

Data warehouse refreshment

- Detect changes in external data sources:

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources
- Extract the changes and integrate into the warehouse.

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources
- Extract the changes and integrate into the warehouse.
- Update indexes, subaggregates and any other additional data structures.

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources
- Extract the changes and integrate into the warehouse.
- Update indexes, subaggregates and any other additional data structures.
- Types of refreshments:

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources
- Extract the changes and integrate into the warehouse.
- Update indexes, subaggregates and any other additional data structures.
- Types of refreshments:
 - ▶ Periodical refreshment (daily or weekly).

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources
- Extract the changes and integrate into the warehouse.
- Update indexes, subaggregates and any other additional data structures.
- Types of refreshments:
 - ▶ Periodical refreshment (daily or weekly).
 - ▶ Immediate refreshment.

Data warehouse refreshment

- Detect changes in external data sources:
 - ▶ Different monitoring techniques: external and intrusive techniques.
 - ▶ Snapshot vs. timestamped sources
 - ▶ Queryable, logged, and replicated sources
 - ▶ Callback and internal action sources
- Extract the changes and integrate into the warehouse.
- Update indexes, subaggregates and any other additional data structures.
- Types of refreshments:
 - ▶ Periodical refreshment (daily or weekly).
 - ▶ Immediate refreshment.
 - ▶ Determined by usage, types of data source, etc.

Outline

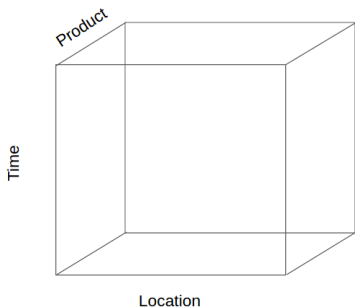
- 1 Motivation
- 2 ETL
- 3 OLAP Systems**
- 4 Summary

OLAP systems

- The next step is to provide solutions for querying and reporting multidimensional analytical data.

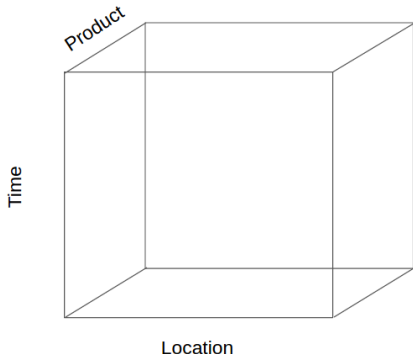
Multidimensional cube

- The proper data model for multidimensional reporting is the multidimensional one.



Operations in multidimensional data model

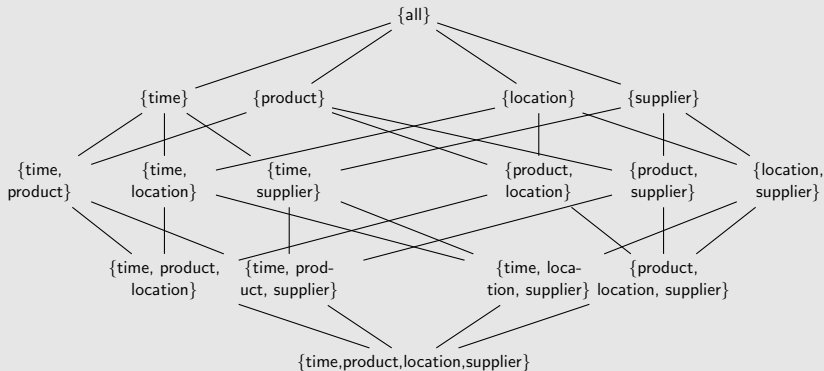
- Roll up – summarize data along a dimension hierarchy.
- Drill down – go from higher level summary to lower level summary or detailed data.
- Slice and dice – corresponds to selection and projection.
- Pivot – reorient cube.
- Raking, Time functions, etc.



Lattice of cuboids

- Different degrees of summarizations are presented as a lattice of cuboids.

Example for dimensions: time, product, location, supplier



Using this structure, one can easily show roll up and drill down operations.

Total number of cuboids

- For an n -dimensional data cube, the total number of cuboids that can be generated is:

$$T = \prod_{i=1}^n (L_i + 1),$$

where L_i is the number of levels associated with dimension i (excluding the virtual top level "all" since generalizing to "all" is equivalent to the removal of a dimension).

- For example, if the cube has 10 dimensions and each dimension has 4 levels, the total number of cuboids that can be generated will be:

$$T = 5^{10} = 9,8 \times 10^6.$$

Total number of cuboids

- **Example:** Consider a simple database with two dimensions:

Total number of cuboids

- **Example:** Consider a simple database with two dimensions:
 - ▶ Columns in Date dimension: day, month, year
 - ▶ Columns in Localization dimension: street, city, country.
 - ▶ Without any information about hierarchies, the number of all possible group-bys is

Total number of cuboids

- **Example:** Consider a simple database with two dimensions:
 - ▶ Columns in Date dimension: day, month, year
 - ▶ Columns in Localization dimension: street, city, country.
 - ▶ Without any information about hierarchies, the number of all possible group-bys is 2^6 :

Total number of cuboids

- **Example:** Consider a simple database with two dimensions:
 - ▶ Columns in Date dimension: day, month, year
 - ▶ Columns in Localization dimension: street, city, country.
 - ▶ Without any information about hierarchies, the number of all possible group-bys is 2^6 :

\emptyset		\emptyset
day		street
month		city
year		country
day, month	⋈	street, city
day, year		street, country
month, year		city, country
day, month, year		street, city, country

Total number of cuboids

- **Example:** Consider the same relations but with defined hierarchies:

Total number of cuboids

- **Example:** Consider the same relations but with defined hierarchies:
 - ▶ `day` → `month` → `year`
 - ▶ `street` → `city` → `country`

Total number of cuboids


- **Example:** Consider the same relations but with defined hierarchies:
 - ▶ `day` → `month` → `year`
 - ▶ `street` → `city` → `country`
 - ▶ Many combinations of columns can be excluded, e.g., group by `day`, `year`, `street`, `country`.
 - ▶ The number of group-bys is then

Total number of cuboids

- **Example:** Consider the same relations but with defined hierarchies:
 - ▶ `day` → `month` → `year`
 - ▶ `street` → `city` → `country`
 - ▶ Many combinations of columns can be excluded, e.g., group by `day`, `year`, `street`, `country`.
 - ▶ The number of group-bys is then 4^2 :

Total number of cuboids

- **Example:** Consider the same relations but with defined hierarchies:
 - ▶ $\text{day} \rightarrow \text{month} \rightarrow \text{year}$
 - ▶ $\text{street} \rightarrow \text{city} \rightarrow \text{country}$
 - ▶ Many combinations of columns can be excluded, e.g., group by day, year, street, country.
 - ▶ The number of group-bys is then 4^2 :



Three types of aggregate functions

- distributive: `count()`, `sum()`, `max()`, `min()`,
- algebraic: `ave()`, `stddev()`,
- holistic: `median()`, `mode()`, `rank()`.

OLAP servers

- Relational OLAP (ROLAP),
- Multidimensional OLAP (MOLAP),
- Hybrid OLAP (HOLAP).

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:
 - ▶ Denormalization,

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:
 - ▶ Denormalization,
 - ▶ Materialized views,

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:
 - ▶ Denormalization,
 - ▶ Materialized views,
 - ▶ Partitioning,

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:
 - ▶ Denormalization,
 - ▶ Materialized views,
 - ▶ Partitioning,
 - ▶ Joins,

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:
 - ▶ Denormalization,
 - ▶ Materialized views,
 - ▶ Partitioning,
 - ▶ Joins,
 - ▶ Indexes (join index, bitmaps),

ROLAP

- **ROLAP servers** use a relational or post-relational database management system to store and manage warehouse data.
- ROLAP systems use SQL and its OLAP extensions.
- Optimization techniques:
 - ▶ Denormalization,
 - ▶ Materialized views,
 - ▶ Partitioning,
 - ▶ Joins,
 - ▶ Indexes (join index, bitmaps),
 - ▶ Query processing.

ROLAP

- Advantages of ROLAP Servers:

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,
 - ▶ Scalable with respect to the size of data,

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,
 - ▶ Scalable with respect to the size of data,
 - ▶ Sparsity is not a problem (fact tables contain only facts),

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,
 - ▶ Scalable with respect to the size of data,
 - ▶ Sparsity is not a problem (fact tables contain only facts),
 - ▶ Mature and well-developed technology.

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,
 - ▶ Scalable with respect to the size of data,
 - ▶ Sparsity is not a problem (fact tables contain only facts),
 - ▶ Mature and well-developed technology.
- Disadvantage of ROLAP Servers:

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,
 - ▶ Scalable with respect to the size of data,
 - ▶ Sparsity is not a problem (fact tables contain only facts),
 - ▶ Mature and well-developed technology.
- Disadvantage of ROLAP Servers:
 - ▶ Worse performance than MOLAP,

ROLAP

- Advantages of ROLAP Servers:
 - ▶ Scalable with respect to the number of dimensions,
 - ▶ Scalable with respect to the size of data,
 - ▶ Sparsity is not a problem (fact tables contain only facts),
 - ▶ Mature and well-developed technology.
- Disadvantage of ROLAP Servers:
 - ▶ Worse performance than MOLAP,
 - ▶ Additional data structures and optimization techniques used to improve the performance.

MOLAP

- MOLAP Servers use array-based multidimensional storage engines.

MOLAP

- MOLAP Servers use array-based multidimensional storage engines.
- Optimization techniques:

MOLAP

- MOLAP Servers use array-based multidimensional storage engines.
- Optimization techniques:
 - ▶ Two-level storage representation: dense cubes are identified and stored as array structures, sparse cubes employ compression techniques,

MOLAP

- MOLAP Servers use array-based multidimensional storage engines.
- Optimization techniques:
 - ▶ Two-level storage representation: dense cubes are identified and stored as array structures, sparse cubes employ compression techniques,
 - ▶ Materialized cubes.

MOLAP

- Advantages of MOLAP Servers:

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:
 - ▶ Scalability problem in the case of larger number of dimensions,

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:
 - ▶ Scalability problem in the case of larger number of dimensions,
 - ▶ Not tailored for sparse data.

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:
 - ▶ Scalability problem in the case of larger number of dimensions,
 - ▶ Not tailored for sparse data.
 - ▶ **Example:**

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:
 - ▶ Scalability problem in the case of larger number of dimensions,
 - ▶ Not tailored for sparse data.
 - ▶ **Example:**
 - Logical model consists of four dimensions: customer, product, location, and day

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:
 - ▶ Scalability problem in the case of larger number of dimensions,
 - ▶ Not tailored for sparse data.
 - ▶ **Example:**
 - Logical model consists of four dimensions: customer, product, location, and day
 - In case of 100 000 customers, 10 000 products, 1 000 locations and 1 000 days, the data cube will contain 1 000 000 000 000 000 cells!

MOLAP

- Advantages of MOLAP Servers:
 - ▶ Multidimensional views are directly mapped to data cube array structures – efficient access to data,
 - ▶ Can easily store subaggregates.
- Disadvantages of MOLAP Servers:
 - ▶ Scalability problem in the case of larger number of dimensions,
 - ▶ Not tailored for sparse data.
 - ▶ **Example:**
 - Logical model consists of four dimensions: customer, product, location, and day
 - In case of 100 000 customers, 10 000 products, 1 000 locations and 1 000 days, the data cube will contain 1 000 000 000 000 000 cells!
 - A huge number of cells is empty: a customer is not able to buy all products in all locations . . .

HOLAP

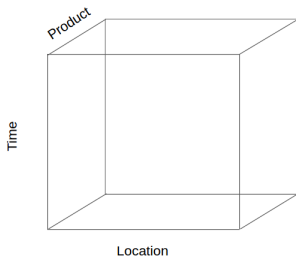
- HOLAP servers are a hybrid approach that combines ROLAP and MOLAP technology.
- HOLAP benefits from the greater scalability of ROLAP and the faster computation of MOLAP.

Multidimensional queries

- We need an intuitive way of expressing analytical (multidimensional) queries:

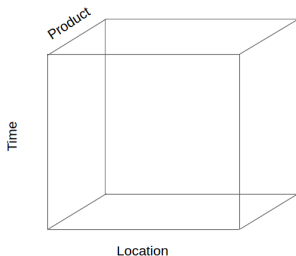
Multidimensional queries

- We need an intuitive way of expressing analytical (multidimensional) queries:
 - ▶ Operations like roll up, drill down, slice and dice, pivoting, ranking, time and window functions, etc.



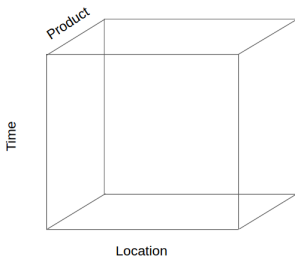
Multidimensional queries

- We need an intuitive way of expressing analytical (multidimensional) queries:
 - ▶ Operations like roll up, drill down, slice and dice, pivoting, ranking, time and window functions, etc.
- Two solutions:



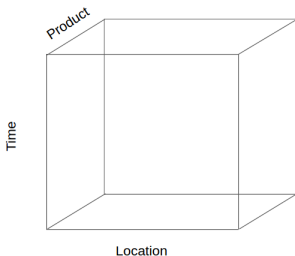
Multidimensional queries

- We need an intuitive way of expressing analytical (multidimensional) queries:
 - ▶ Operations like roll up, drill down, slice and dice, pivoting, ranking, time and window functions, etc.
- Two solutions:
 - ▶ Extending **SQL**, or



Multidimensional queries

- We need an intuitive way of expressing analytical (multidimensional) queries:
 - ▶ Operations like roll up, drill down, slice and dice, pivoting, ranking, time and window functions, etc.
- Two solutions:
 - ▶ Extending **SQL**, or
 - ▶ Inventing a new language (→ **MDX**),



- OLAP extensions in SQL:
 - ▶ GROUP BY CUBE,
 - ▶ GROUP BY ROLLUP,
 - ▶ GROUP BY GROUPING SETS,
 - ▶ GROUPING and DECODE/CASE
 - ▶ OVER and PARTITION BY,
 - ▶ RANK.

MDX

- MDX \longrightarrow Multidimensional expressions.

MDX

- MDX → Multidimensional expressions.
- For OLAP queries, MDX is an alternative to SQL:

MDX

- MDX → Multidimensional expressions.
- For OLAP queries, MDX is an alternative to SQL:

MDX

- MDX → Multidimensional expressions.
- For OLAP queries, MDX is an alternative to SQL:

Academic_year	Instructor	AVG(Grade)
2011/2	Stefanowski	4.2
2011/2	Słowiński	4.1
2012/3	Stefanowski	4.0
2012/3	Słowiński	3.8
2013/4	Stefanowski	3.9
2013/4	Słowiński	3.6
2013/4	Dembczyński	4.8

MDX

- MDX → Multidimensional expressions.
- For OLAP queries, MDX is an alternative to SQL:

Academic_year	Instructor	AVG(Grade)
2011/2	Stefanowski	4.2
2011/2	Słowiński	4.1
2012/3	Stefanowski	4.0
2012/3	Słowiński	3.8
2013/4	Stefanowski	3.9
2013/4	Słowiński	3.6
2013/4	Dembczyński	4.8



AVG(Grade) Name	Academic_year		
	2011/2	2012/3	2013/4
Stefanowski	4.2	4.0	3.9
Słowiński	4.1	3.8	3.6
Dembczyński			4.8

MDX

- MDX query:

```
SELECT {[Time].[1997],[Time].[1998]} ON COLUMNS,  
{[Measures].[Sales],[Measures].[Cost]} ON ROWS  
FROM Warehouse  
WHERE ([Store].[All].[USA])
```

- Seems to be similar to **SQL**, but in fact it is quite **different!**

Outline

- ① Motivation
- ② ETL
- ③ OLAP Systems
- ④ Summary

Summary

- ETL process is a strategic element of data warehousing.
- Main concepts: extraction, transformation and integration, load, data warehouse refreshment and metadata.
- New emerging technology . . .
- OLAP systems: ROLAP, MOLAP and HOLAP.

Bibliography

- J. Han and M. Kamber. *Data Mining: Concepts and Techniques (second edition)*. Morgan Kaufmann Publishers, 2006