
Advanced Database Technology
Anna Östlin Pagh and Rasmus Pagh
IT University of Copenhagen
Spring 2005

**Introduction, practical information,
and recap of database background**

January 31, 2005

Based on
selected parts of chapters 1-10 in GUW
(not curriculum)

— Today's lecture (2-3 hours) —

- Course introduction
- Practical information
- Recap of relevant database background
- Overview of lectures

— What is a database? —

According to Webster's dictionary:

da·ta·base

a usually large collection of data organized especially for rapid search and retrieval (as by a computer)

Remark:

The need for (and the ability to give) rapid answers to a multitude of **queries** about data is increasing. Databases have thus grown to perform much more advanced processing than search and retrieval.

— What is a database management system? —

Database management system (DBMS):

Software system used when implementing databases

more precisely

System for providing **efficient**, **convenient**, and **safe** storage of and **multi-user** access to, possibly **massive**, amounts of **persistent** data.

Problem session: (5 minutes, discuss in groups of 2-4 students)

In this course we focus on massive amounts of data. Think about examples of large data sets for which there is a need for some kind of DBMS (existing or still not existing).

— What the course is about —

The two main goals of the course:

- To understand how relational database systems work and what influence their performance. Focus is on large data sets.
- Provide efficient solutions when classical relational databases don't do it, e.g. when there is text data or geometric data.

— Why this course is interesting/important —

There is a growing need for people that are able to implement efficient methods for handling massive data sets:

- Most common tasks in databases can be done fast when the amount of data is small, while it may be very time consuming for large data sets.
- In many areas the amount of data grows faster than the size of internal memory, e.g. biological data, internet pages. This means that the data to process can not fit into internal memory.

— Focus on algorithmic aspects —

The focus will be on algorithmic aspects of database implementation:

- Efficient data structures for indexes
- Efficient implementation of relational operations
- Models of computation when data is stored on secondary memory

Internal vs. external memory:

- Time to retrieve one word from internal memory: ≈ 100 ns
- Time to retrieve one word from external memory: ≈ 10 ms (= 10^7 ns)
- Time to perform 100 operations: ≈ 50 ns (or even less)

Conclusion: When data must be stored in external memory, the time to retrieve data from disk is most often the bottleneck.

— Database course overview —

Database Systems

Course book: *Modern Database Management*.

Assumes: Introductory programming.

Introduction to Databases

Course book: *Database systems – the complete book* (GUW).

Assumes: Nothing.

Advanced Database Technology

Course book: *Database systems – the complete book*.

Assumes: Introduction to algorithms and an introductory database course.

— Course format —

Lectures and problem sessions: (Mondays 10.00 to around 14-15)

Mix of lectures and short problem sessions without preparation.

Exercises: (about an hour, typically after the lecture)

You are **not** expected to prepare before the exercises. This is just a big problem session, that give you an oportunity to think about the material in the lecture and discuss it with the other students.

Hand-ins: **Satisfactory hand-ins required to enter exam.**

Exam: Written, 4 hours, on June 6, 2005.

— About the mandatory hand-ins —

There will be hand-ins due at 10.00 on 5 Mondays.

Hand-ins must be completed in groups of size at most 2. You are allowed to discuss hand-ins with fellow students, but you must understand and prepare your own solution, within the group.

Every assignment is graded on a scale from A-E (A=very good, B=good, C=acceptable, D=not good enough, E=not handed in/late hand-in), and students must achieve A, B, or C on at least 4 out of 5 assignments.

If you are not able to hand in in time, for some reason, inform the teacher as soon as possible.

— Manning of course and course homepage —

Lecturers:

Rasmus Pagh, pagh@itu.dk, office 3C.07.

Anna Östlin Pagh, annao@itu.dk, office 3C.11.

Homepage:

www.itu.dk/people/pagh/ADBT05/

- News.
- Reading directions for each lecture.
- Lecture slides.
- Problems for hand-ins.

— What we expect from you —

- Basic course in databases, e.g.,
 - Relational data model/ Relational databases
 - Relational algebra
 - SQL
- Basic course in algorithms and data structures, e.g.,
 - Search trees
 - Sorting algorithms
 - Hashing
 - Big-O notation
 - Basic algorithm analysis

— Recap part of the lecture —

- Relations and SQL
 - Basic concepts in relational data model, like attribute, schemas, keys etc.
 - Define relational algebra
 - Basic operations, like set operations, joins, selection etc.
 - Bags (multisets)
 - More operations, e.g., duplicate removal, grouping
- DBMS
 - What it is
 - Indexes
 - Transactions

— What is a relational database? —

All major general purpose DBMSs are based on the so-called **relational data model**. This means that all data is stored in a number of tables (with named columns), such as:

<i>accountNo</i>	<i>balance</i>	<i>type</i>
12345	1000.00	savings
67890	2846.92	checking
32178	-3210.00	loan
...

For historical, mathematical reasons such tables are referred to as **relations**.

— The relational data model —

- The model behind most implementations of databases.
- Data is always represented as relations (2-dim. tables).
- SQL is a query language for relational databases and is based on relational algebra.

Some basic concepts:

- Relation
- Attribute
- Schema
- Tuple
- Key

How many of these do you recognize?

How many of these are you able to define now?

Relation

The way to represent data is through relations.

A **relation** is a two-dimensional table.

The order of rows and columns can be exchanged, and it is still the same relation.

Example:

<i>title</i>	<i>year</i>	<i>length</i>	<i>film Type</i>
Star Wars	1977	124	color
Mighty Ducks	1991	104	color
Wayne's World	1992	95	color

Attribute

An **attribute** is the name of a column in a relation. It usually describes the meaning of the content in the column.

Example:

<i>title</i>	<i>year</i>	<i>length</i>	<i>filmType</i>
Star Wars	1977	124	color
Mighty Ducks	1991	104	color
Wayne's World	1992	95	color

Schema

A **schema** is a description of a class of relations. It consists of the name of the relation and the set of attributes in the relation.

That it is a **set** of attributes means that the attributes are unordered.

Example:

<u>Movies</u>			
<i>title</i>	<i>year</i>	<i>length</i>	<i>filmType</i>
Star Wars	1977	124	color
Mighty Ducks	1991	104	color
Wayne's World	1992	95	color

Schema for the above relation:

Movies(title, year, length, filmType)

A set of schemas for relations are called a **design** or a **database schema**.

Tuple

A **tuple** is a row in a table. The values in the row are called components.

A relation can be seen as a set of tuples.

Example:

<u>Movies</u>			
<i>title</i>	<i>year</i>	<i>length</i>	<i>filmType</i>
Star Wars	1977	124	color
Mighty Ducks	1991	104	color
Wayne's World	1992	95	color

Keys of a relation

A **key** for a relation is a set of its attributes that satisfy:

- **Uniqueness.** The values of the attributes uniquely identify a tuple.
- **Minimality.** No proper subset of the attributes has the uniqueness property.

If uniqueness is satisfied (but not necessarily minimality) the attributes are said to form a **superkey**.

<i>title</i>	<i>year</i>	<i>length</i>	<i>filmType</i>	<i>studioName</i>	<i>starName</i>
Star Wars	1977	124	color	Fox	Carrie Fisher
Star Wars	1977	124	color	Fox	Mark Hamill
Star Wars	1977	124	color	Fox	Harrison Ford
Mighty Ducks	1991	104	color	Disney	Emilio Estevez

Key: { *title*, *year*, *starName* } **Superkey:** { *title*, *year*, *length*, *starName* }

— Relational algebra and SQL examples —

Relational algebra is notation for operations on relations, like constructing new relations and defining queries on relations.

The rest of the lecture is about:

- Basic operations in **relational algebra**:
 - **set operations** (e.g. union)
 - **selection** and **projection**
 - **join**
- **Bags** (multisets). Why they are used and what the consequence is.
- More operations, e.g., duplicate removal, grouping
- **Indexes**
- **Transactions**

How many of the above red words do you recognize?

How many of them are you able to define now?

Set operations

Operations on two sets R and S , where R and S must have the same set of attributes. We have the three set operations:

- Union, $R \cup S$,
- Intersection, $R \cap S$, and
- Difference, $R \setminus S$.

Example:

<u>Movies1</u>			<u>Movies2</u>		
<i>title</i>	<i>year</i>	<i>filmType</i>	<i>title</i>	<i>year</i>	<i>filmType</i>
Star Wars	1977	color	Star Wars	1999	color
Mighty Ducks	1991	color	Mighty Ducks	1991	color
Wayne's World	1992	color	Star Wars	2002	color

Projection

A **projection** of relation R on attributes A_1, \dots, A_n is denoted by

$$\pi_{A_1, \dots, A_n}(R)$$

and is the relation R restricted to columns for attributes A_1, \dots, A_n .

<i>title</i>	<i>year</i>	<i>length</i>	<i>filmType</i>	<i>studioName</i>	<i>starName</i>
Star Wars	1977	124	color	Fox	Carrie Fisher
Star Wars	1977	124	color	Fox	Mark Hamill
Star Wars	1977	124	color	Fox	Harrison Ford
Mighty Ducks	1991	104	color	Disney	Emilio Estevez

	<i>title</i>	<i>length</i>	<i>studioName</i>
$\pi_{title, length, studioName}(Movies) =$	Star Wars	124	Fox
	Mighty Ducks	104	Disney

Selection

A **selection** of tuples satisfying condition C from relation R is denoted by

$$\sigma_C(R)$$

and is the relation R restricted to tuples for which condition C is satisfied.

C can be any boolean expression, i.e. it may involve multiple attributes, constants, AND, OR, and NOT.

Example:

Movies		
<i>title</i>	<i>year</i>	<i>filmType</i>
Star Wars	1977	color
Mighty Ducks	1991	color
Wayne's World	1992	color

$\sigma_{year > 1981}(Movies) =$

<i>title</i>	<i>year</i>	<i>filmType</i>
Mighty Ducks	1991	color
Wayne's World	1992	color

— Join (natural) —

Natural-Join or Inner-Join:

Let R and S be two relations with attributes R_1, \dots, R_n and S_1, \dots, S_m respectively. The join of relations R and S , denoted

$$R \bowtie S$$

has attributes $\{R_1, \dots, R_n\} \cup \{S_1, \dots, S_m\}$.

If $r \in R$ and $s \in S$ agree on attributes $\{R_1, \dots, R_n\} \cap \{S_1, \dots, S_m\}$ then the joint tuple for r and s is in $R \bowtie S$.

There are other types of join, e.g., Theta-Join and Outer-Join.

Join example

<u>Movies</u>				<u>Stars in</u>		
<i>title</i>	<i>year</i>	<i>length</i>	<i>studioN.</i>	<i>title</i>	<i>year</i>	<i>starN.</i>
Star Wars	1977	124	Fox	Star Wars	1977	Carrie F.
Mighty D.	1991	104	Disney	Star Wars	1977	Mark H.
Wayne's W.	1992	95	Param.	Star Wars	1977	Harrison F.
				Mighty D.	1991	Emilio E.
				Wayne's W.	1992	Dana C.
				Wayne's W.	1992	Mike M.

Movies ⋈ Stars in =

<i>title</i>	<i>year</i>	<i>length</i>	<i>studioN.</i>	<i>starN.</i>
Star Wars	1977	124	Fox	Carrie F.
Star Wars	1977	124	Fox	Mark H.
Star Wars	1977	124	Fox	Harrison F.
Mighty D.	1991	104	Disney	Emilio E.
Wayne's W.	1992	95	Param.	Dana C.
Wayne's W.	1992	95	Param.	Mike M.

Bags

Relational algebra is an algebra on sets, but most database systems do not only use sets, they also use bags.

A **bag** or a multiset, is a set where elements may appear more than once. (In a relation there may be two or more identical rows.)

The motivation for using bags instead of sets is that some operations can be implemented faster. E.g.,

- union
- projection

Operations on bags vs. sets

Some examples of the difference between operations on bags and sets.

- $R \cup S$: All rows in R and S , even if they appear in both or if they appear more than once in R or in S .
- $R \cap S$: if tuple t appears n times in R and m times in S , then it appears $\min(n, m)$ times in $R \cap S$.
- $\pi_{A_1, \dots, A_n}(R)$ (projection): All tuples in R also appear in $\pi_{A_1, \dots, A_n}(R)$, even if the rows become identical when some columns are removed.

<u>Movies1</u>			<u>Movies2</u>		
<i>title</i>	<i>year</i>	<i>filmType</i>	<i>title</i>	<i>year</i>	<i>filmType</i>
Star Wars	1977	color	Star Wars	1999	color
Mighty Ducks	1991	color	Mighty Ducks	1991	color
Wayne's World	1992	color	Star Wars	2002	color

— More operations —

Other useful relational operations often used in languages like SQL:

- Duplicate elimination: When bags are used it is useful to be able to get rid of duplicates. $\delta(R)$
- Aggregation operators: E.g., sum, average, maximum in a column. $\gamma_{OP(A)}(R)$, where OP is e.g., max.
- Grouping: Divide a relation up into groups of tuples depending on the values in one or more attributes. Used together with aggregation. $\gamma_{A_1, \dots, A_n, OP(A)}(R)$.
- Extended projection: Creation of new columns from existing columns by performing some kind of computation.

SQL

SQL is a language that can be used for expressing queries on relations. It is based on a “mixture” of relational algebra for sets and bags.

Some SQL examples:

- `SELECT A_1, \dots, A_n FROM R` means $\pi_{A_1, \dots, A_n}(R)$.
- `SELECT * FROM R WHERE C` means $\sigma_C(R)$.
- `R UNION S` means $R \cup S$.
- `R EXCEPT S` means $R \setminus S$.
- `R NATURAL JOIN S` means $R \bowtie S$.
- `SELECT DISTINCT * FROM R` means $\delta(R)$.
- `SELECT $OP(B)$ FROM R GROUP BY A` means $\gamma_{A, OP(B)}(R)$

— SQL update operations —

SQL also supports the creation and modification of relations.

Some SQL examples:

- `CREATE TABLE R (<schema description>)`
- `INSERT INTO R VALUES (v_1, \dots, v_n).`
- `DELETE FROM R WHERE C .`
- `UPDATE R SET $A = v$ WHERE C .`

Selective queries

Consider the selection query:

```
SELECT *  
FROM R  
WHERE <condition>
```

- If we have to report 80% of the tuples in R, it makes sense to do a full table scan.
- On the other hand, if the query is very **selective**, and returns just a small percentage of the tuples we might hope to do better, by using an **index**.

Indexes

To be able to quickly find the first tuple with a specific value for an attribute, the DBMS may build an **index** on that attribute.

A database index is similar to an index in the back of a book:

1. For every piece of data you might be interested in (e.g., the attribute year=1977), the index says where to find it.
2. The index itself is organized such that one can quickly do the lookup.

Indexes

Some indexes are efficient for both **point queries** ($\text{year}=1977$) and **range queries** ($1985 < \text{year} < 1999$), while others only support efficient point queries.

Indexes are also used by the DBMS to speed up other operations, e.g., **join operations** are often considerably faster when the join attributes are indexed.

In most DBMSs we can specify what indexes should be created, e.g.:

- `CREATE INDEX I ON R(A)`

— Transactions

One or more updates in a database can be grouped into something called a **transaction**. This is a way to ensure correct updates of the database.

Ideal transactions are said to meet the **ACID** test:

- **A**tomicity – the all-or-nothing execution of transactions.
- **C**onsistency – transactions preserve database constraints.
- **I**solation – the appearance that transactions are executed one by one.
- **D**urability – the effect of a transaction is never lost once it has completed.

Commercial DBMSs tend to fully implement **C** and **D**, while **A** and **I** are only partially implemented (for efficiency reasons).

— Summary of recap —

This part of the lecture was about:

- the relational data model
- relational algebra
- relational algebra on bags
- some examples in SQL
- Properties of DBMSs

These concepts will underlie much of the first part of the course.