

12. Set-Oriented DB-Interface

Theo Härder
www.haerder.de

Goals

- Derivation of concepts for translation of descriptive queries and their
- Optimization and access module creation

Main references:

Theo Härder, Erhard Rahm: Datenbanksysteme – Konzepte und Techniken der Implementierung, Springer, 2001, Chapter 12.

Mitschang, B.: Anfrageverarbeitung in Datenbanksystemen – Entwurfs- und Implementierungskonzepte, Reihe Datenbanksysteme, Vieweg-Verlag, 1995

Goetz Graefe: Query Evaluation Techniques for Large Databases, ACM Computing Surveys 25:2, June 1993, pp. 73-170.

Realization of Database Systems – SS 2011

Logical Data Structures

Characterization of the mapping

```
SELECT Emp.Eno, Dept.Name
FROM Dept, Emp, Skill
WHERE Emp.Job = 'Programmer' &
      Skill.Sno = Emp.Sno &
      Emp.Dno = Dept.Dno
```

Mapping functions

- | | | |
|--------------------------|-----|-------------------------------------|
| - views | <-> | base relations |
| - relational expressions | <-> | logical access paths |
| - record sets | <-> | single records, currency indicators |


FETCH Skill USING ...
 FETCH NEXT Emp ...
 FETCH OWNER WITHIN ...

Properties of the upper interface

- Access-path-independent (relational) data model
- All facts and relationships are represented by values
- Non-procedural (descriptive) query languages
- Access to record sets

Realization of DBS

- Descriptive DB languages
- Evaluation of DB statements**
- Host language embedding
- Query optimization
- Estimation of execution plans
- Creation of execution plans
- Code creation
- Ad-hoc queries



© 2011 AG DBIS

Examples of Descriptive SQL Queries

- **Simple query**

```
SELECT Eno, EName, Salary/12
FROM   Emp
WHERE  Job = W
      AND Bonus > Salary
```
- Replaced by**


```
DECLARE C1 CURSOR FOR SELECT Eno, EName, Salary/12
INTO :X, :Y, :Z
FROM   Emp
WHERE  Job = :W
      AND Bonus > Salary
```
- With operators**

```
OPEN C1
FETCH C1 INTO :X, :Y, :Z
CLOSE C1
```

12-3

Realization of DBS

- Descriptive DB languages
- Evaluation of DB statements**
- Host language embedding
- Query optimization
- Estimation of execution plans
- Creation of execution plans
- Code creation
- Ad-hoc queries



© 2011 AG DBIS

Evaluation of DB Statements

- **Processing steps for the evaluation of DB statements:**
 - 1. Lexical and syntactical analysis**
 - Creation of a query graph (QG) as reference structure for the subsequent compilation steps
 - Checking for correct syntax (parsing)
 - 2. Semantic analysis**
 - Checking for the existence and validity of referenced tables, views, and attributes
 - Replacement of views in the QG by their view definitions
 - Replacement of external by internal names (name resolution)
 - Conversion of external formats into internal representations
 - 3. Access control and integrity control**

should already be done for performance reasons, if possible, at compile time

 - Access control requires generation of runtime actions in case of value dependencies
 - Execution of simple integrity controls (control of formats and conversion of data types)
 - Generation of runtime actions for more complex controls

12-4

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries



© 2011 AG DBIS

Evaluation of DB Statements (2)

4. Standardization and simplification

serve for more effective compilation and early error detection

- Transformation of the QG into a normal form
- Elimination of redundancies

5. Restructuring and transformation

- Restructuring aims at global improvement of the QG; transformation inserts executable operations
- Application of heuristic rules (**algebraic optimization** for QG restructuring)
- Transformation leads to replacement and aggregation of logical operators by plan operators (**non-algebraic optimization**): In most cases, there are several plan operators available for the implementation of a logical operator
- Determination of **alternative access plans**: Frequently, many execution sequences or access paths can be chosen
- Cost estimation and selection of the cheapest execution plans

→ Steps 4 + 5 summarized as query optimization

6. Code generation

- Generation of a tailor-made program for the given (SQL) statement
- Creation of an executable access module
- Management of access modules in a DBMS library

12-5

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries



© 2011 AG DBIS

Embedding of a Set-Oriented Interface

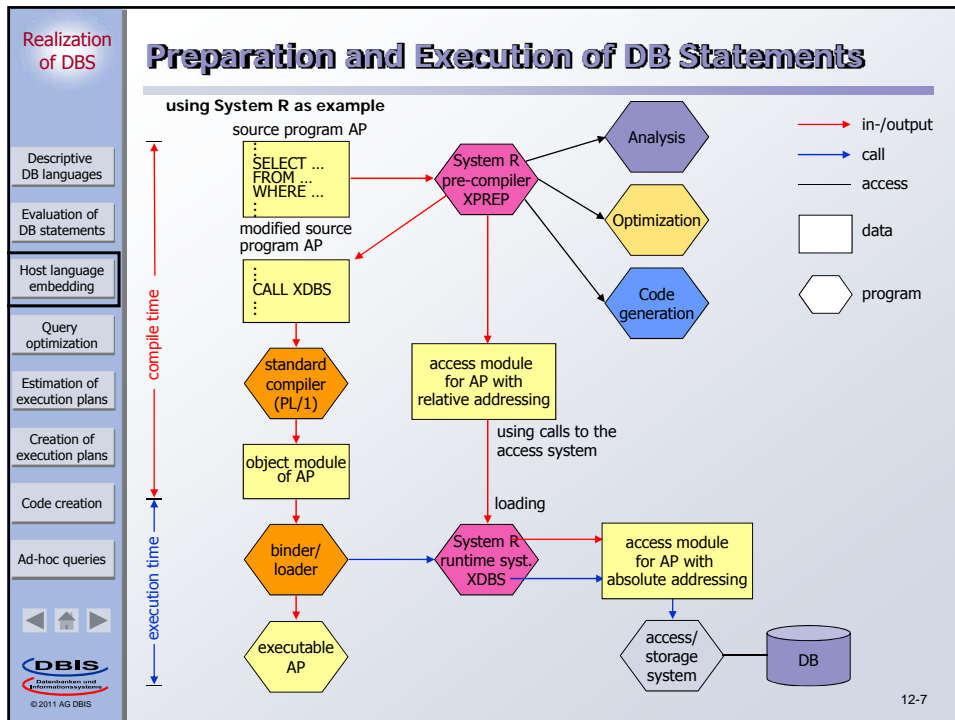
- Specification of the desired record set (qualification operator)**

SELECT Eno, EName, Salary/12 FROM Emp WHERE Job='Operator' AND Bonus>Salary	DECLARE C1 CURSOR FOR SELECT Eno, EName, Salary/12 FROM Emp WHERE Job=:W AND Bonus>Salary
---	---
- Successive provision of qualified records (fetch operator)**

OPEN C1; FETCH C1 INTO :X, :Y, :Z; CLOSE C1;	→ binding of W, e.g. 'Operator', activation of the cursor → fetching of a record → deactivation of the cursor
--	--
- Possible solution: replacement by pre-compiler**

DECLARE C1 ... OPEN C1 FETCH C1 INTO ...	→ comment → DCL T(3) POINTER; T(1) = ADDR(W) CALL XDBS (AM1, 2, OPEN, ADDR(T)) → T(1) = ADDR(X); T(2) = ADDR(Y); T(3) = ADDR(Z); CALL XDBS (AM1, 2, FETCH, ADDR(T));
--	---

12-6



Realization of DBS

- Descriptive DB languages
- Evaluation of DB statements
- Host language embedding
- Query optimization
- Estimation of execution plans
- Creation of execution plans
- Code creation
- Ad-hoc queries

Query Optimization*

- From query (what?) to evaluation (how?)**
 - Goal: cost-effective evaluation plan
- Use of a large number of methods and strategies**
 - Logical transformation of queries
 - Selection of access paths
 - Optimized storage of data on external memory
- Key problem**
 - Exact optimization is not 'computable', in general
 - Lack of accurate statistical information
 - Broad use heuristics (rules of thumb)
- Optimization goal**
 - either maximization of output with given resources
 - or minimization of resource usage for given output
 - throughput maximization?
 - response time minimization for given query language, mix of queries of different types and given system environment!

DBIS

© 2011 AG DBIS

* Jarke, M., Koch, J.: Query Optimization in Database Systems, in: ACM Comp. Surveys 16:2, 1984, pp. 111-152

12-8

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Query Optimization (2)

- Which costs are to be considered?
 - Communication cost (# of messages, set of data to be transmitted)
→ Distributed DBMS!
 - Computation cost (CPU cost, path lengths)
 - I/O cost (# of physical references)
 - Storage cost (temporary storage occupancy in DB buffer and on disk)

→ Variety of costs are not independent of one another

→ In centralized DBMS often "weighted function of computation- and I/O-costs"
- What is the best approach?

Step 1: After compilation, find appropriate internal representation for the query (QG)

Step 2: Apply logical restructuring to the query graph

Step 3: Map restructured query onto alternative sequences of plan operators (transformation) (→ set of execution plans)

Step 4: Compute cost estimates for each QEP and select the cheapest one

12-9

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Query Processing – Cost Issues

How are the costs divided up for transaction processing?
DB system : communication system : application

When do costs incur for the query processing?

Analysis

query
↓
syntactic analysis
semantic analysis (access- and integrity control)

run time
↓
without any preparation

Optimization

query graph
↓
standardization
simplification
query restructuring
query transformation

execution plan
↓
Code Generation
execution control

compilation time
↑
run time
↓
for maximal preparation

Code generation
Execution

query result
↓
allocation in program environment

Provision of result data

query result
↓
allocation in program environment

12-10

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries



© 2011 AG DBIS

Standardization of a Query

- **Standardization**
 - Choice of a normal form
 - E.g., conjunctive normal form
 - $(A_{11} \text{ OR } \dots \text{ OR } A_{1n}) \text{ AND } \dots \text{ AND } (A_{m1} \text{ OR } \dots \text{ OR } A_{mn})$
 - Displacement of quantors
- **Transformation rules for Boolean expressions**

Commutative rules

$A \text{ OR } B \Leftrightarrow B \text{ OR } A$

$A \text{ AND } B \Leftrightarrow B \text{ AND } A$

Associative rules

$(A \text{ OR } B) \text{ OR } C \Leftrightarrow A \text{ OR } (B \text{ OR } C)$

$(A \text{ AND } B) \text{ AND } C \Leftrightarrow A \text{ AND } (B \text{ AND } C)$

Distributive rules

$A \text{ OR } (B \text{ AND } C) \Leftrightarrow (A \text{ OR } B) \text{ AND } (A \text{ OR } C)$

$A \text{ AND } (B \text{ OR } C) \Leftrightarrow (A \text{ AND } B) \text{ OR } (A \text{ AND } C)$

De Morgan rules

$\text{NOT } (A \text{ AND } B) \Leftrightarrow \text{NOT } (A) \text{ OR } \text{NOT } (B)$

$\text{NOT } (A \text{ OR } B) \Leftrightarrow \text{NOT } (A) \text{ AND } \text{NOT } (B)$

Double negation rule

$\text{NOT } (\text{NOT } (A)) \Leftrightarrow A$
- **Itempotence rules for Boolean expressions**

A	OR	A	\Leftrightarrow	A
A	AND	A	\Leftrightarrow	A
A	OR	NOT (A)	\Leftrightarrow	TRUE
A	AND	NOT (A)	\Leftrightarrow	FALSE
A	AND	(A OR B)	\Leftrightarrow	A
A	OR	(A AND B)	\Leftrightarrow	A
A	OR	FALSE	\Leftrightarrow	A
A	OR	TRUE	\Leftrightarrow	TRUE
A	AND	FALSE	\Leftrightarrow	FALSE

12-11

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries



© 2011 AG DBIS

Simplification of a Query

- **Equivalent expressions**
 - Can have a varying degree of **redundancy**
 - Usage/elimination of common sub-expressions
 - $(A_1 = a_{11} \text{ OR } A_1 = a_{12}) \text{ AND } (A_1 = a_{12} \text{ OR } A_1 = a_{11})$
 - $(\text{Age} > 25 \text{ OR } (\text{Age} > 30 \text{ AND } \text{Job} = \text{'Programmer'}))$
 - Simplification of expressions referring to an "empty table"
- **Constant propagation**

$A \text{ op } B \text{ AND } B = \text{const.} \rightarrow A \text{ op const.}$
- **Non-satisfiable expressions**

$A \geq B \text{ AND } B > C \text{ AND } C \geq A \rightarrow A > A \rightarrow \text{false}$
- **Use of integrity constraints (IC)**
 - ICs are true for all records of the related table
 - A is primary key: $\pi_A \rightarrow$ no duplicate elimination required
 - Rule: Family-Status = 'married' AND Tax-Class ≥ 3
 - $\rightarrow \text{Expression: } (\text{Family-Status} = \text{'married'} \text{ AND } \text{Tax-Class} = 1) \rightarrow \text{false}$
- **Improvement of evaluation**
 - Adding of an IC to the WHERE clause does not change its truth value
 - Simpler evaluation structure, however more efficient heuristics needed

12-12

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries



© 2011 AG DBIS

Query Restructuring

- **Most important rules for restructuring and transformation**
 - **Early execution** of Selection (σ) and Projection (π) without duplicate elimination
 - Aggregation of unary operator sequences (as σ and π) to a single operation
 - Evaluation of equal subtrees in the QG only once
 - Binary operator sequences (as \cap , \cup , $-$, \times , \bowtie): minimization of intermediate results

→ **Selective operations (σ , π) before constructive operations (\times , \bowtie)**
- **Aggregation of operator sequences**
 - R1: $\pi_{A_n}(\dots \pi_{A_2}(\pi_{A_1}(\text{Tab}))\dots) \Leftrightarrow \pi_{A_n \dots A_1}(\text{Tab})$
 - R2: $\sigma_{p_n}(\dots \sigma_{p_2}(\sigma_{p_1}(\text{Tab}))\dots) \Leftrightarrow \sigma_{p_1 \text{ AND } p_2 \dots \text{ AND } p_n}(\text{Tab})$
- **Restructuring algorithm**
 - (1) **Decompose complex join predicates** such that they can be assigned to binary joins
 - (2) Divide Selections with several predicate terms in **separate Selections** with a single predicate term each.
 - (3) Execute Selections as early as possible, i.e., **push down Selections** to the leaves of the QG
 - (4) Aggregate simple Selections such that subsequent Selections (of the same table) are **executed at a time**
 - (5) Execute Projections without duplicate elimination as early as possible, i.e., **push down Projections** to the leaves as far as possible
 - (6) Aggregate simple projections (on a table) to a single operation

12-13

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries



© 2011 AG DBIS

Query Transformation

- **Task**
aggregation of logical operators (one- and two-variable expressions) and their replacement by **plan operators**
- **Typical plan operators in relational systems**
 - **On a single table:**
Selection, Projection, Sorting, Aggregation, Update operations (IUD) and ACCESS to base tables
+
Extensions: recursion, grouping . . .
 - **On two tables:**
Join- and set-operations, Cartesian product
- **Adjustments in QG for the effective use of plan operators**
 - (1) *Grouping of adjacent operators (if possible);*
 - (2) *Determination of processing sequence for binary operations; (minimize the size of intermediate tables)*
 - (3) *Detection of common subtrees (compute them only once).*

12-14

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Estimation of Execution Plans – Fundamental Problems

- **Query optimization rests on “fatal” assumptions, in general**
 - (1) All data elements and all attribute values are uniformly distributed
 - (2) Independence of Predicates:
 - $\text{sel}(X = a \text{ and } Y = b) = \text{sel}(X = a) * \text{sel}(Y = b)$
 - (3) Principle of Inclusion
 - $\text{sel}(T1.X = T2.Y) = 1/\max\{|X|, |Y|\}$

➔ These assumptions are wrong (in the general case)
- **Example**

(Family_Status = 'Married') AND (Age < 20)

➔ Linear interpolation, multiplication of probabilities

Although cost estimates are mostly wrong . . .

12-15

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans


Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

What is the Major Problem?



- **Correlation in predicates!**
- **EXAMPLE:**
 - Assume a query with the following WHERE clause:
 WHERE make = 'Honda' AND model = 'Accord',
 suppose
 - 10 makes ==> selectivity(make) = 1/10
 - 100 models ==> selectivity(model) = 1/100
 - So selectivity of both = 1/10 * 1/100 = 1/1000
 - But only Honda makes an Accord model!
 - We assumed the predicates were independent by multiplying the selectivities!
 - In fact, model and make are heavily correlated (predicate on make really adds no information)!
 - Effect: We underestimate cardinality by an order of magnitude!

12-16

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Estimation of Execution Plans – Fundamental Problems

- **Solution?**
 - Improvement of statistics/heuristics
 - **Histograms, Sampling**
- **Random samplings**
 - Speed: large data set, complex algorithms
 - **Example:** estimation "Sales in Europe" in a TPC-H application:
 1%: 8.46 Mio. in 4 sec.
 10%: 8.51 Mio. in 52 sec.
 100%: 8.54 Mio. in 200 sec.
- **Areas of sampling usage**
 - Approximate query evaluation, estimation of response time
 - **Query optimization**
 - Load balancing
 - Data mining
 - Interactive query design

12-17

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Use of Histograms

Equi-width histogram of T.a

Query 1:

```
SELECT *
FROM T
WHERE T.a <= 50
```

Query 2:

```
SELECT *
FROM T
WHERE T.a <= 52
```

- Cardinality of result for Query 1?
sum of bucket values for $T.a \leq 50$: $20 + 30 + 50 + 40 = 140$
- Cardinality of result for Query 2?
 - Cardinality for $T.a \leq 50$: 140
 - Cardinality for $50 < T.a \leq 52$:
uniform distribution assumption within buckets:
 $((52-50) / (60-50)) * 70 = 1/5 * 70 = 14$
 $\Rightarrow 140 + 14 = 154$

12-18

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Cost Model - Problems

- **Optimizer task**
 - Obtains cost estimate for each "promising" execution plan
 - Use of a weighted cost formula:

$$C = \text{\#physical page accesses} + W * (\text{\#calls of the access system})$$
 - Desired: weighted measure for I/O- and CPU utilization
 - W is the cost ratio of AS call to page access
- **Permanent problem**
 - In 1985, SQL was not standardized
 - SQL2 and SQL3 are essentially more complex
 - UDTs
 - Type and table hierarchies
 - Recursion, Constraints, Triggers, ...
- **Compilation and optimization**
 - **Cost-based optimizer**
 - Histograms
 - but: UDTs need their own cost model
 - "Optimizing the XXX optimizer"
 - **Dynamic QEPs**
 - alternative plans depending on resource availability
 - "reduce the braking distance"
 - **Seduction to gambling**

12-19

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Cost Model – Statistical Quantities

- **Statistical quantities for segments**
 - M_S number of data pages of segment S
 - L_S number of empty pages in S
- **Statistical quantities for tables**
 - N_R number of records of table R (Card(R))
 - $T_{R,S}$ number of pages in S with records of R
 - C_R clustering factor (number of records per page)
- **Statistical quantities per index** I on attributes A of a table R:
 - j_I number of attribute values / key values in the index (=Card ($\pi_A(R)$))
 - B_I number of leaf pages (B*-tree)
 - ...

➔ Statistics must be maintained in the DB catalog
- **Updating for each modification very expensive**
 - Additional write- and log operations
 - DB catalog would be the lock bottleneck
- **Alternative**
 - Initialization of statistical quantities at load- or generation time of tables and index structures
 - Periodical re-calculation of statistics by special command/service program (DB2: RUNSTATS)

12-20

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

 DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Cost Model – Computational Basis

- Selectivity Factor SF for**

$A_i = a_i$ $SF = \begin{cases} 1/j_i & \text{if index on } A_i \\ 1/10 & \text{otherwise} \end{cases}$

$A_i = A_k$ $SF = \begin{cases} 1 / \text{Max}(j_i, j_k) & \text{if index on } A_i, A_k \\ 1 / j_i & \text{if index on } A_i \\ 1 / 10^k & \text{otherwise} \end{cases}$

$A_i \geq a_i$ (or $A_i > a_i$) $SF = \begin{cases} (a_{\max} - a_i) / (a_{\max} - a_{\min}) & \text{if index on } A_i \text{ and interpolatable} \\ 1/3 & \text{otherwise} \end{cases}$

$A_i \text{ BETWEEN } a_i \text{ AND } a_k$ $SF = \begin{cases} (a_k - a_i) / (a_{\max} - a_{\min}) & \text{if index on } A_i \text{ and interpolatable} \\ 1/4 & \text{otherwise} \end{cases}$

$A_i \text{ IN } (a_1, a_2, \dots, a_r)$ $SF = \begin{cases} r / j_i & \text{if index on } A_i \text{ and } SF < 0.5 \\ 1/2 & \text{otherwise} \end{cases}$

Selectivity factor SF ($0 \leq SF \leq 1$)
 $\rightarrow \text{Card}(\sigma_p(R)) = SF(p) \cdot \text{Card}(R)$
- Computation of expressions**

 - $SF(p(A) \wedge p(B)) = SF(p(A)) \cdot SF(p(B))$
 - $SF(p(A) \vee p(B)) = SF(p(A)) + SF(p(B)) - SF(p(A)) \cdot SF(p(B))$
 - $SF(\neg p(A)) = 1 - SF(p(A))$
- Join Selectivity Factor (JSF)**

 - $\text{Card}(R \bowtie S) = JSF \cdot \text{Card}(R) \cdot \text{Card}(S)$
 - for (N:1)-joins (loss-free): $\text{Card}(R \bowtie S) = \text{Max}(\text{Card}(R), \text{Card}(S))$

12-21

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding


Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

 DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Creation and Selection of Execution Plans

- Input:**

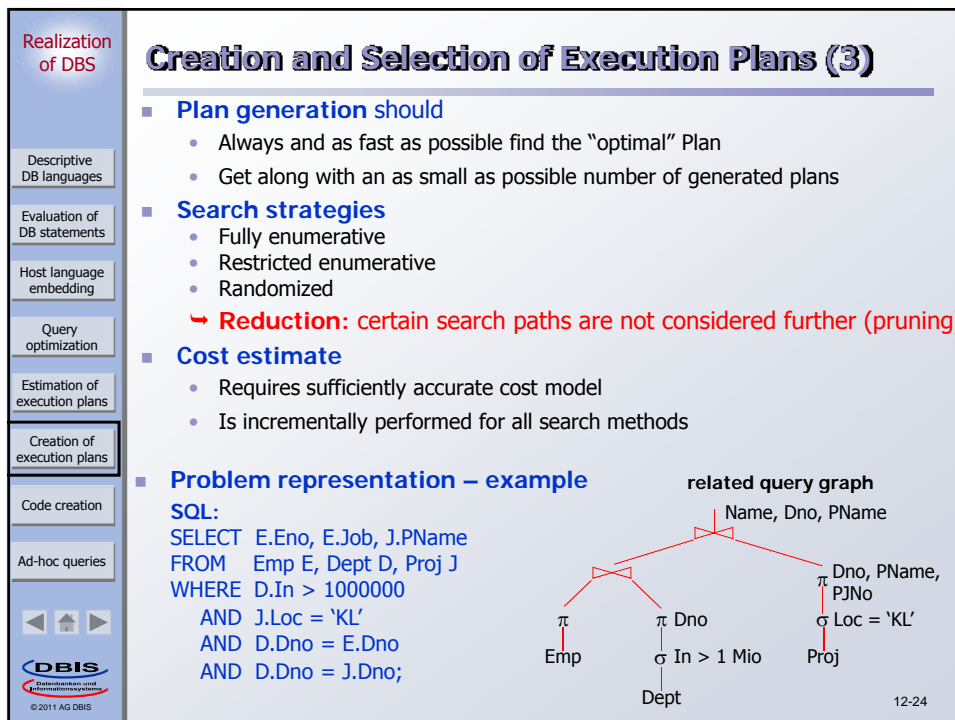
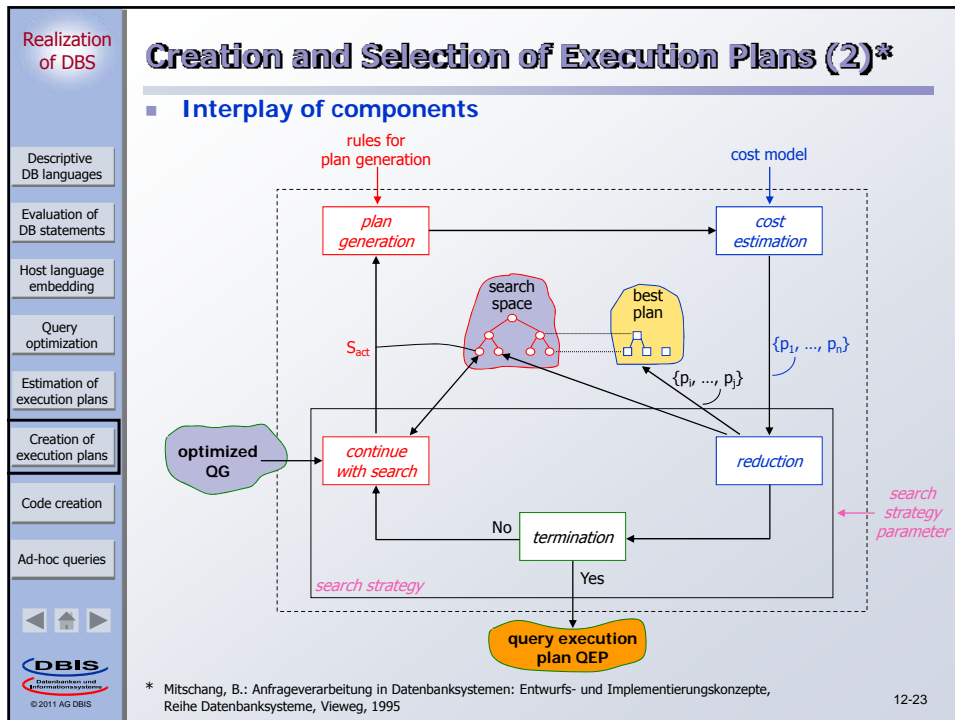
 - Optimized query graph (QG)
 - Existing storage structures and access paths
 - Cost model
- Output: optimal execution plan (or at least "good")**
- Approach:**

 1. **Generate all "reasonable" logical execution plans** for the evaluation of the query
 2. **Make the execution plans** complete by adding information for physical data representation (sort sequence, access path properties, statistical information)
 3. **Select the cheapest execution plan** corresponding to the given cost model


→ Alternative execution plans for a QG primarily emerge, because various methods (implementations) exist for each plan operator and because operation sequences (e.g. in case of multi-joins) can be varied. In case of complex queries, very large search spaces with alternatives are formed (e.g. 10^{70} possible execution plans for a query with 15 joins).
- Generation by the query optimizer**

 - **Small set plans containing the optimal plan**
 - **Confinement by heuristics**
 - Hierarchical generation based on the concept of nesting of SQL
 - Decomposition in a set of subqueries with at most two-variable expressions


12-22



Creation and Selection of Execution Plans (4)




cost estimate: $C(D.Dno) \dots C(E.Dno) \dots$



cost estimate: $C(J.Dno) \dots$


a) Possible access path for the single tables




b) Solution tree for single tables: reduction by pruning of subtrees

12-25

Creation and Selection of Execution Plans (5)



b) Solution tree for single tables: after pruning



c) Extended solution tree for the Nested-Loop join with the second table

cost estimate per path: e.g. by $C(C(D.Dno) + c(E.Dno) + \text{join cost})$

12-26

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Code Creation

- **Optimized query graph**
 - Result of optimization phase
 - Input data structure for code generator
- **Use of the operations of the access system**
 - **Direct operations** (e.g. INSERT <record>)
 - **Scan operations** (example SYSTEM R)
 - CALL RSS (OPEN, SCAN_STRUCTURE, RETURN_CODE)
 - CALL RSS (NEXT, SCAN_STRUCTURE, RETURN_CODE)
 - SCAN_STRUCTURE is complex data structure for the handing-over of in-/output values, search arguments, etc.

➔ **These operations are base operations for the code generation**
- **Classification of SQL statements**
 - Each class is described by a base process (e.g., by means of a cursor)
 - Skeleton of a base process is called 'model'
 - Processing step in the model is called 'fragment' (as code stored in a library)

➔ **Classification happens according to the kind of access actions**
- **Provision of models and fragments**
 - 4 models for simple queries (query blocks)
 - In total: 30 models with 5-10 fragments each (<100 fragments)

12-27

Realization of DBS

Descriptive DB languages

Evaluation of DB statements

Host language embedding

Query optimization

Estimation of execution plans

Creation of execution plans

Code creation

Ad-hoc queries

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

Flow Diagram for an Access Module

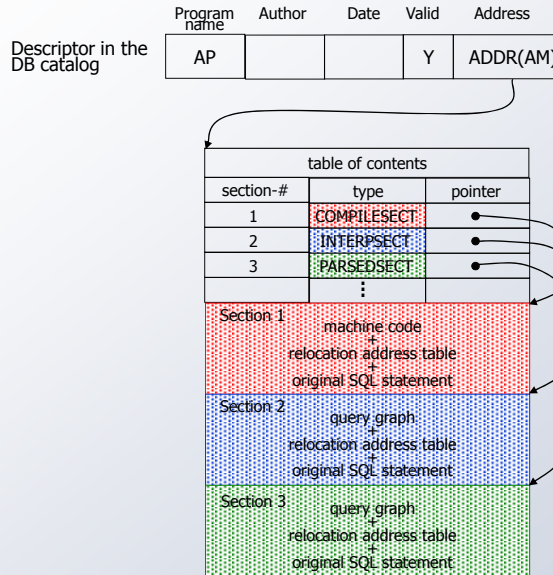
```

graph TD
    Prolog[Prolog] --> Decision1{OPEN or FETCH}
    Decision1 --> OPEN[OPEN]
    OPEN --> Binding[binding of input variables]
    Binding --> RSS_OPEN[RSS call for OPEN]
    RSS_OPEN --> Decision2{OK?}
    Decision2 -- N --> SettingReturn[setting of the RETURN code]
    Decision2 -- Y --> RSS_NEXT[RSS call for NEXT]
    RSS_NEXT --> Decision3{OK?}
    Decision3 -- N --> SettingReturn
    Decision3 -- Y --> EvalWhere[evaluation of the WHERE clause]
    EvalWhere --> Decision4{result}
    Decision4 -- T --> CompRecord[computation of the output record]
    CompRecord --> AssignVars[assignment to output variables]
    AssignVars --> SettingReturn
    Decision4 -- F --> SettingReturn
    SettingReturn --> BackJump[back jump]
    BackJump --> Decision1
  
```

Model for the selection of a record set by means of a cursor

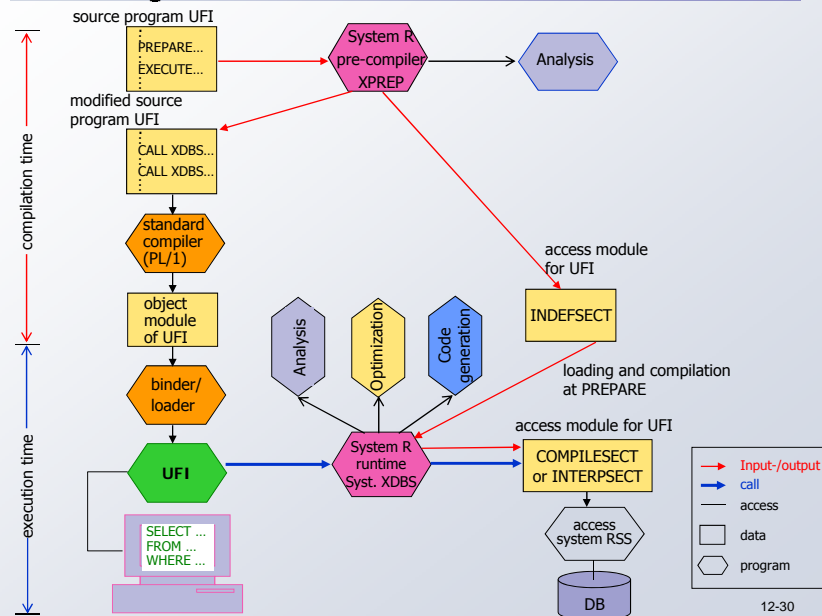
12-28

Structure of an Access Module

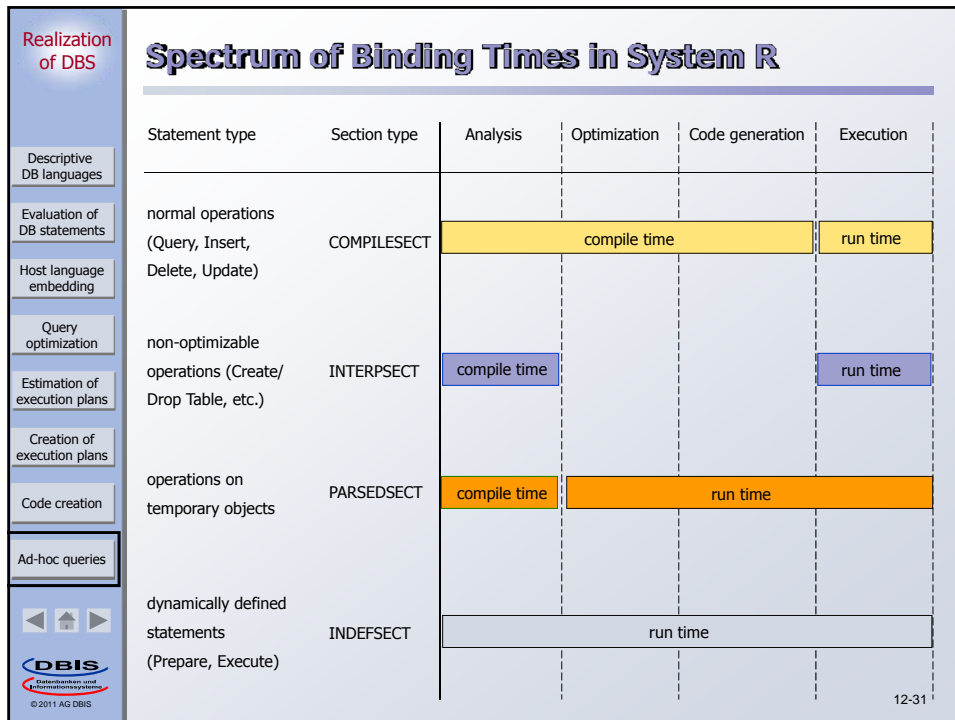


12-29

Preparation, Compilation and Execution of Ad-hoc Queries



12-30



Realization of DBS

Summary

- **Interpretation of a DB statement**
 - General program (Interpreter) accepts statements of the DB language as input and creates result by means of calls to the access system
 - High overhead at run time (esp. for repeated statement executions)
- **Compilation, code creation, and execution of a DB statement**
 - For each DB statement, a tailor-made program is created (compile time) which is evaluated at run time and thereby derives the result by means of calls to the access system
 - Compilation overhead is avoided as far as possible at run time
- **Query optimization: core problem**
of compilation of set-oriented DB languages
 - "Fatal" assumptions
 - uniform distribution of all attribute values
 - independence of predicates, principle of inclusion
 - Cost estimates for execution plans
 - CPU time and I/O overhead
 - no. of messages and data volumes to be transmitted (distributed case)
 - Good heuristics for the selection of execution plans is very important
- **Cost model**
 - Minimization of cost in dependency of the system state
 - Problem: Update of statistical quantities

DBIS
Datenbanken und Informationssysteme
© 2011 AG DBIS

12-32