

Luca Longinotti, 08-920-142

- Internal Schema Physical DI
- Conceptual Schema Logical DI
- External Schema

- Domain constraints (attribute domain)
- Key constraints (primary key)
- Entity constraints (not null)
- Referential integrity (foreign key)

insertion $r \leftarrow r \cup E$

updating $r \leftarrow E, r \leftarrow \pi_{F_1, F_2, \dots}(r)$

deletion $r \leftarrow r - E$

select σ_p , predicate $\wedge \vee \neg$

project Π_A , keep only columns A

union \cup , same schema, union of rows

set difference $-$, union-compatible, row difference

cartesian product \times , disjoint attribute names, else rename!

rename $p_r(A, \dots)$, rename rel. + attr.

set intersection \cap , equal to $r - (r - s)$

theta join \bowtie_θ , θ is bool condition, joins all attrs. from both schemas

natural join \bowtie , all attrs. of R + new ones in S

division \div , "for all", $R - S$

assignment \leftarrow , temp. variables

aggregate functions $G \dots \theta F(A)(E)$
group fns table

outer join, join with no info. loss (use null)

- simplicity
- no redundancy
- no update anomalies
- avoid null values
- avoid spurious tuples

\Rightarrow Functional Dependencies

\rightarrow minimal Rules:

$Y \subseteq X = X \rightarrow Y$

$X \rightarrow Y = XZ \rightarrow YZ$

$X \rightarrow Y, Y \rightarrow Z = X \rightarrow Z$

- 1NF: no composite, multival, nested
- 2NF: partial functionally dependent
- 3NF: BCNF + contained in candidate key
- BCNF: trivial or superkey

Join Evaluation

cost-based \rightarrow statistics

- nested loop
- block nested loop
- indexed nested loop
- merge join
- hash join (only equi.)

heuristic \rightarrow

- selection early
- projection early
- most restrictive select/joins before others

TRC / DRC | tables, columns, rows

PRIMARY KEY (X)

FOREIGN KEY (X) REFERENCES tab(Y)

LIKE % (0-n) or _ (1)

SELECT FROM WHERE GROUP HAVING AS to rename (on groups)

FROM t_1 JOIN t_2 ON cond

NOT/AND/OR GROUP BY

SELECT DISTINCT *

UNION/INTERSECT/EXCEPT

IN, EXISTS, NOT EXISTS

UNIQUE, IS NULL, ORDER BY ASC/DESC

INSERT INTO x VALUES (...)

UPDATE x SET ..., CASE when then else END

Views: CREATE VIEW x AS ...

WITH clause for temporary views

CREATE DOMAIN, CAST x AS type

CREATE FUNCTION x(y)

RETURNS type AS \$\$

[DECLARE]


BEGIN stmt END;

\$\$ language plpgsql;

CREATE TRIGGER x AFTER/BEFORE

WS/UP/DEL ON table

ER diagrams



1:1, 1:N, N:M
 participation:
 - optional
 = required

ER-to Relational

- 1) Regular Entity types (PRIMARY KEYS)
- 2) Weak Entity types (FOREIGN KEYS)
- 3) 1:1 Relations (FOREIGN KEY)
- 4) 1:N Relations (FOREIGN KEY)
- 5) N:M Relations (rel. relation, 2 FOREIGN KEY)
- 6) Multivalued Attributes (relation, FOREIGN KEY)
- 7) N-ary Relationships (rel. relation, N FOREIGN KEY)
- 8) Specialization or Generalization

B+ tree! multi-level index

- path length $\log_{\frac{m}{2}}(keys)$
- node: $m-1$ keys, m pointers
- nodes half to full
- internal nodes between $m/2$ and m
- leaf nodes: $(m-1)/2$ and $m-1$ children
- root at least 2 children (or leaf)
- leaf nodes linked together (last ptr.)
- insertion: split
- deletion: coalesce, redistribute

Static/Dynamic/Extendable Hashing

Transaction States:

- active
- partially committed
- committed
- failed
- aborted (\Rightarrow kill/restart)

ACID properties:

- Atomicity
- Consistency
- Isolation
- Durability

Schedule! seq. of instructions from a set of concurrent transactions

\Rightarrow conflict if there is a write!

Serial schedule: sequential transactions

\Rightarrow swaps can be made:

Serializable schedule: if equivalent to serial schedule

- both are reads
- they refer to different data
- one of them is not a R or W

Conflict equivalence! transform into another using only non-conflicting swaps

Concurrency control

Conflict serializable schedule: is conflict equivalent to a serial schedule

Schedule: - serializable
 - recoverable
 - cascadeless

- Lock-based (X/S -locks)
 $\rightarrow \Delta$ deadlocks, non-serializable

Precedence Graph (if acyclic, conflict serializable)

A directed graph with a node T_i for each transaction and an edge $T_i \rightarrow T_j$ if:

- T_i executes write(Q) before T_j executes read(Q)
- T_i executes read(Q) before T_j executes write(Q)
- T_i executes write(Q) before T_j executes write(Q)

- 2PL (obtain, release)
 - strict (X-only) 2^{eng}
 - rigorous (both) 5^{eng}

- undesirable phenomena
 - dirty read
 - nonrepeatable read
 - phantom read

\Rightarrow isolation levels

- read uncommitted
 - read committed
 - repeatable read / serializable