

Query Processing for Non-traditional Applications

CS848 Spring 2013

Cheriton School of CS

Updating Data

- 1 What are updates (how to understand dynamic aspects of instances)?


Plan

- 1 What are updates (how to understand dynamic aspects of instances)?
- 2 How do we understand updates *in our framework*?

- 1 What are updates (how to understand dynamic aspects of instances)?
- 2 How do we understand updates *in our framework*?
 - updates and logical relations
 - updates and constraints
 - updates and access paths


- 1 What are updates (how to understand dynamic aspects of instances)?
- 2 How do we understand updates *in our framework*?
 - updates and logical relations
 - updates and constraints
 - updates and access paths
- 3 Difficulties on the way
 - sequencing updates
 - value invention

Physical Design and Query Compilation: Overview




UPDATES IN NUTSHELL

Physical Design and Updates: Overview




old instance

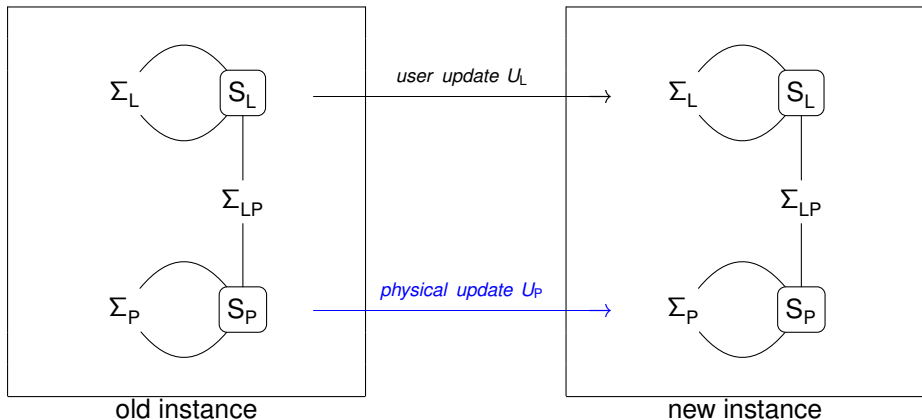


new instance

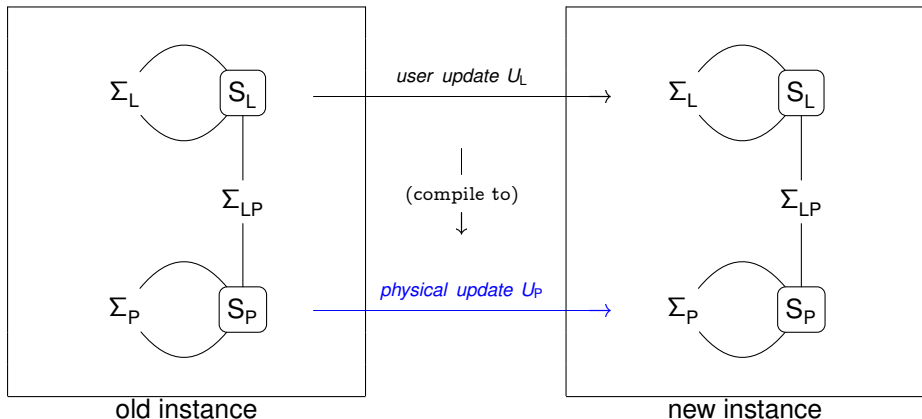
Physical Design and Updates: Overview



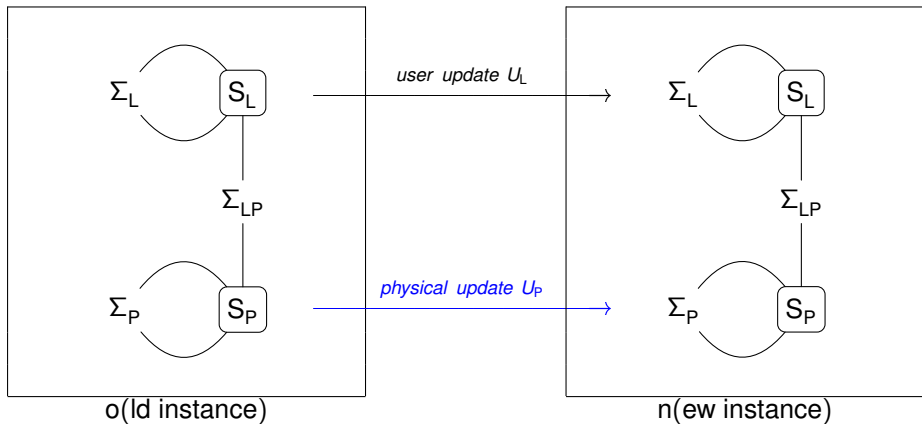
Physical Design and Updates: Overview



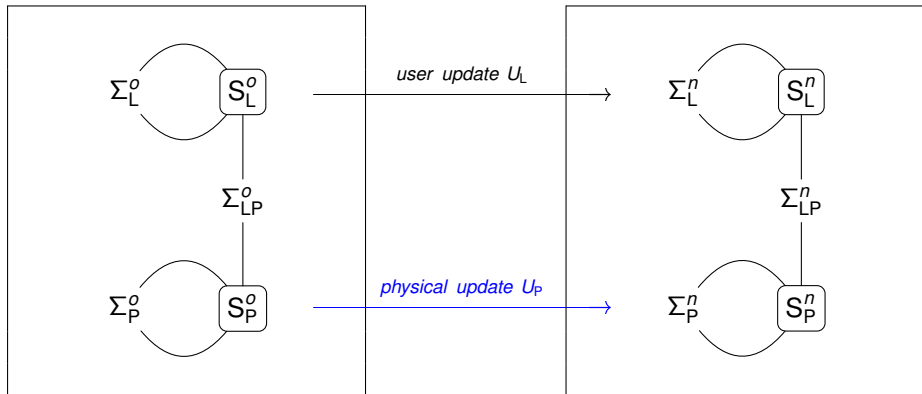
Physical Design and Updates: Overview




Update Schema



Update Schema




Update Schema



$$S_L^\pm = \{P^+, P^- \mid P \in S_L\},$$

$$\Sigma_L^\pm = \{\forall \bar{x}. (P^o(\bar{x}) \vee P^+(\bar{x})) \leftrightarrow (P^n(\bar{x}) \vee P^-(\bar{x})) \mid P \in S_L\}$$

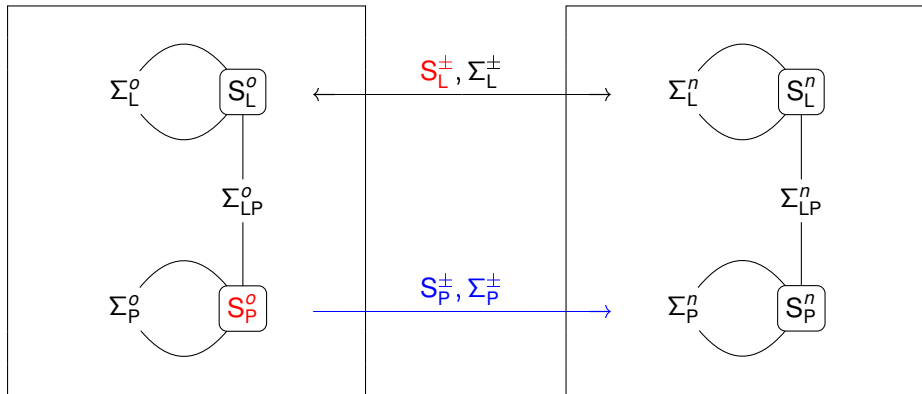
Update Schema




$$S_L^\pm = \{P^+, P^-\mid P \in S_A\},$$

$$\Sigma_L^\pm = \{\forall \bar{x}. (P^o(\bar{x}) \vee P^+(\bar{x})) \leftrightarrow (P^n(\bar{x}) \vee P^-(\bar{x})) \mid P \in S_A\}$$


Update Schema




Update Schema




Physical Design and Update Compilation



Physical Design and Update Compilation




Physical Design and Update Compilation



- U_L is a user query $P^+(\bar{x})$ ($P^-(\bar{x})$) for $P \in S_A$;

Physical Design and Update Compilation



- U_L is a *user query* $P^+(\bar{x})$ ($P^-(\bar{x})$) for $P \in S_A$;
- U_P is a *plan* for the user query $P^+(\bar{x})$ ($P^-(\bar{x})$) for $P \in S_A$
 - \Rightarrow w.r.t. the access paths $S_A \cup S_L^\pm$, and
 - \Rightarrow aux code that inserts (deletes) the result of the plan into (from) P .

Example

Setup: standard relational design for `Employee (id, name, salary)`

- A *base file* `empfile` of `emp` records (organized by `id`)
- An `emp-name` index on employee `names` (links `name` to `id`)

Example

Setup: standard relational design for `Employee(id, name, salary)`

- A *base file* `empfile` of `emp` records (organized by `id`)
- An `emp-name` index on employee names (links `name` to `id`)

Logical Schema:

$$S_L = \{\text{Employee}/3\}, \Sigma_L = \{\text{"id is a key"}\}$$

Physical Schema:

$$S_P = S_A = \{\text{empfile}/3/0, \text{emp-id}/3/1, \text{emp-name}/2/1\}$$
$$\Sigma_{LP} = \left\{ \begin{array}{l} \forall x, y, z. \text{Employee}(x, y, z) \leftrightarrow \text{empfile}(x, y, z) \\ \forall x, y, z. \text{Employee}(x, y, z) \leftrightarrow \text{emp-id}(x, y, z) \\ \forall x, y, z. \text{Employee}(x, y, z) \leftrightarrow \text{emp-name}(y, x) \end{array} \right\}$$

Logical Update Schema: (just the signature)

$$S_L = \{\text{empfile}^+/3, \text{empfile}^-/3, \text{emp-name}^+/2, \text{emp-name}^-/2\}$$

Physical Update Schema:

$$S_P = \{\text{Employee}^+/3, \text{Employee}^-/3, \text{empfile}^0/3, \text{empfile}^0/3, \dots\}$$

Example

Setup: standard relational design for `Employee(id, name, salary)`

- A *base file* `empfile` of `emp` records (organized by `id`)
- An `emp-name` index on employee names (links `name` to `id`)

Logical Update Schema: (just the signature)

$$S_L = \{\text{empfile}^+/3, \text{empfile}^-/3, \text{emp-name}^+/2, \text{emp-name}^-/2\}$$

Physical Update Schema:

$$S_P = \{\text{Employee}^+/3, \text{Employee}^-/3, \text{empfile}^0/3, \text{empfile}^0/3, \dots\}$$
$$\Sigma_{LP} = \{\forall x, y, z. (\text{empfile}^0(x, y, z) \vee \text{empfile}^+(x, y, z)) \\ \leftrightarrow (\text{empfile}^n(x, y, z) \vee \text{empfile}^-(x, y, z)), \dots\}$$
$$\Sigma_P = \{\forall x, y, z. \text{Employee}^+(x, y, z) \wedge \text{Employee}^-(x, y, z) \rightarrow \perp, \dots\}$$

Update Queries:

`empfile+(x, y, z)`

`empfile-(x, y, z)`

Example

Setup: standard relational design for `Employee (id, name, salary)`

- A *base file* `empfile` of `emp` records (organized by `id`)
- An `emp-name` index on employee names (links `name` to `id`)

Logical Update Schema: (just the signature)

$$S_L = \{\text{empfile}^+/3, \text{empfile}^-/3, \text{emp-name}^+/2, \text{emp-name}^-/2\}$$

Physical Update Schema:

$$\begin{aligned} S_P &= \{\text{Employee}^+/3, \text{Employee}^-/3, \text{empfile}^o/3, \text{empfile}^o/3, \dots\} \\ \Sigma_{LP} &= \{\forall x, y, z. (\text{empfile}^o(x, y, z) \vee \text{empfile}^+(x, y, z)) \\ &\quad \leftrightarrow (\text{empfile}^n(x, y, z) \vee \text{empfile}^-(x, y, z)), \dots\} \\ \Sigma_P &= \{\forall x, y, z. \text{Employee}^+(x, y, z) \wedge \text{Employee}^-(x, y, z) \rightarrow \perp, \dots\} \end{aligned}$$

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

Example

Setup: standard relational design for `Employee (id, name, salary)`

- A *base file* `empfile` of `emp` records (organized by `id`)
- An `emp-name` index on employee names (links `name` to `id`)

Logical Update Schema: (just the signature)

$$S_L = \{\text{empfile}^+/3, \text{empfile}^-/3, \text{emp-name}^+/2, \text{emp-name}^-/2\}$$

Physical Update Schema:

$$\begin{aligned} S_P &= \{\text{Employee}^+/3, \text{Employee}^-/3, \text{empfile}^o/3, \text{empfile}^o/3, \dots\} \\ \Sigma_{LP} &= \{\forall x, y, z. (\text{empfile}^o(x, y, z) \vee \text{empfile}^+(x, y, z)) \\ &\quad \leftrightarrow (\text{empfile}^n(x, y, z) \vee \text{empfile}^-(x, y, z)), \dots\} \\ \Sigma_P &= \{\forall x, y, z. \text{Employee}^+(x, y, z) \wedge \text{Employee}^-(x, y, z) \rightarrow \perp, \dots\} \end{aligned}$$

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

... similar for `emp-name`, no-op for `emp-id` (why?)

Transaction Types

Transactions

A user update (expressed as diffs on *logical* symbols) that transforms an consistent instance to another consistent instance.

Transaction Types

Transactions

A user update (expressed as diffs on *logical* symbols) that transforms an consistent instance to another consistent instance.

Additional information about transaction behaviour?

- 1 transaction only adds tuples to a certain relation,
- 2 transaction only modifies certain relations,
- 3 ...

Transaction Types

Transactions

A user update (expressed as diffs on *logical* symbols) that transforms an consistent instance to another consistent instance.

Additional information about transaction behaviour?

- 1 transaction only adds tuples to a certain relation,
- 2 transaction only modifies certain relations,
- 3 ...

Additional information \Rightarrow additional constraints:

- 1 $P^- = \emptyset$ for the “insert-only” relation P ,
- 2 $P^+ = P^- = \emptyset$ for unmodified relations.
- 3 ...

The View Update Problem

Classical View Update Problem

Given a relational view

$$\forall \bar{x}. V(\bar{x}) \leftrightarrow Q(\bar{x})$$

with Q expressed over S_L , is it possible to **update the content** of V by appropriately modifying the interpretation of the S_L symbols?

\Rightarrow *insertable, deletable, and updatable* views

The View Update Problem

Classical View Update Problem

Given a relational view

$$\forall \bar{x}. V(\bar{x}) \leftrightarrow Q(\bar{x})$$

with Q expressed over S_L , is it possible to **update the content** of V by appropriately modifying the interpretation of the S_L symbols?

\Rightarrow *insertable, deletable, and updatable* views

Answer

Define *update schema* for V and S_L (where every symbol is also an access path). Then V is

- *insertable* if P^n is definable w.r.t. the update design with $V^- = \emptyset$,
- *deletable* if P^n is definable w.r.t. the update design with $V^+ = \emptyset$, and
- *updatable* if P^n and V^- are definable w.r.t. the update design

for all $P \in S_L$.

\Rightarrow when the answer is positive, we construct a corresponding *update* queries.

ADVANCED ISSUES IN UPDATE COMPILATION

Progressive Updates

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

Progressive Updates

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

This doesn't quite work:

after *executing* the 1st update query **we no longer have** empfile^o !

Progressive Updates

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

This doesn't quite work:

after *executing* the 1st update query **we no longer have** empfile^o !

Possible Solutions:

- 1 simultaneous *relational* assignment:
 - \Rightarrow compute all deltas and store results in temporary storage,
 - \Rightarrow *only then* apply all deltas to S_A ;

Progressive Updates

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

This doesn't quite work:

after *executing* the 1st update query **we no longer have** empfile^o !

Possible Solutions:

- 1 simultaneous *relational* assignment:
 - \Rightarrow compute all deltas and store results in temporary storage,
 - \Rightarrow *only then* apply all deltas to S_A ;
- 2 using independent deltas:
 - \Rightarrow add constraints to avoid the problem (e.g., $P^- \subseteq P^o$);

Progressive Updates

Update Queries:

$$\begin{aligned} \text{empfile}^+(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^+(x, y, z) \wedge \neg \text{empfile}^o(x, y, z) \\ \text{empfile}^-(x, y, z) &\xrightarrow{\text{compiles}} \text{Employee}^-(x, y, z) \wedge \text{empfile}^o(x, y, z) \end{aligned}$$

This doesn't quite work:

after *executing* the 1st update query **we no longer have** empfile^o !

Possible Solutions:

- 1 simultaneous *relational* assignment:
 - \Rightarrow compute all deltas and store results in temporary storage,
 - \Rightarrow *only then* apply all deltas to S_A ;
- 2 using independent deltas:
 - \Rightarrow add constraints to avoid the problem (e.g., $P^- \subseteq P^o$);
- 3 evolving physical schema one AP at a time
 - \Rightarrow sequence of update schemas with a subset of S_A “updated”,
 - \Rightarrow subsequent updates compiled w.r.t. partially updated schema.

Value Invention

Setup: advanced relational design for `Employee(id, name, salary)`

- A *base file* `empfile(r, x, y, z)` of `emp` records *with RIds "r"*
- An `emp-name(y, r)` index on employee names (links name to RIds)

Value Invention

Setup: advanced relational design for `Employee(id, name, salary)`

- A *base file* `empfile(r, x, y, z)` of emp records *with RIds "r"*
- An `emp-name(y, r)` index on employee names (links name to RIds)
 - ⇒ no update query, e.g., for `empfile+(r, x, y, z)`: no “source” of RIds!
(due to: $\forall x, y, z. \text{Employee}(x, y, z) \leftrightarrow (\exists r. \text{empfile}(r, x, y, z))$)

Value Invention

Setup: advanced relational design for `Employee(id, name, salary)`

- A *base file* `empfile(r, x, y, z)` of `emp` records *with RIds “r”*
- An `emp-name(y, r)` index on employee names (links name to RIds)
⇒ no update query, e.g., for `empfile+(r, x, y, z)`: no “source” of RIds!

IDEA (Constant Complement [Bancilhon and Spyratos])

An *oracle access path* that provides the required value
given the values of remaining attributes as parameters.

Value Invention

Setup: advanced relational design for `Employee(id, name, salary)`

- A *base file* `empfile(r, x, y, z)` of `emp` records *with RIds “r”*
- An `emp-name(y, r)` index on employee names (links name to RIds)
⇒ no update query, e.g., for `empfile+(r, x, y, z)`: no “source” of RIds!

IDEA (Constant Complement [Bancilhon and Spyratos])

An *oracle access path* that provides the required value
given the values of remaining attributes as parameters.

In practice: a record allocation mechanism

(e.g., `malloc+code` that initializes fields of the allocated record)

Value Invention

Setup: advanced relational design for `Employee(id, name, salary)`

- A *base file* `empfile(r, x, y, z)` of `emp` records *with RIds “r”*
- An `emp-name(y, r)` index on employee names (links name to RIds)
⇒ no update query, e.g., for `empfile+(r, x, y, z)`: no “source” of RIds!

IDEA (Constant Complement [Bancilhon and Spyratos])

An *oracle access path* that provides the required value
given the values of remaining attributes as parameters.

In practice: a record allocation mechanism

(e.g., `malloc+code` that initializes fields of the allocated record)

- a separate access path (may need to “remember” all allocated records!)
- a part of the record insertion code (AP^+ doesn't have the attribute)
⇒ update query for `emp-name+` must execute *after* `empfile+`.

Value Invention and Schematic Cycles

Can we *always* schedule the updates of record IDs before using these as values (e.g., in an index)?

Value Invention and Schematic Cycles

Can we *always* schedule the updates of record IDs before using these as values (e.g., in an index)?

NO: recall our Employee-Works-Department physical schema in which

- emp records have a pointer to a dept record (for the Works relationship),
- dept records have a pointer to an emp record (to the “manager”).

Value Invention and Schematic Cycles

Can we *always* schedule the updates of record IDs before using these as values (e.g., in an index)?

NO: recall our Employee-Works-Department physical schema in which

- emp records have a pointer to a dept record (for the Works relationship),
- dept records have a pointer to an emp record (to the “manager”).

⇒ impossible to insert the 1st employee and 1st department!

Value Invention and Schematic Cycles

Can we *always* schedule the updates of record IDs before using these as values (e.g., in an index)?

NO: recall our Employee-Works-Department physical schema in which

- emp records have a pointer to a dept record (for the Works relationship),
- dept records have a pointer to an emp record (to the “manager”).

⇒ impossible to insert the 1st employee and 1st department!

IDEA: reify (one of) the AP (we have done that already in our example) and then interleave updates to the reified relations.

Value Invention and Schematic Cycles

Can we *always* schedule the updates of record IDs before using these as values (e.g., in an index)?

NO: recall our Employee-Works-Department physical schema in which

- `emp` records have a pointer to a `dept` record (for the `Works` relationship),
- `dept` records have a pointer to an `emp` record (to the “manager”).

⇒ impossible to insert the 1st employee and 1st department!

IDEA: reify (one of) the AP (we have done that already in our example) and then interleave updates to the reified relations.

- 1 insert an employee's `Id` into `emp-id` AP (yields address of `emp`);
- 2 insert department record (the above value used for the manager field; yields address of `dept`);
- 3 insert the same employee into `emp-dept` AP using the `dept` address.

Additional Issues

- How to know when an *constant complement* is needed?

Additional Issues

- How to know when an *constant complement* is needed?
- How to determine the *ordering* of the individual AP updates?

Additional Issues

- How to know when an *constant complement* is needed?
- How to determine the *ordering* of the individual AP updates?
- How to identify when *reification* is needed?

Additional Issues

- How to know when an *constant complement* is needed?
- How to determine the *ordering* of the individual AP updates?
- How to identify when *reification* is needed?
- How to determine if the user update preserves *consistency*?
 - ⇒ guaranteed by the user (e.g., extra user queries to make sure)
 - ⇒ system-generated checks—HARD!