A tutorial on

Temporal Databases

by

Jan Chomicki
Monmouth University

David Toman
University of Toronto

List of Slides

Temporal databases

Databases with a **time** dimension.

Examples:

- various kinds of records:
  - credit
  - personnel
  - financial
  - judicial
  - cadastral (property)
  - medical
- monitoring and measurement results
- $\cdots$

As current data grows stale, it becomes **historical** data.

Plan of the tutorial

- Models of time
- Abstract Temporal Data Models and Query Languages
- Practical Temporal Models and Query Languages
- More Powerful Languages
- Incomplete Information in Temporal Databases
- Temporal Integrity Constraints
- *Research Problems*

The structure of time

Temporal **ontology**:

- databases: points (instants)
- AI: intervals (periods)

Temporal **domain** T:

- a linearly-ordered set
- typically one of the well known mathematical domains: $\mathbf{N}$ (natural numbers, $\mathbf{Z}$ (integers), $\mathbf{Q}$ (rationals), $\mathbf{R}$ (reals)

There is also a data domain D of uninterpreted constants.

Ontology important because it determines the kind of temporal object facts are associated with.

Domain properties (discreteness, density, ...) also important.

Skipped in this tutorial

Nonlinear temporal domains:

- branching time has potential database applications in versioning and workflows [Attie et al., 1993]
- proposals too preliminary

Multiple time granularities:

- important practical issue
- many possible approaches [Ladkin, 1986, Wang et al., 1993, Bettini et al., 1995]

Multiple temporal dimensions

To model multiple kinds of time:

- valid time vs. transaction time

To represent intervals using pairs of points.

To represent multiple temporal attributes in query results.

Plan

- Models of time
- **Abstract Temporal Data Models and Query Languages**
  - **Data models**
  - **Query languages**
  - **Expressive power**
- Practical Temporal Models and Query Languages
- More Powerful Languages
- Incomplete Information in Temporal Databases
- Temporal Integrity Constraints

Abstract Temporal Databases

Representation-independent semantics of temporal databases.

Advantages:

- allows high-level, declarative query languages
- provides a formal framework to solve outstanding problems in temporal databases:
  - interoperability of different data models
  - functional dependencies and normal forms

Explicitly introduced in [Chomicki, 1994] but implicit in many recent papers (*snapshot-equivalence*).

Basic Building Blocks

- Temporal domain $(T, <)$ (linearly ordered, unbounded)
- Relational databases $DB(D, \rho)$ over $(D, =)$ and schema ρ:

$$\Big(\quad \underbrace{D, =,}_{\text{domain of uninterpreted constants}} \quad \underbrace{r_1, \ldots, r_n}_{\text{a finite instance of } \rho \text{ over } D} \quad \Big)$$

- Major approaches of putting these together:

  1. the snapshot model
  2. the timestamp model

Relational databases serve as descriptions of *worlds* at every single time instant.

This is the major difference from the propositional uses of TL (e.g., in the area of specification and verification).

There are several ways of adding time (a single temporal dimension) to the relational model.

The Snapshot Model

Temporal database is a function $T \to DB(D, \rho)$

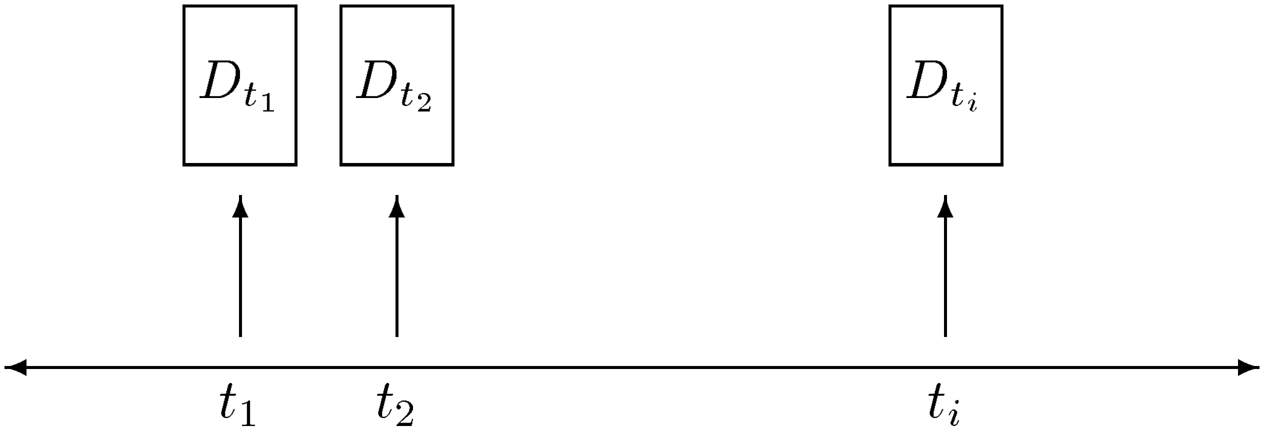

Data type of relations $T \to (D^n \to \text{bool})$

This approach is the closest to Propositional TL: every database describes the state of the world at a particular time instant; the order relation $<$ on T then defines the flow of time (the access relation).

Snapshots: Example

Year	Snapshot
. . .	
1990	{Work(John, IBM), Work(Mary,IBM), Work(Steve,HP)}
1991	. . .
1992	. . .
1993	{ Work(John,Microsoft),Work(Mary,IBM), Work(Steve,HP)}
. . .	

Histories

Relational database updates map states to states.

History:

- a sequence of database states
- thus also a **snapshot temporal database**

The Timestamp Model

For every $r_i \in \rho$ we define a relation R_i such that

$$R_i = \{(t, a_1, \ldots, a_k) : (a_1, \ldots, a_k) \in r_i \text{ in } D_t\}$$

A *Timestamp Temporal Database* is a structure

$$\Big(\underbrace{D, =,}_{\text{domain of uninterpreted constants}} \quad \underbrace{T, <,}_{\text{temporal domain}} \quad \underbrace{R_1, \ldots, R_n}_{\text{an instance of } \rho \text{ over } T \text{ and } D} \Big)$$

Data type of relations $(T \times D^n) \to \text{bool}$

The snapshot model is not very suitable for queries of the form

$$\{t : DB \models \varphi(t)\}$$

(all time instants such that φ holds in DB).

The alternative *timestamp model* makes time to be just another data type (however, we must remember that it comes with an interpreted order relation).

The two models are equivalent (type-isomorphic).

Timestamp Example

Work		
Name	Company	Year
John	IBM	1990
John	IBM	1991
John	MicroSoft	1993
John	MicroSoft	...
Mary	DEC	1984
Mary	DEC	1985
Mary	IBM	1990
Mary	IBM	...
Steve	HP	1990
Steve	HP	...

Finiteness not an issue *at this level*.

Different temporal data models - different representations.

Query Languages

Start from *Relational Calculus* over $\rho \cup \{=\}$

$$L ::= r_i(x_1, \ldots, x_k) | x_i = x_j | L \wedge L | \neg L | \exists x.L$$

Example: *list everyone who works for IBM.*

$$\{x : \exists y.\mathrm{Works}(x, y) \wedge y = \mathrm{IBM}\}$$

Two principal temporal extensions:

- implicit references to time (temporal connectives)
- explicit references to time (variables and quantifiers)

First-order Temporal Connectives

Temporal Connectives:

$$
\begin{array}{lllll}
O & ::= & X_i & - & \text{placeholders} \\
 & & O \wedge O & \multirow{2}{*}{\}} & \multirow{2}{*}{propositional connectives} \\
 & & \neg O & & \\
 & & t_i < t_j & \multirow{2}{*}{\}} & \multirow{2}{*}{temporal variables} \\
 & & \exists t_i . O & &
\end{array}
$$

Definition. A *(k-ary) temporal connective* is an O-formula with exactly one free variable t_0 and k free predicate variables $X_1, \ldots, X_k$.

The O-formulas define the *truth tables* for the new connectives.

At this point we only allow *first-order* connectives; therefore the language of O-formulas is just FO logic over the signature of the temporal domain $(T, <)$.

Examples of Temporal Connectives

The standard connectives **since** and **until**:

$$X_1 \textbf{ until } X_2 \triangleq \exists t_2 . t_0 < t_2 \wedge X_2 \wedge \forall t_1 (t_0 < t_1 < t_2 \rightarrow X_1)$$
$$X_1 \textbf{ since } X_2 \triangleq \exists t_2 . t_0 > t_2 \wedge X_2 \wedge \forall t_1 (t_0 > t_1 > t_2 \rightarrow X_1)$$

The derived connectives:

$$\Diamond X_1 \triangleq \text{true } \textbf{until } X_1 \qquad \Box X_1 \triangleq \neg \Diamond \neg X_1$$
$$\blacklozenge X_1 \triangleq \text{true } \textbf{since } X_1 \qquad \blacksquare X_1 \triangleq \neg \Diamond \neg X_1$$

For discrete linear order we also define the $\bigcirc$ (next) and $\bullet$ (previous) operators as

$$\bigcirc X_1 \triangleq \exists t_1 . t_1 = t_0 + 1 \wedge X_1 \qquad \bullet X_1 \triangleq \exists t_1 . t_1 + 1 = t_0 \wedge X_1$$

Note that all the connectives are first-order definable.

Propositional Temporal Logic

Propositional TL(**since**, **until**)

is expressively equivalent to monadic FOL($<$)
over complete linear orders [Kamp, 1968]

Propositional TL(**until**)

is equivalent to TL(**since**, **until**)
on (discrete) linear orders bounded in the past.

Formulas in Propositional TL(**since**, **until**)

can be separated to pure past, present and future parts:

$$\varphi \equiv \varphi_{\text{past}} \vee \varphi_{\text{present}} \vee \varphi_{\text{future}}$$

First-order Temporal Logic: syntax

Let Ω be a finite set of temporal connectives.

$$
\begin{array}{lll}
F ::= & r_i(x_{i_1}, \ldots, x_{i_k}) & - \quad \text{database schema} \\
& F \wedge F & \\
& & \} \quad \text{propositional connectives} \\
& \neg F & \\
& x_i = x_j & \\
& & \} \quad \text{data variables} \\
& \exists x_i . F & \\
& \boxed{\omega(F_1, \ldots, F_k)} & - \quad \text{temporal connectives}
\end{array}
$$

Features: encapsulation, same schema as the original relational database,...

Temporal Logic: semantics

$$
\begin{array}{ll}
DB, \theta, t \models r_j(x_{i_1}, \ldots, x_{i_k}) & \text{if } r_j \in \rho, (\theta(x_{i_1}), \ldots, \theta(x_{i_k})) \in ! r_j^{DB(t)} \\
DB, \theta, t \models x_i = x_j & \text{if } \theta(x_i) = \theta(x_j) \\
DB, \theta, t \models \varphi \wedge \psi & \text{if } DB, \theta, t \models \varphi \text{ and } DB, \theta, t \models \psi \\
DB, \theta, t \models \neg \varphi & \text{if not } DB, \theta, t \models \varphi \\
DB, \theta, t \models \exists x_i . \varphi & \text{if there is } a \in D \text{ and } DB, \theta[x_i \mapsto a], t \models \varphi \\
DB, \theta, t \models \omega(F_1, \ldots, F_k) & \text{if } T_P, [t_0 \mapsto t] \models \omega^* \text{ where} \\
 & \qquad T_P, \delta \models X_i \text{ iff } DB, \theta, \delta(t_i) \models F_i
\end{array}
$$

Answer: $\varphi(DB) = \{t, \theta : DB, \theta, t \models \varphi\}$.

Examples

- *list John's work history*

$$\exists x.\mathrm{Works}(x, y) \wedge x = \mathrm{John}$$

- *list all people who were rehired by the same company*

$$\exists y.\mathrm{Works}(x, y) \wedge \blacklozenge(\neg\mathrm{Works}(x, y) \wedge \blacklozenge\mathrm{Works}(x, y))$$

- *list all people who have been unemployed between two jobs*

$$\exists y.\mathrm{Works}(x, y) \wedge \blacklozenge(\neg\exists y.\mathrm{Works}(x, y) \wedge \blacklozenge\exists y.\mathrm{Works}(x, y))$$

- *list all people who continuously worked at IBM since John left*

$$\mathrm{Works}(x, \mathrm{IBM})\ \mathbf{since}\ (\mathrm{Works}(\mathrm{John}, \mathrm{IBM}) \wedge \bigcirc\Box\neg\mathrm{Works}(\mathrm{John}, \mathrm{IBM}))$$

Temporal Relational Calculus: syntax

$$
\begin{array}{lllll}
M & ::= & \boxed{R_i(t_j, x_{i_1}, \ldots, x_{i_k})} & - & \text{extended database schema} \\
 & & M \wedge M & \} & \text{propositional connectives} \\
 & & \neg M & & \\
 & & x_i = x_j & \} & \text{data variables} \\
 & & \exists x_i . M & & \\
 & & \boxed{\begin{array}{l} t_i < t_j \\ \exists t_i . M \end{array}} & \} & \text{temporal variables}
\end{array}
$$

$\Rightarrow$ essentially a two-sorted first-order logic (2-FOL)

Temporal Relational Calculus: semantics

$$
\begin{array}{ll}
DB, \theta \models R_j(t_i, x_{i_1}, \ldots, x_{i_k}) & \text{if } R_j \in \rho, (\theta(t_i), \theta(x_{i_1}), \ldots, \theta(x_{i_k})) \in R_j^{DB} \\
DB, \theta \models t_i < t_j & \text{if } \theta(t_i) < \theta(t_j) \\
DB, \theta \models x_i = x_j & \text{if } \theta(x_i) = \theta(x_j) \\
DB, \theta \models \varphi \wedge \psi & \text{if } DB, \theta \models \varphi \text{ and } DB, \theta \models \psi \\
DB, \theta \models \neg\varphi & \text{if not } DB, \theta \models \varphi \\
DB, \theta \models \exists t_i . \varphi & \text{if there is } s \in T_P \text{ and } DB, \theta[t_i \mapsto s] \models \varphi \\
DB, \theta \models \exists x_i . \varphi & \text{if there is } a \in D \text{ and } DB, \theta[x_i \mapsto a] \models \varphi
\end{array}
$$

Answer: $\varphi(DB) = \{\theta : DB, \theta \models \varphi\}$.

Examples

- *list John's work history*

$$\exists x.\text{Works}(t_0, x, y) \wedge x = \text{John}$$

- *list all people who were rehired by the same company*

$$\exists y.\text{Works}(t_0, x, y) \wedge \exists t_1 (t_1 < t_0 \wedge \neg\text{Works}(t_1, x, y) \wedge \\ \exists t_2.t_2 < t_1 \wedge \text{Works}(t_2, x, y))$$

- *list all people who have been unemployed between two jobs*

$$\exists y.\text{Works}(t_0, x, y) \wedge \exists t_1 (t_1 < t_0 \wedge \neg\exists y.\text{Works}(t_1, x, y) \wedge \\ \exists t_2.t_2 < t_1 \wedge \exists y.\text{Works}(t_2, x, y))$$

- *list all people who continuously worked at IBM since John left*

$$\exists t_1 < t_0 \, (\text{Works}(t_1, \text{John}, \text{IBM}) \wedge \\ \forall t_2.t_2 > t_1 \to \neg\text{Works}(t_2, \text{John}, \text{IBM}) \wedge \\ \forall t_2.t_1 < t_2 \leq t_0 \to \text{Works}(t_2, x, \text{IBM}))$$

Expressive Power

$$
\begin{aligned}
\mathrm{Embed}(r_i(x_1, \ldots, x_{v_i})) &= R_i(t_0, x_1, \ldots, x_{v_i}) \\
\mathrm{Embed}(x_i = x_j) &= x_i = x_j \\
\mathrm{Embed}(F_1 \wedge F_2) &= \mathrm{Embed}(F_1) \wedge \mathrm{Embed}(F_2) \\
\mathrm{Embed}(\neg F) &= \neg\, \mathrm{Embed}(F) \\
\mathrm{Embed}(\exists x. F) &= \exists x.\, \mathrm{Embed}(F) \\
\mathrm{Embed}(\omega(F_1, \ldots, F_k)) &= \\
&\omega^*(\mathrm{Embed}(F_1)[t_0/t_1], \ldots, \mathrm{Embed}(F_k)[t_0/t_k])
\end{aligned}
$$

Expressive Power (cont.)

1. FOTL is strictly less expressive than 2-FOL
   for arbitrary finite set of connectives.

$$\exists t_1 < t < t_2 . \forall x (P(t_1, x) \iff R(t_2, x))$$

2. Future FOTL is less expressive than FOTL
   for linear order with left endpoint

$$\exists t > 0 . \forall x (P(0, x) \iff R(t, x))$$

3. FOTL = 2-FOL
   if the data domain D has fixed size

Why: limited scope of temporal variables

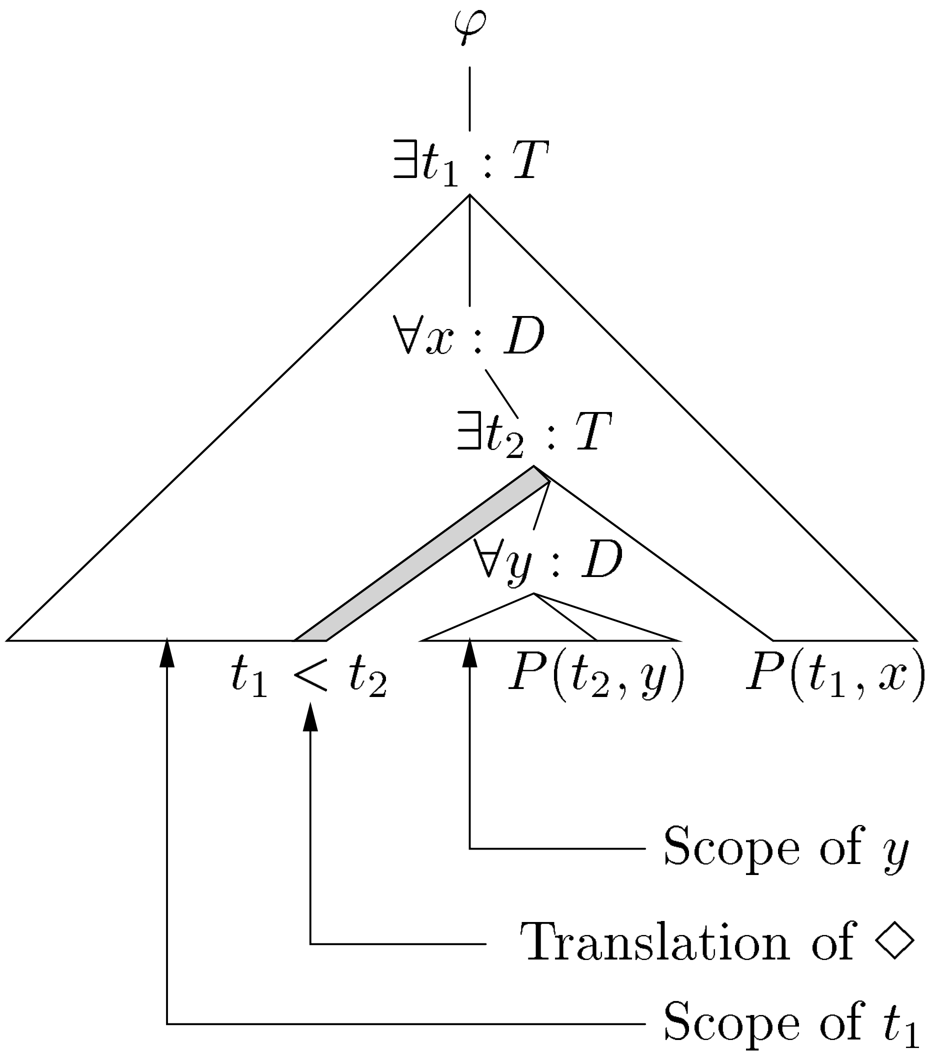

Temporal Relational Algebra

A Temporal Relational Algebra (TRA) consists of:

- Universe: single-dimensional temporal relations
- Operations: a finite set of operations of the type

$$2^{T \times D^{n_1}} \times \ldots \times 2^{T \times D^{n_k}} \to 2^{T \times D^{n}}$$

- First-order operations: can be defined by a first-order formula.

All operations preserve closure over the universe

$\Rightarrow$ can *not* express all FO queries.

The problem is not in the operators, but in the *closure* requirement—the universe contains only single-dimensional temporal relations.

Major problem from the implementation point of view: a limitation to single-dimensional temporal relations prevents expressing all FO queries.

However, the majority of practical temporal query languages ignore this fact!

TRA: example

- Universe: temporal relations of type $T \times D^k$.
- Operations:

$$
\begin{aligned}
\pi_V(R) &= \{t, \theta_{|V} : DB, \theta, t \models R\} \\
\sigma_F(R) &= \{t, \theta_{|FV(R)} : DB, \theta, t \models R \wedge F\} \\
R \bowtie S &= \{t, \theta_{|FV(R) \cup FV(S)} : DB, \theta, t \models R \wedge S\} \\
R \cup S &= \{t, \theta_{|FV(R) \cup FV(S)} : DB, \theta, t \models R \vee S\} \\
R - S &= \{t, \theta_{|FV(R) \cup FV(S)} : DB, \theta, t \models R \wedge \neg S\} \\
\mathcal{S}(R, S) &= \{t, \theta_{|FV(R) \cup FV(S)} : DB, \theta, t \models R \ \textbf{since} \ S\} \\
\mathcal{U}(R, S) &= \{t, \theta_{|FV(R) \cup FV(S)} : DB, \theta, t \models R \ \textbf{until} \ S\}
\end{aligned}
$$

Temporal Logic TL(FO)

FOTL does not have the separation property of the propositional logic:

$$\forall x. \blacklozenge p(x) \iff \Diamond p(x)$$

Weaker two-layer logic TL(FO) [Gabbay et al., 1994]:

$\Rightarrow$ allows temporal connectives only outside of the scopes of $\forall$ and $\exists$.

$\Rightarrow$ has separation property

Plan

- Models of time
- Abstract Temporal Data Models and Query Languages
- **Practical Temporal Models and Query Languages**
  - **Compact encoding of temporal databases**
  - **Query languages over the encodings**
  - **Query languages and translations**
  - **Efficient relational operations**
- More Powerful Languages
- Incomplete Information in Temporal Databases
- Temporal Integrity Constraints

Concrete Temporal Databases

Grouping by non-temporal attributes and

+

finite encodings of the resulting sets of time instants

$$
\underset{T \times D^k}{\boxed{\begin{array}{cc} & \ldots \\ t_1, & a_1, \ldots, a_k \\ t_2, & a_1, \ldots, a_k \\ \vdots & \vdots \\ t_k, & a_1, \ldots, a_k \\ & \ldots \end{array}}}
\longrightarrow
\underset{2^T \times D^k}{\boxed{\begin{array}{cc} & \ldots \\ \{t_1, t_2, \ldots, t_k\}, & a_1, \ldots, a_k \\ & \ldots \end{array}}}
$$

Implicit functional dependency $A_1 \ldots A_k \to T$ in $2^T \times D^k$.

Observation:
a single (data) tuple is usually related to a large number of time instants

Finite Encoding using Constraints

- sets of time instants described by their characteristic formulas
- the language of constraints admits quantifier elimination
- fixed-size tuples = quantifier-free formulas in CNF

$$\boxed{\begin{array}{c} \ldots \\ \{t_1, t_2, \ldots, t_k\}, \quad a_1, \ldots, a_k \\ \ldots \end{array}} \longrightarrow \boxed{\begin{array}{c} \ldots \\ \varphi(t), \quad a_1, \ldots, a_k \\ \ldots \end{array}} \longrightarrow \boxed{\begin{array}{cc} \multicolumn{2}{c}{\ldots} \\ \varphi_1, & a_1, \ldots, a_k \\ \vdots & \vdots \\ \varphi_l, & a_1, \ldots, a_k \\ \multicolumn{2}{c}{\ldots} \end{array}}$$

Interval Encoding

Let T_P be the temporal domain. We define the set

$$I(T) = \{(a,b) : a \leq b, a \in T \cup \{-\infty\}, b \in T \cup \{\infty\}\}$$

Relations on the elements of $I(T)$:

$$\begin{array}{ll} ([a,b] <_{--} [a',b']) \iff a < a' & ([a,b] <_{+-} [a',b']) \iff b < a' \\ ([a,b] <_{-+} [a',b']) \iff a < b' & ([a,b] <_{++} [a',b']) \iff b < b' \end{array}$$

The structure $T_I = (I(T), <_{--}, <_{+-}, <_{-+}, <_{++})$ is the *Interval-based Temporal Domain (corresponding to T_P).*

Most prominent example of constraint encoding: quantifier free formulas in the theory of (linear) order with constants.

Interval Encoding (cont.)

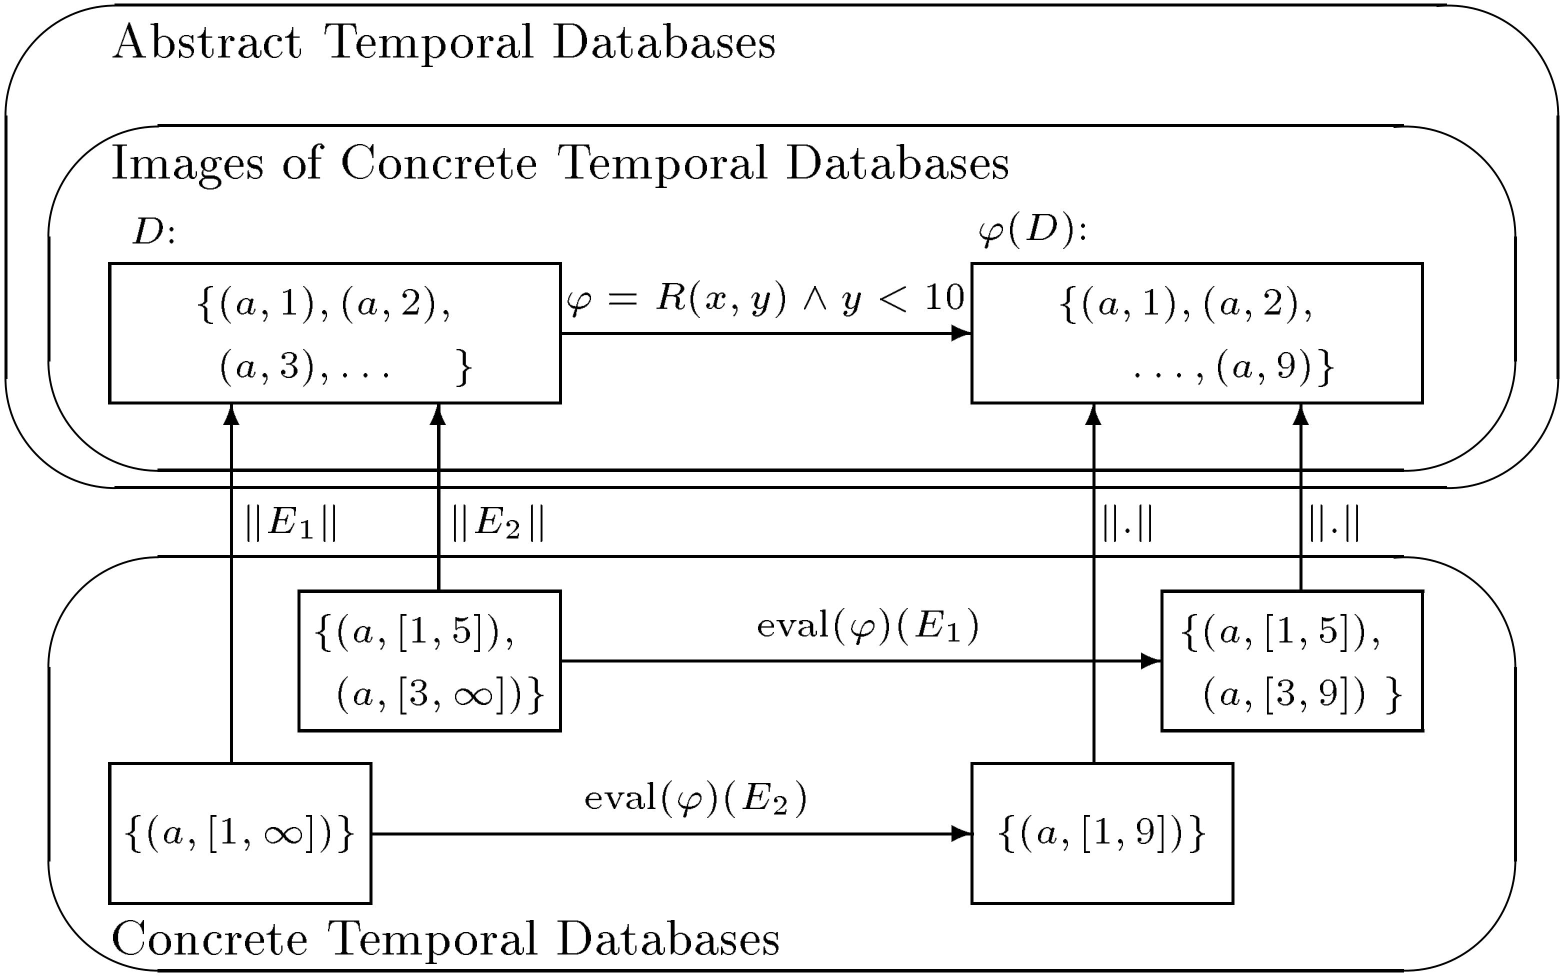

Example

Work		
Name	Company	Year
John	IBM	[1990, 1991]
John	MicroSoft	[1993, ∞]
Mary	DEC	[1984, 1985]
Mary	IBM	[1990, ∞]
Steve	HP	[1990, ∞]

Interval Queries

$$
\begin{array}{lll}
M \quad ::= & R_i(I_j, x_{i_1}, \ldots, x_{i_k}) & - \quad \text{extended database schema} \\
& M \wedge M & \\
& \neg M & \Big\} \quad \text{propositional connectives} \\
& x_i = x_j & \\
& \exists x_i.M & \Big\} \quad \text{data variables} \\
& I_i^* < I_j^* & \\
& \exists I_i.M & \Big\} \quad \text{temporal variables}
\end{array}
$$

L^I is the language of M-formulas.

Note that I's range over $I(T)$ – INTERVALS.

Genericity

Definition. $\varphi \in L_I$ is $\|.\|$-*generic* if

$$\|D_1\| = \|D_2\| \supset \|\varphi D_1\| = \|\varphi D_2\|$$

for all D_1, D_2 concrete temporal databases.

Let D_1, D_2 be two concrete temporal databases such that:

$$R^{D_1} = \{([0,3], a)\}$$
$$R^{D_2} = \{([0,2], a), ([1,3], a)\}$$

Then $\exists I, J. \exists x (R(I,x) \wedge R(J,x) \wedge I \neq J)$ is true in D_2 but false in D_1.

$\Rightarrow$ not generic.

Coalescing

A single-dimensional temporal relation is *coalesced* if every fact is associated only with maximal non-overlapping intervals.

- coalescing has to be enforced using non-logical operation on relations
- relational operations do not preserve coalescing in general e.g., projection, union, or set difference.

A way to curb non-genericity in the single-dimensional case.

TQUEL [Snodgrass, 1987]

- a "minimal" extension of QUEL
- valid and transaction time
- interval encoding and coalescing
- only PSJ queries directly in syntax

TSQL2 [Snodgrass, 1995]

- similar extension to SQL
- valid and transaction time:
  $\Rightarrow$ temporal elements as timestamps
  $\Rightarrow$ tuples can not be bounded by their arity
- no formal semantics
  $\Rightarrow$ unclear expressive power

Example

list all people who were rehired by the same company

TSQL2:
```
select  r.name
from    Works r, Works s
where   r.name=s.name
  and   r.company=s.company
  and   validtime(r) precedes validtime(s)
```

Does not work in general:

⇒ `indep` has to be coalesced and can not contain duplicates.

Example (cont.)

list all people who were unemployed (between jobs):

TSQL2:
```
select   r.name
from     Works r, Works s
where    r.name=s.name
  and    validtime(r) precedes validtime(s)
```

Does not work:

⇒ `indep` has to be recoalesced **after projecting out** company.

$$\begin{aligned} \text{Works} = \{ \quad & (\text{Joe}, \text{Amdahl}, \quad [1985, 1990]), \\ & (\text{Joe}, \text{IBM}, \qquad [1990, 1993]), \\ & (\text{Joe}, \text{MicroSoft}, [1993, \infty]) \quad \} \end{aligned}$$

⇒ to get correct answer we have to use subqueries/views.

Other Proposals

HRDM [Clifford and Croker, 1987]

time-dependent attributes (functions of time instants)

timestamps form the temporal ontology of *lifespans*

IXRM [Lorentzos, 1993]

interval, multidimensional approach

introduces EXPSPACE operation (unfold)

TSQL [Navathe and Ahmed, 1993]

interval encoding of 1-dimensional time line

based on temporal relational algebra

TSQL2's Successors

1. ATSQL
2. SQL/Temporal

Common characteristics:

- *interval ontology* of time rather than point based ontology.
- still 1 or 2 dimensional
- allow coercion of temporal attributes to data ones and vice versa,
- levels of compatibility with SQL

Intervals vs. True Intervals

- Intervals as encodings of sets of time instants:

```
king("Charles IV",  "Czech Kingdom", [1347, 1378])
king("Casimir III", "Poland",        [1333, 1370])
```

- Intervals as points in two-dimensional space:

```
electricity("Jones A.", 40, 05/15/96, 06/15/96)
electricity("Smith J.", 35, 05/01/96, 06/01/96)
```

Moral: don't mix apples and oranges!

Temporal Connectives in L_I

- Coalescing simplifies select-project-join queries
- However, it is of little help even for elementary *temporal* operations:

$$\begin{aligned}
A \textbf{ since } B &\mapsto [\max(A^-, succ(B^-)), A^+] \\
&\quad \text{for } \max(A^-, succ(B^-)) \leq A^+ \text{ and } B^+ \geq A^- \\
A \textbf{ until } B &\mapsto [A^-, \min(A^+, pred(B^+))] \\
&\quad \text{for } A^- \leq \min(A^+, pred(B^+)) \text{ and } A^+ \geq B^-
\end{aligned}$$

For details see [Böhlen et al., 1996].

It is still fairly difficult to state even elementary temporal queries.

Failure of Coalescing

Known facts:

- there is no unique coalescing for temporal dimension > 1
- FO-complete query languages can not avoid arbitrarily large dimensionality in queries
  [Abiteboul et al., 1996, Toman and Niwinski, 1996]

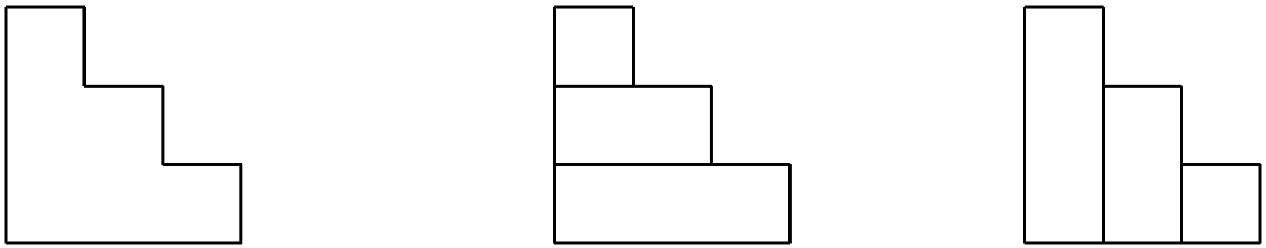

- queries depend on the representation

  $\Rightarrow$ non-$\|.\|$-generic queries

Translations

Another approach:

- use an *abstract* query language
- devise an evaluation procedure over encoded databases

In our case we consider the following two pairs:

1. FOTL + interval encoding [Böhlen et al., 1996]
2. 2-FOL + interval encoding [Toman, 1996, Toman, 1997]

The target language is L^I (in the first case with coalescing).

SQL/TP [Toman, 1997]

$\boxed{\text{Temporal attributes range over individual time instants}}$

- Syntax and Semantics:
  $\Rightarrow$ an extension of SQL by a new data type
  $\Rightarrow$ supports finite duplication

- Encoding of relations:
  $\Rightarrow$ compact representation of temporal relations using intervals
  $\Rightarrow$ syntactic restrictions to maintain closure (safety)

- Query evaluation:
  $\Rightarrow$ efficiency: depends on the (size of the) encoding only
  $\Rightarrow$ allows compilation to SQL/92

The goals:

(1) The syntax and semantics of standard SQL has to be extended in a *natural way*: by adding a new data type `time` that behaves (almost) identically as the existing data types, and

(2) The syntax and semantics are *independent* of the chosen encoding of timestamps.

The chosen language is derived from calculus with finite duplicates.

The choice of encoding is not critical: we aim on subsuming TSQL, so we have chosen TSQL's encoding of temporal databases.

Efficient of query evaluation means that it uses only elements in the active domain of the encoding (plus some small neighborhoods of `Time` values).

SQL/TP: syntax

```
<exp>     ::= <ddl_exp> | <dql_exp> | <dml_exp>

<ddl_exp> ::= create table <rid> ( <signat> ) |
              create view  <rid> ( <dql_exp> )

<dql_exp> ::= select [all] <s_list> [from <f_list>]
                     [where <cond>] [group by <g_list>] |
              ( <dql_exp> ) union [all] ( <dql exp> ) |
              ( <dql_exp> ) except [all] ( <dql exp> ) |
              ( <dql_exp> ) intersect [all] ( <dql exp> )

<dml_exp> ::= insert [all] into <rid> ( <dql_exp> ) |
              delete [all] from <rid> ( <dql_exp> )
```

Standard SQL, all nesting of queries is realized in the `from` clause.

The `having` clause and nesting in the `where` clause are easily definable using nesting in the `from` clause and thus are considered to be just a syntactic sugar.

Where is the "temporal extension"? In the data definition language: when we define an attribute of type `time` (next slide)...

SQL/TP: encoding of time

```
<type>    ::= time [using <enc>] | integer | char(N) | ...
```

Modifiers `<enc>`:

- `points`
- `bounded intervals`
- `unbounded intervals`

Other classes of constraints can be readily incorporated `points` is for atomic events (e.g., transaction time).

`bounded intervals` is for facts with bounded "duration".

`unbounded intervals` is for facts with unbounded duration (e.g., x holds from now on)

Easy to introduce other constraint classes (back-door for the CDB folks): If you can't beat them (SQL), join them.

Duplicate Semantics

- standard SQL semantics

  ⇒ select-project-join, aggregation, difference, union

  ⇒ infinite bags w/finite duplication

- no duplicate-preserving projection of unbounded attributes:

  ⇒ `select all` name `from indep`

  Result:

$$\{\ldots, \underbrace{(\text{Poland}), \ldots, (\text{Poland})}_{\infty}, \ldots\}$$

- duplicate preserving projection of *bounded* types may be allowed

The semantics of the new language is exactly SQL semantics on the $\|.\|$-images (the semantics is naturally extended to infinite relations).

We need to restrict projections to maintain finite duplication (infinite duplication is not very interesting nor useful).

Note the difference between infinite and unbounded (we can "measure" infinite bounded sets, e.g., in dense case; we may get non-integral "counts", e.g., 1.5 years)

Doesn't work after coalescing; aggregation often based on the encoding!

Query Processing

Efficient execution of queries on the chosen encoding.

⇒ General constraint encoding [Kanellakis et al., 1993].

⇒ Interval encoding:

- equality on time instants maps into interval intersection

  ⇒ indices for intervals

- set operations on points map to

  ⇒ normalization (join-based) followed by the original operation

  ⇒ on-the-fly multidimensional indices

Intervals (or hypercubes) are important in general: more complex encodings base indexing on *bounding boxes*.

Indexing in Temporal Databases

Problem:

given an interval I, retrieve all tuples whose timestamp intersects I.

Additional knowledge:

- intervals are nonempty (thus not all pairs of time instants form an interval)
- "current" state may be accessed by far more queries than the other states of the database
- temporal databases are often append-only

Leads to unbalanced access structures.

Indexing (cont.)

- Time Index [Elmasri et al., 1990, Ramaswamy, 1997]
- R-Tree [Guttman, 1984]

  general multi-dimensional index
  used in POSTGRESS [Stonebraker and Kemnitz, 1991]
- AP-Tree

  append-only index
  combination of ISAM and B^+-tree
- ...

Survey of indexing methods: [Saltzberg and Tsotras, 1994]

Indexing (example)

Time Index [Elmasri et al., 1990]:

- at every interval start- or end-point store the list of all intervals containing this point
- build a B-tree over the above points

Queries in $O(\log n + s/B)$.

Disadvantage: $O(n^2/B)$ space requirement.

$\Rightarrow$ [Ramaswamy, 1997] improves this method to space $O(n/B)$.

Updates

Depends on choice of semantics: sets vs. bags.

Insertions

$\Rightarrow$ duplicate tuples and overlapping timestamps

Deletions

$\Rightarrow$ may have to delete only part of a concrete tuple

$\Rightarrow$ set-based deletions

Updates

$\Rightarrow$ similar to deletes: parts of concrete tuples

$\Rightarrow$ problem: updates *in-place*

Plan

- Models of time
- Abstract Temporal Data Models and Query Languages
- Practical Temporal Models and Query Languages
- **More Powerful Languages**
  - **multidimensional connectives**
  - **higher-order features**
- Incomplete Information in Temporal Databases
- Temporal Integrity Constraints

More Powerful Languages

Additional expressive power:

- More temporal dimensions
- Higher-order temporal connectives
- Higher-order features in the underlying language

Multidimensional Temporal Logics

- A *k-ary m-dimensional temporal connective* is a formula in the first-order language of the temporal domain T_P with exactly m free variable $t_0^1, \ldots, t_0^m$ and k free predicate variables $X_1, \ldots, X_k$
- The language $L^{\Omega(m)}$ is a first-order logic extended with a finite set $\Omega(m)$ of m-dimensional temporal connectives.
- Semantics defined using satisfaction relation with m evaluation points:

$$DB, \theta, t_1, \ldots, t_m \models \varphi$$

Multidimensional Temporal Logics (example)

- Two-dimensional logic

$$\begin{aligned} \Diamond \; &= \exists t_1 . X_1 \\ \uparrow\downarrow &= \exists t_1 . (t_1[1] = t_0[2] \wedge t_1[2] = t_0[1] \wedge X_1) \end{aligned}$$

- other examples:
  - the temporal logic with the *now* operator,
  - the Vlach and Aqvist system, and
  - most of the interval logics.
- However, k-FOTL$\sqsubset$ $(k+1)$-FOTL [Toman and Niwinski, 1996]

Higher-order features

1. in the temporal dimension:
   - ETL (connectives use regular expressions) [Wolper, 1983]
   - fixed point extensions (μTL) [Vardi, 1988]
   - second order logic (in $<$) for defining temporal connectives
2. in the data dimension (stratified w.r.t. temporal connectives):
   - Datalog($^\neg$)
   - fixed-point extensions of relational calculus
3. in both:
   - Datalog$_{1S}$ and TempLog [Baudinet et al., 1993]

Datalog$_{1S}$ and Templog

"Find all the persons at risk where "being at risk" is defined in the following way: a person is at risk *at a given time if she has been earlier infected or she has been in contact with someone already at risk.*

$$atRisk(X, T+1) \leftarrow infected(X, T).$$
$$atRisk(X, T+1) \leftarrow atRisk(X, T).$$
$$atRisk(X, T+1) \leftarrow contacted(X, Y, T), atRisk(Y, T).$$

This is Datalog$_{1S}$, the Templog formulation is similar.

A query that is inductive on both *time* and *data*.

New application area of Datalog$_{1S}$ and extensions: operational semantics of active databases [Motakis and Zaniolo, 1997].

Plan

- Models of time
- Abstract Temporal Data Models and Query Languages
- Practical Temporal Models and Query Languages
- More Powerful Languages
- **Incomplete Information in Temporal Databases**
  - **Motivation**
  - **Marked Nulls and Constraints**
  - **Queries**
- Temporal Integrity Constraints

Incomplete temporal information

Partial information:

- *Sue stopped working for Microsoft and started to work for IBM before 1992*
- *John worked for IBM before working for Microsoft*

Different granularities:

- *The Beatles broke up in the sixties*

More natural here than in many other database applications.

Null values

SQL-92 NULL is not enough:

- stands for *"value unknown"*
- lack of consistent formal interpretation in queries
- unable to express order, succession and other logical conditions

Marked nulls

The model of Imieliński and Lipski:

- marked nulls
- local conditions associated with tuples
- global conditions associated with tables

Semantics of a table:

- set of relations (possible worlds)

Constraint databases with marked nulls

[Koubarakis, 1994]:

- two kinds of variables: universally and existentially quantified
- conditions: *(gap)-order constraints.*

Two-stage formal semantics:

- substitute time values for nulls in such a way that the global condition is satisfied, obtaining a set of generalized relations $rep(R)$
- for every element of $rep(R)$ take the corresponding abstract temporal database

This works for an arbitrary number of temporal dimensions.

One dimension: one can use *indefinite intervals*, together with global conditions.

Example

Work		
Name	Company	Year
John	IBM	$1990 \leq Year \leq c_1$
John	Microsoft	$c_2 \leq Year$
Sue	Microsoft	$Year \leq c_3$
Sue	IBM	$c_3 < Year$
Global condition		
$1990 < c_1 < c_2 \wedge c_3 < 1992$		

One possible world

Work		
Name	Company	Year
John	IBM	1990
John	Microsoft	1991, 1992, ...
Sue	Microsoft	..., 1991
Sue	IBM	1992, 1993, ...

Another possible world

Work		
Name	Company	Year
John	IBM	1990
John	Microsoft	1993, 1994, ...
Sue	Microsoft	..., 1989
Sue	IBM	1990, 1991, ...

There are more possible worlds.

Query languages

Relational languages:

- (extended) relational algebra
- (extended) relational calculus

Modal operators:

- possibility: *"Is it possible that John and Sue overlapped at IBM?"*
- certainty: *"Is it certain that John worked at IBM in 1990?"*

Both calculus and algebra are extended with the operators.

Query evaluation

Relational calculus:

- translation to relational algebra
- quantifier elimination (for modal queries)

Complexity results for the general model [Koubarakis, 1997]:

- not worse than for incomplete *relational* databases

Tractability results for a restricted model [van der Meyden, 1997]

- local conditions = equalities
- any practical temporal domain
- crucially depends on the *number of temporal dimensions*:
  - PTIME for one dimension,
  - co-NP-complete for more than one dimension.

Other approaches

Temporal databases:

- point and interval relationships [Chaudhuri, 1988]
- quantitative indeterminacy [Dyreson and Snodgrass, 1993]
- least and upper bounds on an interval [Gadia et al., 1992]

Artificial intelligence:

- 13 basic kinds of relationships between intervals, algebra of relationships [Allen, 1983]
- disjunctive information
- metric information [Meiri, 1991], [Kautz and Ladkin, 1991]

Example: $I_1\{\mathbf{before}, \mathbf{after}\}I_2$ equivalent to $I_1^+ < I_2^- \vee I_2^+ < I_1^-$

The *qualitative* part of these can be recast using Koubarakis' approach.

Plan

- Models of time
- Abstract Temporal Data Models and Query Languages
- Practical Temporal Models and Query Languages
- More Powerful Languages
- Incomplete Information in Temporal Databases
- **Temporal Integrity Constraints**
  - **Functional Dependencies and Normal Forms**
  - **Constraint Dependencies**
  - **Temporal Integrity Constraints in Relational Databases**

Temporal integrity constraints

Relational databases: integrity constraints are closed first-order logic formulas.

Temporal databases: integrity constraints are closed formulas in a first-order temporal query language.

Why first order: evaluated over *first order structures*.

Applications of integrity constraints

Integrity constraints capture the semantics of a database application.

Enforcing database integrity:

- only meaningful information is stored.

Database design:

- normal forms (criteria for *good*, anomaly-free schemas)
- *good* decompositions

Functional dependencies (FDs)

The most popular class of dependencies, essential for defining **keys** and **normal forms**.

Example: relation schema `Emp(SSN, Name,Salary)`.

The FD $\texttt{SSN} \rightarrow \texttt{Salary}$ holds in `Emp` if for every instance r of `Emp`, corresponding to a possible state of the real world, and for every pair of tuples $t_1, t_2 \in r$:

if $t_1[\texttt{SSN}] = t_2[\texttt{SSN}]$, then $t_1[\texttt{Salary}] = t_2[\texttt{Salary}]$.

$$\forall s_1, s_2, n_1, n_2, d_1, d_2 . \texttt{Emp}(s_1, n_1, d_1) \wedge \texttt{Emp}(s_2, n_2, d_2) \wedge s_1 = s_2 \Rightarrow d_1 = d_2$$

Can be generalized to *sets* of attributes in both sides.

FDs in database design

A set of attributes X is a **key** of a relation schema U w.r.t. a given set of FDs F if:

1. for every attribute $A \in U$, the FD $X \rightarrow A$ is implied by F, and
2. no proper subset of X has property 1.

A relation schema U is in **Boyce-Codd Normal Form** (BCNF) w.r.t. a set of FDs F if for every dependency $X \rightarrow A$ implied by F, X is a superset of a key of U.

Functional dependencies in temporal databases

Snapshot functional dependencies: should hold in every **snapshot**.

Enough to consider temporal extensions of relation schemas with one (or more) temporal attributes.

Example: relation schema `Emp(SSN,Name,Salary,Year)`

Snapshot dependency is a classical FD:

$$\texttt{SSN Year} \rightarrow \texttt{Salary}$$

Compare with:

$$\texttt{SSN} \rightarrow \texttt{Salary}$$

which means that the salary never changes.

The classical relational database design theory is still applicable.

Keys, normal forms, decompositions etc. - as in relational databases.

Temporal functional dependencies

Many attempts to define temporal FDs for concrete temporal databases:

- latest and most general in [Jensen et al., 1996]
- no logical formulation for dependencies
- technical problems [Chomicki, 1997]

The approach of Jensen, Snodgrass and Soo is subsumed, generalized, and clarified by the classical approach [Chomicki, 1997].

Constraint dependencies

Functional dependencies are **equality-generating.** In temporal databases it is natural to consider more general **constraint-generating** dependencies [Baudinet et al., 1995].

Example: *"The past never changes."*

$$\forall x, x', t_v, t'_v, t_t, t'_t.\ p(x, t_v, t_t) \wedge p(x', t'_v, t'_t) \wedge t_v < t'_v \Rightarrow t_t < t'_t$$

Valid time: t_v, transaction time t_t.

Many examples in [Jensen and Snodgrass, 1994].

Complexity results for the implication problem [Baudinet et al., 1995].

Not clear how to use CGDs for database design.

Further generalizations: tuple-generating constraint dependencies [Maher and Srivastav

Temporal integrity constraints in relational databases

Integrity constraints on **histories.**

Examples:

"salaries of employees should never decrease"

"employees should not be hired if they have been fired in the past"

"a recalled book should be returned within a week"

Constraint satisfaction

A history H **satisfies** a constraint C if C is true in every state of H.

A finite history H **potentially satisfies** C if it can be extended to an infinite history that satisfies C.

Implementation: after each update check whether C is potentially satisfied in the current history.

More flexible enforcement: temporal triggers.

Temporal logic constraints

Many variants proposed and studied.

"employees can not be hired if they have been fired in the past"

$$\neg(\exists x)(hired(x) \land \blacklozenge fired(x))$$

"employees can not be hired if they have been fired in the past and not subsequently reinstated"

$$\neg(\exists x)(hired(x) \land (\neg reinstated(x)\ \mathbf{since}\ fired(x)))$$

"employees can be reinstated at most once"

$$\neg(\exists x)(reinstated(x) \land \blacklozenge reinstated(x))$$

Biquantified formulas

[Lipeck and Saake, 1987]

- only **future** connectives
- quantifiers either **external** (not in the scope of any temporal connective) or **internal** (no temporal connective in their scope)
- no internal quantifiers: potential constraint satisfaction decidable in exponential time [Chomicki and Niwinski, 1995]
- one internal quantifier: undecidable.

Not surprising, in view of the fact that first-order temporal logic is highly undecidable.

Past formulas

[Chomicki, 1995]:

- only past connectives
- arbitrary quantifiers
- potential constraint satisfaction undecidable
- a practical method that approximates potential constraint satisfaction:
  - check whether the constraint is satisfied in the current state

Also *real-time* temporal logic.

History encoding

Encoding in [Chomicki, 1995] :

- the database schema is augmented by *auxiliary* relations
- every database state is extended with the instances of auxiliary relations to form an *extended* state
- the extended state H'_k encodes the whole current history $(H_0, \ldots, H_k)$.

Properties of the encoding

Incrementally computable.

Lossy.

Space-efficient.

Auxiliary relations

Auxiliary relations for a constraint C:

- one auxiliary relation r_α for each temporal subformula α of C (i.e., of the form $\bullet A$ or A **since** B),
- the arity of r_α is equal to the number of free variables in α,
- auxiliary constraint relation r_C (0-ary).

Auxiliary relations are defined inductively (induction on time). The definitions are automatically derived from the constraints. The relations are implemented as materialized views.

The constraint C is satisfied iff the constraint relation r_C contains the empty tuple $()$.

Example

"employees can not be hired if they have been fired in the past and not subsequently reinstated"

$$\neg(\exists x)(\mathit{hired}(x) \land (\neg \mathit{reinstated}(x) \textbf{ since } \mathit{fired}(x)))$$

Auxiliary relations:

$$\begin{aligned} r_\alpha(x) &\stackrel{df}{=} \neg \mathit{reinstated}(x) \textbf{ since } \mathit{fired}(x) \\ r_C &\stackrel{df}{=} \neg(\exists x)(\mathit{hired}(x) \land r_\alpha(x)) \end{aligned}$$

Inductive definition of r_α:

$$\begin{aligned} r_\alpha^0(x) &\stackrel{df}{=} \mathit{False} \\ r_\alpha^{k+1}(x) &\stackrel{df}{=} (r_\alpha^k(x) \lor \mathit{fired}^k(x)) \land \neg \mathit{reinstated}^{k+1}(x) \end{aligned}$$

The superscript is the consecutive number of the state.

The definitions do not depend on the specific value of the superscript, provided it is greater than 0.

The definitions are implemented as materialized views, using an active DBMS [Chomicki and Toman, 1995].

Space and time requirements

Assumptions:

- a fixed finite set of integrity constraints IC,
- a fixed database schema.

The active domain $adom(D)$ of a history $D = (D_0, \ldots, D_k)$ is the set of domain values that appear in D.

A history encoding is *bounded* if:

1. the extended schema is fixed and finite.
2. for every $n \in N$ there is an $m \in N$ such that for every history $H = (H_0, \ldots, H_k)$, the following condition is satisfied:

   if $|adom(H)| < n$, then $size(H'_k) < m$.

Space efficiency

The encoding of [Chomicki, 1995] is (polynomially) bounded.

Other published encodings:

- [Lipeck and Saake, 1987]: bounded
- [Sistla and Wolfson, 1995]: unbounded

Time requirements.

Of course, not only databases over fixed active domains! However, the space savings of the encoding are the most spectacular when the same domain values appear in many different database states.

Bounded vs. unbounded encoding

Example.

	0	1	2	3	$\cdots$
$\{x : p(x)\}$	$\{a\}$	$\{\}$	$\{a\}$	$\{\}$	$\cdots$
$\{x : \blacklozenge p(x)\}$	$\{\}$	$\{a\}$	$\{a\}$	$\{a\}$	$\cdots$

The auxiliary relation requires constant space.

	0	1	2	3	$\cdots$
$\{x : p(x)\}$	$\{a_0\}$	$\{a_1\}$	$\{a_2\}$	$\{a_3\}$	$\cdots$
$\{x : \blacklozenge p(x)\}$	$\{\}$	$\{a_0\}$	$\{a_0, a_1\}$	$\{a_0, a_1, a_2\}$	$\cdots$

The auxiliary relation requires unbounded space because the active domain is unbounded.

References

[Abiteboul et al., 1996] Abiteboul, S., Herr, L., and den Bussche, J. V. (1996). Temporal Versus First-Order Logic to Query Temporal Databases. In *ACM Symposium on Principles of Database Systems*, Montreal, Canada.

[Allen, 1983] Allen, J. (1983). Maintaining Knowledge about Temporal Intervals. *Communications of the ACM*, 26(11):832–843.

[Attie et al., 1993] Attie, P. C., Singh, M., Sheth, A., and Rusinkiewicz, M. (1993). Specifying and Enforcing Intertask Dependencies. In *International Conference on Very Large Data Bases*, pages 134–145.

[Baudinet et al., 1993] Baudinet, M., Chomicki, J., and Wolper, P. (1993). Temporal Deductive Databases. In [Tansel et al., 1993], pages 294–320.

[Baudinet et al., 1995] Baudinet, M., Chomicki, J., and Wolper, P. (1995). Constraint-Generating Dependencies. In *International Conference on Database Theory*, Prague, Czech Republic. Springer-Verlag. Short version in: Proc. 2nd Workshop on Principles and Practice of Constraint Programming, 1994.

[Bettini et al., 1995] Bettini, C., Wang, X., Bertino, E., and Jajodia, S. (1995). Semantic Assumptions and Query Evaluation in Temporal Databases. In *ACM SIGMOD International Conference on Management of Data*, pages 257–268, San Jose, California.

[Böhlen et al., 1996] Böhlen, M., Chomicki, J., Snodgrass, R., and Toman, D. (1996). Querying TSQL2 Databases with Temporal Logic. In *International Conference on Extending Database Technology*, Avignon, France. Springer Verlag, LNCS 1057.

[Chaudhuri, 1988] Chaudhuri, S. (1988). Temporal Relationships in Databases. In *International Conference on Very Large Data Bases.*

[Chomicki, 1994] Chomicki, J. (1994). Temporal Query Languages: A Survey. In Gabbay, D. and Ohlbach, H., editors, *Temporal Logic, First International Conference*, pages 506–534. Springer-Verlag, LNAI 827.

[Chomicki, 1995] Chomicki, J. (1995). Efficient Checking of Temporal Integrity Constraints Using Bounded History Encoding. *ACM Transactions on Database Systems*, 20(2):149–186.

[Chomicki, 1997] Chomicki, J. (1997). "Temporal" Considered Harmful in Temporal Database Design. In preparation.

[Chomicki and Niwinski, 1995] Chomicki, J. and Niwinski, D. (1995). On the Feasibility of Checking Temporal Integrity Constraints. *Journal of Computer and System Sciences*, 51(3):523–535.

[Chomicki and Toman, 1995] Chomicki, J. and Toman, D. (1995). Implementing Temporal Integrity Constraints Using an Active DBMS. *IEEE Transactions on Knowledge and Data Engineering*, 7(4):566–582.

[Clifford and Croker, 1987] Clifford, J. and Croker, A. (1987). The Historical Relational Data Model (HRDM) and Algebra Based on Lifespans. In *IEEE International Conference on Data Engineering*.

[Dyreson and Snodgrass, 1993] Dyreson, C. E. and Snodgrass, R. (1993). Historical Indeterminacy. In *IEEE International Conference on Data Engineering*.

[Elmasri et al., 1990] Elmasri, R., Wuu, G. T., and Kim, Y. G. (1990). The time index : An access structure for temporal data. In *International Conference On Very Large Data Bases*, pages 1–12, Palo Alto, Ca., USA. Morgan Kaufmann Publishers, Inc.

[Gabbay et al., 1994] Gabbay, D., Hodkinson, I., and Reynolds, M. (1994). *Temporal Logic: Mathematical Foundations and Computational Aspects*. Oxford University Press.

[Gadia et al., 1992] Gadia, S., Nair, S., and Poon, Y. (1992). Incomplete Information in Relational Temporal Databases. In *International Conference on Very Large Data Bases*.

[Guttman, 1984] Guttman, A. (1984). R-trees: A dynamic index structure for spatial searching. In *ACM SIGMOD International Conference on Management of Data*, pages 47–57.

[Jensen and Snodgrass, 1994] Jensen, C. and Snodgrass, R. (1994). Temporal Specialization and Generalization. *IEEE Transactions on Knowledge and Data Engineering*, 6(6):954–974.

[Jensen et al., 1996] Jensen, C., Snodgrass, R., and Soo, M. (1996). Extending Existing Dependency Theory to Temporal Databases. *IEEE Transactions on Knowledge and Data Engineering*, 8(4).

[Kamp, 1968] Kamp, J. (1968). *Tense Logic and the Theory of Linear Order*. PhD thesis, University of California, Los Angeles.

[Kanellakis et al., 1993] Kanellakis, P., Ramaswamy, S., Vengroff, D., and Vitter, J. (1993). Indexing for Data Models with Constraints and Classes. In *ACM Symposium on Principles of Database Systems*.

[Kautz and Ladkin, 1991] Kautz, H. and Ladkin, P. (1991). Integrating Metric and Qualitative Temporal Reasoning. In *National Conference on Artificial Intelligence.*

[Koubarakis, 1994] Koubarakis, M. (1994). Database Models for Infinite and Indefinite Temporal Information. *Information Systems*, 19(2):141–174.

[Koubarakis, 1997] Koubarakis, M. (1997). The Complexity of Query Evaluation in Indefinite Temporal Constraint Databases. *Theoretical Computer Science*, 171(1):89–112.

[Ladkin, 1986] Ladkin, P. (1986). Primitives and Units for Time Specification. In *National Conference on Artificial Intelligence*, pages 354–359.

[Lipeck and Saake, 1987] Lipeck, U. and Saake, G. (1987). Monitoring Dynamic Integrity Constraints Based on Temporal Logic. *Information Systems*, 12(3):255–269.

[Lorentzos, 1993] Lorentzos, N. A. (1993). The Interval-extended Relational Model and Its Application to Valid-time Databases. In [Tansel et al., 1993], pages 67–91.

[Maher and Srivastava, 1996] Maher, M. J. and Srivastava, D. (1996). Chasing Constrained Tuple-Generating Dependencies. In *ACM Symposium on Principles of Database Systems*, pages 128–138.

[Meiri, 1991] Meiri, I. (1991). Combining Qualitative and Quantitative Constraints in Temporal Reasoning. In *National Conference on Artificial Intelligence*.

[Motakis and Zaniolo, 1997] Motakis, I. and Zaniolo, C. (1997). Temporal Aggregation in Active Database Rules. In *ACM SIGMOD International Conference on Management of Data*.

[Navathe and Ahmed, 1993] Navathe, S. and Ahmed, R. (1993). Temporal Extensions to the Relational Model and SQL. In [Tansel et al., 1993], pages 92–109.

[Ramaswamy, 1997] Ramaswamy, S. (1997). Efficient indexing for constraint and temporal databases. In *International Conference on Database Theory*, pages 419–431. Springer-Verlag.

[Saltzberg and Tsotras, 1994] Saltzberg, B. and Tsotras, V. J. (1994). A Comparison of Access Methods for Time Evolving Data. Technical Report NU-CCS-94-21, College of Comp. Sci., Northeastern University.

[Sistla and Wolfson, 1995] Sistla, A. and Wolfson, O. (1995). Temporal Triggers in Active Databases. *IEEE Transactions on Knowledge and Data Engineering*, 7(3):471–486.

[Snodgrass, 1987] Snodgrass, R. (1987). The Temporal Query Language TQuel. *ACM Transactions on Database Systems*, 12(2):247–298.

[Snodgrass, 1995] Snodgrass, R., editor (1995). *The TSQL2 Temporal Query Language*. Kluwer Academic Publishers.

[Stonebraker and Kemnitz, 1991] Stonebraker, M. and Kemnitz, G. (1991). The POSTGRES Next-Generation Database Management System. *Communications of the ACM*, 34(10):78–92.

[Tansel et al., 1993] Tansel, A., Clifford, J., Gadia, S., Jajodia, S., Segev, A., and Snodgrass, R., editors (1993). *Temporal Databases: Theory, Design, and Implementation*. Benjamin/Cummings.

[Toman, 1996] Toman, D. (1996). Point vs. Interval-based Query Languages for Temporal Databases. In *ACM Symposium on Principles of Database Systems*, Montreal, Canada.

[Toman, 1997] Toman, D. (1997). Point-based temporal extensions of SQL. Technical Report TR-CS-363, University of Toronto.

[Toman and Niwinski, 1996] Toman, D. and Niwinski, D. (1996). First-order queries over temporal databases inexpressible in temporal logic. In *International Conference on Extending Database Technology*, Avignon, France.

[van der Meyden, 1997] van der Meyden, R. (1997). The Complexity of Querying Indefinite Data about Linearly Ordered Domains. *Journal of Computer and System Sciences*, 54(1):113–135.

[Vardi, 1988] Vardi, M. Y. (1988). A Temporal Fixpoint Calculus. In *ACM Symposium on Principles of Programming Languages*.

[Wang et al., 1993] Wang, X., Jajodia, S., and Subrahmanian, V. (1993). Temporal Modules: An Approach Toward Federated Temporal Databases. In *ACM SIGMOD International Conference on Management of Data*, pages 227–236.

[Wolper, 1983] Wolper, P. (1983). Temporal Logic Can Be More Expressive. *Information and Control*, 56:72–99.