

Logical Data Expiration

A Tutorial

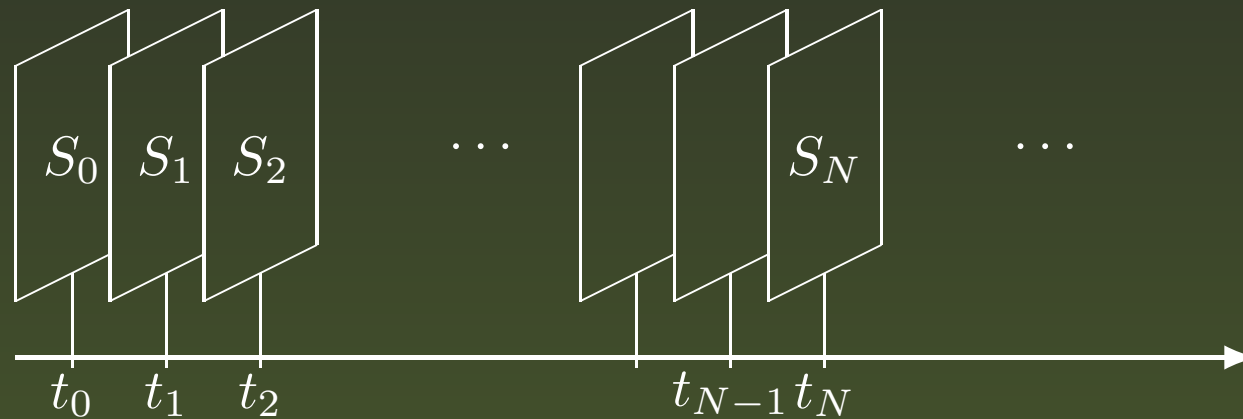
David Toman

david@uwaterloo.ca

School of Computer Science, University of Waterloo, Canada

Data Evolution and Histories

Changes of data are captured (conceptually) by *histories*:



- states S_i describe *system state*
- transitions $S_i \rightarrow S_{i+1}$ represent system evolution
 \Rightarrow **append only histories (new states at the end)**

Data Access and Queries

Data is accessed using *queries*

- simple value look-ups vs. complex query languages
- current state only vs. access to *past states*
 - analysis of data warehouse evolution
 - enforcement of temporal integrity constraints
 - monitoring applications
 - ...

Expiration

Question: What data do we need to keep?

Expiration

Question: What data do we need to keep?

1. *Policy-driven* expiration
2. *Query-driven* (logical) expiration

Expiration

Question: What data do we need to keep?

1. *Policy*-driven expiration
2. *Query*-driven (logical) expiration

Data to be expired is determined by
the (class of) queries we are allowed to
ask w.r.t. **all possible extensions** of a history

Examples

- Record keeping/business rules:
 - ⇒ tax forms must be kept 5 years back
- Enforcing *dynamic integrity constraints*:
 - ⇒ don't hire anyone you've fired in the past
- Caching policy management:
 - ⇒ what data should be moved to backup storage?
- Moving window queries, etc...

Outline of the Talk

- Temporal Database Primer
- Expiration Operators
 - how good is an expiration operator?*
- Administrative Approaches to Expiration
 - ⇒ materialized views and queries
- Query-driven Expiration
 - ⇒ Temporal Logic and Materialized Views
 - ⇒ First-order Queries and Partial Evaluation
 - space limits for expiration operators*
- Infinite Extensions of Histories
 - ⇒ Certain/Potential Answers

TDB Primer

Temporal Databases and Histories

System states: Relational structures (fixed schema)

Time: discrete (integer-like) $\{0, \dots, N, \dots\}$

1. Snapshot Temporal Database:

⇒ time-indexed sequence of relational structures

⇒ **append-only**: $H; D_{N+1}$

2. Timestamp Temporal Database:

⇒ time-indexed tuples (using a *temporal attribute*)

Choices 1 and 2 equivalent [Chomicki and Toman, 1998]

Example

Information about TA and courses by semester:

Snapshot

0	{(John, CS448)}
1	{(John, CS448), (Sue, CS234)}
2	{(John, CS448)}
3	{(Sue, CS234)}

Timestamp

{ (0, John, CS448), (1, John, CS448), (1, Sue, CS234), (2, John, CS448), (3, Sue, CS234) }
--

Temporal Queries

Queries: first-order formulas (over a fixed schema)

1. Temporal logic (FOTL)

⇒ modal (temporal) connectives

⇒ implicit references to time

2. Temporal Relational Calculus (2-FOL):

⇒ temporal variables/attributes/quantifiers

⇒ explicit access to time and ordering of time

Proposition 1 *FOTL cannot express all 2-FOL queries.*

[Abiteboul et al., 1996, Toman and Niwinski, 1996, Toman, 2003b]

Examples

Students who TA'ed at least one class twice:

- in (past) FOTL:

$$\{x : \blacklozenge(\exists y. \text{TA}(x, y) \wedge \bullet\blacklozenge\text{TA}(x, y))\}$$

- in 2-FOL:

$$\{x : \exists t_1, t_2. t_1 < t_2 \wedge \exists y. \text{TA}(t_1, x, y) \wedge \text{TA}(t_2, x, y)\}$$

Examples

Students who TA'ed at least one class twice:

- in (past) FOTL:

$$\{x : \blacklozenge(\exists y. \text{TA}(x, y) \wedge \bullet\blacklozenge\text{TA}(x, y))\}$$

- in 2-FOL:

$$\{x : \exists t_1, t_2. t_1 < t_2 \wedge \exists y. \text{TA}(t_1, x, y) \wedge \text{TA}(t_2, x, y)\}$$

$$\exists t_1, t_2. t_1 < t_2 \wedge \forall x, y. \text{TA}(t_1, x, y) \iff \text{TA}(t_2, x, y)$$

cannot be expressed in FOTL

Finite vs. Infinite Histories

Semantics of queries defined w.r.t:

- current (finite) history
 - ⇒ query evaluation on a finite temporal database
- a completion of current history
 - ⇒ hypothetical reasoning

Data Expiration

Expiration Operator

An *expiration operator* is a triple $(0^\mathcal{E}, \Delta^\mathcal{E}, Q^\mathcal{E})$ s.t.:

1. it provides an *inductive definition*

$$\mathcal{E}(\langle \rangle) = 0^\mathcal{E} \quad (\text{initial state})$$

$$\mathcal{E}(H; D) = \Delta^\mathcal{E}(\mathcal{E}(H), D) \quad (\text{extension maintenance})$$

2. it maintains the following invariant:

$$Q(H) = Q^\mathcal{E}(\mathcal{E}(H)) \quad (\text{answer preservation})$$

Examples

- the *identity* operator:

$$\begin{aligned} 0^{\mathcal{E}_{\text{id}}} &= \langle \rangle & Q^{\mathcal{E}_{\text{id}}} &= Q \\ \Delta^{\mathcal{E}_{\text{id}}} &= \lambda H \lambda S. H; S \end{aligned}$$

- the *current* operator:

$$\begin{aligned} 0^{\mathcal{E}_{\text{now}}} &= \langle \rangle & Q^{\mathcal{E}_{\text{now}}} &= Q \\ \Delta^{\mathcal{E}_{\text{now}}} &= \lambda H \lambda S. \langle S \rangle \end{aligned}$$

the queries *preserved* are *different*...

Another Example

- (lossless) *compression* based operators:

$$0^{\mathcal{E}_{\text{compress}}} = \text{compress}(\langle \rangle)$$

$$\Delta^{\mathcal{E}_{\text{compress}}} = \lambda H \lambda S. \text{compress}(\text{decompress}(H); S)$$

$$Q^{\mathcal{E}_{\text{compress}}} = \lambda H. Q(\text{decompress}(H))$$

⇒ compress and decompress are *lossless*

... *no reduction*; $|H| \sim |\mathcal{E}_{\text{compress}}(H)|$

⇒ special case: *interval timestamps*.

Expiration vs. Queries Revisited

- Given an *expiration operator*
 - for what class of queries it preserves answers?
⇒ can these be characterized syntactically?
- Given a *set* of temporal queries:
 - is there an expiration operator that
⇒ maintains answers to these queries?
⇒ can be found algorithmically?
 - for what *query languages*?

How Good is It ?

What is the space needed by $\mathcal{E}(H)$ in terms of

- size of the original history, $|H|$,
- length of H (number of states, $|\text{dom}_T|$),
- the size of the *active data domain* of H (number of constants that have appeared in H , $|\text{dom}_D|$),
- size of the answer $Q(H)$,
- size of the queries.

How Good is It ?

What is the space needed by $\mathcal{E}(H)$ in terms of

- size of the original history, $|H|$,
- length of H (number of states, $|\text{dom}_T|$),
- the size of the *active data domain* of H (number of constants that have appeared in H , $|\text{dom}_D|$),
- size of the answer $Q(H)$,
- size of the queries.

General Goal: make $|\mathcal{E}(H)|$ independent of length of H . \Rightarrow **bounded expiration operator**

Example

Proposition 2 \mathcal{E}_{now} *is bounded.*

Proposition 3 $\mathcal{E}_{\text{compress}}$ *cannot be bounded for lossless compression schemes.*

\mathcal{E}_Q for a temporal query Q in a language \mathcal{L} ?

... depends on (expressive power) of \mathcal{L}

Administrative Expiration Policies

Administrative Approaches

Query-independent expiration policies.

- characterize queries whose answers are not affected

Administrative Approaches

Query-independent expiration policies.

- characterize queries whose answers are not affected
 - expiration operator = a view of the history
 - ⇒ the view must be *self-maintainable*
 - query reformulation = query over the view
 - ⇒ *answering queries over views* problem

Administrative Approaches

Query-independent expiration policies.

- characterize queries whose answers are not affected
 - expiration operator = a view of the history
 - ⇒ the view must be *self-maintainable*
 - query reformulation = query over the view
 - ⇒ *answering queries over views* problem
- detect attempts to access the *missing data*
 - ⇒ at run-time

Cutoff Points

Common approach: *history truncation* or *cutoff point*

1. policies based on fixed *absolute cutoff point*, or
2. policies based on *now-relative cutoff point*.

⇒ generalization of the \mathcal{E}^{id} and the \mathcal{E}^{now} operators

Cutoff Points

Common approach: *history truncation* or *cutoff point*

1. policies based on fixed *absolute cutoff point*, or
2. policies based on *now-relative cutoff point*.

⇒ generalization of the \mathcal{E}^{id} and the \mathcal{E}^{now} operators

Example: Vacuuming [Jensen, 1995]:

- $\rho(R) : e$ (a *remove* specification), and
- $\kappa(R) : e$ (a *keep* specification).

R is a temporal relation; e a selection condition

- absolute/*now*-relative specifications

A Query-driven Expiration: Finite Histories

Query Driven Expiration

Expiration for queries in a temporal query language

- ⇒ Past FOTL (and variants)
- ⇒ Future FOTL
- ⇒ 2-FOL (temporal relational calculus)

Query Driven Expiration

Expiration for queries in a temporal query language

⇒ Past FOTL (and variants)

⇒ Future FOTL

⇒ 2-FOL (temporal relational calculus)

- *Finite relational structures can be completely characterized by first-order queries.*

⇒ best expiration operator for a fixed query Q .

Query Driven Expiration

Expiration for queries in a temporal query language

⇒ Past FOTL (and variants)

⇒ Future FOTL

⇒ 2-FOL (temporal relational calculus)

■ *Finite relational structures can be completely characterized by first-order queries.*

⇒ best expiration operator for a fixed query Q .

■ *Optimal expiration operator cannot exist.*

⇒ we look for a **bounded expiration operator**.

Query Driven Approaches

1. Removal of “old” states (expiration)

- ⇒ removes a **subset** of existing states
- ⇒ no other changes (history \rightarrow history)

2. Auxiliary (non-temporal) view maintenance

- ⇒ maintains **auxiliary** relations
- ⇒ maps a history to *a single extended state*

3. Specialization of queries

- ⇒ specialize a query w.r.t. the known prefix H .

Past Temporal Logic

- Syntax: First-order logic past temporal operators

$$Q ::= R(\mathbf{x}) \mid F \mid Q \wedge Q \mid \neg Q \mid \exists x.Q \mid \bullet Q \mid Q \text{ since } Q$$

- Semantics:

$$Q(H) = \{\theta : H, \theta, n \models Q\}$$

\Rightarrow queries over unbounded past: $\{x : \blacklozenge R(x)\}$

Unfolding and Materialized Views

- Crux of the approach:

$$Q_1 \text{ since } Q_2 \equiv Q_1 \wedge \bullet(Q_2 \vee (Q_1 \text{ since } Q_2))$$

- *auxiliary views* for temporal subformulas
⇒ only the “previous” state needed

Unfolding and Materialized Views

- Crux of the approach:

$$Q_1 \text{ since } Q_2 \equiv Q_1 \wedge \bullet(Q_2 \vee (Q_1 \text{ since } Q_2))$$

- *auxiliary views* for temporal subformulas
 \Rightarrow only the “previous” state needed
- recurrent definitions to maintain the views.

α	R_α^0	R_α^n
$\bullet Q$	false	Q^{n-1}
$Q_1 \text{ since } Q_2$	false	$Q_1^n \wedge (Q_2^{n-1} \vee R_\alpha^{n-1})$

Example

- Query: *Students that TA'ed at least one class twice.*

$$\{x : \blacklozenge(\exists y. \text{TA}(x, y) \wedge \bullet\blacklozenge\text{TA}(x, y))\}$$

- Temporal subqueries:

$$\alpha_1 = \blacklozenge\text{TA}(x, y) \text{ and}$$

$$\alpha_2 = \bullet\blacklozenge\text{TA}(x, y) \text{ and}$$

$$\alpha_3 = \blacklozenge\exists y. \text{TA}(x, y) \wedge \bullet\blacklozenge\text{TA}(x, y).$$

Example (cont.)

Inductive maintenance of views:

$R_{\alpha_1}(x, y)$	$R_{\alpha_2}(x, y)$	$R_{\alpha_3}(x)$
{ (John, CS448) }	{ }	{ }
{ (John, CS448), (Sue, CS234) }	{ (John, CS448) }	{ John }
{ (John, CS448), (Sue, CS234) }	{ (John, CS448), (Sue, CS234) }	{ John }
{ (John, CS448), (Sue, CS234) }	{ (John, CS448), (Sue, CS234) }	{ John, Sue }

Space Utilization

- $Q = \diamond(p(x_1) \wedge \dots \wedge p(x_k))$
- $H = \langle \{a_1\}, \{a_2\}, \{a_3\}, \dots, \{a_n\} \rangle$.

For $\alpha = \diamond(p(x_1) \wedge \dots \wedge p(x_k))$:

$$|R_\alpha| = (n - 1)^k$$

... the same holds for every prefix of H .

Full details: [Chomicki, 1995],

\Rightarrow subsumes approaches based on TRA

[Yang and Widom, 1998, Yang and Widom, 2000].

Adding Fixpoints

- Syntax:

$$Q ::= R(\mathbf{x}) \mid F \mid Q \wedge Q \mid \neg Q \mid \exists x.Q \mid \bullet Q \mid \mu X.Q.$$

- Unfolding a fixpoint: $\mu X.Q \equiv Q(\mu X.Q)$
- Inductive maintenance of auxiliary relation:

α	R_{α}^0	R_{α}^n
$\bullet Q$	false	Q^{n-1}

\Rightarrow careful definition of Q^{n-1} .

Full details: [Toman, 2003a]

Metric Temporal Logic

- access to *real time* time instants
 - ⇒ a clk constant in each state (current *real time*)
 - ⇒ not part of the active data domain
- additional temporal operators

$$Q ::= \dots \mid \text{since}_{\sim c} \mid \bullet_{\sim c}$$

- ⇒ semantics respects $\sim c$ distances
- materialized views now contain *distance* values
 - ⇒ bounded by c
 - ⇒ bounded expiration if $\text{clk}^i - \text{clk}^{i-1} \geq \epsilon > 0$

Future Temporal Logic

- Syntax:

$Q ::= R(\mathbf{x}) \mid F \mid Q \wedge Q \mid \neg Q \mid \exists x.Q \mid \circ Q \mid Q \text{ until } Q$

- Semantics: $Q(H) = \{\theta : H, \theta, 0 \models Q\}$

\Rightarrow *still active domain semantics*

Future Temporal Logic

- Syntax:

$$Q ::= R(\mathbf{x}) \mid F \mid Q \wedge Q \mid \neg Q \mid \exists x.Q \mid \circ Q \mid Q \text{ until } Q$$

- Semantics: $Q(H) = \{\theta : H, \theta, 0 \models Q\}$

\Rightarrow *still active domain semantics*

- Unfolding rule (similarly to PastTL):

$$Q_1 \text{ until } Q_2 \equiv Q_1 \wedge (\circ Q_2 \vee \circ(Q_1 \text{ until } Q_2))$$

\Rightarrow now we need to represent a formula with *holes* to be substituted when the history is extended.

Biquantified Formulas

- [Lipeck and Saake, 1987, Lipeck et al., 1994]
 - ⇒ restrictions to Future FOTL syntax: 3 layers
 1. FO formulas (evaluated in a *state*),
 2. TL(FO) formulas: temporal outside (1),
 3. Universal quantifiers on top of (2)
- Automata-based approach
 - ⇒ designed in the *propositional* setting
 - ⇒ mix quantifiers and temporal connectives?
 - ... bounded expiration based on an automaton for (2) implemented by *triggers*

Two-sorted First-order Language

- Temporal Relational Calculus (2-FOL)

$$L ::= R(t, \mathbf{x}) \mid x = x' \mid t < t' \mid L \wedge L \mid L \wedge \neg L \mid L \vee L \mid \exists x.L \mid \exists t.L$$

for $R(t, \mathbf{x})$ true in H iff $R(\mathbf{x})$ is true in D_t

A bounded expiration operator for 2-FOL?

\Rightarrow conjectured that it does NOT exist

Expiration Revisited

Idea: remove those states that

1. do not contribute to query answer (due to \wedge)
2. contribute duplicate information (due to \exists)

Expiration Revisited

Idea: remove those states that

1. do not contribute to query answer (due to \wedge)
2. contribute duplicate information (due to \exists)

Easy for a fixed history:

- \Rightarrow compute answer to Q bottom-up
- \Rightarrow propagate “back” to remove redundant data

NOTE: 2-FOL queries with *unbounded answers* cannot have bounded expiration operator

\Rightarrow consider only *bounded queries*

Handling History Extensions

Atomic formulas:

- $\begin{bmatrix} x \\ a \end{bmatrix} \equiv \begin{cases} x = a & a \in \text{dom}_D \\ \forall a \in \text{dom}_D. x \neq a & a = \bullet \end{cases}$
- $\begin{bmatrix} t \\ s \end{bmatrix} \equiv \begin{cases} t = s & s \in \text{dom}_T \\ t > \text{maxtime}(\text{dom}_T) & s = \bullet \end{cases}$

Specialization of base relations *and their extensions*:

$$R(t, \mathbf{x}) \equiv \left(\bigvee_{\mathbf{a} \in R_{D_s}} \text{true}_{\begin{bmatrix} t\mathbf{x} \\ s\mathbf{a} \end{bmatrix}} \right) \vee \left(\bigvee_{\mathbf{a} \in \text{dom}_D \cup \{\bullet\}} R(t, \mathbf{x})_{\begin{bmatrix} t\mathbf{x} \\ \bullet\mathbf{a} \end{bmatrix}} \right)$$

\Rightarrow depends **only** on the future extensions of history

Query Specialization

$$\text{PE}_H(Q) = \left\{ \begin{array}{l}
 \{\text{true}_{[s\mathbf{a}]}^{t\mathbf{x}} : R(s, \mathbf{a}) \in D\} \\
 \cup \{R(t, \mathbf{x})_{[\bullet\mathbf{a}]}^{t\mathbf{x}} : \mathbf{a} \in (\text{dom}_D \cup \{\bullet\})^{|\mathbf{x}|}\} \quad Q \equiv R(t, \mathbf{x}) \\
 \\
 \{Q'_1[\mathbf{a}] : Q'_1[\mathbf{a}] \in \text{PE}_H(Q_1), \models [\mathbf{a}] \wedge F\} \quad Q \equiv Q_1 \wedge F \\
 \\
 \{Q'_1 \wedge Q'_2[\mathbf{ab}] : Q'_1[\mathbf{a}] \in \text{PE}_H(Q_1), Q'_2[\mathbf{b}] \in \text{PE}_H(Q_2), \models [\mathbf{ab}]\} \quad Q \equiv Q_1 \wedge Q_2 \\
 \\
 \{(\exists y. \bigvee_{Q'_1[\mathbf{ab}] \in \text{PE}_H(Q_1)} Q'_1[\mathbf{a}] : \exists b. Q''_1[\mathbf{ab}] \in \text{PE}_H(Q_1)\} \quad Q \equiv \exists y. Q_1 \\
 \\
 \{(\exists t. \bigvee_{Q'_1[\mathbf{as}] \in \text{PE}_H(Q_1)} Q'_1[\mathbf{a}] : \exists s. Q''_1[\mathbf{as}] \in \text{PE}_H(Q_1)\} \quad Q \equiv \exists t. Q_1 \\
 \\
 \{Q'_1 \wedge \neg Q'_2[\mathbf{a}] : Q'_1[\mathbf{a}] \in \text{PE}_H(Q_1), Q'_2[\mathbf{a}] \in \text{PE}_H(Q_2)\} \\
 \cup \{Q'_1[\mathbf{a}] : Q'_1[\mathbf{a}] \in \text{PE}_H(Q_1), Q'_2[\mathbf{a}] \notin \text{PE}_H(Q_2)\} \quad Q \equiv Q_1 \wedge \neg Q_2 \\
 \\
 \{Q'_1 \vee Q'_2[\mathbf{a}] : Q'_1 \in \text{PE}_H(Q_1)[\mathbf{a}], Q'_2[\mathbf{a}] \in \text{PE}_H(Q_2)\} \\
 \cup \{Q'_1[\mathbf{a}] : Q'_1[\mathbf{a}] \in \text{PE}_H(Q_1), Q'_2[\mathbf{a}] \notin \text{PE}_H(Q_2)\} \\
 \cup \{Q'_2[\mathbf{a}] : Q'_1[\mathbf{a}] \notin \text{PE}_H(Q_1), Q'_2[\mathbf{a}] \in \text{PE}_H(Q_2)\} \quad Q \equiv Q_1 \vee Q_2
 \end{array} \right.$$

Duplicate Information Removal

IDEA: Modify the PE_H for quantification over time:

$$(\exists t. \bigvee_{\substack{Q'_1[\mathbf{x}t] \in PE_H(Q_1) \\ s \in TB_a(t)}} Q'_1) [\mathbf{x}]$$

where $Q''_1[\mathbf{x}t] \in PE_H(Q_1)$ for some s

Duplicate Information Removal

IDEA: Modify the PE_H for quantification over time:

$$(\exists t. \bigvee_{\substack{Q'_1[\mathbf{x}t] \in PE_H(Q_1) \\ s \in TB_a(t)}} Q'_1) [\mathbf{x}]_a$$

← what is this??

where $Q''_1[\mathbf{x}t]_a \in PE_H(Q_1)$ for some s

Equivalence w.r.t. History Extensions

Definition 1: Let $Q_1[\frac{\mathbf{x}t}{\mathbf{a}s_1}]$, $Q_2[\frac{\mathbf{x}t}{\mathbf{a}s_2}] \in \text{PE}_H(Q)$ for $s_1 \neq s_2$.

We define $[\frac{\mathbf{x}t}{\mathbf{a}s_1}] \sim_Q^H [\frac{\mathbf{x}t}{\mathbf{a}s_2}]$ iff for any extension H' of H

$$(\mathbf{a}, s_1) \in Q(H; H') \iff (\mathbf{a}, s_2) \in Q(H; H')$$

Definition 2: $\text{TB}_a(t)$ is the set of representatives of the

$[\frac{\mathbf{x}t}{\mathbf{a}s_1}] \sim_Q^H [\frac{\mathbf{x}t}{\mathbf{a}s_2}]$ equivalence classes [e.g., min in $<$].

Residual History Reconstruction

- *Specialization-based* expiration:

$$\begin{aligned}Q(H) &= \text{PE}_H(Q)(\emptyset) \\ \text{PE}_{H;H'}(Q) &\equiv \text{PE}_{H'}(\text{PE}_H(Q))\end{aligned}$$

- We use $\text{PE}_H(Q)$ to *construct* $\mathcal{E}(H)$
 \Rightarrow temporal variable $t_i \rightarrow$ a unary relation

$$T_i(t) = \bigcup_a \text{TB}_a(t).$$

- \Rightarrow each quantifier $\exists t_i.Q'$ in Q is restricted to $T_i(t)$
- \Rightarrow a state $D_j \in H$ expires if $j \notin \bigcup_i T_i$.

Properties

- $Q(H; H') = Q(\mathcal{E}_Q(H); H')$
for all H, H' histories and Q FO query
- $|\mathcal{E}_Q(H)| \leq f(|\mathbf{dom}_D|, |Q|)$,
 f is an exponential tower in number of nested $\neg\exists$
- $|\mathcal{E}_Q(H)| \leq |H| + |\mathbf{dom}_T||Q|$

... and can be implemented by FO queries/updates.

Space: Lower Bounds

Example: $\exists t_1, t_2. t_1 < t_2 \wedge \forall x. R(t_1, x) \iff R(t_2, x)$

- Potentially we need to keep all states for which R contains distinct subsets of dom_D
 - \Rightarrow potentially all subsets of dom_D
 - \Rightarrow any residual history is exponential in $|\text{dom}_D|$.
- *sequences of states* yield more exponents.
- non-elementary blowup even when translating monadic FO to propositional TL

General Lower Bounds

Limits of Bounded Encoding

Clearly, this cannot work for all possible queries:

Example 1: Query $\{t : R(t)\}$.

$$\text{answer} \sim |\text{dom}_T H|$$

Example 2: Query $\{t : R(t) \wedge \forall t'. R(t') \rightarrow t \geq t'\}$.

$$\text{answer} \sim \log(|\text{dom}_T H|)$$

Counting

Example: “is the number of states containing a greater than the number of states containing b ?”

\Rightarrow we need $\Omega(\log(|\mathbf{dom}_T|))$ space for counter(s)

$\Rightarrow \mathcal{O}(\log(|\mathbf{dom}_T|))$ is sufficient.

Conjecture: we can use the above technique (but remember counts of the expired values) to answer queries with counting

$\Rightarrow |\mathcal{E}_Q(H)| \leq \text{POLY}(\log(|\mathbf{dom}_T|))$

Duplicates

Example (in SQL-style syntax):

```
( select '1'
  from R
  where R.x='a' ) except all ( select '1'
                              from R
                              where R.x='b' )
```

... is nonempty iff

the number of states containing a is greater than
the number of states containing b .

\Rightarrow just like counting ...

Retroactive Updates

Example:

while $\exists t.R(t, a) \wedge \exists t.R(t, b)$ **do** { while both a and b exist in R }

delete $R(t, a)$

where $\forall t'.R(t', a) \supset t' > t;$ { delete (chronologically) first a }

delete $R(t, b)$

where $\forall t'.R(t', b) \supset t' > t;$ { delete (chronologically) first b }

return $\exists t.R(t, a)$ { return true if R contains an a }

\Rightarrow we need $\Omega(\log(|\text{dom}_T|))$ space for counter(s)

\Rightarrow just like for counting ...

Full Future μ TL

No bounded operator can exist:

\Rightarrow [Toman, 2003a] shows $\Omega(|\text{dom}_T|)$ lower bound

$$\varphi = \exists x, y. \diamond(Q(x, y) \wedge \mu X. R(x, y) \vee \circ \exists z. X(x, z) \wedge X(z, y))$$

State	Database instance
-------	-------------------

0	$\{\}$ or $\{Q(a, b)\}$
---	-------------------------

\vdots	\vdots
----------	----------

n	$\{\}$ or $\{Q(a, b)\}$
-----	-------------------------

$n + 1$	$\{R(a, c_1), \dots, R(c_i, c_{i+1}), \dots, R(c_k, b)\}$
---------	---

Certain and Potential Answers

Infinite Histories

Definition 4 *Let H be a finite history, Q a query (in an appropriate query language), and θ a substitution.*

- *θ is a potential answer for Q with respect to H if there is an infinite completion H' of H such that $H', \theta \models Q$.*
- *θ is a certain answer for Q with respect to H if for all infinite completions H' of H we have $H', \theta \models Q$.*

potential answer: a direct generalization of of potential constraint satisfaction [Chomicki, 1995].

Infinite Histories (cont.)

Proposition 5 ([Gabbay et al., 1994]) *Satisfaction for two dimensional propositional temporal logic over natural numbers-based time domain is not decidable.*

Proposition 6 ([Chomicki, 1995]) *For past formulas potential constraint satisfaction is undecidable.*

Proposition 7 ([Chomicki and Niwinski, 1995])
For biquantified formulas

⇒ no internal quantifiers: EXPTIME,

⇒ a single internal quantifier: undecidable.

... data expiration is a moot point

Related Problems

- garbage collection in programming languages
⇒ *navigational* query language (ala IMS)
based on *reachability queries*
- data streams and streaming queries

data stream	=	history
synopse(is)	=	residual history
streaming query	=	temporal query
standing query	=	fixed query
	...	

Open Problems

FutureTL

⇒ expiration operator for full FOETL

Rich Temporal Domains (more than linear \leq)

⇒ constraint DB techniques [Libkin et al., 2000]

Space Bounds For Aggregate Queries

⇒ a weaker bound, e.g., $|\mathcal{E}H| \in O(\log(|\text{dom}_T H|))$?

Decidable Certain/Potential Answers

⇒ Decidable languages (monodic TL)

⇒ Optimal Expiration Operators?

Acknowledgment

- A chapter in *Logics for Emerging Applications of Databases*, J. Chomicki, G. Saake, and R. van der Mayden (eds.), Springer 2003.
(<http://db.uwaterloo.ca/~david/book-lead.ps>)
- Part of this research was done while at  BRICS
Centre for Basic Research in Computer Science
funded by the Danish National Science Foundation.
- The research was supported by the National Sciences and Engineering Research Council of Canada (NSERC).

Advertisement

Call For Papers

The 12th International Symposium on TEMPORAL REPRESENTATION AND REASONING (TIME 2005)

Burlington, Vermont, USA, June 23-25, 2005

URL: <http://time2005.cse.buffalo.edu/>

Papar submission: 11 pages on January 22, 2005.

References

- [Abiteboul et al., 1996] Abiteboul, S., Herr, L., and Van den Bussche, J. (1996). Temporal Versus First-Order Logic to Query Temporal Databases. In *ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems*, pages 49–57.
- [Chomicki, 1995] Chomicki, J. (1995). Efficient Checking of Temporal Integrity Constraints Using Bounded History Encoding. *TODS*, 20(2):149–186.
- [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, 1998] Chomicki, J. and Toman, D. (1998). Temporal Logic in Information Systems. In Chomicki, J. and Saake, G., editors, *Logics for Databases and Information Systems*, pages 31–70. Kluwer.
- [Gabbay et al., 1994] Gabbay, D. M., Hodkinson, I. M., and Reynolds, M. (1994). *Temporal Logic: Mathematical Foundations and Computational Aspects*. Oxford University Press.
- [Jensen, 1995] Jensen, C. S. (1995). Vacuuming. In Snodgrass, R. T., editor, *The TSQL2 Temporal Query Language*, pages 447–460.
- [Libkin et al., 2000] Libkin, L., Kuper, G., and Paredaens, J., editors (2000). *Constraint Databases*. Springer.
- [Lipeck et al., 1994] Lipeck, U. W., Gertz, M., and Saake, G. (1994). Transitional Monitoring of Dynamic Integrity Constraints. *IEEE Data Engineering Bulletin*.
- [Lipeck and Saake, 1987] Lipeck, U. W. and Saake, G. (1987). Monitoring Dynamic Integrity Constraints Based on Temporal Logic. *Information Systems*, 12(3):255–269.

- [Toman, 2003a] Toman, D. (2003a). Logical Data Expiration for Fixpoint Extensions of Temporal Logics. In *International Symposium on Spatial and Temporal Databases*, page to appear. Springer LNCS.
- [Toman, 2003b] Toman, D. (2003b). On Incompleteness of Multi-dimensional First-order Temporal Logics. In *International Symposium on Temporal Representation and Reasoning*, page to appear. IEEE Press.
- [Toman and Niwinski, 1996] Toman, D. and Niwinski, D. (1996). First-Order Queries over Temporal Databases Inexpressible in Temporal Logic. In *Advances in Database Technology, EDBT'96*, volume 1057, pages 307–324. Springer.
- [Yang and Widom, 1998] Yang, J. and Widom, J. (1998). Maintaining Temporal Views over Non-Temporal Information Sources for Data Warehousing. In *Advances in Database Technology, EDBT'98*, pages 389–403.
- [Yang and Widom, 2000] Yang, J. and Widom, J. (2000). Temporal View Self-Maintenance. In *Advances in Database Technology, EDBT'00*, pages 395–412.