

Secure XML Database Access with Views

SecReT'09

Benoit Groz

(joint work with Anne-Cécile Caron, Yves Roos, Sławek Staworko,
Sophie Tison)

Mostraré

10 juillet 2009

Securing databases with views

Many ways to enforce access control for XML. Among others:

- Checking the queries:

- ▶ statically ⇒ may reject proper queries and access
[Oasis project: XACML]
- ▶ dynamically ⇒ incurs costly runtime security check
[Murata et al. CCS'03]

Securing databases with views

Many ways to enforce access control for XML. Among others:

- Checking the queries:
 - ▶ statically ⇒ may reject proper queries and access
[Oasis project: XACML]
 - ▶ dynamically ⇒ incurs costly runtime security check
[Murata et al. CCS'03]
- Annotating the data:
 - ▶ annotating the data, or materializing the view ⇒ expensive maintenance [Damiani et al. EDBT'00, Cho et al. VLDB'02]
 - ▶ **annotating the DTD with *Non-materialized view***
Rewriting queries from the view to the document
[Fan et al. SIGMOD'04, Vercammen et al, Rassadko et al ...]

Outline

1 Non-materialized views and query rewriting

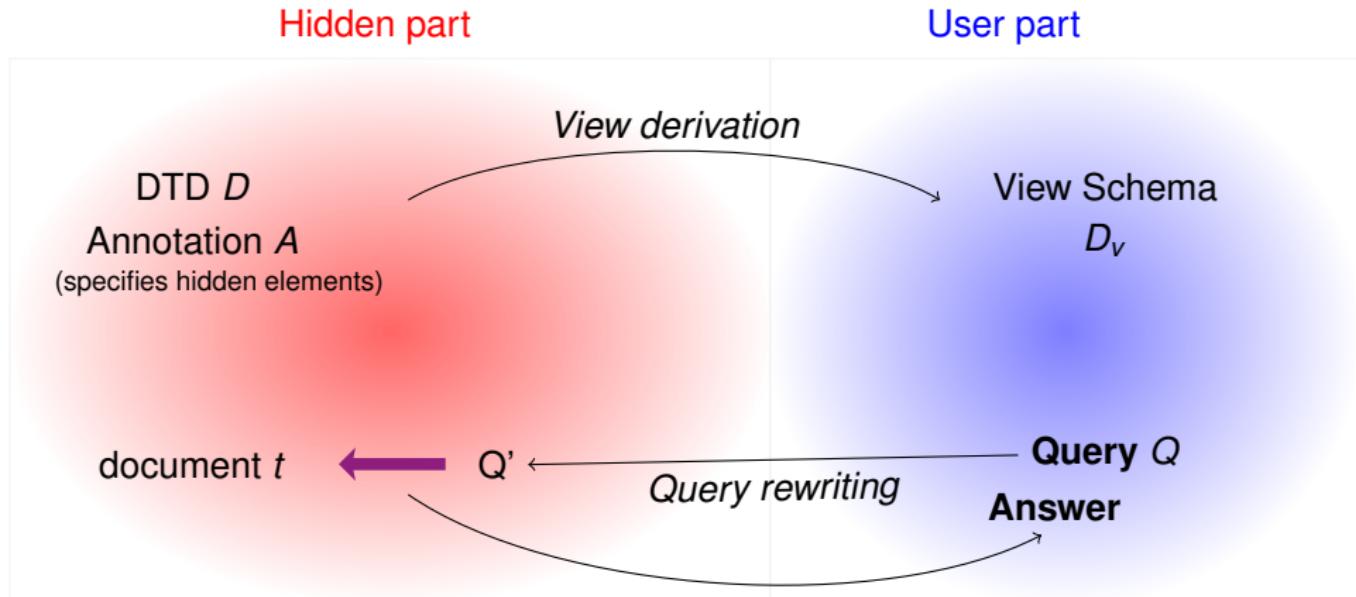
2 Comparing Access Control Policies

Visibilium omnium... et invisibilium

“Whoever wishes to keep a secret
must hide the fact that he possesses one”.

attributed to Johann Wolfgang von Goethe

Overview

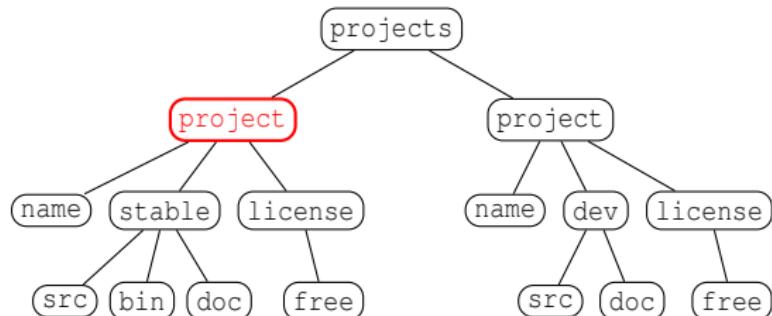


Answer to query Q = evaluation of Q'
on the original document t

Framework: XML

- XML document=tree.
- No data-values.

```
<projects>
  <project>
    <name>
    </name>
    ...
  </project>
  <project>
    ...
  </project>
</projects>
```

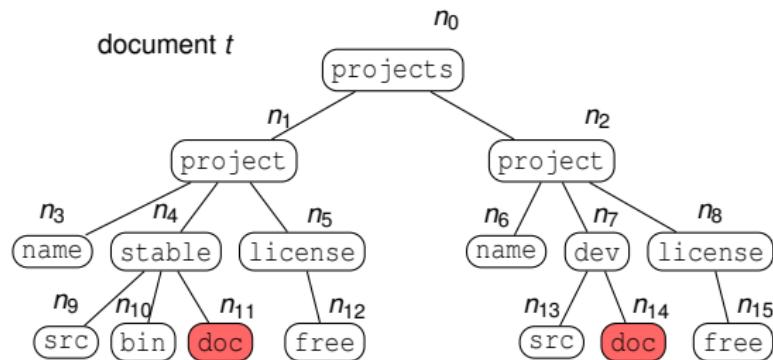


RegularXPath

- We use *Regular XPath queries*

Query $q_1 = \Downarrow^*/\Downarrow :: doc$

Ans(q_1, t) = { n_{11}, n_{14} }



“get all documentations”

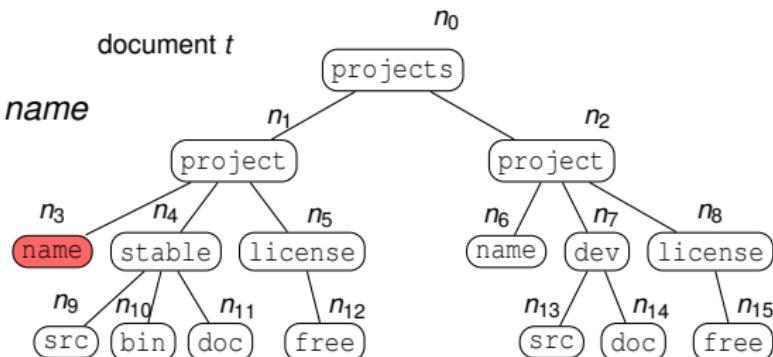
Regular XPath

- We use *Regular XPath* queries

Query $q_2 =$

$\Downarrow :: project[\Downarrow :: stable]/\Downarrow :: name$

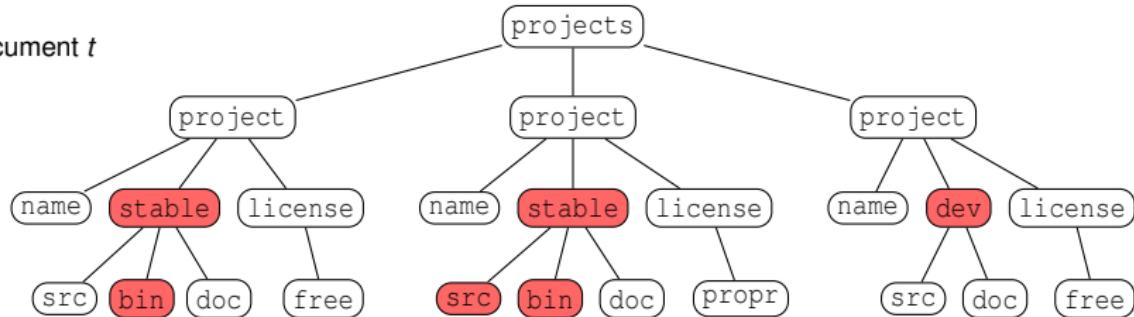
$\text{Ans}(q_2, t) = \{n_3\}$



“get names of stable projects”

Access control for XML

document t

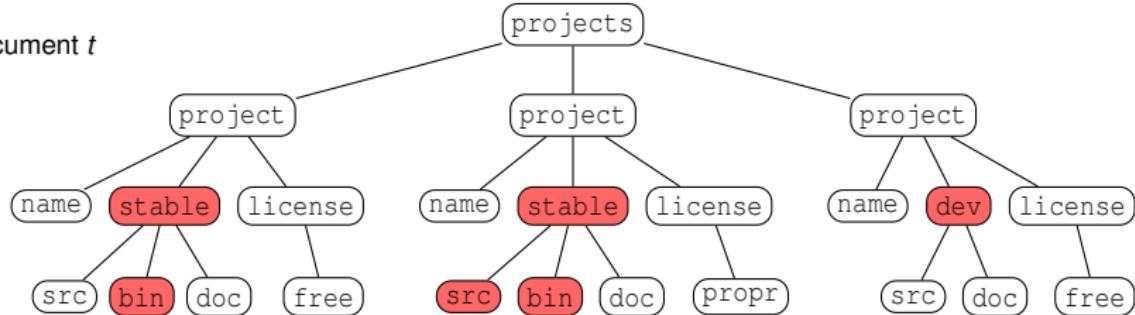


We wish to hide:

- whether a project is *stable* or *in-development*
- the *binaries*
- the *sources* for non-free projects

DTD and Annotation

document t



Example

$\text{projects} \rightarrow \text{project}^*$

$\text{project} \rightarrow \text{name}, (\text{stable} \mid \text{dev}), \text{license}$

$A_0(\text{project}, \text{stable}) = \text{false}$

$A_0(\text{project}, \text{dev}) = \text{false}$

$\text{license} \rightarrow \text{free} \mid \text{proper}$

$\text{stable} \rightarrow \text{src}, \text{bin}, \text{doc}$

$A_0(\text{stable}, \text{src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

$A_0(\text{stable}, \text{doc}) = \text{true}$

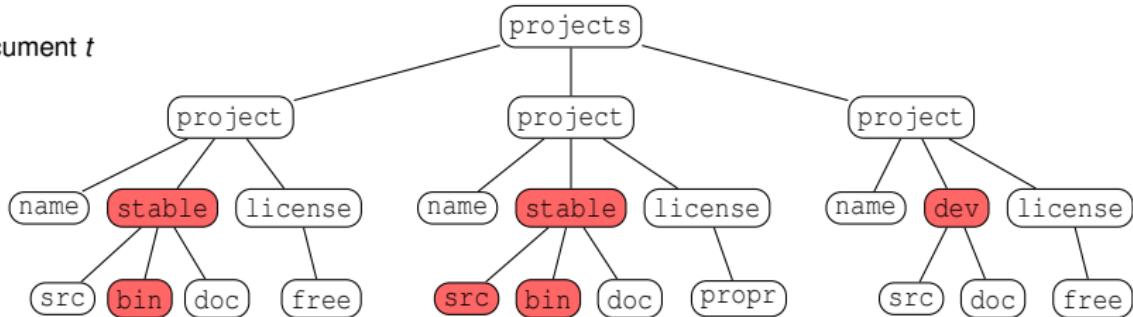
$\text{dev} \rightarrow \text{src}, \text{doc}$

$A_0(\text{dev}, \text{src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

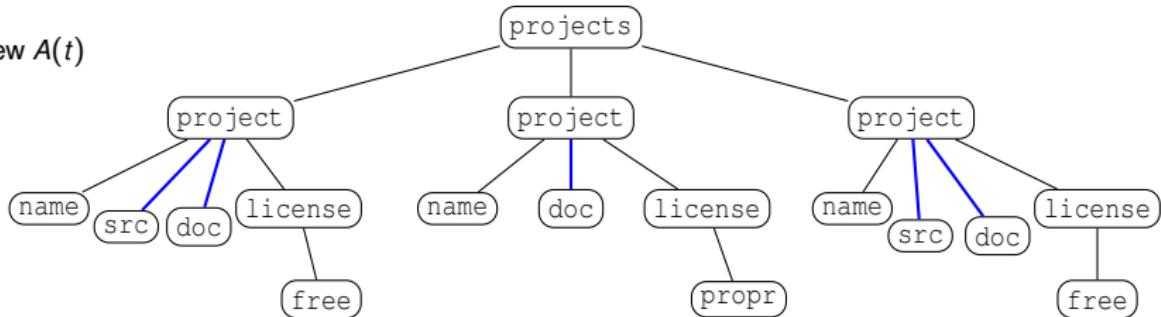
$A_0(\text{dev}, \text{doc}) = \text{true}$

The security view

document t



View $A(t)$



Annotating the DTDs

▷ annotation as a function $A : \Sigma \times \Sigma \rightarrow \{\text{true}, \text{false}, [f]\}$.

Example

projects → project*

project → name, (stable | dev), license

$A_0(\text{project, stable}) = \text{false}$

$A_0(\text{project, dev}) = \text{false}$

license → free | propr

stable → src, bin, doc

$A_0(\text{stable, src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

$A_0(\text{stable, doc}) = \text{true}$

dev → src, doc

$A_0(\text{dev, src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

$A_0(\text{dev, doc}) = \text{true}$

Annotating the DTDs

▷ annotation as a function $A : \Sigma \times \Sigma \rightarrow \{\text{true}, \text{false}, [f]\}$.

Example

projects → project*

project → name, (stable | dev), license

$A_0(\text{project, stable}) = \text{false}$

$A_0(\text{project, dev}) = \text{false}$

license → free | propr

stable → src, bin, doc

$A_0(\text{stable, src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

$A_0(\text{stable, doc}) = \text{true}$

dev → src, doc

$A_0(\text{dev, src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

$A_0(\text{dev, doc}) = \text{true}$

Proposition

This model of annotation is equivalent to defining accessible elements with a $XReg$ filter f_{acc}^A such that :

$\forall n \in N_t. n \text{ accessible wrt. } A \iff (t, n) \models f_{acc}^A$

Rewriting Queries

Theorem: *Regular XPath* is closed under query rewriting

There exists a function Rewrite such that :

$$\forall t. \text{Ans}(Q, A(t)) = \text{Ans}(\text{Rewrite}(Q, A), t)$$

Moreover, $\text{Rewrite}(Q, A)$ is computable in time $O(|A| * |Q|)$.

Rewriting Queries

Theorem: *Regular XPath* is closed under query rewriting

There exists a function Rewrite such that :

$$\forall t. \text{Ans}(Q, A(t)) = \text{Ans}(\text{Rewrite}(Q, A), t)$$

Moreover, $\text{Rewrite}(Q, A)$ is computable in time $O(|A| * |Q|)$.

Proof.

Translate the base axes using f_{acc}^A :

$$\text{Rewrite}(\uparrow, A) = \text{self}[f_{acc}^A]/(\uparrow[\neg f_{acc}^A])^*/\text{self}[f_{acc}^A].$$

Rewrite the query inductively.



Rewriting Queries

Hidden part

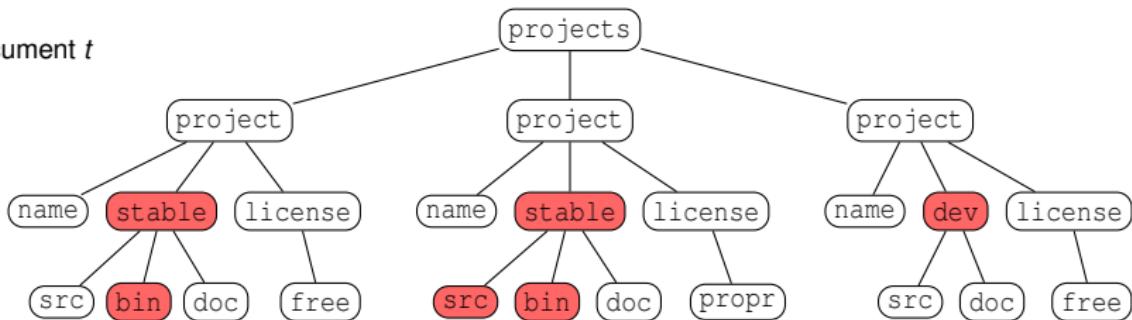
User part

document t

$Q' = \text{Rewrite}(Q, A)$

Query Q
Answer

document t



Rewriting Queries

Hidden part

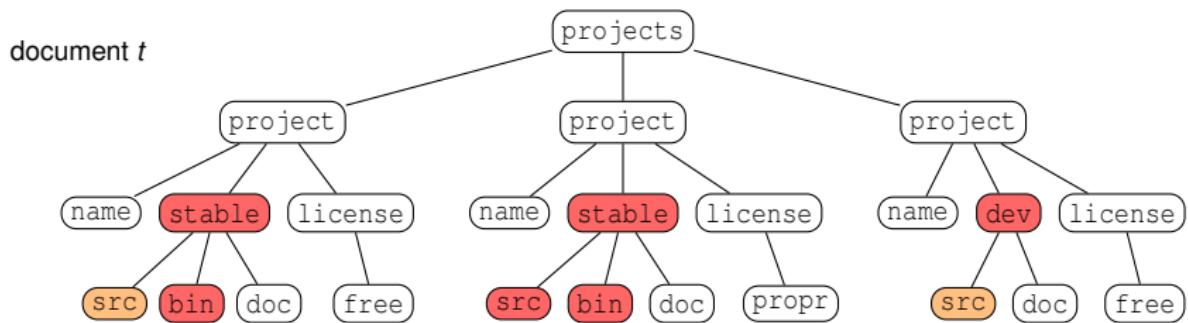
User part

$$Q' = \Downarrow :: project[license/free]/\Downarrow :: * / \Downarrow :: src$$
$$Q = \Downarrow :: project / \Downarrow :: src$$

document t

$$Q' = \text{Rewrite}(Q, A)$$

Query Q
Answer



Outline

1 Non-materialized views and query rewriting

2 Comparing Access Control Policies

Comparing access control policies

Definition

Two annotations A_1 and A_2 over DTD D are *equivalent* iff they hide the same nodes:

$$A_1 \equiv^D A_2 \text{ iff } \forall t \in L(D). A_1(t) = A_2(t)$$

Comparing access control policies

Definition

Two annotations A_1 and A_2 over DTD D are *equivalent* iff they hide the same nodes:

$$A_1 \equiv^D A_2 \text{ iff } \forall t \in L(D). A_1(t) = A_2(t)$$

Proposition

Testing equivalence of annotations is *EXPTIME*-complete.

Proof.

This problem is polynomially equivalent to the problem of equivalence of \mathcal{XReg} filters over a DTD. □

Comparing Access control policies

Definition

A_1 and A_2 annotations over DTD D . A_1 is *1-restriction* of A_2 in the presence of D , denoted

$$A_1 \preccurlyeq_1^D A_2 \text{ iff } \forall t \in L(D). N_{A_1(t)} \subseteq N_{A_2(t)}$$

Intuition:

The simplest way for comparing two annotations:

A_1 is more “restrictive” than A_2 if it shows no element hidden by A_2 .

Proposition

Testing 1-restriction is *EXPTIME*-complete.

Does \Leftarrow_1 ensure the properties we expect?

Example

projects → project*

project → name, (stable | dev), license

$A_0(\text{project}, \text{stable}) = \text{false}$

$A_0(\text{project}, \text{dev}) = \text{false}$

license → free | prop

stable → src, bin, doc

$A_0(\text{stable}, \text{src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

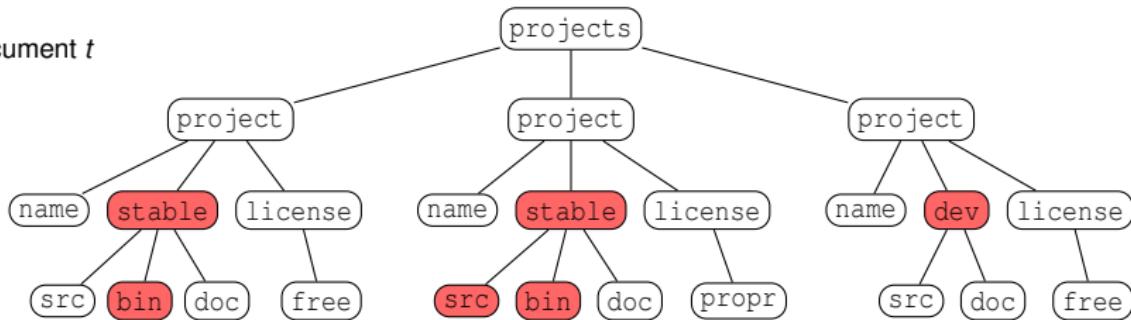
$A_0(\text{stable}, \text{doc}) = \text{true}$

dev → src, doc

$A_0(\text{dev}, \text{src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

$A_0(\text{dev}, \text{doc}) = \text{true}$

document t



Does \Leftarrow_1 ensure the properties we expect?

Example

projects → project*

project → name, (stable | dev), license

$A_0(\text{project, stable}) = \text{false}$

$A_0(\text{project, dev}) = \text{false}$

license → free | prop

stable → src, bin, doc

$A_0(\text{stable, src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

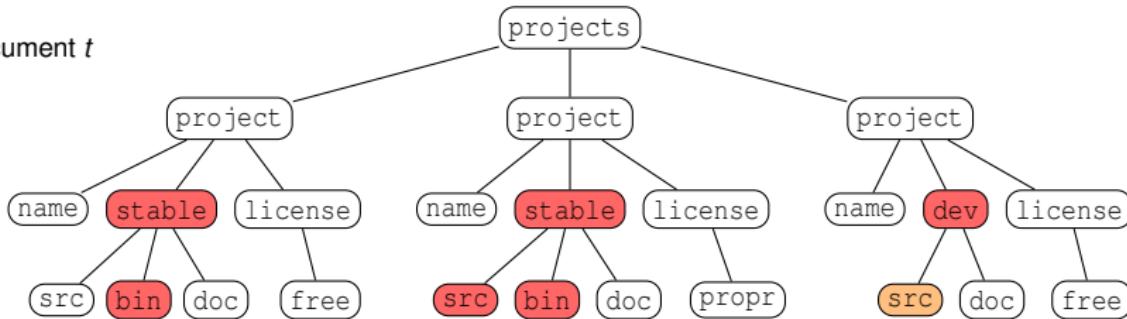
$A_0(\text{stable, doc}) = \text{true}$

dev → src, doc

$A_0(\text{dev, src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}] \text{false}$

$A_0(\text{dev, doc}) = \text{true}$

document t



Does \Leftarrow_1 ensure the properties we expect?

Example

projects → project*

project → name, (stable | dev), license

$A_0(\text{project}, \text{stable}) = \text{false}$

$A_0(\text{project}, \text{dev}) = \text{false}$

license → free | prop

stable → src, bin, doc

$A_0(\text{stable}, \text{src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}]$

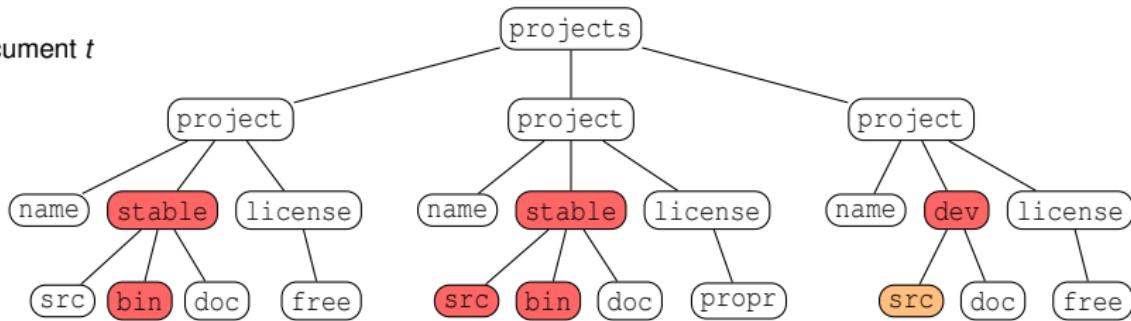
$A_0(\text{stable}, \text{doc}) = \text{true}$

dev → src, doc

$A_0(\text{dev}, \text{src}) = [\uparrow^*::\text{project} / \downarrow^*::\text{free}] \text{false}$

$A_0(\text{dev}, \text{doc}) = \text{true}$

document t



⇒ User can select all projects under free license that are not stable !

An information-oriented comparison

Argument

A_1 should be more “restrictive” than A_2 if every information inferred from A_1 can be inferred from A_2 .

Definition

A_1 and A_2 annotations over DTD D . A_1 is *2-restriction* of A_2 in the presence of D , denoted

$$A_1 \preceq_2^D A_2 \text{ iff } \forall Q_1 \exists Q_2. \quad \forall t \in L(D). \quad Ans(Q_1, A_1(t)) = Ans(Q_2, A_2(t))$$

An information-oriented comparison

Definition

A_1 and A_2 annotations over DTD D . A_1 is *2-restriction* of A_2 in the presence of D , denoted

$$A_1 \preceq_2^D A_2 \text{ iff } \forall Q_1 \exists Q_2. \quad \forall t \in L(D). \quad Ans(Q_1, A_1(t)) = Ans(Q_2, A_2(t))$$

Theorem

This property is undecidable.

An information-oriented comparison

Definition

A_1 and A_2 annotations over DTD D . A_1 is *2-restriction* of A_2 in the presence of D , denoted

$$A_1 \preccurlyeq_2^D A_2 \text{ iff } \forall Q_1 \exists Q_2. \forall t \in L(D). Ans(Q_1, A_1(t)) = Ans(Q_2, A_2(t))$$

Theorem

This property is undecidable.

Alternative characterization

$A_1 \preccurlyeq_2^D A_2$ if and only if

$$\exists f. \forall t \models D \forall n \in N_{A_2}(t). (n, A_2(t)) \models f \iff n \in N_{A_1}(t)$$

\implies : if filter f is provided, then one can verify the property in EXPTIME

An information-oriented comparison

Definition

A_1 and A_2 annotations over DTD D . A_1 is *2-restriction* of A_2 in the presence of D , denoted

$$A_1 \preceq_2^D A_2 \text{ iff } \forall Q_1 \exists Q_2. \quad \forall t \in L(D). \quad Ans(Q_1, A_1(t)) = Ans(Q_2, A_2(t))$$

Theorem

This property is undecidable.

Theorem

However, for non-recursive DTDs, 2-restriction can be tested in
EXPTIME

Further work

- implementation
- update propagation
- richer schema and query language
- other view formalisms