

XML Data Exchange

Relational Data Exchange Settings

Data Exchange Setting: (σ, τ, Σ)

σ : Source schema.

τ : Target schema.

Σ : Set of rules that specify relationship between the target and the source (source-to-target dependencies).

- Source-to-target dependency:

$$\psi_{\tau}(\bar{x}, \bar{z}) :- \varphi_{\sigma}(\bar{x}, \bar{y}).$$

- $\varphi_{\sigma}(\bar{x}, \bar{y})$: conjunction of atomic formulas over σ .
- $\psi_{\tau}(\bar{x}, \bar{z})$: conjunction of atomic formulas over τ .

Example: Relational Data Exchange Setting

- $\sigma = \text{Book}(\text{Title}, \text{AName}, \text{Aff})$
- $\tau = \text{Writer}(\text{Name}, \text{BTitle}, \text{Year})$
- $\Sigma = \text{Writer}(x_2, x_1, z_1) :- \text{Book}(x_1, x_2, y_1).$

Relational Data Exchange Problem

- Given a source instance S , find a target instance T such that (S, T) satisfies Σ .
 - (S, T) satisfies $\psi_\tau(\bar{x}, \bar{z}) :- \varphi_\sigma(\bar{x}, \bar{y})$ if whenever S satisfies $\varphi_\sigma(\bar{a}, \bar{b})$, there is a tuple \bar{c} such that T satisfies $\psi_\tau(\bar{a}, \bar{c})$.
 - T is called a **solution** for S .
- Previous example:

	<i>Book</i>	<i>Title</i>	<i>AName</i>	<i>Aff</i>
S :		Algebra	Hungerford	U. Washington
		Real Analysis	Royden	Stanford

Relational Data Exchange Problem

Possible solutions:

<i>Writer</i>	<i>Name</i>	<i>BTitle</i>	<i>Year</i>
$T_1:$	Hungerford	Algebra	1974
	Royden	Real Analysis	1988

<i>Writer</i>	<i>Name</i>	<i>BTitle</i>	<i>Year</i>
$T_2:$	Hungerford	Algebra	\perp_1
	Royden	Real Analysis	\perp_2

Query Answering

- Q is a query over target schema.

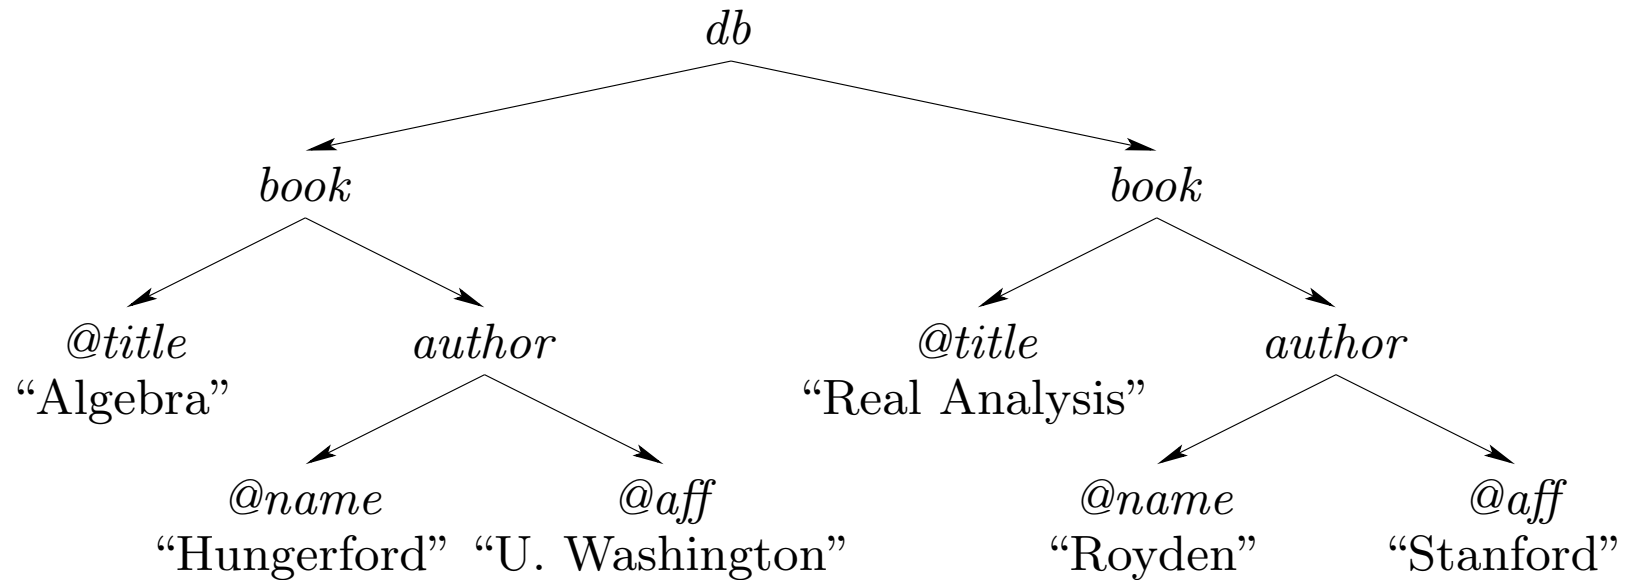
What does it mean to answer Q ?

$$\underline{\text{certain}}(Q, S) = \bigcap_{T \text{ is a solution for } S} \bigcap_{R \in \text{POSS}(T)} Q(R)$$


- Previous example:

- $\underline{\text{certain}}(\exists y \exists z \text{Writer}(x, y, z), I) = \{\text{Hungerford, Royden}\}$

XML Documents



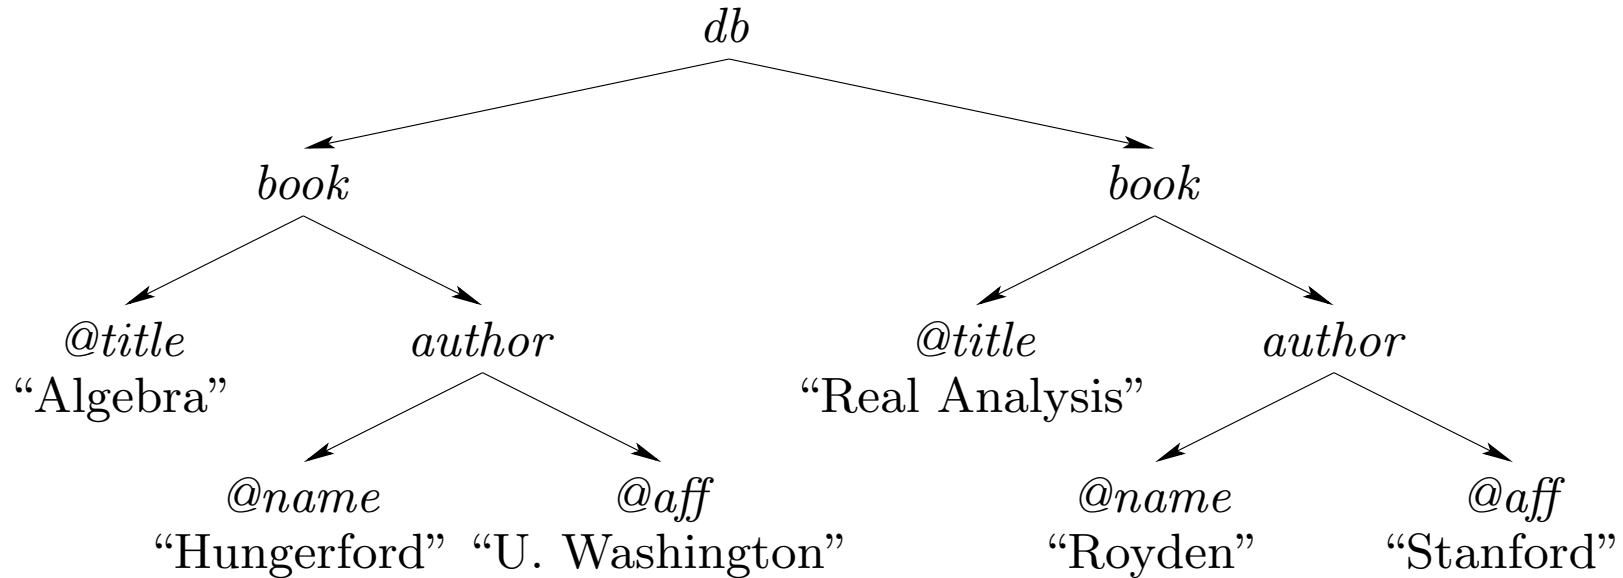
XML Documents



DTD :

<i>db</i>	→	<i>book</i> ⁺
<i>book</i>	→	<i>author</i> ⁺
<i>author</i>	→	ε

XML Documents



DTD :

<i>db</i>	→	<i>book</i> ⁺	<i>book</i>	→	<i>@title</i>
<i>book</i>	→	<i>author</i> ⁺	<i>author</i>	→	<i>@name, @aff</i>
<i>author</i>	→	ϵ			


XML Data Exchange Settings

- Instead of source and target relational schemas, we have **source** and **target DTDs**.
- But what are the source-to-target dependencies?

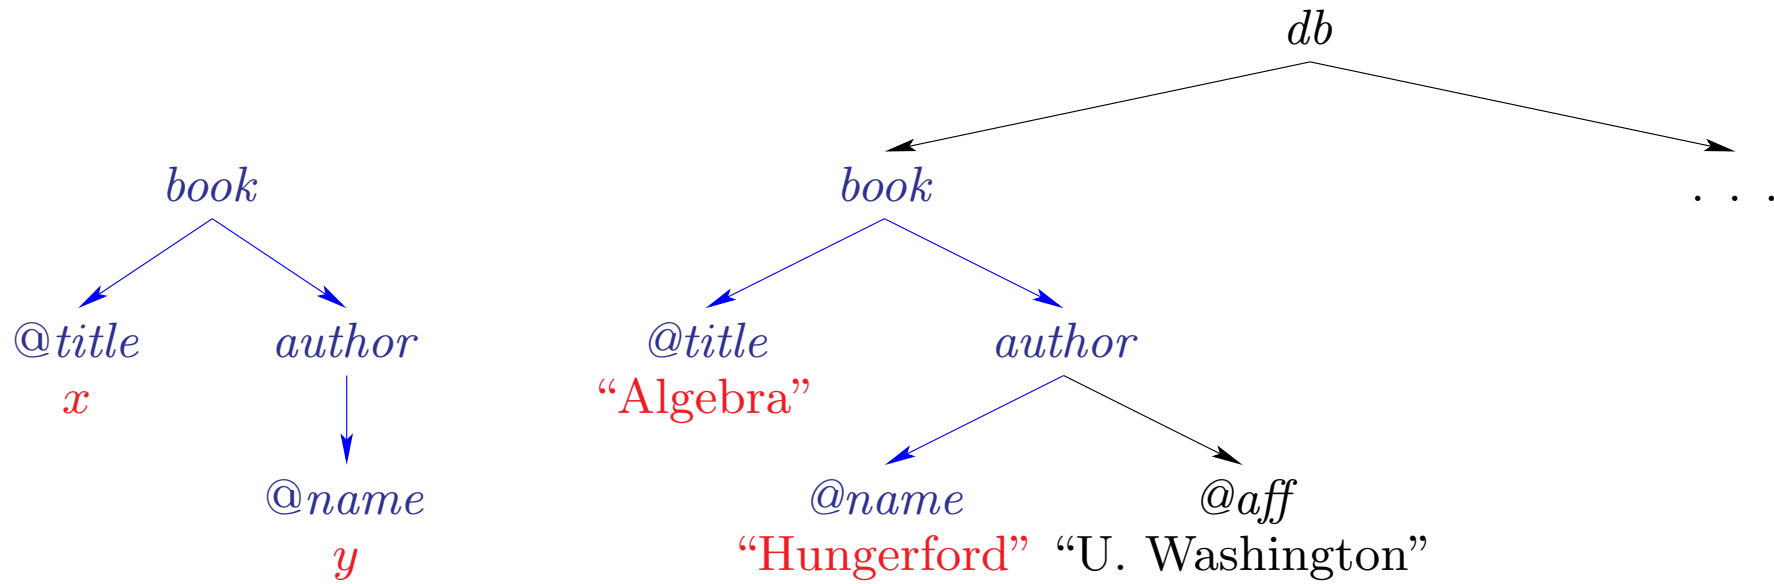
To define them, we use tree patterns.

If a certain pattern is found in the source, another pattern has to be found in the target.


Tree Patterns: Example




Tree Patterns: Example



Tree Patterns: Example



Tree Patterns: Example



Collect tuples (x, y) : (Algebra, Hungerford), (Real Analysis, Royden)

Tree Patterns

- Example: $book(@title = x)[author(@name = y)]$.
- Language also includes wildcard `_` (matching more than one symbol) and descendant operator `//`.


XML Source-to-target Dependencies

- Source-to-target dependency (STD):

$$\psi_{\tau}(\bar{x}, \bar{z}) \text{ :- } \varphi_{\sigma}(\bar{x}, \bar{y}),$$

where $\varphi_{\sigma}(\bar{x}, \bar{y})$ and $\psi_{\tau}(\bar{x}, \bar{z})$ are tree-patterns over the source and target DTDs, resp.

- Example:



XML Data Exchange Settings

XML Data Exchange Setting: $(D_\sigma, D_\tau, \Sigma)$

D_σ : Source DTD.

D_τ : Target DTD.

Σ : Set of XML source-to-target dependencies.

Each constraint in Σ is of the form $\psi_\tau(\bar{x}, \bar{z}) :- \varphi_\sigma(\bar{x}, \bar{y})$.

- $\varphi_\sigma(\bar{x}, \bar{y})$: tree-pattern over D_σ .
- $\psi_\tau(\bar{x}, \bar{z})$: tree-pattern over D_τ .

Example: XML Data Exchange Setting

- Source DTD:

db	\rightarrow	$book^+$		
$book$	\rightarrow	$author^+$	$book$	\rightarrow $@title$
$author$	\rightarrow	ε	$author$	\rightarrow $@name, @aff$

- Target DTD:

bib	\rightarrow	$writer^+$		
$writer$	\rightarrow	$work^+$	$writer$	\rightarrow $@name$
$work$	\rightarrow	ε	$work$	\rightarrow $@title, @year$

- Σ :

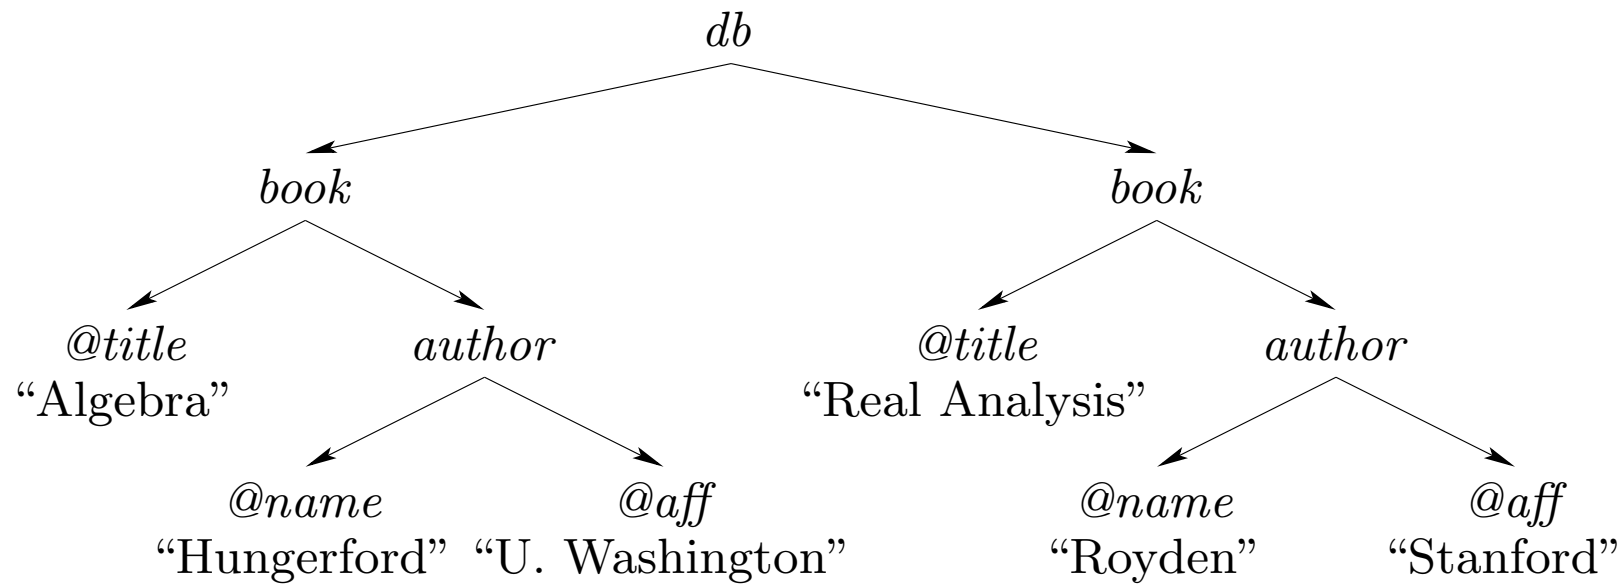
$$writer(@name = y)[work(@title = x, @year = z)] :-$$
$$book(@title = x)[author(@name = y)].$$

XML Data Exchange Problem

- Given a source tree T , find a target tree T' such that (T, T') satisfies Σ .
 - (T, T') satisfies $\psi_\tau(\bar{x}, \bar{z}) :- \varphi_\sigma(\bar{x}, \bar{y})$ if whenever T satisfies $\varphi_\sigma(\bar{a}, \bar{b})$, there is a tuple \bar{c} such that T' satisfies $\psi_\tau(\bar{a}, \bar{c})$.
 - T' is called a **solution** for T .


XML Data Exchange Problem

Let T be our original tree:



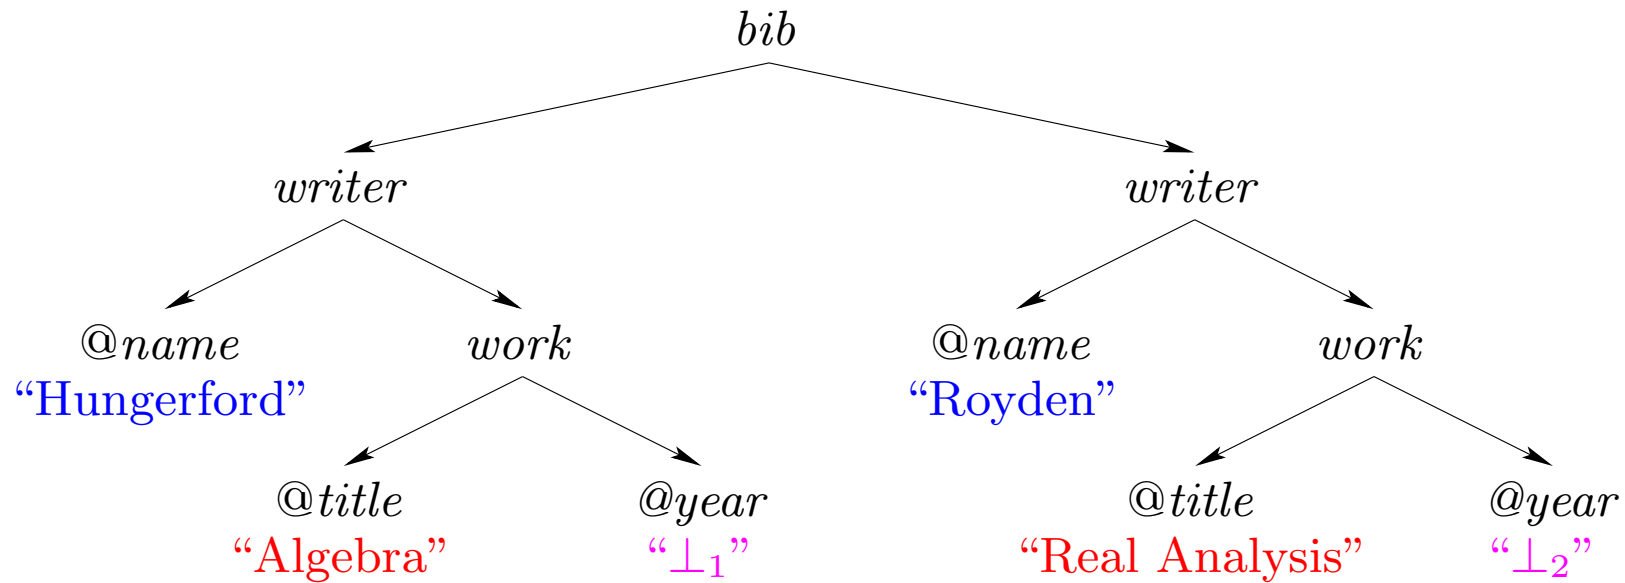
XML Data Exchange Problem

A solution for T :



XML Data Exchange Problem

Another solution for T :



Consistency of XML Data Exchange Settings

- What if we have target DTD

bib	\rightarrow	$writer^+$		
$writer$	\rightarrow	$novelist^*, poet^*$	$writer$	\rightarrow $@name$
$novelist$	\rightarrow	$work^+$		
$poet$	\rightarrow	$work^+$		
$work$	\rightarrow	ε	$work$	\rightarrow $@title, @year$

in our previous example?

- The setting becomes **inconsistent!**
 - There are no T conforming to D_σ and T' conforming to D_τ such that (T, T') satisfies Σ .

Consistency of XML Data Exchange Settings

- An XML data exchange setting is **inconsistent** if it does not admit solutions for any given source tree. Otherwise it is **consistent**.
- A relational data exchange setting is always consistent.
- An XML data exchange setting is not always consistent.
 - What is the complexity of checking whether a setting is consistent?

Bad News: General Case

Theorem Checking if an XML data exchange setting is consistent necessarily takes **exponential time**.

Complexity-theoretic statement: EXPTIME-complete.

But the parameter is the size of the DTDs and constraints – typically not very large. Hence $2^{O(n)}$ is not too bad.

Good News: Consistency for Commonly used DTDs

DTDs that commonly occur in practice tend to be simple. In fact more than 50% of regular expressions are of this form:

$$l \rightarrow \hat{l}_1, \dots, \hat{l}_m,$$

where all the l_i 's are distinct, and \hat{l} is one of the following: l , or l^* , or l^+ , or $l?$

For example, $\text{book} \rightarrow \text{title}, \text{author}^+, \text{chapter}^*, \text{publisher}?$

Theorem For non-recursive DTDs that only have these rules, checking if an XML data exchange setting is consistent is solvable in time $O((\|D_\sigma\| + \|D_\tau\|) \cdot \|\Sigma\|^2)$.

Query Answering in XML Data Exchange

- Decision to make: what is our query language?
- XML query languages such as XQuery take XML trees and produce XML trees.
 - This makes it hard to talk about **certain answers**.
- For now we use a query language that produces **tuples of values**.

Conjunctive Tree Queries

- Query language *CTQ//* is defined by

$$Q \quad := \quad \varphi \quad | \quad Q \wedge Q \quad | \quad \exists x Q,$$

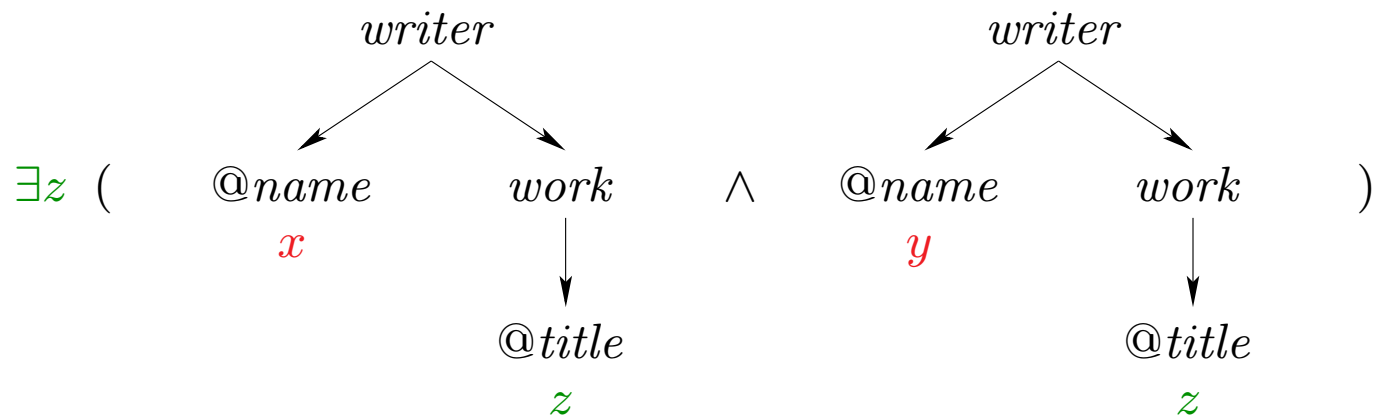
where φ ranges over tree-patterns.

- By disallowing descendant // we obtain restriction *CTQ*.

Example: Conjunctive Tree Query

List all pairs of authors that have written articles with the same title.

$Q(x, y) :=$



Computing Certain Answers

- Semantics: as in the relational case.

$$\underline{\text{certain}}(Q, T) = \bigcap_{T' \text{ is a solution for } T} Q(T').$$

- Given data exchange setting $(D_\sigma, D_\tau, \Sigma)$ and query Q :

PROBLEM: CERTANSW(Q).

INPUT: Tree T conforming to D_σ and tuple \bar{a} .

QUESTION: Is $\bar{a} \in \underline{\text{certain}}(Q, T)$?

Computing Certain Answers: General Picture

It is not even clear if the problem is solvable.

Good news For every XML data exchange setting and $CTQ//$ -query Q , the problem $CERTANSW(Q)$ is solvable in exponential time.

More precisely, it is in $coNP$.

Not so good news Sometimes exponential time is “unavoidable”: There exist an XML data exchange setting and a $CTQ//$ -query Q such that $CERTANSW(Q)$ is $coNP$ -complete.

We want to find cases that admit fast algorithms.

Computing Certain Answers: Eliminating bad cases

Suppose one of the following is allowed in tree patterns over the target in STDs:

- descendant operator $//$, or
- wildcard $_$, or
- patterns that **do not start at the root**.

Then one can find source and target DTDs (in fact, very simple DTDs) and a CTQ -query Q such that $CERTANSW(Q)$ must take exponential time.

A more precise statement: is **coNP-complete**.

Fully specified constraints

- We disallow the three features that make query answering hard.
- This gives us **fully-specified STDs**:

We impose restrictions on tree patterns over target DTDs:

- no descendant relation `//`; and
- no wildcard `_`; and
- all patterns **start at the root**.

No restrictions imposed on tree patterns over source DTDs.

- Subsume non-relational data exchange handled by IBM.

An efficient case

- Recall relational data exchange and conjunctive queries:
then $\underline{\text{certain}}(Q, S) = \text{certain}(Q, \text{CANSOL}(S))$.
- Idea: given a source tree T , compute a solution T^* for T
such that

$$\underline{\text{certain}}(Q, T) = \text{remove_null_tuples}(Q(T^*)).$$

- T^* is a **canonical** solution for T .
- We compute T^* in two steps:
 - We use STDs to compute a canonical pre-solution $\text{cps}(T)$ from T .
 - Then we use target DTD to compute T^* from $\text{cps}(T)$.

Example: XML Data Exchange Setting

- Source DTD:

$$\begin{array}{ll} r & \rightarrow A^*, B^* \\ A & \rightarrow \varepsilon \qquad A \rightarrow @l \\ B & \rightarrow \varepsilon \qquad B \rightarrow @l \end{array}$$


- Target DTD:

$$\begin{array}{ll} r & \rightarrow (C, D)^* \\ C & \rightarrow \varepsilon \qquad C \rightarrow @m \\ D & \rightarrow E \\ E & \rightarrow \varepsilon \qquad E \rightarrow @n \end{array}$$


- Σ :

$$\begin{array}{ll} r[C(@m = x)] & :- A(@l = x), \\ r[C(@m = x)] & :- B(@l = x). \end{array}$$


Example: Computing Canonical Pre-solution





Example: Computing Canonical Pre-solution




Example: Computing Canonical Pre-solution




Example: Computing Canonical Pre-solution





Example: Computing Canonical Pre-solution





Example: Computing Canonical Pre-solution





Example: Computing Canonical Pre-solution



Example: Computing Canonical Pre-solution



Example: Computing Canonical Pre-solution




Example: Computing Canonical Pre-solution

r
↓
 C
↓
 $@m$
“1”

r
↓
 C
↓
 $@m$
“2”


Example: Computing Canonical Pre-solution

Canonical pre-solution:




Not yet a solution: it does not conform to the target DTD.

Example: Computing Canonical Solution




Example: Computing Canonical Solution




$$r \rightarrow (C, D)^*$$

Example: Computing Canonical Solution




$$r \rightarrow (C, D)^*$$

Example: Computing Canonical Solution




$$D \rightarrow E$$

Example: Computing Canonical Solution




$$D \rightarrow E$$

Example: Computing Canonical Solution




$$E \rightarrow @n$$

Example: Computing Canonical Solution




$$E \rightarrow @n$$

Example: Computing Canonical Solution




$$D \rightarrow E$$

Example: Computing Canonical Solution




$$D \rightarrow E$$

Example: Computing Canonical Solution




$$E \rightarrow @n$$

Example: Computing Canonical Solution



$$E \rightarrow @n$$

Example: Computing Canonical Solution



Does this always work?

Depends on regular expressions in target DTDs.

- class of **good** regular expressions.
 - Examples: $(A|B)^*$, $A, B^+, C^*, D?$, $(A^*|B^*)$, $(C, D)^*$.
 - bad: $A, (B|C)$.
 - exact definition: quite involved.

Does this always work? cont'd

- For target DTDs only using good regular expressions:
 - There exists a solution for a tree T iff there exists a canonical solution T^* for T .
 - Previous algorithm computes canonical solution T^* for T in polynomial time.
 - $\text{certain}(Q, T) = \text{remove_null_tuples}(Q(T^*))$, for every $CTQ//$ -query.
- Complexity: **polynomial time**.