Declarative Inconsistency Handling in Relational and Semi-structured Databases

Ph.D. Dissertation Defense

Sławek Staworko

Department of Computer Science and Engineering University at Buffalo, SUNY

9 May 2007

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

Motivation

Schema

Emp(<u>Name</u>, Dept)

Q: who works in IT

 $Q(X) := \operatorname{Emp}(X, IT)$. Answers are: {*Mary*, *John*} Are we sure about John ?



▲ロト ▲帰ト ▲ヨト ▲ヨト 三日 - の々ぐ

Motivation

Schema

Emp(<u>Name</u>, Dept)

Q: who works in IT

 $Q(X) := \operatorname{Emp}(X, IT)$. Answers are: {*Mary*, *John*} Are we sure about John ?



◆□▶ ◆圖▶ ◆臣▶ ◆臣▶ ─ 臣

What about data cleaning?

- all conflicting tuples are removed
- "clean" instance $r' = \{Emp(Mary, IT)\}$
- may cause loss of information
- sources may be autonomous
- inconsistencies may be results of updates in progress, long running transactions, manual entry, exceptions from the constraints

Motivation

Schema

Emp(<u>Name</u>, Dept)

Q: who works in IT

 $Q(X) := \operatorname{Emp}(X, IT)$. Answers are: {*Mary*, *John*} Are we sure about John ?



Repairs [Arenas et al. PODS'99]:

 $r_1 = \{ Emp(John, IT), Emp(Mary, IT) \}$ $r_2 = \{ Emp(John, PR), Emp(Mary, IT) \}$

Consistent answers are:

 $\{Mary\}$ John is not a consistent answer because of r_2

- 1. Analysis of computational complexity of consistent query answers (CQA) in the presence of universal constraints.
- 2. Building a practical system for computing CQA.
- 3. Developing a general framework for preference-based conflict resolution.
- 4. Adapting the framework of CQA to semi-structured databases.

1. universal constraints

 $\forall^* \neg [R_1(\bar{x}_1) \land \ldots \land R_n(\bar{x}_n) \land \neg P_1(\bar{y}_1) \land \ldots \land \neg P_m(\bar{y}_m) \land \rho]$

where $\bar{y} \subseteq \bar{x}$ (for safety [Abiteboul et al.])

2. Horn constraints

$$R_1(\bar{x}_1) \wedge \ldots \wedge R_n(\bar{x}_n) \wedge \rho \Rightarrow P(\bar{y})$$

3. denial constraints

$$\forall^* \neg [R_1(\bar{x}_1) \land \ldots \land R_n(\bar{x}_n) \land \rho]$$

4. functional dependencies (FD)

$$R: X \to Y$$

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > <

1. relational algebra (a subclass of SQL)

 $E ::= R \mid \sigma_{\varphi}(E) \mid \pi_{X}(E) \mid E_{1} \times E_{2} \mid E_{1} \cup E_{2} \mid E_{1} \setminus E_{2}.$

2. first-order queries

 $\varphi ::= R(\bar{x}) \mid \varphi_1 \lor \varphi_2 \mid \varphi_1 \land \varphi_2 \mid \neg \varphi \mid \exists x. \varphi(x) \mid \forall x. \varphi(x).$

3. conjunctive queries

 $Q(\bar{x}) = \exists \bar{y}. R_1(\bar{z}_1) \land \ldots \land R_k(\bar{z}_k).$

Repairs and Consistent Query Answers [PODS'99]

Repair:

a consistent instance minimally different from the database



There can be an exponential number of repairs

Consistent Query Answers:

answers present in every repair.

Computational complexity for universal constraints

Data complexity

- query and IC assumed to be fixed
- input size depends on the DB size only
- standard measure in relational databases [Vardi, STOC'82]

Constraints	Repair	Consistent Answers to	
	Checking	$\{\forall, \exists\}$ -free queries	conjunctive queries
Universal	coNP-c	П <u></u> 2-с	П <u></u> 2-с
Horn	PTIME	coNP-c	coNP-c
Acyclic Horn	PTIME	PTIME	coNP-c
Denial	PTIME	PTIME	coNP-c

[Chomicki and Marcinkowski, IC'05]

Conflict hypergraph [Chomicki and Marcinkowski, IC'05]

Denial constraint (*Emp* : *Name* \rightarrow *Dept*)

 $\forall n, d_1, d_2. \neg [Emp(n, d_1) \land Emp(n, d_2) \land d_1 \neq d_2]$

Conflict: {*Emp*(*John*, *IT*), *Emp*(*John*, *PR*)}



Conflict graph

- vertices = tuples from DB
- edges connect conflicting tuples
- ▶ repairs ≅ maximal independent sets

used to effectively compute CQA

Extended conflict hypergraph

Universal constraint ($Dept(DeptName, MgrName) \Rightarrow Emp[MgrName, DeptName]$)

 $\forall n_1, n_2, d_1, d_2. \neg [Dept(d, m) \land \neg Emp(m, d)]$

Conflict: { $Dept(HR, Bob), \neg Emp(Bob, HR)$ }



Extended conflict graph:

- vertices = tuples present and missing in DB
- edges connect conflicting tuples
- repairs \cong independent sets minimally different from DB

Computing consistent answers to π -free queries

System Hippo:

- π -free queries
- denial constraints
- implemented in Java
- JDBC front-end

Efficiency

- ▶ for {σ, ×}-queries no overhead
- For {σ, ×, ∪, \}-queries twice the time of standard evaluation

Handling projection

co-NP-complete problem



- ロ ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 回 ト - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4 □ - 4

Related work: Computing consistent query answers

- 1. Query rewriting [Arenas et al., PODS'99]
 - ► A query *Q* is rewritten to a query *Q*' defining consistent answers to *Q*.
 - easy to incorporate into already existing applications
 - relatively small overhead
 - applicable to a limited class of queries and constraints
 - ConQuer: Select-Project-Join SQL queries in the presence of key dependencies [Fuxman et al., SIGMOD'05].
- 2. Logic programs
 - logic programs specify all repairs
 - evaluation with existing systems like dlv or smodels
 - very general: arbitrary first order queries and universal constraints (even some referential constraints)
 - ► computationally expensive: Π^p₂-completeness of logic programs

 INFOMIX: incorporates various optimizations [Eiter et al., SIGMOD'05]

- 1. Adding projection (π)
- 2. Analysis of inconsistencies encountered in practice
- 3. Benchmarking CQA systems
- 4. Universal constraints without the safety restriction

 $\forall x. \neg [\neg P(x) \land x < 100] \equiv \forall x < 100. P(x)$

Priorities, Preferences, and Cleaning

Priority \succ

- an acyclic orientation of the conflict graph
- ► ≻ is called total when all edges are oriented

 $Emp(John, PR) \succ Emp(John, IT)$

▲ロト ▲帰ト ▲ヨト ▲ヨト 三日 - の々ぐ

Priorities, Preferences, and Cleaning

Priority \succ

- an acyclic orientation of the conflict graph
- ► ≻ is called total when all edges are oriented

Preferred CQA

- ► use ≻ to define a family of preferred repairs Rep(≻)
- ► preferred consistent answers w.r.t. > are the answers present in every preferred repair w.r.t >



 $Emp(John, PR) \succ Emp(John, IT)$

Priorities, Preferences, and Cleaning

Priority \succ

- an acyclic orientation of the conflict graph
- ► ≻ is called total when all edges are oriented

 $Emp(John, PR) \succ Emp(John, IT)$

$$\omega_{\succ}(r) = \{t \in r | \neg \exists t' \in r.t' \succ t\}$$

Preferred CQA

- ► use ≻ to define a family of preferred repairs Rep(≻)
- ► preferred consistent answers w.r.t. > are the answers present in every preferred repair w.r.t >

Database cleaning with a total \succ

► *r*′ := Ø

- while $\omega_{\succ}(r) \neq \emptyset$ do
 - 1. choose any $x \in \omega_{\succ}(r)$
 - 2. add x to r'
 - 3. remove x from r with neighbors

▶ return r'

Desirable Properties of Preferred Repairs

$(\mathcal{P}1)$ Non-emptiness

 $Rep(\succ) \neq \varnothing$

 $(\mathcal{P}2)$ Monotonicity

$$\succ_1 \subseteq \succ_2 \Rightarrow \operatorname{Rep}(\succ_2) \subseteq \operatorname{Rep}(\succ_1)$$

 $(\mathcal{P}3)$ Non-discrimination

 $Rep(\varnothing) = Rep$

 $(\mathcal{P}4)$ Categoricity

 \succ is total $\Rightarrow |Rep(\succ)| = 1$

Optimal Use of Priorities



Common optimal repairs

 $r' \in C\text{-}Rep(\succ)$ iff r' is a result of cleaning the database with \succ .

Properties

- *C-Rep* satisfies $\mathcal{P}1 \mathcal{P}4$
- $\blacktriangleright \ C\text{-}Rep \subseteq \mathcal{G}\text{-}Rep$
- C-Rep = G-Rep for priorities that cannot be extended to a cyclic orientation

CQA: co-NP-complete

- 1. Conditioned active integrity constraints [Flesca, et all., PPDP'04] allow to specify how to resolve some conflicts
 - repairs may be not optimal
 - $\mathcal{P}1$ and $\mathcal{P}2$ are satisfied
 - $\mathcal{P}3$ and $\mathcal{P}4$ are violated
- 2. Repair constraints can be used to restrict considered repairs [Greco and Lembo, ER'04]
 - $\mathcal{P}3$ is satisfied, but $\mathcal{P}1$ is violated
 - \blacktriangleright in a weaker version: $\mathcal{P}1$ is satisfied, but $\mathcal{P}3$ is violated

- 1. Experimental evaluation
 - $\blacktriangleright \ one \ FD \rightarrow PTIME$
- 2. Cyclic priorities (what about monotonicity)
- 3. Generalization to denial and universal constraints

XML and DTD

XML Documents

- ordered finite trees
- Iabels from finite Σ
- PCDATA for text nodes
- text labels from Γ



C(A(d),B(e),B)

DTD is a function that:

- takes the label of a node
- returns a regular expression to which the labels of the children must comply

Example of DTD $C \rightarrow (A,B)^*$ $A \rightarrow EMPTY | PCDATA$ $B \rightarrow EMPTY$

C(A(d),B(e),B) doesn't satisfy D, but C(A(d),B()) does.











◆□▶ ◆□▶ ◆ □▶ ◆ □▶ ○ □ ○ ○ ○ ○





Edit distance, repairs, and valid query answers

Distance between documents

dist(T, S) is the minimal cost of transforming T into S



Edit distance, repairs, and valid query answers

Distance between documents

dist(T, S) is the minimal cost of transforming T into S

Distance to a DTD

dist(T, D) is the minimal cost of repairing T w.r.t D i.e.,

 $min\{dist(T, S)|S valid w.r.t D\}$

DTD

 $\begin{array}{l} \mathsf{C} \rightarrow (\mathsf{A},\mathsf{B})^{*} \\ \mathsf{A} \rightarrow \mathsf{EMPTY} \mid \mathsf{PCDATA} \\ \mathsf{B} \rightarrow \mathsf{EMPTY} \end{array}$





Edit distance, repairs, and valid query answers

Distance between documents

dist(T, S) is the minimal cost of transforming T into S

Distance to a DTD

dist(T, D) is the minimal cost of repairing T w.r.t D i.e.,

 $min\{dist(T, S)|S valid w.r.t D\}$

DTD

 $\begin{array}{l} \mathsf{C} \rightarrow (\mathsf{A},\mathsf{B})^{*} \\ \mathsf{A} \rightarrow \mathsf{EMPTY} \mid \mathsf{PCDATA} \\ \mathsf{B} \rightarrow \mathsf{EMPTY} \end{array}$

Repair

T' is a repair of T w.r.t D iff dist(T', T) = dist(T, D)





Valid Query Answers

x is a valid answer to query Q in T w.r.t. D iff x is an answer to Q in every repair of T w.r.t. D.











▲□▶ ▲圖▶ ★ 国▶ ★ 国▶ - 国 - のへで







▲□▶ ▲圖▶ ★ 国▶ ★ 国▶ - 国 - のへで





▲□▶ ▲圖▶ ★ 国▶ ★ 国▶ - 国 - のへで





◆□▶ ◆□▶ ◆三▶ ◆三▶ 三三 のへぐ





◆□ > ◆□ > ◆豆 > ◆豆 > ̄豆 = のへ⊙





Compact representation of all repairs



Repairing Paths:

- ▶ (Read, Read, Del)
- ► (Read, Read, Ins A, Read)
- ▶ (Read, Del, Read)

Complexity of VQA

XPath queries

sequences of steps (with filter conditions) to navigate through the tree.

Combined complexity

- both the document and the query are parts of the input
- standard evaluation in PTIME [Gottlob et al., PODS'03]

Positive results

- VQA in PTIME for simple XPath queries (only descending axes, no negation, no union/disjunction, and no join conditions)
- Rhino: uses trace graphs to effectively compute VQA

Negative results

VQA is coNP-complete for other XPath queries.

Related work: Consistent query answers for XML

- 1. [Flesca et. all, Xsym'03] queries valid documents that violate FDs
 - two operations: marking nodes as unreliable and setting null value
 - polynomial time algorithm for queries of form $A_1/\ldots/A_n$
 - no experimental evaluation
- 2. [Flesca et. all, WISE'05] deals with invalid documents and FDs
 - edit operations similar to ours, but
 - ➤ a different (set-theoretic) definition of repairs: C(A(d), B(e), B) has only one repair C(A(d), B, A, B).

- polynomial algorithm for a restricted class of queries (no union) and DTDs
- no experimental evaluation

1. Other editing operations

- ► General Deletion/Insertion: dist(T, D) computed in O(|T|⁵) time [Suzuki, SAC'05].
- Moving nodes: intractable for string edit distance [Comrode et al., SODA'02] and edit distance on unordered trees [de Rougement, ISIP'03]

- 2. Incorporation of semantic constraints (e.g., key dependencies)
- 3. Data complexity

 CQA: foundation of enlivened research in the ares of inconsistent databases (more than 100 papers over the last 10 years)

- practical systems have been constructed
- tractability has been studied
- framework has been adapted to XML databases