Slide 0

Data Integration

View-based Query Processing

Maurizio Lenzerini, *Giuseppe De Giacomo*, *Diego Calvanese*

Dipartimento di Informatica e Sistemistica
Università di Roma "La Sapienza"

Dottorato in Ingegneria Informatica
Dicembre, 2000 — Roma

Slide 1

View-based query processing

Basic Idea: Computing the answer to a query based on a set of views, rather than on the raw data in the database.

Significance: Relevant problem in query optimization, query answering with incomplete information, data warehousing, data integration, etc.

Two approaches:

- View-based query rewriting (indirect)
- View-based query answering (direct)

Slide 2

Part 1

View-based query processing: Computing the answer to a query based on a set of views, rather than on the raw data in the database.

- View-based query rewriting (indirect)
- View-based query answering (direct)
- Relationship between query answering and query rewriting

We study view-based query processing within the **relational data model**, focusing on **conjunctive queries** (CQs).

Slide 3

LAV: formal framework

The integrated database (DB) *DB* is simply a set of structures (relations, in the relational model), one for each symbol in an alphabet $\mathcal{A}^{\mathcal{G}}$

- the structure of the global view is specified in the schema language $\mathcal{L}_{\mathcal{G}}$ over $\mathcal{A}^{\mathcal{G}}$
- each source structure is modeled as a view over the global view, expressed in the view language $\mathcal{L}_{\mathcal{V}}$ over $\mathcal{A}^{\mathcal{G}}$
- queries are issued over the global view, and are expressed in the query language $\mathcal{L}_{\mathcal{Q}}$ over $\mathcal{A}^{\mathcal{G}}$

LAV: formal framework

Slide 4

- The global view $\mathcal{G}$ is specified as a set of constraints in $\mathcal{L}_{\mathcal{G}}$, and associated to each source structure we have a set (its extension)
- We have a view V_i for each source structure, with
  - extension $ext(V_i)$,
  - definition $def(V_i)$, i.e., a query $V_i(\vec{\mathbf{x}})$:- $v_i(\vec{\mathbf{x}}, \vec{\mathbf{y}})$, where $v_i(\vec{\mathbf{x}}, \vec{\mathbf{y}})$ is expressed in the language $\mathcal{L}_{\mathcal{V}}$ over $\mathcal{A}^{\mathcal{G}}$
  - assumption $as(V_i)$, i.e., how to interpret $ext(V_i)$ wrt the tuples satisfying V_i
- A query Q is expressed in the language $\mathcal{L}_Q$ over $\mathcal{A}^{\mathcal{G}}$. If DB satisfies $\mathcal{G}$, $ans(Q, DB)$ is the set of objects in DB that satisfy Q

LAV: formal framework

Slide 5

The specification $as(V_i)$ determines how accurate is the extension of the the view with respect to the specification $def(V_i)$

- **Sound Views**: a database DB is *coherent with the sound view* V_i, if $ext(V_i) \subseteq ans(def(V_i), DB)$
- **Complete Views**: a database DB is *coherent with the complete view* V_i, if $ext(V_i) \supseteq ans(def(V_i), DB)$
- **Exact Views**: a database DB is *coherent with the exact view* V_i, if $ext(V_i) = ans(def(V_i), DB)$

LAV: formal framework

Slide 6

Suppose we have the extensions of the source structures. Let $\mathcal{G}$ be the specification (or, schema) of the global view, Q a query of arity n, and $\vec{\mathbf{d}} = (d_1, \ldots, d_n)$ a tuple of constants

Query answering is defined as follows:

$\vec{\mathbf{d}} \in cert(Q, \mathcal{V})$ iff $(d_1, \ldots, d_n) \in ans(Q, DB)$, for each DB such that:

- DB satisfies $\mathcal{G}$
- DB is coherent with $V_1, \ldots, V_m$

View-based query answering

Slide 7

Basic Idea: Given a query Q, a set of views $\mathcal{V}$ with definitions $def(\mathcal{V})$ and extensions $ext(\mathcal{V})$, compute the tuples t which are in the answer to Q in all databases consistent with the views (certain tuples).

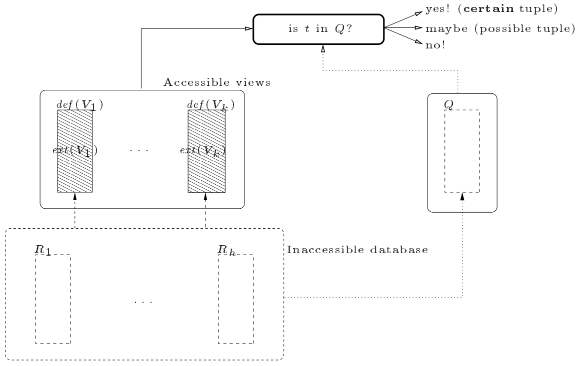

View-based query rewriting

Slide 8

Basic Idea: Given a query Q and a set of views $\mathcal{V}$ with definitions $def(\mathcal{V})$, reformulate Q into a new query Q' expressed in some language on the alphabet of the view names $\mathcal{V}$.

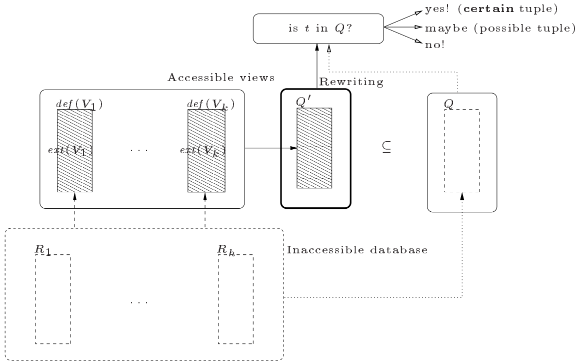

Query rewriting

Slide 9

There are many results on view-based query rewriting.

- Conjunctive queries [Levy et al 95], [Rajaraman et al 95], [Levy et al 97];
- Limitations on binding patterns [Rajaraman et al 95], [Duschka 97]
- Queries with aggregates [Srivastava et al 96], [Cohen et al 99], [Grumbach et al 99];
- Under constraints (Functional dependencies, Inclusions dependencies, etc.) [Duschka 97], [Gryz 98]
- Description Logics queries [Beeri et al 97];
- Recursive queries [Duschka et al 97], [Afrati et al 99];
- Queries for semi-structured data [Fernandez et al 98], [Milo et al 99], [Papakonstantinou et al 99], [Calvanese et al, PODS'99].

Query rewriting: the setting studied here

Slide 10

- **View definitions are CQs** over the global schema
- **View are all sound**
- **Queries are CQs**, or also **UCQs**, i.e., union of conjunctive queries, over the global schema
- **Rewritings are UCQs**, i.e., finite sets of CQs, in the alphabet of the views.

Query rewriting

Slide 11

Query answering by rewriting:

- Rewrite the query in the alphabet of the view names;
- Evaluate the rewriting on the view extension.

Typically people are interested in **rewritings** that are:

- **Contained** in the original query;
- Expressed in a given **query language**;
- **Maximal** for the given class of queries;
- **Exact**(??), i.e., rewritings that are logically equivalent to the query, if possible (observe that such rewritings may not exists).

Exact rewriting: example

Slide 12

Global schema:

```
nonstop(Airline, Number, From, To).
```

Views:

```
flights_by_United(From, To) -: nonstop('UA', Number, From, To)
flights_from_SFO(Airline, Number, To) -: nonstop(Airline, Number, 'SFO', To)
```

Query:

```
q(Airline, Number) :- nonstop(Airline, Number, 'SFO', 'LAX')
```

An **exact rewriting** exists! Namely:

```
q(Airline, Number) :- flight_from_SFO(Airline, Number, 'LAX')
```

Maximal rewriting: example

Slide 13

Global schema:

```
nonstop(Airline, Number, From, To).
```

Views:

```
flights_by_United(From, To) -: nonstop('UA', Number, From, To)
flights_from_SFO(From, To) -: nonstop(Flight, Number, 'SFO', To)
```

Query:

```
q(Airline, Number) :- nonstop(Airline, Number, 'LAX', 'PHX')
```

A **maximal rewriting** (wrt UCQs) is:

```
q('UA', Number) :- flight_by_UA(Number, 'LAX', 'PHX')
```


Rewriting: extensions are not considered!

Slide 14

Observe that in computing the rewriting we are taking into account only the **view definitions**, not the view extensions!

How should such definitions be interpreted wrt to the possible extensions?

- A view definition gives properties that the **tuples** produced by the view must have.
- The view definition is **not** a guarantee that all such tuples are provided by the view (i.e., views are only **sound** in general).
- There is not even guarantee that the results produced by two views are consistent (they may need to be reconciled).

Example: a richer domain

Slide 15

Global schema:

`emp(E)` E is an employee
`phone(E,P)` P is E's phone
`office(E,O)` O is E's office
`mgr(E,M)` M is E's manager
`dept(E,D)` D is E's department

Views:

```
v1(E,P,M) -: emp(E), phone(E,P), mgr(E,M)
v2(E,O,D) -: emp(E), office(E,O), dept(E,D)
v3(E,P) -: emp(E), phone(E,P), dept(E,'ToyDept')
```

- `v1` gives information about employees, their phones and managers.
- `v2` gives information about employees, their offices and department.
- `v3` gives information about employees and their phones but only of the Toy Department.

Slide 16

Example: a richer domain (cont.)

Query: "What are Sally's phone and office?"

```
q(P,O) :- phone('Sally',P), office('Sally',O)
```

Maximal rewriting (wrt UCQs):

```
q(P,O) :- v1('Sally',P,M), v2('Sally',O,D)
q(P,O) :- v3('Sally',P), v2('Sally',O,D)
```

Slide 17

Example: view expansion (cont.)

If we **expand the views** in the rewriting above, we can compare the rewriting with the original query:

```
q(P,O) :- emp('Sally'), phone('Sally',P), mgr('Sally',M),
          emp('Sally'), office('Sally',O), dept('Sally',D)
q(P,O) :- emp('Sally'), phone('Sally',P), dept('Sally','ToyDept'),
          emp('Sally'), office('Sally',O), dept('Sally',D);
```

Observe:

- Both CQs are **contained** in the original query.
- The original query is contained in neither of them, nor in their union, i.e., the rewriting is **not exact**.
- These are the CQs that come closest to the original query while still constructable from the views (their union is a **maximal rewriting** wrt UCQs).

Slide 18

Rewriting algorithms

Given a query and a collection of views that are both **CQs**, there is a **maximal rewriting** made up of finite set of CQs (i.e., a UCQ)!

We can use two algorithms to compute maximal rewritings:

- **Bucket algorithm**
- **Inverse-rules algorithm**

Slide 19

Bucket algorithm: basis

Theorem 1 *If R is a CQ rewriting for a query Q, and R has more atoms then Q, then there exists a CQ rewriting R' such that $R \subseteq R'$.*

Proof. $R^{expd} \subseteq Q$ (rewriting), hence there is a containment mapping μ from Q to R^{expd}.

If R has more atoms that Q, then there is an atom α such that no atoms of Q is mapped to any atom that comes form α^{expd}.

If we delete α from R we get a new rewriting R' ($R'^{expd} \subseteq Q$, since μ is a containment mapping from Q to R'^{expd}).

Moreover $R \subseteq R'$ (identity mapping on atoms gives us the containment mapping). $\square$

Bucket algorithm: raw form

Slide 20

`for all` (CQs on the alphabet of the views with a number of atoms
that is less or equal to than those in Q) {

`if` ($R^{expd} \subseteq Q$) add R to the returned solution, `else` discard R

}

Bucket algorithm: refined form

Slide 21

`/* bucket initialization */`

Create a bucket for each atom α in the query Q, that will contain views that are relevant to answering the subgoal

Put a view V in the bucket for α if the definition of V contains an atom β that unifies with α

`/* solution generation */`

`for all` (CQs R on the alphabet of the views
formed by taking one atom from each bucket) {

`if` ($R^{expd} \subseteq Q$) add R to the returned solution, `else` discard R

}

Bucket algorithm: example

Slide 22

Views:

```
v1(E,P,M) -: emp(E), phone(E,P), mgr(E,M)
v2(E,O,D) -: emp(E), office(E,O), dept(E,D)
v3(E,P) -: emp(E), phone(E,P), dept(E,'ToyDept')
```

Query: "What are Sally's phone and office?"

```
q(P,O) :- phone('Sally',P), office('Sally',O)
```

Buckets:

B1	B2
`v1('Sally',P,M)`	`v2('Sally',O,D)`
`v3('Sally', P)`	

Rewriting:

```
q(P,O) :- v1('Sally',P,M), v2('Sally',O,D)
q(P,O) :- v3('Sally',P), v2('Sally',O,D)
```

Bucket algorithm: main result

Slide 23

Theorem 2 *The rewriting generated by the bucket algorithm is the* **maximal rewriting wrt UCQs**.

Proof. For the raw form, immediate.

For the refined form [Grahne&Mendelzon, ICDT'99] □

Inverse-rule algorithm: basis

Slide 24

Obtain **inverse rules** by:

1. Replace existential variable in the body of each view definition by a **Skolem function**.

   Recall `v(X) -: a1(X,Y), a2(X,Y)` stands for:
   $\forall x.v(x) \supset \exists y.a_1(x,y) \wedge a_2(x,y)$.
   By applying Skolemization we get $\forall x.v(x) \supset a_1(x, f(x)) \wedge a_2(x, f(x))$.

2. **Split** the body of the rules.

   Recall that $\forall x.v(x) \supset a_1(x, f(x)) \wedge a_2(x, f(x))$ is equivalent to
   $(\forall x.v(x) \supset a_1(x, f(x))) \ \wedge \ (\forall x.v(x) \supset a_2(x, f(x)))$, thus we obtain:

   ```
   a1(X,f(X)) :- v(X)
   a2(X,f(X)) :- v(X)
   ```

Evaluate the query considering **inverse rules as IDB** and **views as EDB**.

Inverse-rule algorithm: observations

Slide 25

- Because all functions symbols are in the head of the inverse rules, we never introduce a function symbol within a function symbol, leading to a finite process.
- Bottom-up evaluation can produce tuples with function symbols, but these cannot be real answers to the query and need to be discarded.
- The algorithm works for datalog (recursive) queries (but CQ views) as well.
- For non recursive queries we can easily get an equivalent UCQs by evaluating the query a la prolog.
- If required, functional symbols can be polynomially eliminated by adding new predicates.

Inverse-rule algorithm: example

Slide 26

Views:

```
v1(E,P,M) -: emp(E), phone(E,P), mgr(E,M)
v2(E,O,D) -: emp(E), office(E,O), dept(E,D)
v3(E,P) -: emp(E), phone(E,P), dept(E,'ToyDept')
```

Inverse rules:

```
emp(E) :- v1(E,P,M)                mgr(E,M) :- v1(E,P,M)
emp(E) :- v2(E,O,D)
emp(E) :- v3(E,P)                  office(E,O) :- v2(E,O,D)

phone(E,P) :- v1(E,P,M)            dept(E,D) :- v2(E, O, D)
phone(E,P) :- v3(E,P)              dept(E, 'ToyDept') :- v3(E,P)
```

Query: "What are Sally's phone and office?"

```
q(P,O) :- phone('Sally',P), office('Sally',O)
```

Inverse-rule algorithm: example (cont.)

Slide 27

Observe, if we unfold the query

```
q(P,O) :- phone('Sally',P), office('Sally',O)
```

using the inverse rules

```
emp(E) :- v1(E,P,M)                mgr(E,M) :- v1(E,P,M)
emp(E) :- v2(E,O,D)
emp(E) :- v3(E,P)                  office(E,O) :- v2(E,O,D)

phone(E,P) :- v1(E,P,M)            dept(E,D) :- v2(E, O, D)
phone(E,P) :- v3(E,P)              dept(E, 'ToyDept') :- v3(E,P)
```

we get

```
q(P,O) :- v1('Sally',P,M), v2('Sally',O,D)
q(P,O) :- v3('Sally',P), v2('Sally',O,D)
```

Inverse-rule algorithm: main result

Slide 28

Theorem 3 *For CQs (and UCQs) the rewriting generated by the inverse-rule algorithm is the* **maximal rewriting wrt UCQs**.

For datalog (recursive) queries the rewriting generated by the inverse-rule algorithm is the maximal rewriting wrt datalog queries.

Proof. See [Duschka&Genesereth, PODS'97]. □

Maximal rewriting vs certain answers

Slide 29

Theorem 4 *Every answer generated by evaluating the* **maximal rewriting** $Rmax_{UCQs}(Q, \mathcal{V})$ *is a* **certain answer** *in* $cert(Q, \mathcal{V})$.

Proof. Trivial. □

Maximal rewriting vs certain answers (cont.)

Slide 30

Theorem 5 *Every* **certain answer** *in* $cert(Q, \mathcal{V})$ *is generated by evaluating the* **maximal rewriting** $Rmax_{UCQs}(Q, \mathcal{V})$.

Proof. Suppose not. Let $t \in cert(Q, \mathcal{V})$ such that $t \notin ans(Rmax_{UCQs}(Q, \mathcal{V}), ext(\mathcal{V}))$. Consider the CQ C_t on the alphabet of the views defined as:

$$c(t) \text{ :- } V_1(t_{11}), \cdots, V_1(t_{1k_1}), \cdots, V_n(t_{n1}), \cdots, V_n(t_{nk_n})$$

where $V_1, \cdots V_n$ are all the views in $\mathcal{V}$ and $t_{i,1} \cdots, t_{ik_i}$ is $ext(v_i)$.

(cont...)

Slide 31

Then, $C_t^{expd} \subseteq Q$, indeed

- for all DB conforming with the views $ans(C_t^{expd}, DB) = \{t\} \subseteq ans(Q, DB)$
- while for those DB' not conforming with the views $ans(C_t^{expd}, DB') = \emptyset \subseteq ans(Q, DB')$.

Being a CQ on the alphabet of the views and being $C_t^{expd} \subseteq Q$ it follows $C_t \subseteq Rmax_{UCQs}(Q, \mathcal{V})$. Contradiction! □

Slide 32

Part 1: conclusions

For the setting considered here, i.e,

- **View definitions are CQs** over the global schema
- **View are all sound**
- **Queries are CQs**, or also **UCQs**, i.e., union of conjunctive queries, over the global schema
- **Rewritings are UCQs**, i.e., finite sets of CQs, in the alphabet of the views

everything works fine:

- We can focus on **maximal rewriting** wrt UCQs.
- Maximal rewriting computes exactly the **certain answers**.
- Maximal rewriting is expressible in a simple (PTIME in data complexity) query language.

Does these nice results extend to more general settings? See Part 2 and Part 3.

Slide 33

Part 2

View-based query processing: Computing the answer to a query based on a set of views, rather than on the raw data in the database.

- View-based query rewriting (indirect)
- View-based query answering (direct)
- Relationship between query answering and query rewriting

We are studying view-based query processing within the **relational data model**.

Slide 34

View-based query answering

Basic Idea: Given a query Q, a set of views $\mathcal{V}$ with definitions $def(\mathcal{V})$ and extensions $ext(\mathcal{V})$, compute the tuples t which are in the answer to Q in all databases consistent with the views (certain tuples).

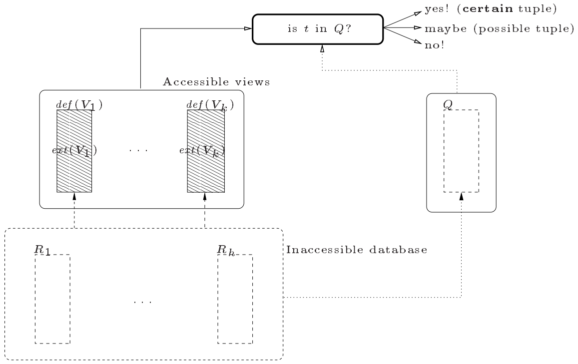

Slide 35

View-based query rewriting

Basic Idea: Given a query Q and a set of views $\mathcal{V}$ with definitions $def(\mathcal{V})$, reformulate Q into a new query Q' expressed in some language on the alphabet of the view names $\mathcal{V}$.

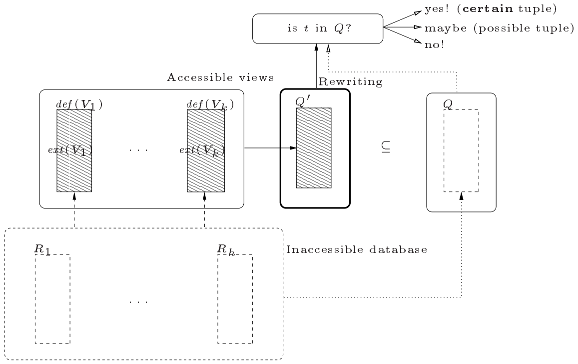

View-based query answering

Slide 36

There are several results also on view-based query answering.

- Complexity for several query and view languages [Abiteboul et al 98], [Grahne et al 99];
- Description Logics queries [Calvanese et al, AAAI'00];
- Regular path queries [Calvanese et al, ICDE'00, PODS'00].

View-based query answering

Slide 37

Query answering (QA) is defined as follows:

$\vec{\mathbf{d}} \in cert(Q, \mathcal{V})$ **iff** $(d_1, \ldots, d_n) \in ans(Q, DB)$**, for each** DB **such that:**

- DB **satisfies** $\mathcal{G}$
- DB **is coherent with** $V_1, \ldots, V_m$

View-based query answering and query containment

Slide 38

There is a strong relationship between QA and QC when the views are **sound**.

For query languages at least as powerful as CQs (CQs, UCQs, datalog, FOL, etc.).

QA and QC are mutually reducible one into the other in polynomial time!

Reduction to query containment

Slide 39

Consider the query Q' on the alphabet of the views defined as:

$$Q'(t) \text{ :- } V_1(t_{11}), \cdots, V_1(t_{1k_1}), \cdots, V_n(t_{n1}), \cdots, V_n(t_{nk_n})$$

where $V_1, \cdots V_n$ are all the views in $\mathcal{V}$ and $t_{i,1} \cdots, t_{ik_i}$ is $ext(V_i)$.

Theorem 6 $t \in cert(Q, \mathcal{V})$ *iff* $Q'^{expd} \subseteq Q$.

Slide 40

Proof.
$\Rightarrow$ If $t \in cert(Q, \mathcal{V})$ then

- for all DB conforming with the views $ans(Q'^{expd}, DB) = \{t\} \subseteq ans(Q, DB)$
- while for those DB' not conforming with the views $ans(Q'^{expd}, DB') = \emptyset \subseteq ans(Q, DB')$.

hence $Q'^{expd} \subseteq Q$.

$\Leftarrow$ Assume $Q'^{expd} \subseteq Q$. For all DB conforming with the views $ans(Q'^{expd}, DB) = \{t\}$ and since $Q'^{expd} \subseteq Q$, we have $t \in ans(Q, DB)$. $\square$

Slide 41

Reduction from query containment

Let Q_1 and Q_2 be two queries (x, y are tuples of variables):

$$Q_1(x) \text{ -: } body_1(x, y)$$
$$Q_2(x) \text{ -: } body_2(x, y)$$

Consider a single view V:

- $def(V)$: $V(c)$ -: $body_1(x, y), p(x)$ (p is a new predicate)
- $ext(V) = \{c\}$
- $as(V) = sound$

and the query Q defined as

$$Q(c) \text{ :- } body_2(x, y), p(x)$$

Theorem 7 *$Q_1 \subseteq Q_2$ iff $c \in cert(Q, \{V\})$.*

Slide 42

Proof.
$\Rightarrow$

If $c \notin cert(Q, \mathcal{V})$, then there exists a database DB conforming with the views such that $c \notin ans(Q, DB)$. This implies that there exists a tuple in $t \in ans(Q_1, DB)$ but $t \notin ans(Q_2, DB)$, i.e., $Q_1 \not\subseteq Q_2$.

$\Leftarrow$ Assume $Q_1 \not\subseteq Q_2$. Then there exists a database DB and a tuple t such that $t \in ans(Q_1, DB)$ but $t \notin ans(Q_2, DB)$. We can extend DB by assigning to new predicate p the interpretation $ans(p, DB) = \{t\}$. But then we have that DB conforms to the view V while $ans(Q, DB) = \emptyset$. Hence $c \notin cert(Q, \mathcal{V})$. $\square$

Slide 43

QA and QC: observations

- We can transfer **upper-bounds** from QC to QA, using the reduction form QA to QC.
- We can transfer **lower-bounds** from QC to QA, using the reduction from QC to QA.
- We can use algorithms for query containment to get **algorithms** for query answering.
- *What kind of complexity are we characterizing for QA?* **Combined complexity.**

Slide 44

Combined complexity of QA for sound views

$QC[Q_1, Q_2]$	CQs	PQs*	datalog	FOL
CQs	NP	NP	EXPTIME	undec.
PQ	Π_2^p	Π_2^p	EXPTIME	undec.
datalog	2EXPTIME	2EXPTIME	undec.	undec.
FOL	undec.	undec.	undec.	undec.
$QA[\mathcal{V}, Q]$	CQs	PQs	datalog	FOL
CQs	**NP**	**NP**	**EXPTIME**	**undec.**
PQs	Π_2^p	Π_2^p	**EXPTIME**	**undec.**
datalog	**2EXPTIME**	**2EXPTIME**	**undec.**	**undec.**
FOL	**undec.**	**undec.**	**undec.**	**undec.**

* PQs, i.e., are UCQs that allow to nest disjunctions in conjunctions.

Slide 45

What about data complexity?

We want to refine the complexity analysis to take into account that data in the database are orders of magnitude bigger that the size of the queries.

In other words we would like to characterize the **data complexity** of QA!

Slide 46

Complexity of view based query answering

Can be measured in three different ways:

data complexity: as a function of the size of the view extensions $ext(V_1) \cup \cdots \cup ext(V_k)$

expression complexity: as a function of the size of the query Q and of the view definitions $def(V_1), \ldots, def(V_k)$

combined complexity: as a function of the size of both $ext(V_1) \cup \cdots \cup ext(V_k)$ and the expressions Q, $def(V_1), \ldots, def(V_k)$

Slide 47

QA coNP-hard for UCQs views

Theorem 8 *Let $\mathcal{L}_\mathcal{Q}$ and $\mathcal{L}_\mathcal{V}$ be UCQs, and let view be sound. Then QA is coNP-hard in data complexity.*

Proof. Reduction from graph 3-colorability.

Views:
$$\begin{array}{lll} V_s(x,y) & \text{:-} & R_s(x,y) \\ V_G(x,y) & \text{:-} & R_{rg}(x,y) \cup R_{gr}(x,y) \cup R_{rb}(x,y) \cup R_{br}(x,y) \cup R_{gb}(x,y) \cup R_{bg}(x,y) \\ V_f(x,y) & \text{:-} & R_f(x,y) \end{array}$$

Query:
$$Q(x,y) \;\text{:-}\; \bigcup_{\beta \neq \gamma}(R_s(x,z), R_{\alpha\beta}(z,v), R_{\gamma\delta}(v,w), R_f(w,y))$$

Only the view extensions depend on graph $G = (N, E)$

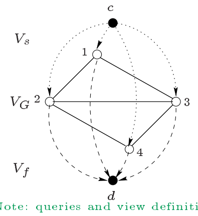

Extensions:
$$\begin{array}{lll} ext(V_s) & = & \{(c,a) \mid a \in N\} \\ ext(V_G) & = & \{(a,b),(b,a) \mid (a,b) \in E\} \\ ext(V_f) & = & \{(a,d) \mid a \in N\} \end{array}$$

G is 3-colorable iff (c,d) is not a certain answer of Q.

Note: queries and view definitions used in this proof are UPQs! See Part 3.

QA for PQs views: algorithm

Slide 48

Let n be the number of tuples in $ext(\mathcal{V})$ and k the number of atoms in the longest (wrt atoms) view definition.

Theorem 9 *Let $\mathcal{L}_{\mathcal{V}}$ be PQs and $\mathcal{L}_Q$ be datalog, then $t \notin cert(Q, \mathcal{V})$ iff there exists a database DB' of* **size** nk *conforming with the views such that $t \notin ans(Q, DB)$. (For both sound and exact views).*

Proof. Since $t \notin cert(Q, \mathcal{V})$, there exists a database DB conforming with the views such that $t \notin ans(Q, DB)$.

Consider the database $DB' \subseteq DB$ having only the nk tuples required by the views. DB' still conforms with the views, moreover still have $t \notin ans(Q, DB')$ (it is sufficient only that Q is monotone!). □

QA for PQs views: algorithm (cont.)

Slide 49

Theorem 10 *Let $\mathcal{L}_{\mathcal{V}}$ be PQs and $\mathcal{L}_Q$ be datalog, then QA is in* **coNP**. *(For both sound and exact views.)*

Algorithm:

```
bool certain (tuple t, query Q, views V)
{
   guess(a database DB' of size nk) {   (nondet.)
      verify whether DB' conforms to the views   (poly.)
      verify whether t ∉ ans(Q, DB')   (poly.)
      if (both test positive ) return false
   }
   return true;
}
```


Data complexity for sound views

Slide 50

Sound Views	CQ	PQ	datalog	FOL
CQ	**PTIME***	**PTIME***	**PTIME***	**undec.**
PQ	**coNP**	**coNP**	**coNP**	**undec.**
datalog	**coNP**	**coNP**	**undec.**	**undec.**
FOL	**undec.**	**undec.**	**undec.**	**undec.**

* This is shown, by proving that the maximal rewriting wrt a PTIME query language (namely CQs, PQs, datalog) computes exactly the certain answers. See Part 1.

What about exact views?

Slide 51

Till now we have focused on sound views. What happen when we consider **exact views***?*

QA becomes coNP-hard even for views defined by CQs! See below.

QA coNP-hard for exact CQs views

Slide 52

Theorem 11 *Let $\mathcal{L}^{\mathcal{V}}$ and $\mathcal{L}^{Q}$ be CQs, and let $as(\mathcal{V}) = exact$. Then verifying whether $c \in cert(Q, \mathcal{V})$ is* **coNP-hard** *in data complexity.*

Proof. Reduction from graph 3-colorability. Let $G = (N, E)$ be an arbitrary graph. Consider three exact views $\mathcal{V} = \{V_1, V_2, V_3\}$:

definitions	extensions
$V_1(x)$:- $color(x, y)$	$ext(V_1) = N$
$V_2(y)$:- $color(x, y)$	$ext(V_2) = \{red, green, blue\}$
$V_3(y)$:- $edge(x, y)$	$ext(V_3) = E$

and the query Q:

$$Q(c) \text{ :- } edge(x, y), color(x, z), color(y, z)$$

Then $c \in cert(Q, \mathcal{V})$ iff the graph G is not 3-colorable. □

Data complexity for exact views

Slide 53

Complete Views	CQ	PQ	datalog	FOL
CQ	**coNP**	**coNP**	**coNP**	**undec.**
PQ	**coNP**	**coNP**	**coNP**	**undec.**
datalog	**undec.**	**undec.**	**undec.**	**undec.**
FOL	**undec.**	**undec.**	**undec.**	**undec.**

Part 2: conclusions

Slide 54

We have gathered a lot of results on query rewriting (Part 1) and on query answering (Part 2).

What can we say now about the relationships between these two ways of performing view-based query processing? See Part 3.

Part 3

Slide 55

View-based query processing: Computing the answer to a query based on a set of views, rather than on the raw data in the database.

- View-based query rewriting (indirect)
- View-based query answering (direct)
- Relationship between query answering and query rewriting

We are studying view-based query processing within the **relational data model**.

Slide 56

View-based query answering

Basic Idea: Given a query Q, a set of views $\mathcal{V}$ with definitions $def(\mathcal{V})$ and extensions $ext(\mathcal{V})$, compute the tuples t which are in the answer to Q in all databases consistent with the views (certain tuples).

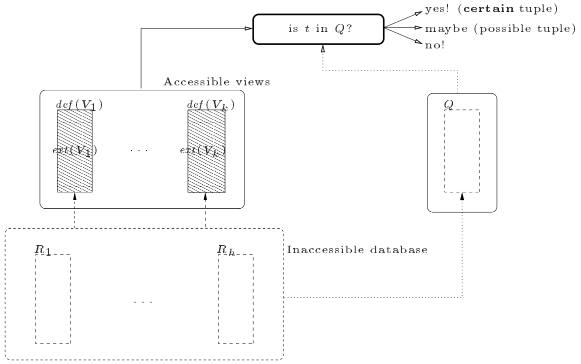

Slide 57

View-based query rewriting

Basic Idea: Given a query Q and a set of views $\mathcal{V}$ with definitions $def(\mathcal{V})$, reformulate Q into a new query Q' expressed in some language on the alphabet of the view names $\mathcal{V}$.

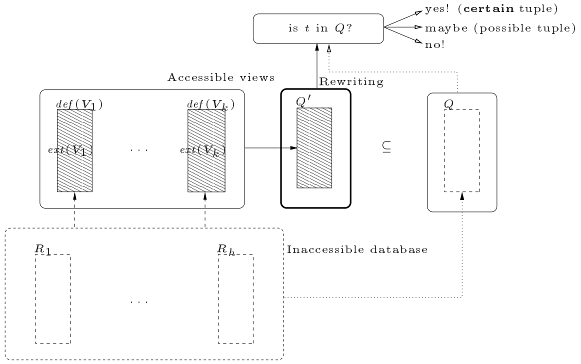

Slide 58

Previous results

There are many previous results on view-based query processing.

View-based query rewriting:

- Conjunctive queries
- Queries with aggregates
- Under constraints (inclusion and functional dependencies)
- Recursive queries
- Description Logics queries
- Queries for semi-structured data

View-based query answering:

- Complexity for several query and view languages
- Under Description Logics constraints
- Regular path queries

But: *No clear understanding of the relationships between query answering and query rewriting!!!*

Slide 59

Query rewriting vs. query answering

Query answering by rewriting:

- Rewrite the query in the alphabet of the view names;
- Evaluate the rewriting on the view extension.

Typically people are interested in **rewritings** that are:

- Contained in the original query
- Expressed in a given **query language**
- **Maximal** for the given class of queries
- Exact(??), i.e., rewritings that are logically equivalent to the query, if possible (observe that such rewritings may not exists).

But:

- **Q1:** *When is the rewriting also* **complete** *– i.e., computes all certain tuples?*
- **Q2:** *What do we gain or lose by focusing on a given class of queries?*

Perfect rewriting

Slide 60

Query answering: Let $cert(Q, def(\mathcal{V}), ext(\mathcal{V}))$ be the function that computes the certain tuples for Q wrt $def(\mathcal{V})$ and $ext(\mathcal{V})$.

Perfect rewriting:
Define $cert_{[Q,def(\mathcal{V})]}$ to be the **Currying** of $cert$ wrt Q and $def(\mathcal{V})$.

$\Rightarrow$
- $cert_{[Q,def(\mathcal{V})]}$ is a query on the alphabet of the view names that given $ext(\mathcal{V})$ returns the certain tuples for the query Q wrt $def(\mathcal{V})$ and $ext(\mathcal{V})$;
- $cert_{[Q,def(\mathcal{V})]}$ is a (*sound*) **rewriting** of Q wrt $def(\mathcal{V})$;
- $cert_{[Q,def(\mathcal{V})]}$ is **complete** (no better rewritings may exist);
- $cert_{[Q,def(\mathcal{V})]}$ is called the **perfect rewriting** of Q wrt $def(\mathcal{V})$.

Comparing with the perfect rewriting

Slide 61

- **Q1:** *Can we express the perfect rewriting in a certain query language?*
- **Q2:** *How does maximal rewriting for a given class of queries compare with the perfect rewriting?*
  - *From a semantical point of view?*
  - *From a computational point of view?*
- *Which is the computational complexity of the perfect rewriting?*

The case of conjunctive queries

Slide 62

Rewriting: ([Levy et al 95], [Duschka at al 97], [Abiteboul et al 98])
Let Q and $def(\mathcal{V})$ be CQs, and let Q' be the **union of all maximal rewritings for the class of CQs**. Then:

- Q' is the maximal rewriting for the class of unions of conjunctive queries (UCQs);
- **Query answering:** generate Q', evaluate Q' on $ext(\mathcal{V})$;
- **Q' is the perfect rewriting**;
- Q' is a PTIME query (in fact, NLOGSPACE).
- Q' is an exact rewriting, if an exact rewriting exists.

Q: *Does this "ideal situation" carry on to cases where Q and $def(\mathcal{V})$ allow for union?*

Unions of path queries (UPQs)

Slide 63

Very simple query language defined as follows:

$$\begin{array}{lcl} Q & \longrightarrow & P \mid Q_1 \cup Q_2 \\ P & \longrightarrow & R \mid P_1 \circ P_2 \end{array}$$

where R denotes a **binary database relation**, P denotes a **path query**, which is a chaining of database relations, and Q denotes a **union of path queries**.

Observe: UPQs are a simple form of:

- Unions of conjunctive queries;
- Regular path queries.

View-based query processing for UPQs

Slide 64

Thm [Calvanese et al, ICDE'00]: View-based query answering for UPQs is coNP-complete in data complexity.

In other words: $cert(Q, def(\mathcal{V}), ext(\mathcal{V}))$ with Q and $def(\mathcal{V})$ fixed is coNP-complete.

$\Rightarrow$ **The perfect rewriting $cert_{[Q,def(\mathcal{V})]}$ is a coNP-complete query.**

Bad news: *For query languages that include UPQs the perfect rewriting is coNP-hard!*

PTIME perfect rewritings

Slide 65

Typically we are interested in **PTIME queries**.

Program: *Isolate those UPQs Q and $def(\mathcal{V})$ for which the perfect rewriting is PTIME (assuming P$\neq$NP).*

We tackle this issue exhibiting a tight correspondence between view-based query processing and constraint satisfaction problems.

Constraint satisfaction problems

Slide 66

Uniform constraint satisfaction problem CSP: given two relational structures A and B over the same vocabulary, determine if $CSP(A, B)$ is satisfiable – i.e., if there exists an homomorphism $h : A \to B$.

Non-uniform constraint satisfaction problem CSP_B: is the Currying of CSP wrt to B, i.e., for a fixed relational structure B, given a relational structure A over the same vocabulary as B, determine if $CSP_B(A)$ is satisfiable – i.e., if there exists an homomorph. $h : A \to B$.

Complexity:

- CSP is NP-complete;
- CSP_B can be in PTIME for certain B;
- But also, CSP_B can be NP-hard for certain B.

CSP and query processing for UPQs

Slide 67

Intimate connection: between CSP and view-based query processing

Thm [Calvanese et al, LICS'00]

- **Unform CSP and QA for UPQa are mutually reducible!**
- **Nonuniform CSP and evaluating the perfect rewriting for UPQs are mutually reducible!!**

CSP and query processing for UPQs (cont.)

Slide 68

Long standing open problem: ([Feder&Vardi 93,99],[Jeavons et al 95,97])

- *Characterize the structures B such that CSP_B is in PTIME.*

Consequence:

$\Rightarrow$ Characterizing PTIME $cert_{Q,def(\mathcal{V})}$ for UPQs – and hence query languages that includes UPQs – and characterizing PTIME CSP_B are closely related.

Bad news:

- *Characterizing PTIME $cert_{Q,def(\mathcal{V})}$ for query languages including UPQs is going to be difficult;*
- *Testing PTIME rewritings for perfectness in general is going to be difficult in general.*

Rewritings in PTIME query languages

Slide 69

Fall back Program: *Fix the language of the rewriting, choosing a PTIME query language.*

Observe: This is exactly what is done in most papers on rewriting!!!

Important problem (open for most languages):

- How can we test the rewriting obtained for perfectness?

What about QA under constraints?

Slide 70

We have considered QA (and QR) over a global (relational) schema without any form of constraint. In other words, we have not specified any relationship between the (virtual) relations of the global schema.

It would be important to understand QA under constraints (see Lenzerini's slides).

In [Calvanese et al, AAAI'00] QA under constraints is studied in the setting of a global schema expressed in the description logics $\mathcal{DLR}$. See AAAI'00 slides.