

# Ontology and Database Systems: Knowledge Representation and Ontologies

## Part 4: Ontology Based Data Access

*Diego Calvanese*

Faculty of Computer Science  
European Master in Computational Logic

A.Y. 2017/2018



Fakultät für Informatik  
Facoltà di Scienze e Tecnologie informatiche  
Faculty of Computer Science

# Part 4

## Ontology-based data access

# Outline of Part 4

- 1 The *DL-Lite* family of tractable Description Logics
  - Basic features of *DL-Lite*
  - Syntax and semantics of *DL-Lite*
  - Identification assertions in *DL-Lite*
  - Members of the *DL-Lite* family
  - Properties of *DL-Lite*
  
- 2 Linking ontologies to relational data
  - The impedance mismatch problem
  - Ontology-Based Data Access systems
  - Query answering in Ontology-Based Data Access
  - The ONTOP framework for Ontology-Based Data Access

# Outline of Part 4

## 1 The *DL-Lite* family of tractable Description Logics

- Basic features of *DL-Lite*
- Syntax and semantics of *DL-Lite*
- Identification assertions in *DL-Lite*
- Members of the *DL-Lite* family
- Properties of *DL-Lite*

## 2 Linking ontologies to relational data

# Outline of Part 4

## 1 The *DL-Lite* family of tractable Description Logics

- Basic features of *DL-Lite*
  - Syntax and semantics of *DL-Lite*
  - Identification assertions in *DL-Lite*
  - Members of the *DL-Lite* family
  - Properties of *DL-Lite*

## 2 Linking ontologies to relational data

# The *DL-Lite* family

[Calvanese, De Giacomo, Lembo, Lenzerini, and Rosati 2007a; Artale et al. 2009; Calvanese, De Giacomo, Lembo, Lenzerini, Poggi, et al. 2009]

- A family of DLs optimized according to the tradeoff between expressive power and **complexity** of query answering, with emphasis on **data**.
- Carefully designed to have nice computational properties for answering UCQs (i.e., computing certain answers):
  - The same data complexity as relational databases.
  - In fact, query answering can be delegated to a relational DB engine.
  - The DLs of the *DL-Lite* family are essentially the maximally expressive ontology languages enjoying these nice computational properties.
- Captures conceptual modeling formalism.

The *DL-Lite* family provides new foundations for Ontology-Based Data Access.

# Basic features of $DL\text{-}Lite_{\mathcal{A}}$

$DL\text{-}Lite_{\mathcal{A}}$  is an expressive member of the  $DL\text{-}Lite$  family [Calvanese, De Giacomo, Lembo, Lenzerini, Poggi, et al. 2009].

- Takes into account the distinction between **objects** and **values**:
  - Objects are elements of an abstract interpretation domain.
  - Values are elements of concrete data types, such as integers, strings, ecc.
  - Values are connected to objects through **attributes** (rather than roles).
- Is equipped with identification assertions.
- Captures most of UML class diagrams and Extended ER diagrams.
- Enjoys nice computational properties, both w.r.t. the traditional reasoning tasks, and w.r.t. query answering (see later).

# Outline of Part 4

## 1 The *DL-Lite* family of tractable Description Logics

- Basic features of *DL-Lite*
- Syntax and semantics of *DL-Lite*
- Identification assertions in *DL-Lite*
- Members of the *DL-Lite* family
- Properties of *DL-Lite*

## 2 Linking ontologies to relational data



# Syntax of the $DL\text{-}Lite_A$ description language

- Role expressions:

- atomic role:  $P$
- basic role:  $Q ::= P \mid P^-$
- arbitrary role:  $R ::= Q \mid \neg Q$  (to express disjointness)

- Concept expressions:

- atomic concept:  $A$
- basic concept:  $B ::= A \mid \exists Q \mid \delta(U)$
- arbitrary concept:  $C ::= \top_C \mid B \mid \neg B$  (to express disjointness)

- Attribute expressions:

- atomic attribute:  $U$
- arbitrary attribute:  $V ::= U \mid \neg U$  (to express disjointness)

- Value-domain expressions:

- attribute range:  $\rho(U)$
- RDF datatypes:  $T_i$
- top domain:  $\top_D$

# Semantics of $DL\text{-}Lite_{\mathcal{A}}$ – Objects vs. values

	Objects	Values
Interpretation domain $\Delta^{\mathcal{I}}$	Domain of objects $\Delta_O^{\mathcal{I}}$	Domain of values $\Delta_V^{\mathcal{I}}$
Alphabet $\Gamma$ of constants	Object constants $\Gamma_O$ $c^{\mathcal{I}} \in \Delta_O^{\mathcal{I}}$	Value constants $\Gamma_V$ $d^{\mathcal{I}} = \text{val}(d)$ given a priori
Unary predicates	Concept $C$ $C^{\mathcal{I}} \subseteq \Delta_O^{\mathcal{I}}$	RDF datatype $T_i$ $T_i^{\mathcal{I}} \subseteq \Delta_V^{\mathcal{I}}$ given a priori
Binary predicates	Role $R$ $R^{\mathcal{I}} \subseteq \Delta_O^{\mathcal{I}} \times \Delta_O^{\mathcal{I}}$	Attribute $V$ $V^{\mathcal{I}} \subseteq \Delta_O^{\mathcal{I}} \times \Delta_V^{\mathcal{I}}$

# Semantics of the $DL\text{-}Lite_A$ constructs

Construct	Syntax	Example	Semantics
atomic role	$P$	<i>child</i>	$P^I \subseteq \Delta_O^I \times \Delta_O^I$
inverse role	$P^-$	<i>child</i> <sup>-</sup>	$\{(o, o') \mid (o', o) \in P^I\}$
role negation	$\neg Q$	$\neg$ <i>manages</i>	$(\Delta_O^I \times \Delta_O^I) \setminus Q^I$
atomic concept	$A$	<i>Doctor</i>	$A^I \subseteq \Delta_O^I$
existential restriction	$\exists Q$	$\exists$ <i>child</i> <sup>-</sup>	$\{o \mid \exists o'. (o, o') \in Q^I\}$
concept negation	$\neg B$	$\neg \exists$ <i>child</i>	$\Delta^I \setminus B^I$
attribute domain	$\delta(U)$	$\delta$ ( <i>salary</i> )	$\{o \mid \exists v. (o, v) \in U^I\}$
top concept	$\top_C$		$\top_C^I = \Delta_O^I$
atomic attribute	$U$	<i>salary</i>	$U^I \subseteq \Delta_O^I \times \Delta_V^I$
attribute negation	$\neg U$	$\neg$ <i>salary</i>	$(\Delta_O^I \times \Delta_V^I) \setminus U^I$
top domain	$\top_D$		$\top_D^I = \Delta_V^I$
datatype	$T_i$	<i>xsd:int</i>	$T_i^I \subseteq \Delta_V^I$ (predefined)
attribute range	$\rho(U)$	$\rho$ ( <i>salary</i> )	$\{v \mid \exists o. (o, v) \in U^I\}$
object constant	$c$	<i>john</i>	$c^I \in \Delta_O^I$
value constant	$d$	' <i>john</i> '	$val(d) \in \Delta_V^I$ (predefined)

# DL-Lite<sub>A</sub> assertions

**TBox** assertions can have the following forms:

- Inclusion assertions (also called positive inclusions):

$$\begin{array}{ll} B_1 \sqsubseteq B_2 & \text{concept inclusion} \\ Q_1 \sqsubseteq Q_2 & \text{role inclusion} \end{array} \qquad \begin{array}{ll} \rho(U) \sqsubseteq T_i & \text{value-domain inclusion} \\ U_1 \sqsubseteq U_2 & \text{attribute inclusion} \end{array}$$

- Disjointness assertions (also called negative inclusions):

$$\begin{array}{ll} B_1 \sqsubseteq \neg B_2 & \text{concept disjointness} \\ Q_1 \sqsubseteq \neg Q_2 & \text{role disjointness} \end{array} \qquad U_1 \sqsubseteq \neg U_2 \quad \text{attribute disjointness}$$

- Functionality assertions:

$$(\text{funct } Q) \quad \text{role functionality} \qquad (\text{funct } U) \quad \text{attribute functionality}$$

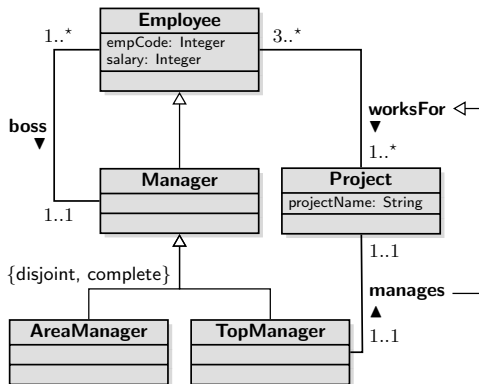
- Identification assertions:  $(\text{id } B \ I_1, \dots, I_n)$   
where each  $I_j$  is a role, an inverse role, or an attribute

**ABox** assertions:  $A(c), \ P(c, c'), \ U(c, d),$   
where  $c, c'$  are object constants and  $d$  is a value constant

# Semantics of the $DL\text{-}Lite_{\mathcal{A}}$ assertions

Assertion	Syntax	Example	Semantics
conc. incl.	$B_1 \sqsubseteq B_2$	$Father \sqsubseteq \exists child$	$B_1^{\mathcal{I}} \subseteq B_2^{\mathcal{I}}$
role incl.	$Q_1 \sqsubseteq Q_2$	$father \sqsubseteq anc$	$Q_1^{\mathcal{I}} \subseteq Q_2^{\mathcal{I}}$
v.dom. incl.	$\rho(U) \sqsubseteq T_i$	$\rho(age) \sqsubseteq xsd:int$	$\rho(U)^{\mathcal{I}} \subseteq T_i^{\mathcal{I}}$
attr. incl.	$U_1 \sqsubseteq U_2$	$offPhone \sqsubseteq phone$	$U_1^{\mathcal{I}} \subseteq U_2^{\mathcal{I}}$
conc. disj.	$B_1 \sqsubseteq \neg B_2$	$Person \sqsubseteq \neg Course$	$B_1^{\mathcal{I}} \cap B_2^{\mathcal{I}} = \emptyset$
role disj.	$Q_1 \sqsubseteq \neg Q_2$	$sibling \sqsubseteq \neg cousin$	$Q_1^{\mathcal{I}} \cap Q_2^{\mathcal{I}} = \emptyset$
attr. disj.	$U_1 \sqsubseteq \neg U_2$	$offPhn \sqsubseteq \neg homePhn$	$U_1^{\mathcal{I}} \cap U_2^{\mathcal{I}} = \emptyset$
role funct.	( <b>funct</b> $Q$ )	( <b>funct</b> $father$ )	$\forall o, o_1, o_2. (o, o_1) \in Q^{\mathcal{I}} \wedge (o, o_2) \in Q^{\mathcal{I}} \rightarrow o_1 = o_2$
att. funct.	( <b>funct</b> $U$ )	( <b>funct</b> $ssn$ )	$\forall o, v, v'. (o, v) \in U^{\mathcal{I}} \wedge (o, v') \in U^{\mathcal{I}} \rightarrow v = v'$
id const.	( <b>id</b> $B \ I_1, \dots, I_n$ )	( <b>id</b> $Person \ name, dob$ )	$I_1, \dots, I_n$ identify instances of $B$
mem. asser.	$A(c)$	$Father(bob)$	$c^{\mathcal{I}} \in A^{\mathcal{I}}$
mem. asser.	$P(c_1, c_2)$	$child(bob, ann)$	$(c_1^{\mathcal{I}}, c_2^{\mathcal{I}}) \in P^{\mathcal{I}}$
mem. asser.	$U(c, d)$	$phone(bob, '2345')$	$(c^{\mathcal{I}}, val(d)) \in U^{\mathcal{I}}$

# DL-Lite<sub>A</sub> – Example



$Manager \sqsubseteq Employee$   
 $AreaManager \sqsubseteq Manager$   
 $TopManager \sqsubseteq Manager$   
 $AreaManager \sqsubseteq \neg TopManager$   
 $Employee \sqsubseteq \delta(empCode)$   
 $\delta(empCode) \sqsubseteq Employee$   
 $\rho(empCode) \sqsubseteq xsd:int$   
 $(\text{funcnt } empCode)$   
 $(\text{id } Employee \ empCode)$   
 $\exists worksFor \sqsubseteq Employee$   
 $\exists worksFor^- \sqsubseteq Project$   
 $Employee \sqsubseteq \exists worksFor$   
 $Project \sqsubseteq \exists worksFor^-$   
 $(\text{funcnt } manages)$   
 $(\text{funcnt } manages^-)$   
 $manages \sqsubseteq worksFor$   
 $\vdots$

**Note:** DL-Lite<sub>A</sub> cannot capture completeness of a hierarchy. This would require **disjunction** (i.e., **OR**).


# Outline of Part 4

## 1 The *DL-Lite* family of tractable Description Logics

- Basic features of *DL-Lite*
- Syntax and semantics of *DL-Lite*
- Identification assertions in *DL-Lite*
- Members of the *DL-Lite* family
- Properties of *DL-Lite*

## 2 Linking ontologies to relational data

# Identification assertions – Example



What we would like to additionally capture:

- ① No two leagues with the same year and the same nation exist
- ② Within a certain league, the code associated to a round is unique
- ③ Every match is identified by its code within its round
- ④ Every referee can umpire at most one match in the same round
- ⑤ No team can be the home team of more than one match per round
- ⑥ No team can be the host team of more than one match per round




# Identification assertions – Example (cont'd)

$League \sqsubseteq \exists of$	$PlayedMatch \sqsubseteq Match$	
$\exists of \sqsubseteq League$	$\delta(code_M) \sqsubseteq Match$	$Match \sqsubseteq \delta(code_M)$
$\exists of^- \sqsubseteq Nation$	$\delta(code_R) \sqsubseteq Round$	$Round \sqsubseteq \delta(code_R)$
$Round \sqsubseteq \exists belongsTo$	$\delta(playedOn_P) \sqsubseteq PlayedMatch$	$\dots$
$\exists belongsTo \sqsubseteq Round$	$\dots$	
$\exists belongsTo^- \sqsubseteq League$	$\rho(code_M) \sqsubseteq \text{xsd:int}$	
$Match \sqsubseteq \exists playedIn$	$\rho(playedOn_P) \sqsubseteq \text{xsd:date}$	
$\dots$	$\dots$	

( <b>funct</b> <i>of</i> )	( <b>funct</b> <i>hostTeam</i> )	( <b>funct</b> <i>homeGoals</i> )
( <b>funct</b> <i>belongsTo</i> )	( <b>funct</b> <i>umpiredBy</i> )	( <b>funct</b> <i>hostGoals</i> )
( <b>funct</b> <i>playedIn</i> )	( <b>funct</b> <i>code</i> )	( <b>funct</b> <i>playedOn</i> )
( <b>funct</b> <i>homeTeam</i> )	( <b>funct</b> <i>year</i> )	

# Identification assertions – Example (cont'd)



- ① No two leagues with the same year and the same nation exist
- ② Within a certain league, the code associated to a round is unique
- ③ Every match is identified by its code within its round
- ④ Every referee can umpire at most one match in the same round
- ⑤ No team can be the home team of more than one match per round
- ⑥ No team can be the host team of more than one match per round

(id *League of*,  $year_L$ )

(id *Round belongsTo*,  $code_R$ )

(id *Match playedIn*,  $code_M$ )

(id *Match umpiredBy*, *playedIn*)

(id *Match homeTeam*, *playedIn*)

(id *Match hostTeam*, *playedIn*)

# Semantics of identification assertions

Let  $(\text{id } B \ I_1, \dots, I_n)$  be an identification assertion in a  $DL\text{-}Lite_{\mathcal{A}}$  TBox.

An interpretation  $\mathcal{I}$  satisfies such an assertion if for all  $o_1, o_2 \in B^{\mathcal{I}}$ , if there exist objects or values  $u_1, \dots, u_n$  such that

$$(o_1, u_j) \in I_j^{\mathcal{I}} \text{ and } (o_2, u_j) \in I_j^{\mathcal{I}}, \text{ for } j \in \{1, \dots, n\},$$

then  $o_1 = o_2$ .

In other words, the instance  $o_i$  of  $B$  is identified by the tuple  $(u_1, \dots, u_n)$  of objects or values to which it is connected via  $I_1, \dots, I_n$ , respectively.

**Note:** the roles or attributes  $I_j$  are not required to be functional or mandatory.

The above definition of semantics implies that, in the case where an instance  $o \in B^{\mathcal{I}}$  is connected by means of  $I_j^{\mathcal{I}}$  to a set  $u_j^1, \dots, u_j^k$  of objects (or values), it is each single  $u_j^h$  that contributes to the identification of  $o$ , and not the whole set  $\{u_j^1, \dots, u_j^k\}$ .

# Outline of Part 4

## 1 The *DL-Lite* family of tractable Description Logics

- Basic features of *DL-Lite*
- Syntax and semantics of *DL-Lite*
- Identification assertions in *DL-Lite*
- **Members of the *DL-Lite* family**
- Properties of *DL-Lite*

## 2 Linking ontologies to relational data

# Restriction on TBox assertions in $DL\text{-}Lite_{\mathcal{A}}$ ontologies

We will see that, to ensure the good computational properties that we aim at, we have to impose a **restriction** on the use of functionality and role/attribute inclusions.

## Restriction on $DL\text{-}Lite_{\mathcal{A}}$ TBoxes

**No functional or identifying role or attribute can be specialized** by using it in the right-hand side of a role or attribute inclusion assertion.

Formally:

- If  $(\mathbf{funct}\ P)$ ,  $(\mathbf{funct}\ P^-)$ ,  $(\mathbf{id}\ B\ \dots, P, \dots)$ , or  $(\mathbf{id}\ B\ \dots, P^-, \dots)$  is in  $\mathcal{T}$ , then  $Q \sqsubseteq P$  and  $Q \sqsubseteq P^-$  are **not in**  $\mathcal{T}$ .
- If  $(\mathbf{funct}\ U)$  or  $(\mathbf{id}\ B\ \dots, U, \dots)$  is in  $\mathcal{T}$ , then  $U' \sqsubseteq U$  is **not in**  $\mathcal{T}$ .

# $DL-Lite_{\mathcal{F}}$ and $DL-Lite_{\mathcal{R}}$

We consider also two sub-languages of  $DL-Lite_{\mathcal{A}}$  (that trivially obey the previous restriction):

- $DL-Lite_{\mathcal{F}}$ : Allows for functionality assertions, but does not allow for role inclusion assertions.
- $DL-Lite_{\mathcal{R}}$ : Allows for role inclusion assertions, but does not allow for functionality assertions.

In both  $DL-Lite_{\mathcal{F}}$  and  $DL-Lite_{\mathcal{R}}$  we do not consider data values (and hence drop value domains and attributes).

*Note:* We simply use  $DL-Lite$  to refer to any of the logics of the  $DL-Lite$  family.

# Outline of Part 4

## 1 The *DL-Lite* family of tractable Description Logics

- Basic features of *DL-Lite*
- Syntax and semantics of *DL-Lite*
- Identification assertions in *DL-Lite*
- Members of the *DL-Lite* family
- Properties of *DL-Lite*

## 2 Linking ontologies to relational data

# Capturing basic ontology constructs in $DL\text{-}Lite_{\mathcal{A}}$

ISA between classes	$A_1 \sqsubseteq A_2$
Disjointness between classes	$A_1 \sqsubseteq \neg A_2$
Mandatory participation to relations	$A_1 \sqsubseteq \exists P \quad A_2 \sqsubseteq \exists P^-$
Domain and range of relations	$\exists P \sqsubseteq A_1 \quad \exists P^- \sqsubseteq A_2$
Functionality of relations	$(\mathbf{funct} \ P) \quad (\mathbf{funct} \ P^-)$
ISA between relations	$Q_1 \sqsubseteq Q_2$
Disjointness between relations	$Q_1 \sqsubseteq \neg Q_2$
Domain and range of attributes	$\delta(U) \sqsubseteq A \quad \rho(U) \sqsubseteq T_i$
Mandatory and functional attributes	$A \sqsubseteq \delta(U) \quad (\mathbf{funct} \ U)$
Identification constraints	$(\mathbf{id} \ A \ P, \dots, P'^-, \dots, U, \dots)$



# Properties of *DL-Lite*

- The TBox may contain **cyclic dependencies** (which typically increase the computational complexity of reasoning).

Example:  $A \sqsubseteq \exists P$ ,  $\exists P^- \sqsubseteq A$

- In the syntax, we have not included  $\sqcap$  on the right hand-side of inclusion assertions, but it can trivially be added, since

$$B \sqsubseteq C_1 \sqcap C_2 \text{ is equivalent to } \begin{array}{l} B \sqsubseteq C_1 \\ B \sqsubseteq C_2 \end{array}$$

- A domain assertion on role  $P$  has the form:  $\exists P \sqsubseteq A_1$   
A range assertion on role  $P$  has the form:  $\exists P^- \sqsubseteq A_2$

# Properties of $DL\text{-}Lite_{\mathcal{F}}$

$DL\text{-}Lite_{\mathcal{F}}$  does **not** enjoy the **finite model property**.

## Example

TBox  $\mathcal{T}$ :  $Nat \sqsubseteq \exists succ$                        $\exists succ^- \sqsubseteq Nat$   
                $Zero \sqsubseteq Nat \sqcap \neg \exists succ^-$             (**funct**  $succ^-$ )

ABox  $\mathcal{A}$ :  $Zero(0)$

$\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$  admits only infinite models.

Hence, it is satisfiable, but **not finitely satisfiable**.

Hence, reasoning w.r.t. arbitrary models is different from reasoning w.r.t. finite models only.

# Properties of $DL\text{-}Lite_{\mathcal{R}}$

- $DL\text{-}Lite_{\mathcal{R}}$  **does enjoy the finite model property**. Hence, reasoning w.r.t. finite models is the same as reasoning w.r.t. arbitrary models.
- With role inclusion assertions, we can simulate **qualified existential quantification** in the rhs of an inclusion assertion  $A_1 \sqsubseteq \exists Q.A_2$ .

To do so, we introduce a new role  $Q_{A_2}$  and:

- the role inclusion assertion  $Q_{A_2} \sqsubseteq Q$
- the concept inclusion assertions:
 
$$\begin{array}{lcl} A_1 & \sqsubseteq & \exists Q_{A_2} \\ \exists Q_{A_2}^- & \sqsubseteq & A_2 \end{array}$$

In this way, we can consider  $\exists Q.A$  in the right-hand side of an inclusion assertion as an abbreviation.

# Observations on $DL-Lite_{\mathcal{A}}$

- Captures all the basic constructs of **UML Class Diagrams** and of the **ER Model** ...
- ... **except covering constraints** in generalizations.
- Extends (the DL fragment of) the ontology language **RDFS**.
- Is completely symmetric w.r.t. **direct and inverse properties**.
- Is at the basis of the **OWL 2 QL** profile of OWL 2.

# The OWL 2 QL Profile

OWL 2 defines three **profiles**: OWL 2 QL, OWL 2 EL, OWL 2 RL.

- Each profile corresponds to a syntactic fragment (i.e., a sub-language) of OWL 2 DL that is targeted towards a specific use.
- The restrictions in each profile guarantee better computational properties than those of OWL 2 DL.

The **OWL 2 QL** profile is derived from the DLs of the *DL-Lite* family:

- “[It] includes most of the main features of conceptual models such as UML class diagrams and ER diagrams.”
- “[It] is aimed at applications that use very large volumes of instance data, and where query answering is the most important reasoning task. In OWL 2 QL, conjunctive query answering can be implemented using conventional relational database systems.”

# Complexity of reasoning in $DL\text{-}Lite_{\mathcal{A}}$

- ① We have seen that  $DL\text{-}Lite_{\mathcal{A}}$  can capture the essential features of prominent conceptual modeling formalisms.
- ② In the following, we will analyze reasoning in  $DL\text{-}Lite$ , and establish the following characterization of its computational properties:
  - **Ontology satisfiability** and all classical DL reasoning tasks are:
    - Efficiently tractable in the size of the **TBox** (i.e., **P**TIME).
    - Very efficiently tractable in the size of the **ABox** (i.e., **AC**<sup>0</sup>).
  - **Query answering** for CQs and UCQs is:
    - **P**TIME in the size of the **TBox**.
    - **AC**<sup>0</sup> in the size of the **ABox**.
    - Exponential in the size of the **query** (**NP-complete**).Bad? ... not really, this is exactly as in relational DBs.
- ③ We will also see that  $DL\text{-}Lite$  is essentially the maximal DL enjoying these nice computational properties.

From (1), (2), and (3) we get that:

$DL\text{-}Lite$  is a representation formalism that is very well suited to underlie ontology-based data management systems.

# Outline of Part 4

- 1 The *DL-Lite* family of tractable Description Logics
- 2 Linking ontologies to relational data
  - The impedance mismatch problem
  - Ontology-Based Data Access systems
  - Query answering in Ontology-Based Data Access
  - The ONTOP framework for Ontology-Based Data Access

# Outline of Part 4

- 1 The *DL-Lite* family of tractable Description Logics
- 2 Linking ontologies to relational data
  - The impedance mismatch problem
    - Ontology-Based Data Access systems
    - Query answering in Ontology-Based Data Access
    - The ONTOP framework for Ontology-Based Data Access



# Managing ABoxes

In the traditional DL setting, it is assumed that the data is maintained in the **ABox** of the ontology:

- The ABox is perfectly compatible with the TBox:
  - the vocabulary of concepts, roles, and attributes is the one used in the TBox.
  - The ABox “stores” abstract objects, and these objects and their properties are those returned by queries over the ontology.
- There may be different ways to manage the ABox from a physical point of view:
  - Description Logics reasoners maintain the ABox as main-memory data structures.
  - When an ABox becomes large, managing it in secondary storage may be required, but this is again handled directly by the reasoner.

# Data in external sources

There are several situations where the assumptions of having the data in an ABox managed directly by the ontology system (e.g., a Description Logics reasoner) is not feasible or realistic:

- When the ABox is very large, so that it requires relational database technology.
- When we have no direct control over the data since it belongs to some external organization, which controls the access to it.
- When multiple data sources need to be accessed, such as in Information Integration.


We would like to deal with such a situation by keeping the data in the external (relational) storage, and performing **query answering** by leveraging the capabilities of the **relational engine**.

# Ontology-based data access: Architecture

[Poggi et al. 2008]

The architecture of an OBDA system is based on three main components:

- **Ontology**: provides a unified, conceptual view of the managed information.
- **Data source(s)**: are external and independent (possibly multiple and heterogeneous).
- **Mappings**: semantically link data at the sources with the ontology.



# The impedance mismatch problem

We have to deal with the **impedance mismatch problem**:

- Sources store data, which is constituted by values taken from concrete domains, such as strings, integers, codes, ...
- Instead, instances of concepts and relations in an ontology are (abstract) objects.

## Solution:

- We need to specify how to construct from the data values in the relational sources the (abstract) objects that populate the ABox of the ontology.
- This specification is embedded in the mappings between the data sources and the ontology.


*Note:* the **ABox** is only **virtual**, and the objects are not materialized.

# Solution to the impedance mismatch problem

We need to define a **mapping language** that allows for specifying how to transform data into abstract objects:

- Each mapping assertion maps:
  - a query that retrieves values from a data source to ...
  - a set of atoms specified over the ontology.
- Basic idea: use **Skolem functions** in the atoms over the ontology to “generate” the objects from the data values.
- Semantics of mappings:
  - Objects are denoted by terms (of exactly one level of nesting).
  - Different terms denote different objects (i.e., we make the unique name assumption on terms).

# Impedance mismatch – Example



Actual data is stored in a DB:

- An employee is identified by her SSN.
- A project is identified by its name.

$D_1[SSN: String, PrName: String]$

Employees and projects they work for

$D_2[Code: String, Salary: Int]$

Employee's code with salary

$D_3[Code: String, SSN: String]$

Employee's Code with SSN

...

Intuitively:

- An employee should be created from her SSN: **pers**(SSN)
- A project should be created from its name: **proj**(PrName)

# Creating object identifiers

We need to associate to the data in the tables objects in the ontology.

- We introduce an alphabet  $\Lambda$  of **function symbols**, each with an associated arity.
- To denote values, we use value constants from an alphabet  $\Gamma_V$ .
- To denote objects, we use **object terms** instead of object constants.  
An object term has the form  $\mathbf{f}(d_1, \dots, d_n)$ , with  $\mathbf{f} \in \Lambda$ , and each  $d_i$  a value constant in  $\Gamma_V$ .

## Example

- If a person is identified by her *SSN*, we can introduce a function symbol **pers/1**. If *VRD56B25* is a *SSN*, then **pers**(*VRD56B25*) denotes a person.
- If a person is identified by her *name* and *dateOfBirth*, we can introduce a function symbol **pers/2**. Then **pers**(*Vardi*, *25/2/56*) denotes a person.

# Mapping assertions

Mapping assertions are used to extract the data from the DB to populate the ontology.

We make use of **variable terms**, which are like object terms, but with variables instead of values as arguments of the functions.

Def.: A **mapping assertion** between a database  $\mathcal{D}$  and a TBox  $\mathcal{T}$  has the form


$$\Phi(\vec{x}) \rightsquigarrow \Psi(\vec{t}, \vec{y})$$

where

- $\Phi$  is an arbitrary SQL query of arity  $n > 0$  over  $\mathcal{D}$ ;
- $\Psi$  is a conjunction of atoms whose predicates are atomic concepts and roles of  $\mathcal{T}$ ;
- $\vec{x}, \vec{y}$  are variables, with  $\vec{y} \subseteq \vec{x}$ ;
- $\vec{t}$  are variable terms of the form  $\mathbf{f}(\vec{z})$ , with  $\mathbf{f} \in \Lambda$  and  $\vec{z} \subseteq \vec{x}$ .



# Mapping assertions – Example



$D_1[SSN: String, PrName: String]$

Employees and Projects they work for

$D_2[SSN: String, Code: String]$

Employee's SSN with code

$D_3[Code: String, Salary: Int]$

Employee's code with salary

...

$m_1$ :	SELECT SSN, PrName FROM $D_1$	$\rightsquigarrow$	Employee( <b>pers</b> (SSN)), Project( <b>proj</b> (PrName)), projectName( <b>proj</b> (PrName), PrName), worksFor( <b>pers</b> (SSN), <b>proj</b> (PrName))
$m_2$ :	SELECT SSN, Code FROM $D_2$	$\rightsquigarrow$	Employee( <b>pers</b> (SSN)), empCode( <b>pers</b> (SSN), Code)
$m_3$ :	SELECT SSN, Salary FROM $D_2$ , $D_3$ WHERE $D_2.Code = D_3.Code$	$\rightsquigarrow$	Employee( <b>pers</b> (SSN)), salary( <b>pers</b> (SSN), Salary)

# Concrete mapping languages

Several proposals for concrete languages to map a relational DB to an ontology:

- They assume that the ontology is populated in terms of RDF triples.
- Some template mechanism is used to specify the triples to instantiate.

Examples: D2RQ<sup>1</sup>, SML<sup>2</sup>, Ontop<sup>3</sup>

## R2RML

- Most popular RDB to RDF mapping language
- W3C Recommendation 27 Sep. 2012, <http://www.w3.org/TR/r2rml/>
- R2RML mappings are themselves expressed as RDF graphs and written in Turtle syntax.

---

<sup>1</sup><http://d2rq.org/d2rq-language>


<sup>2</sup>[http://sparqlify.org/wiki/Sparqlification\\_mapping\\_language](http://sparqlify.org/wiki/Sparqlification_mapping_language)

<sup>3</sup><https://github.com/ontop/ontop/wiki/ontopOBDAModel>

# Outline of Part 4

- 1 The *DL-Lite* family of tractable Description Logics
- 2 Linking ontologies to relational data
  - The impedance mismatch problem
  - **Ontology-Based Data Access systems**
  - Query answering in Ontology-Based Data Access
  - The ONTOP framework for Ontology-Based Data Access

# Ontology-based data access: Formalization



To formalize OBDA, we distinguish between the intensional and the extensional level information.

Def.: An **OBDA specification** is a triple  $\mathcal{P} = \langle \mathcal{T}, \mathcal{M}, \mathcal{S} \rangle$ , where:


- $\mathcal{T}$  is a DL TBox providing the intensional level of an ontology.
- $\mathcal{S}$  is a (possibly federated) **relational database schema** for the data sources, possibly with constraints;
- $\mathcal{M}$  is a set of **mapping assertions** between  $\mathcal{T}$  and  $\mathcal{S}$ .

Def.: An **OBDA instance** is a pair  $\mathcal{O} = \langle \mathcal{P}, \mathcal{D} \rangle$ , where

- $\mathcal{P} = \langle \mathcal{T}, \mathcal{M}, \mathcal{S} \rangle$  is an OBDA specification, and
- $\mathcal{D}$  is a relational database compliant with  $\mathcal{S}$ .

# Semantics of an OBDA instance: Intuition

Given an OBDA instance, the **mapping**  $\mathcal{M}$  encodes how the data  $\mathcal{D}$  in the source(s)  $\mathcal{S}$  should be used to populate the elements of the TBox  $\mathcal{T}$ .



The data  $\mathcal{D}$  and the mapping  $\mathcal{M}$  define a **virtual data layer**  $\mathcal{V}$ , which behaves like a (virtual) ABox.

- Queries are answered w.r.t.  $\mathcal{T}$  and  $\mathcal{V}$ .
- One aim is to avoid materializing the data of  $\mathcal{V}$ .
- Instead, the intensional information in  $\mathcal{T}$  and  $\mathcal{M}$  is used to translate queries over  $\mathcal{T}$  into queries formulated over  $\mathcal{S}$ .

## OBDA vs. Ontology Based Query Answering (OBQA)

OBDA relies on OBQA to process queries w.r.t. the TBox  $\mathcal{T}$ , but in addition is concerned with efficiently dealing with the mapping  $\mathcal{M}$ .

**OBDA should not be confused with OBQA.**

# Semantics of mappings

To define the semantics of an OBDA instance  $\mathcal{O} = \langle \mathcal{P}, \mathcal{D} \rangle$ , with  $\mathcal{P} = \langle \mathcal{T}, \mathcal{M}, \mathcal{S} \rangle$ , we first need to define the semantics of mappings.

**Def.:** Satisfaction of a mapping assertion with respect to a database

An interpretation  $\mathcal{I}$  **satisfies** a mapping assertion  $\Phi(\vec{x}) \rightsquigarrow \Psi(\vec{t}, \vec{y})$  in  $\mathcal{M}$  **with respect to a database**  $\mathcal{D}$ , if for each tuple of values  $\vec{v} \in \text{Eval}(\Phi, \mathcal{D})$ , and for each ground atom in  $\Psi[\vec{x}/\vec{v}]$ , we have that:


- if the ground atom is  $A(s)$ , then  $s^{\mathcal{I}} \in A^{\mathcal{I}}$ .
- if the ground atom is  $P(s_1, s_2)$ , then  $(s_1^{\mathcal{I}}, s_2^{\mathcal{I}}) \in P^{\mathcal{I}}$ .

Intuitively,  $\mathcal{I}$  **satisfies**  $\Phi \rightsquigarrow \Psi$  w.r.t.  $\mathcal{D}$  if all facts obtained by evaluating  $\Phi$  over  $\mathcal{D}$  and then propagating the answers to  $\Psi$ , hold in  $\mathcal{I}$ .

**Note:**  $\text{Eval}(\Phi, \mathcal{D})$  denotes the result of evaluating  $\Phi$  over the database  $\mathcal{D}$ .

$\Psi[\vec{x}/\vec{v}]$  denotes  $\Psi$  where each  $x_i$  has been substituted with  $v_i$ .

## Semantics of mappings – Example



$D_1$ :	<i>SSN</i>	<i>PrName</i>
	23AB	optique
	...	...

$D_2$ :	<i>SSN</i>	<i>Code</i>
	23AB	e23
	...	...

$D_3$ :	<i>Code</i>	<i>Salary</i>
	e23	15000
	...	...

The following interpretation  $\mathcal{I}$  satisfies the mapping assertions  $m_1$  and  $m_3$  (we ignore  $m_2$ ) with respect to the above database:

Note that we have directly used object terms as domain elements.

 $\mathcal{I}:$ 

$$\Delta_O^{\mathcal{I}} = \{\mathbf{pers}(23\text{AB}), \mathbf{proj}(\text{optique}), \dots\}, \quad \Delta_V^{\mathcal{I}} = \{\text{optique}, 15000, \dots\}$$

$$Employee^I = \{\mathbf{pers}(23AB), \dots\}, \quad Project^I = \{\mathbf{proj}(\text{optique}), \dots\},$$

$$projectName^I = \{(\mathbf{proj}(\text{optique}), \text{optique}), \dots\},$$

$$worksFor^I = \{(\mathbf{pers}(23AB), \mathbf{proj}(\text{optique})), \dots\},$$

$$salary^I = \{(\mathbf{pers}(23AB), 15000), \dots\}$$

$m_1$ : `SELECT SSN, PrName`  $\rightsquigarrow$  `Employee(pers(SSN)),`  
`FROM D1` `Project(proj(PrName)),`  
`projectName(proj(PrName), PrName),`  
`worksFor(pers(SSN), proj(PrName))`


```
m3:  SELECT SSN, Salary      ~→ Employee(pers(SSN)),
      FROM D2, D3           salary(pers(SSN), Salary)
      WHERE D2.Code = D3.Code
```

# Semantics of an OBDA instance

## Model of an OBDA instance

An interpretation  $\mathcal{I}$  is a **model** of  $\mathcal{O} = \langle \mathcal{P}, \mathcal{D} \rangle$ , with  $\mathcal{P} = \langle \mathcal{T}, \mathcal{M}, \mathcal{S} \rangle$ , if:

- $\mathcal{I}$  is a model of  $\mathcal{T}$ , and
- $\mathcal{I}$  satisfies  $\mathcal{M}$  w.r.t.  $\mathcal{D}$ , i.e.,  
 $\mathcal{I}$  satisfies every assertion in  $\mathcal{M}$  w.r.t.  $\mathcal{D}$ .



An OBDA instance  $\mathcal{O}$  is **satisfiable** if it admits at least one model.



# Outline of Part 4

- 1 The *DL-Lite* family of tractable Description Logics
- 2 Linking ontologies to relational data
  - The impedance mismatch problem
  - Ontology-Based Data Access systems
  - Query answering in Ontology-Based Data Access
  - The ONTOP framework for Ontology-Based Data Access

# Answering queries over an OBDA instance

Given an OBDA instance  $\mathcal{O} = \langle \mathcal{P}, \mathcal{D} \rangle$ , with  $\mathcal{P} = \langle \mathcal{T}, \mathcal{M}, \mathcal{S} \rangle$ :

- Queries are posed over the TBox  $\mathcal{T}$ .
- The data needed to answer queries is stored in the database  $\mathcal{D}$ , which is compliant to  $\mathcal{S}$ .
- The mapping  $\mathcal{M}$  is used to bridge the gap between  $\mathcal{T}$  and  $\mathcal{D}$ .

Two approaches to exploit the mapping:

- bottom-up approach: simpler, but typically less efficient
- top-down approach: more sophisticated, but also more efficient

*Note:* Both approaches require to first **split** the TBox queries in the mapping assertions into their constituent atoms.

# Splitting of mappings

A mapping assertion  $\Phi \rightsquigarrow \Psi$ , where the TBox query  $\Psi$  is constituted by the atoms  $X_1, \dots, X_k$ , can be split into several mapping assertions:

$$\Phi \rightsquigarrow X_1 \quad \dots \quad \Phi \rightsquigarrow X_k$$

This is possible, since  $\Psi$  does not contain non-distinguished variables.

## Example

$m_1$ : `SELECT SSN, PrName FROM D1`  $\rightsquigarrow$  `Employee(pers(SSN)),`  
`Project(proj(PrName)),`  
`projectName(proj(PrName), PrName),`  
`worksFor(pers(SSN), proj(PrName))`

is split into

$m_1^1$ : `SELECT SSN, PrName FROM D1`  $\rightsquigarrow$  `Employee(pers(SSN))`  
 $m_1^2$ : `SELECT SSN, PrName FROM D1`  $\rightsquigarrow$  `Project(proj(PrName))`  
 $m_1^3$ : `SELECT SSN, PrName FROM D1`  $\rightsquigarrow$  `projectName(proj(PrName), PrName)`  
 $m_1^4$ : `SELECT SSN, PrName FROM D1`  $\rightsquigarrow$  `worksFor(pers(SSN), proj(PrName))`

# Bottom-up approach to query answering

Consists in a straightforward application of the mappings:

- 1 Propagate the data from  $\mathcal{D}$  through  $\mathcal{M}$ , materializing an ABox  $\mathcal{A}_{\mathcal{M},\mathcal{D}}$  (the constants in such an ABox are values and object terms).
- 2 Apply to  $\mathcal{A}_{\mathcal{M},\mathcal{D}}$  and to the TBox  $\mathcal{T}$ , the satisfiability and query answering algorithms developed for  $DL-Lite_{\mathcal{A}}$ .

This approach has several drawbacks:

- The technique is no more  $AC^0$  in the data, since the ABox  $\mathcal{A}_{\mathcal{M},\mathcal{D}}$  to materialize is in general polynomial in the size of the data.
- $\mathcal{A}_{\mathcal{M},\mathcal{D}}$  may be very large, and thus it may be infeasible to actually materialize it.
- Freshness of  $\mathcal{A}_{\mathcal{M},\mathcal{D}}$  with respect to the underlying data source(s) may be an issue, and one would need to propagate source updates (cf. Data Warehousing).

# Top-down approach to query answering

Consists of three steps:


- ① **Reformulation:** Compute the perfect rewriting (or reformulation)  $q_{pr} = \text{PerfectRef}(q, \mathcal{T})$  of the original query  $q$ , using the inclusion assertions of the TBox  $\mathcal{T}$  (see later).
  - The perfect rewriting  $q_{pr}$  is such that  $\text{cert}(q, \langle \mathcal{T}, \mathcal{A} \rangle) = \text{Eval}_{\text{CWA}}(q_{pr}, \mathcal{A})$ , for each ABox  $\mathcal{A}$ .
- ② **Unfolding:** Compute from  $q_{pr}$  a new query  $q_{unf}$  by unfolding  $q_{pr}$  using (the split version of) the mappings  $\mathcal{M}$ .
  - Essentially, each atom in  $q_{pr}$  that unifies with an atom in  $\Psi$  is substituted with the corresponding query  $\Phi$  over the database.
  - The unfolded query  $q_{unf}$  is such that  $\text{Eval}(q_{unf}, \mathcal{D}) = \text{Eval}_{\text{CWA}}(q_{pr}, \mathcal{A}_{\mathcal{M}, \mathcal{D}})$ , for each database  $\mathcal{D}$ .
- ③ **Evaluation:** Delegate the evaluation of  $q_{unf}$  to the relational DBMS managing  $\mathcal{D}$ .

# Unfolding

To unfold a query  $q_{pr}$  with respect to a set of mapping assertions:

- ① For each non-split mapping assertion  $\Phi_i(\vec{x}) \rightsquigarrow \Psi_i(\vec{t}, \vec{y})$ :
  - ① Introduce a **view symbol**  $Aux_i$  of arity equal to that of  $\Phi_i$ .
  - ② Add a **view definition**  $Aux_i(\vec{x}) \leftarrow \Phi_i(\vec{x})$ .
- ② For each split version  $\Phi_i(\vec{x}) \rightsquigarrow X_j(\vec{t}, \vec{y})$  of a mapping assertion, introduce a **clause**  $X_j(\vec{t}, \vec{y}) \leftarrow Aux_i(\vec{x})$ .
- ③ Obtain from  $q_{pr}$  in all possible ways queries  $q_{aux}$  defined over the view symbols  $Aux_i$  as follows:
  - ① Find a most general unifier  $\vartheta$  that unifies each atom  $X(\vec{z})$  in the body of  $q_{pr}$  with the head of a clause  $X(\vec{t}, \vec{y}) \leftarrow Aux_i(\vec{x})$ .
  - ② Substitute each atom  $X(\vec{z})$  with  $\vartheta(Aux_i(\vec{x}))$ , i.e., with the body the unified clause to which the unifier  $\vartheta$  is applied.
- ④ The unfolded query  $q_{unf}$  is the **union** of all queries  $q_{aux}$ , together with the view definitions for the predicates  $Aux_i$  appearing in  $q_{aux}$ .

# Unfolding – Example



$m_1$ : SELECT SSN, PrName  
FROM  $D_1$

$\rightsquigarrow$  Employee(**pers**(SSN)),  
Project(**proj**(PrName)),  
projectName(**proj**(PrName), PrName),  
worksFor(**pers**(SSN), **proj**(PrName))

$m_2$ : SELECT SSN, Salary  
FROM  $D_2, D_3$   
WHERE  $D_2.Code = D_3.Code$

$\rightsquigarrow$  Employee(**pers**(SSN)),  
salary(**pers**(SSN), Salary)

We define a view  $Aux_i$  for the source query of each mapping  $m_i$ .

For each (split) mapping assertion, we introduce a clause:

Employee(**pers**(SSN))  $\leftarrow$   $Aux_1(SSN, PrName)$   
 projectName(**proj**(PrName), PrName)  $\leftarrow$   $Aux_1(SSN, PrName)$   
 Project(**proj**(PrName))  $\leftarrow$   $Aux_1(SSN, PrName)$   
 worksFor(**pers**(SSN), **proj**(PrName))  $\leftarrow$   $Aux_1(SSN, PrName)$   
 Employee(**pers**(SSN))  $\leftarrow$   $Aux_2(SSN, Salary)$   
 salary(**pers**(SSN), Salary)  $\leftarrow$   $Aux_2(SSN, Salary)$

# Unfolding – Example (cont'd)

Query over ontology: employees who work for optique and their salary:

$q(e, s) \leftarrow \text{Employee}(e), \text{salary}(e, s), \text{worksFor}(e, p), \text{projectName}(p, \text{optique})$

A unifier  $\vartheta$  between the atoms in  $q$  and the clause heads is:

$\vartheta(e) = \text{pers}(\text{SSN})$

$\vartheta(s) = \text{Salary}$

$\vartheta(\text{PrName}) = \text{optique}$

$\vartheta(p) = \text{proj}(\text{optique})$

After applying  $\vartheta$  to  $q$ , we obtain:

$q(\text{pers}(\text{SSN}), \text{Salary}) \leftarrow \text{Employee}(\text{pers}(\text{SSN})), \text{salary}(\text{pers}(\text{SSN}), \text{Salary}),$   
 $\text{worksFor}(\text{pers}(\text{SSN}), \text{proj}(\text{optique})),$   
 $\text{projectName}(\text{proj}(\text{optique}), \text{optique})$

Substituting the atoms with the bodies of the unified clauses, we obtain:

$q(\text{pers}(\text{SSN}), \text{Salary}) \leftarrow \text{Aux}_1(\text{SSN}, \text{optique}), \text{Aux}_2(\text{SSN}, \text{Salary}),$   
 $\text{Aux}_1(\text{SSN}, \text{optique}), \text{Aux}_1(\text{SSN}, \text{optique})$



# Exponential blowup in the unfolding

When there are multiple mapping assertions for each atom, the unfolded query may be exponential in the original one.

Consider a query:  $q(y) \leftarrow A_1(y), A_2(y), \dots, A_n(y)$

and the mappings:  $m_i^1: \Phi_i^1(x) \rightsquigarrow A_i(\mathbf{f}(x))$  (for  $i \in \{1, \dots, n\}$ )  
 $m_i^2: \Phi_i^2(x) \rightsquigarrow A_i(\mathbf{f}(x))$

We add the view definitions:  $Aux_i^j(x) \leftarrow \Phi_i^j(x)$

and introduce the clauses:  $A_i(\mathbf{f}(x)) \leftarrow Aux_i^j(x)$  (for  $i \in \{1, \dots, n\}, j \in \{1, 2\}$ ).

There is a single unifier, namely  $\vartheta(y) = \mathbf{f}(x)$ , but each atom  $A_i(y)$  in the query unifies with the head of two clauses.

Hence, we obtain **one unfolded query**

$$q(\mathbf{f}(x)) \leftarrow Aux_1^{j_1}(x), Aux_2^{j_2}(x), \dots, Aux_n^{j_n}(x)$$

for each possible combination of  $j_i \in \{1, 2\}$ , for  $i \in \{1, \dots, n\}$ .

Hence, we obtain  **$2^n$  unfolded queries**.

# Computational complexity of query answering

From the top-down approach to query answering, and the complexity results for *DL-Lite*, we obtain the following result.

## Theorem

In a *DL-Lite* OBDA instance  $\mathcal{O} = \langle \mathcal{P}, \mathcal{D} \rangle$ , with  $\mathcal{P} = \langle \mathcal{T}, \mathcal{M}, \mathcal{S} \rangle$ , **query answering** is

- ① **NP-complete** in the size of the query.
- ② **P<sub>TIME</sub>** in the size of the **TBox**  $\mathcal{T}$  and the **mappings**  $\mathcal{M}$ .
- ③ **AC<sup>0</sup>** in the size of the **database**  $\mathcal{D}$ .

*Note:* The **AC<sup>0</sup>** result is a consequence of the fact that query answering in such a setting can be reduced to evaluating an SQL query over the relational database  $\mathcal{D}$ .

# Implementation of top-down approach to query answering

To implement the top-down approach, we need to generate an SQL query.

We can follow different strategies:

- ① Substitute each view predicate in the unfolded queries with the corresponding SQL query over the source:
  - + joins are performed on the DB attributes;
  - + does not generate doubly nested queries;
  - the number of unfolded queries may be exponential.
- ② Construct for each atom in the original query a new view. This view takes the union of all SQL queries corresponding to the view predicates, and constructs also the Skolem terms:
  - + avoids exponential blow-up of the resulting query, since the union (of the queries coming from multiple mappings) is done before the joins;
  - joins are performed on Skolem terms;
  - generates doubly nested queries.

Which method is better, depends on various parameters.

Experiments have shown that (1) behaves better in most cases.

# Towards answering arbitrary SQL queries

- We have seen that answering full SQL (i.e., FOL) queries is undecidable.
- However, we can treat the answers to an UCQ, as “knowledge”, and perform further computations on that knowledge.
- This corresponds to applying a knowledge operator to UCQs that are embedded into an arbitrary SQL query (EQL queries) [Calvanese, De Giacomo, Lembo, Lenzerini, and Rosati 2007b]
  - The UCQs are answered according to the certain answer semantics.
  - The SQL query is evaluated on the facts returned by the UCQs.
- The approach can be implemented by rewriting the UCQs and embedding the rewritten UCQs into SQL.
- The user “sees” arbitrary SQL queries, but these SQL queries are evaluated according to a weakened semantics.

# Outline of Part 4

- 1 The *DL-Lite* family of tractable Description Logics
- 2 Linking ontologies to relational data
  - The impedance mismatch problem
  - Ontology-Based Data Access systems
  - Query answering in Ontology-Based Data Access
  - The ONTOP framework for Ontology-Based Data Access

# The ONTOP framework

- ONTOP is a framework providing advanced functionalities for representing and reasoning over ontologies of the *DL-Lite* family.
- The basic functionality it offers is query answering of UCQs expressed in SPARQL syntax.
- Query answering is also at the basis of
  - ontology satisfiability;
  - intensional reasoning services: concept/role subsumption and disjunction, concept/role satisfiability.
- Reasoning services are highly optimized.
- Can be used with internal and external DBMS (includes drivers for various commercial and non-commercial DBMSs).
- Implemented in Java as an open source project under the Apache 2 licence.

# References I

- [1] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. “Tractable Reasoning and Efficient Query Answering in Description Logics: The *DL-Lite* Family”. In: *J. of Automated Reasoning* 39.3 (2007), pp. 385–429.
- [2] Alessandro Artale, Diego Calvanese, Roman Kontchakov, and Michael Zakharyashev. “The *DL-Lite* Family and Relations”. In: *J. of Artificial Intelligence Research* 36 (2009), pp. 1–69.
- [3] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, Antonella Poggi, Mariano Rodriguez-Muro, and Riccardo Rosati. “Ontologies and Databases: The *DL-Lite* Approach”. In: *Reasoning Web. Semantic Technologies for Informations Systems – 5th Int. Summer School Tutorial Lectures (RW)*. Ed. by Sergio Tessaris and Enrico Franconi. Vol. 5689. Lecture Notes in Computer Science. Springer, 2009, pp. 255–356.

# References II

- [4] Antonella Poggi, Domenico Lembo, Diego Calvanese, Giuseppe De Giacomo, Maurizio Lenzerini, and Riccardo Rosati. “Linking Data to Ontologies”. In: *J. on Data Semantics X* (2008), pp. 133–173. DOI: 10.1007/978-3-540-77688-8\_5.
- [5] Diego Calvanese, Giuseppe De Giacomo, Domenico Lembo, Maurizio Lenzerini, and Riccardo Rosati. “EQL-Lite: Effective First-Order Query Processing in Description Logics”. In: *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI)*. 2007, pp. 274–279.