

# Ontology and Database Systems: Ontology-based Systems

## Part 4: Ontology Based Data Access

*Diego Calvanese*

Faculty of Computer Science  
Master of Science in Computer Science

A.Y. 2013/2014



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 systems
  - 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

- 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.

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

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

- Identification assertions:  $(\mathbf{id} \ B \ I_1, \dots, I_n)$

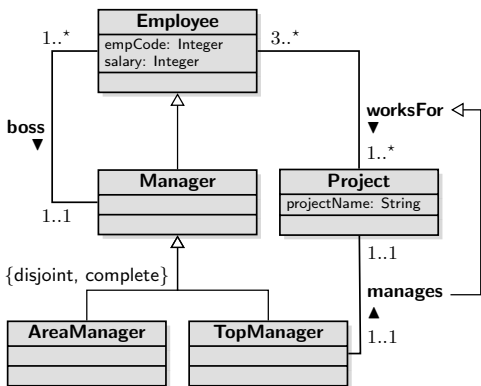
where each  $I_j$  is a role, an inverse role, or an attribute

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 $\sqsubseteq \exists\text{child}$	$B_1^{\mathcal{I}} \subseteq B_2^{\mathcal{I}}$
role incl.	$Q_1 \sqsubseteq Q_2$	father $\sqsubseteq \text{anc}$	$Q_1^{\mathcal{I}} \subseteq Q_2^{\mathcal{I}}$
v.dom. incl.	$\rho(U) \sqsubseteq T_i$	$\rho(\text{age}) \sqsubseteq \text{xsd:int}$	$\rho(U)^{\mathcal{I}} \subseteq T_i^{\mathcal{I}}$
attr. incl.	$U_1 \sqsubseteq U_2$	offPhone $\sqsubseteq \text{phone}$	$U_1^{\mathcal{I}} \subseteq U_2^{\mathcal{I}}$
conc. disj.	$B_1 \sqsubseteq \neg B_2$	Person $\sqsubseteq \neg\text{Course}$	$B_1^{\mathcal{I}} \subseteq (\neg B_2)^{\mathcal{I}}$
role disj.	$Q_1 \sqsubseteq \neg Q_2$	sibling $\sqsubseteq \neg\text{cousin}$	$Q_1^{\mathcal{I}} \subseteq (\neg Q_2)^{\mathcal{I}}$
attr. disj.	$U_1 \sqsubseteq \neg U_2$	offPhn $\sqsubseteq \neg\text{homePhn}$	$U_1^{\mathcal{I}} \subseteq (\neg U_2)^{\mathcal{I}}$
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')	$(c^{\mathcal{I}}, \text{val}(d)) \in U^{\mathcal{I}}$

# DL-Lite<sub>A</sub> – Example




- Manager  $\sqsubseteq$  Employee
- AreaManager  $\sqsubseteq$  Manager
- TopManager  $\sqsubseteq$  Manager
- AreaManager  $\sqsubseteq$   $\neg$ TopManager
- Employee  $\sqsubseteq$   $\delta(\text{empCode})$
- $\delta(\text{empCode}) \sqsubseteq$  Employee
- $\rho(\text{empCode}) \sqsubseteq$  xsd:int
- (**func** empCode)
- (**id** Employee empCode)
- $\exists \text{worksFor} \sqsubseteq$  Employee
- $\exists \text{worksFor}^- \sqsubseteq$  Project
- Employee  $\sqsubseteq$   $\exists \text{worksFor}$
- Project  $\sqsubseteq$   $\exists \text{worksFor}^-$
- (**func** manages)
- (**func** manages<sup>-</sup>)
- manages  $\sqsubseteq$  worksFor
- ⋮

*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)

$\text{League} \sqsubseteq \exists \text{of}$   
 $\exists \text{of} \sqsubseteq \text{League}$   
 $\exists \text{of}^- \sqsubseteq \text{Nation}$   
 $\text{Round} \sqsubseteq \exists \text{belongsTo}$   
 $\exists \text{belongsTo} \sqsubseteq \text{Round}$   
 $\exists \text{belongsTo}^- \sqsubseteq \text{League}$   
 $\text{Match} \sqsubseteq \exists \text{playedIn}$   
 ...


$\text{PlayedMatch} \sqsubseteq \text{Match}$   
 $\text{Match} \sqsubseteq \delta(\text{code})$   
 $\text{Round} \sqsubseteq \delta(\text{code})$   
 $\text{PlayedMatch} \sqsubseteq \delta(\text{playedOn})$   
 ...  
 $\rho(\text{playedOn}) \sqsubseteq \text{xsd:date}$   
 $\rho(\text{code}) \sqsubseteq \text{xsd:int}$   
 ...

**(funct of)**  
**(funct belongsTo)**  
**(funct playedIn)**  
**(funct homeTeam)**

**(funct hostTeam)**  
**(funct umpiredBy)**  
**(funct code)**  
**(funct year)**

**(funct homeGoals)**  
**(funct hostGoals)**  
**(funct playedOn)**

# 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)	(id Match umpiredBy, playedIn)
(id Round belongsTo, code)	(id Match homeTeam, playedIn)
(id Match playedIn, code)	(id Match hostTeam, playedIn)

# Semantics of identification assertions

Let  $(\mathbf{id} \ B \ I_1, \dots, I_n)$  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 **(*funct*  $P$ )**, **(*funct*  $P^-$ )**, **(*id*  $B \dots, P, \dots$ )**, or **(*id*  $B \dots, P^-, \dots$ )** is in  $\mathcal{T}$ , then  $Q \sqsubseteq P$  and  $Q \sqsubseteq P^-$  are **not in  $\mathcal{T}$** .
- If **(*funct*  $U$ )** or **(*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

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

- 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  $\mathcal{T}$ :  $\text{Nat} \sqsubseteq \exists \text{succ}$                        $\exists \text{succ}^- \sqsubseteq \text{Nat}$   
                    $\text{Zero} \sqsubseteq \text{Nat} \sqcap \neg \exists \text{succ}^-$         (**funct succ**<sup>-</sup>)

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

$\mathcal{O} = \langle \mathcal{T}, \mathcal{A} \rangle$  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:
 
$$\begin{array}{l} A_1 \sqsubseteq \exists Q_{A_2} \\ \exists Q_{A_2}^- \sqsubseteq A_2 \end{array}$$

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<sub>A</sub>*

- 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 $DL-Lite_{\mathcal{A}}$

- ① We have seen that  $DL-Lite_{\mathcal{A}}$  can capture the essential features of prominent conceptual modeling formalisms.
- ② In the following, we will analyze reasoning in  $DL-Lite$ , 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  $DL-Lite$  is essentially the maximal DL enjoying these nice computational properties.

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

$DL-Lite$  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 systems
  - 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 systems
    - 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**.

# 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 conjunctive query over  $\mathcal{T}$  of arity  $n' > 0$  **without non-distinguished variables**;
- $\vec{x}, \vec{y}$  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

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

Employee's code with salary

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

Employee's code with SSN

...

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

$m_2$ : `SELECT SSN, Salary`  $\rightsquigarrow$  `Employee(pers(SSN)),`  
`FROM D2, D3` `salary(pers(SSN), Salary)`  
`WHERE D2.Code = D3.Code`



# 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 systems
  - The ONTOP framework for Ontology-Based Data Access

# Ontology-Based Data Access System

The mapping assertions are a crucial part of an Ontology-Based Data Access System.

## Def.: **Ontology-Based Data Access System**

is a triple  $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$ , where

- $\mathcal{T}$  is a TBox.
- $\mathcal{D}$  is a relational database.
- $\mathcal{M}$  is a set of mapping assertions between  $\mathcal{T}$  and  $\mathcal{D}$ .

# Semantics of mappings

To define the semantics of an OBDA system  $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$ , 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


 $D_1:$ 

SSN	PrName
23AB	tones
...	...

 $D_2:$ 

Code	Salary
e23	15000
...	...

 $D_3:$ 

Code	SSN
e23	23AB
...	...

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

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

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

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

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

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

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

$$m_1: \quad \text{SELECT SSN, PrName} \quad \rightsquigarrow \quad \text{Employee}(\mathbf{pers}(\text{SSN})),$$

$$\text{FROM } D_1 \quad \quad \quad \quad \quad \quad \quad \quad \text{Project}(\mathbf{proj}(\text{PrName})),$$

$$\quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \text{projectName}(\mathbf{proj}(\text{PrName}), \text{PrName}),$$

$$\quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \quad \text{worksFor}(\mathbf{pers}(\text{SSN}), \mathbf{proj}(\text{PrName}))$$

$$m_2: \quad \text{SELECT SSN, Salary} \quad \rightsquigarrow \quad \text{Employee}(\mathbf{pers}(\text{SSN})),$$

$$\text{FROM } D_2, D_3 \quad \quad \quad \quad \quad \quad \quad \quad \text{salary}(\mathbf{pers}(\text{SSN}), \text{Salary})$$

$$\text{WHERE } D_2.\text{Code} = D_3.\text{Code}$$

# Semantics of an OBDA system

Def.: **Model** of an OBDA system

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

- $\mathcal{I}$  is a model of  $\mathcal{T}$ ;
- $\mathcal{I}$  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 system  $\mathcal{O}$  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 systems
  - The ONTOP framework for Ontology-Based Data Access

# Answering queries over an OBDA system

In an OBDA system  $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$

- Queries are posed over the TBox  $\mathcal{T}$ .
- The data needed to answer queries is stored in the database  $\mathcal{D}$ .
- 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 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 (hence is only theoretical):

- 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 reformulation  $q_{pr} = \text{PerfectRef}(q, \mathcal{T}_P)$  of the original query  $q$ , using the inclusion assertions of the TBox  $\mathcal{T}$  (see later).
- 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 is such that  $Eval(q_{unf}, \mathcal{D}) = Eval(q_{pr}, \mathcal{A}_{\mathcal{M}, \mathcal{D}})$ .
- 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 `tones` and their salary:

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

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

$$\begin{aligned} \vartheta(e) &= \mathbf{pers}(SSN) & \vartheta(s) &= \mathbf{Salary} \\ \vartheta(PrName) &= \mathbf{tones} & \vartheta(p) &= \mathbf{proj(tones)} \end{aligned}$$

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

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

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

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

# 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:  $\text{Aux}_i^j(x) \leftarrow \Phi_i^j(x)$

and introduce the clauses:  $A_i(\mathbf{f}(x)) \leftarrow \text{Aux}_i^j(x)$  (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

$$q(\mathbf{f}(x)) \leftarrow \text{Aux}_1^{j_1}(x), \text{Aux}_2^{j_2}(x), \dots, \text{Aux}_n^{j_n}(x)$$

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

**Query answering** in a *DL-Lite* OBDM system  $\mathcal{O} = \langle \mathcal{T}, \mathcal{M}, \mathcal{D} \rangle$  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  $AC^0$  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.

# 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 *et al.*, 2007]
  - 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 systems
  - 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

[Calvanese *et al.*, 2007] 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)*, pages 274–279, 2007.