

# SEMANTIC WEB REASONING FOR THE INTELLIGENT INTEGRATION OF WEB RESOURCES

Liviu Badea  
[badea@ici.ro](mailto:badea@ici.ro)

Doina Țilivea  
[doina@ici.ro](mailto:doina@ici.ro)

Anca Hotăran  
[ahotaran@ici.ro](mailto:ahotaran@ici.ro)

*ICI - National Institute for Research and Development in Informatics, Bucharest*

**Abstract:** The Semantic Web should enhance the current World Wide Web with reasoning capabilities for enabling automated processing of possibly distributed information. In this paper we describe an architecture for Semantic Web reasoning and query answering in a very general setting involving several heterogeneous information sources, as well as domain ontologies needed for offering a uniform and source-independent view on the data. Since querying a Web source is very costly in terms of response time, we focus mainly on the query planner of such a system, as it may allow avoiding the access to query-irrelevant sources or combinations of sources based on knowledge about the domain and the sources. This work is part of a larger research project, ROTEL, aimed at combining Semantic Web and Natural Language technology.

**Keywords:** Semantic Web, constraint reasoning, intelligent information integration, ontologies, query planning.

In the ROTEL project we aim at developing an intelligent information framework based on Semantic Web and Natural Language technology. In this paper we concentrate on the Semantic Web query processing subsystem.

Taking advantage of the huge amount of knowledge implicit and distributed on the Web is a significant challenge. The main obstacle is due to the fact that most Web pages were designed for human-centred browsing rather than being machine-processable. In addition to static HTML pages the Web currently offers online access to a large number information resources, such as databases with a Web interface. But real-life applications frequently require combining the information from several such resources, which may not have been developed with this interoperability requirement in mind. Thus, a large amount of knowledge is implicit, heterogeneously distributed among various resources and thus hard to process automatically.

The recent developments towards a “Semantic Web” should help address these problems. Being able to explicitly represent domain-specific knowledge in the form of ontologies, should allow reasoning about such machine-processable Web pages.

The emergence of standards for data markup and interchange such as XML and for representing information about resources and their semantics (such as RDF and RDF Schema) can be seen as a first step in the transition towards a Semantic Web. However, the vast majority of Web pages still conform to the HTML standard, which only controls their visual aspects rather than their informational content. Extracting the informational content from such pages which essentially contain free text is a difficult practical problem. The Resource Description Framework (RDF) has been designed to complement such human-oriented text with machine-processable annotations. A large number of prototype systems able to read and reason about such annotations have been developed (TRIPLE [7], Metalog [20], SiLRI [8], Ontobroker [9]). However, currently only a very small minority of Web pages have RDF annotations. Moreover, existing annotations tend to refer to basic features such as document author, creation date, etc., but do not duplicate the information content of the page.

On the other hand, a large number of information sources have a Web interface and could be the building blocks of complex applications, were it not for the unavoidable semantic mismatch between such resources developed by different people. Such Web interfaces produce pages with a partially stable structure, so that their content can be automatically extracted using wrappers, thus replacing human annotation (which is a significant bottleneck in practice). Dealing with information sources rather than a fixed set of Web pages may pose additional problems. For example, systems like TRIPLE read *all* the relevant (RDF) annotations *before* reasoning about them. In the case of large data sources however, it is obviously impossible to retrieve the *entire* content of such sources before starting reasoning. Also, if additional knowledge is available about the sources, some source accesses may be avoided altogether. Therefore, dealing with information sources requires a certain form of *query planning*, i.e. the ability of constructing and reasoning about alternative sequences of source accesses (plans) before actually

querying these sources. Also, *streaming the query responses* may allow starting processing before the entire content of the information source is retrieved.

In this paper we present an approach to such more complex Semantic Web scenarios, involving the integration of heterogeneous resources using rules and ontologies.

The most significant problem faced when trying to combine several resources is related to their heterogeneity. This heterogeneity can be either *structural* (different schemas), or *semantic* (the same entity can be represented using different terms from different vocabularies). Integrating such resources can be achieved by mapping them (both their schemas and their content) to a *common* “knowledge” level, at which their interoperation is straight-forward. This common level involves not just a common (domain-specific) vocabulary, but also formal (machine processable) descriptions of the terms of this vocabulary, as well as the relationships between them, which form a so-called *ontology*. A *mediator architecture* [21] can be used for query answering. In the following we concentrate mainly on the query planning component of such a mediator-based Semantic Web architecture – other important issues such as wrappers and especially ontologies deserve a more detailed discussion, but which is outside the scope of this paper.

## 1. The architecture: public and private levels

The distinctive feature of the Semantic Web is reasoning. However, the various W3C standards related to the (Semantic) Web are not easy to use by a reasoner, especially due to their heterogeneity (XML, RDF, RuleML, etc.). A *uniform* internal level would be much more appropriate for supporting inference and reasoning. The architecture of our system therefore separates a so-called “*public*” level from the *internal level*. The public level refers to the data, knowledge and models exchanged on the Web and between applications and must conform to the current and emerging Web standards such as XML, RDF(S), RuleML, etc.

### 1.1 The internal representation

We use *F-logic* [12] for describing the content of information sources as well as the domain ontology for several important reasons. First, a logic-based language is not only declarative, but also offers support for reasoning. However, a Prolog-like syntax using predicates with a fixed number of arguments and a position-oriented argument access is not well suited for dealing with the semi-structured data available on the Web. On the other hand, F-logic combines the logical features of Prolog with the frame-oriented features of object-oriented languages, while offering a more powerful query language (allowing e.g. aggregation and meta-level reasoning about the schema). Last but not least, F-logic is widely used in the Semantic Web community [7,18,8,9].

In the following (and throughout the whole paper) we use the term “*predicate*” to denote an F-logic molecule, such as  $X:c[a\bar{1} \rightarrow YI, \dots]$ .

While the data content of an information source can be represented by a number of so-called *source predicates*  $s$ , the user might prefer to have a *uniform* interface to *all* the sources instead of directly querying the sources  $s$ . In fact, she may not even *want* to be aware of the structure of the information sources. Therefore, the mediator uses a uniform knowledge representation language in terms of so-called “*model predicates*”, which represent the user's perspective on the given domain and refer to concepts and terms of an associated domain ontology.

### 1.2 Source predicates, model predicates and description rules

The content of the various Web information sources is represented by so-called *source predicates*. For a uniform and integrated access to the various sources these are described (at the mediator level) in terms of so-called *model predicates*. More precisely, we distinguish between *content* predicates and *constraint* predicates. *Content predicates* (denoted in the following with  $p, q$ ) are predicates which directly or indirectly represent the *content* of information sources. They can be either *source predicates* or *model predicates*. *Source predicates* (denoted with  $s$ ) directly represent the content of (part of) an information source (for example, the content of a Web page or the answer returned by a Web service). *Model predicates* (denoted with  $b$ ) are used to describe the “global model”, including the domain ontology and

the information content of the sources in terms of this ontology. As opposed to content predicates, *constraint predicates* (denoted by  $c$ ) are used to express specific constraints on the content predicate descriptions. For example, a source  $s$  containing information about underpaid employees (with a salary below 1000), would be described as:

$$E:s[\text{name} \rightarrow N, \text{salary} \rightarrow S] \rightarrow E:\text{employee}[\text{name} \rightarrow N, \text{salary} \rightarrow S], S < 1000.$$

Constraint predicates can be either *internal* (treatable internally by the query engine of the source), or *external* (constraints that can only be verified at the mediator level, for example by the built-in constraint solvers of the host CLP environment). Constraints treatable *internally* by the sources can be *verified* at the source level (by the query engines of the sources), but they are also *propagated* at the mediator level. Constraints treatable only *externally* need to be both verified and propagated at the mediator level.

A *complete* description of the source predicates in terms of model predicates is neither possible nor useful, since in general there are too many details of the functioning of the application. Thus, instead of *complete* (iff) descriptions, we shall specify only approximate (necessary) definitions of the source predicates in terms of model predicates (thus, only the relevant features of the sources will be encoded). In the following, we use a *uniform* notation for the domain and source *description rules*:

*Antec, Constra*  $\rightarrow$  *Conseq, Constrc* (dr)

where *Antec* and *Conseq* are conjunctions of content predicates, while *Constra* and *Constrc* are conjunctions of constraints. Variables occurring in the consequent but not in the antecedent are implicitly existentially quantified. As in Prolog and F-logic, all other variables are implicitly universally quantified. *Description rules* are necessary definitions of (combinations of) source or model predicates in terms of model predicates and constraints. (*Source descriptions* are special cases of description rules where the antecedent contains only source predicates. *Integrity constraints* are description rules with only constraints – typically ‘fail’ – in the consequent.)

### 1.3 Querying the sources: wrappers

While the traditional Web was designed mainly for browsing, the Semantic Web should enable automated processing of Web sources. A Semantic Web system should therefore be able to deal not just with *static* Web pages (in various formats such as HTML, XML), but also with *dynamically generated* pages as well as Web services.

While the semantics of static Web pages is encoded using static annotations (e.g. in RDF) the information content of dynamically generated Web pages needs to be extracted by automatic *wrappers*. These have to be quite flexible, i.e. able to deal with XML files, or even possible non-well-formed HTML. Due to its wide spread use worldwide, we currently employ *XQuery* for implementing our wrappers (we also use *tidy*<sup>1</sup> to preprocess non-well-formed HTML sources).

## 2. Ontologies and mapping rules

A key component in a SW application is a *domain ontology*, used by a certain community to enable sharing and exchange of knowledge in a particular domain. Due to the heterogeneous nature of the sources, mapping rules from the various sources to the common ontology are needed for obtaining a unified view of these sources (see Section 0 below for examples).

We currently employ Protégé<sup>2</sup>, a widely used Ontology development tool. Ontologies are exported in RDF(S) format and automatically converted to the internal (F-logic) format. As discussed above, additional rules may be needed to extend the expressiveness of RDF(S).

---

<sup>1</sup> <http://tidy.sourceforge.net/>

<sup>2</sup> <http://protege.stanford.edu>

### 3. Reasoning and Query Planning

Since querying a Web source is extremely costly in terms of response time, every possible optimization that might avoid accessing irrelevant sources should be attempted. The knowledge available about a particular source or combination of sources may imply the inconsistency of some potential query plan even *before* accessing the sources. The purpose of the query planner is to avoid executing such inconsistent plans and possibly to use additional cost measures for ranking plans.

The majority of Semantic Web implementations using F-logic [7, 18, 8, 9] retrieve the content of *all* the sources before query answering. However, in applications in which there are alternative plans for a given query, some of these plans may be inconsistent and -for efficiency -the associated sources should not be accessed. This cannot be achieved simply by the normal F-logic query answering mechanism and requires some form of meta-interpretation.

*Query planning* amounts to unfolding the query in terms of sources and propagating the relevant constraints to eliminate the inconsistent plans. Since this type of reasoning is performed *before* accessing the sources, it can be viewed as an *abductive procedure* [14] in which the source predicates play the role of abducibles.

Due to their flexibility and declarative nature, *Constraint Handling Rules* (CHRs) [10] represent an ideal framework for implementing the reasoning mechanism of the query planner. *CHRs* represent a flexible approach to developing user-defined constraint solvers in a declarative language. As opposed to typical constraint solvers, which are black boxes, CHRs represent a 'no-box' approach to CLP. CHRs can be either *simplification* or *propagation* rules.

A *simplification* rule  $Head \Leftrightarrow Guard \mid Body$  replaces the head constraints by the body provided the guard is true (the *Head* may contain multiple CHR constraint atoms).

*Propagation* rules  $Head \Rightarrow Guard \mid Body$  add the body constraints to the constraint store without deleting the head constraints (whenever the guard is true). A third, hybrid type of rules, *simpagation* rules  $Head1 \setminus Head2 \Leftrightarrow Guard \mid Body$  replace *Head2* by *Body* (while preserving *Head1*) if *Guard* is true. (Guards are optional in all types of rules).

CHRs allow us to combine in an elegant manner the backward reasoning necessary for implementing query unfolding with the forward propagation of abducibles and constraints.

A description rule of the form (dr) will be encoded in CHR using

- *goal regression rules*: for reducing queries given in terms of model predicates to queries in terms of source predicates,

$$G \text{ Conseq} :- \neg G \text{ Antec}, \text{Constra} \tag{b}$$

- *propagation rules*: for completing (intermediate) descriptions in order to allow the discovery of potential inconsistencies.

$$H \text{ Antec} \Rightarrow \text{Constra} \mid H \text{ Conseq}, \text{Constrc} \tag{f}$$

In our CHR embedding of the *abductive procedure* we use two types of constraints,  $Gp$  and  $Hp$ , for each predicate  $p$ . While  $Hp$  represents facts explicitly propagated (abduced),  $Gp$  refers to the *current closure* of the predicate  $p$  (i.e. the explicit definition of  $p$  together with the explicitly abduced literals  $Hp$ ). While propagating  $Hp$  amounts to simply assuming  $p$  to hold (abduction), propagating  $Gp$  amounts to trying to prove  $p$  either by using its definition  $def(p)$ , or by reusing an already abduced fact  $Hp$ . In fact, our description rules can be viewed as a generalization of abductive logic programs with integrity constraints interpreted w.r.t. the 'propertyhood view'<sup>3</sup>.

A set of subgoals in terms of source predicates (induced by backward chaining from the user query using the goal regression rules) may not necessarily be consistent. Applying forward propagation rules ensures the completion ("saturation") of the (partial) query plan and enables detecting potential conflicts *before* actually accessing the sources.

<sup>3</sup> The detailed presentation of the semantics is outside the scope of this paper. Briefly, the goal regression rules make up the program  $P$ , the sources are regarded as (temporary) abducibles  $A$ , while the forward propagation rules play the role of the ALP integrity constraints  $I$ . Our notion of integrity constraints is however more general than that used in ALP.

*Constra* and *Constrc* in (f) represent constraint predicate calls, for which the associated constraints solvers are assumed to be available. The occurrence of *Constra* in the guard ensures that the consequent is propagated only if *Constra* hold.

Source and domain models are described using rules of the form (dr), which are then automatically translated by the system into CHR goal regression and propagation rules (b) and (f). The following *additional problem-independent rules* (gh) and (re) are used for obtaining the complete CHR encoding of a model.

- a CHR simpagation rule<sup>4</sup> for matching a goal  $Gp$  with any existing abduced facts  $Hp$ :  $Hp(X1) \setminus Gp(X2) \Leftrightarrow X1=X2 ; X1 \neq X2, Gp(X2)$ . (re)

This rule should have a higher priority than the unfolding rule in order to avoid re-achieving an already achieved goal. Note that, for completeness, we are leaving open the possibility of achieving  $Gp(X2)$  using its definition or reusing other abduced facts.

- a rule taking care of the consistency between goals and facts that simply hold (since goals will be achieved eventually, they also have to hold):

$Gp \Rightarrow Hp$  (gh)

We have already mentioned the fact that we have to distinguish between goals involving a given predicate  $p$  (which will be treated by a mechanism similar to the normal Prolog backward chaining mechanism) and the instances  $Hp$  of  $p$ , which trigger forward propagation rules. Operationally speaking, while goals  $Gp$  are "consumed" during goal regression, the fact that  $p$  holds should persist even after the goal has been achieved, to enable the activation of the forward propagation rules of  $p$ . Note that rule (gh) should have a higher priority than (b) to allow goals  $p$  to propagate  $Hp$  before applying the goal regression rules for  $p$ .

Certain forward propagation rules (f) may propagate facts that may never be involved (directly or indirectly) in conflicts. These facts may be very useful to the end user in certain applications, especially whenever the user would like to see all the known facts about the instances returned by the query. However, in other applications, we may wish to refrain from propagating such facts that may never lead to an inconsistency (detected by an integrity constraint). In this case, we perform a static dependency analysis of the rules, and generate forward propagation rules (f) only for predicates that may propagate an inconsistency.

## 4. An example

In this paper, we consider a hardware configuration problem as a typical example. A component integrator selling customized computer configurations may use components from several component providers, while trying to satisfy a set of compatibility and user constraints. This problem has many typical features of the Semantic Web:

- it involves distributed and dynamically changing information sources (the component
- catalogs of the different providers available via Web interfaces)
- the information is semantically heterogeneous, requiring a domain ontology for a uniform access
- the domain involves complex compatibility constraints between components, requiring constraint propagation during query planning -there are several alternative sources for some given component, which makes query planning necessary.

Here we consider just a fragment of this hardware configuration domain, in order to emphasize the main features of our architecture. Assume we have two main component providers, or vendors, *flamingo* and *oktal*. (Since their catalogs are accessible via the Web in HTML format, a wrapper is used for extracting the relevant information from these Web pages.) A user query may ask for a system with an upper bound on the total price. Here we just concentrate on a "simple system" containing just a motherboard and a processor. Of course, the two components must satisfy a compatibility constraint (e.g. one should not attempt to use an AMD processor on an Intel-Pentium compatible motherboard).

<sup>4</sup> The rule is more complicated in practice, due to implementation details.

Since the sources are heterogeneous, we first use so-called “mapping rules” to describe their content in terms of the ontology. For example, motherboards sold by *flamingo* are mapped onto the ontology concept *motherboard* (note that here we have a trivial mapping of all slots, with an additional *m\_vendor* slot recording the name of the vendor):

$$Mb:fl\_motherboard[\bar{A} \rightarrow \bar{X}] \rightarrow Mb:motherboard[\bar{A} \rightarrow \bar{X}, m\_vendor \rightarrow flamingo] \quad (r1)$$

Any additional knowledge about the content of the sources may be very useful during query planning for discarding inconsistent plans even before accessing the sources. For example, we may know a lower bound for the price of a *flamingo* motherboard. We may also know that *flamingo* distributes only *intel* and *msi* boards. This is specified using the following source description rules:

$$Mb:fl\_motherboard[m\_price \rightarrow \bar{P}] \rightarrow P \geq 70. \quad (ic1)$$

$$Mb:fl\_motherboard[brand \rightarrow \bar{Br}] \rightarrow member(Br, [intel, msi]). \quad (ic2)$$

We have similar rules for *oktal* motherboards, as well as for processors (available also from both *flamingo* and *oktal*):

$$Mb:okt\_motherboard[\bar{A} \rightarrow \bar{X}] \rightarrow Mb:motherboard[\bar{A} \rightarrow \bar{X}, m\_vendor \rightarrow oktal]. \quad (r2)$$

$$Mb:okt\_motherboard[m\_price \rightarrow \bar{P}] \rightarrow P \geq 70 \quad (ic3)$$

$$Mb:okt\_motherboard[brand \rightarrow \bar{Br}] \rightarrow member(Br, [gigabyte, msi]). \quad (ic4)$$

$$Pr:fl\_processor[\bar{A} \rightarrow \bar{X}] \rightarrow Pr:processor[\bar{A} \rightarrow \bar{X}, p\_vendor \rightarrow flamingo]. \quad (r3)$$

$$Pr:fl\_processor[p\_price \rightarrow \bar{P}] \rightarrow P \geq 150 \quad (ic5)$$

$$Pr:okt\_processor[\bar{A} \rightarrow \bar{X}] \rightarrow Pr:processor[\bar{A} \rightarrow \bar{X}, p\_vendor \rightarrow oktal]. \quad (r4)$$

$$Pr:okt\_processor[p\_price \rightarrow \bar{P}] \rightarrow P \geq 150 \quad (ic6)$$

Note that the source descriptions (ic1), (ic2), (ic3), (ic4), (ic5) and (ic6) express knowledge about the content of the sources. (These descriptions could either be provided by the knowledge engineer or could be automatically retrieved from previous source accesses.)

However, answering queries frequently requires knowledge about the sources that cannot be inferred from their content alone. For example, *oktal* offers discounts for purchases over a given threshold, while *flamingo* does not:

$$oktal:vendor[name \rightarrow oktal, discount \rightarrow 0.1, discount\_threshold \rightarrow 200]. \quad (r5)$$

$$flamingo:vendor[name \rightarrow flamingo, discount \rightarrow 0, discount\_threshold \rightarrow 0]. \quad (r6)$$

The domain ontology encodes our knowledge about this particular domain, such as the concept hierarchy, descriptions of slots/attributes (type, cardinality constraints, etc.), and any additional description rules involving these concepts. For example, *motherboards* and *processors* are sub-concepts of *component*:

*motherboard* :: *component*.

*processor* :: *component*.

For an example of an ontology description rule, we may describe a simple *system* as a set of matching components, such as a *motherboard* and a *processor*. (Note the compatibility constraint between *motherboard* and *processor*: ‘*Pr.id* = *Mb.supported\_CPU*’).

$$Mb:motherboard, Pr:processor, Pr.id = Mb.supported\_CPU \rightarrow \quad (r7)$$

*X:system[motherboard \rightarrow Mb, processor \rightarrow Pr]*.

To show the functioning of the query planner, we consider the following user query which asks for a

system with an upper bound of 210 on the (possibly discounted) price:

$$\begin{aligned} ?-S:\text{system}[\text{motherboard} \bar{\square} Mb[\text{brand} = \text{gigabyte}], \text{processor} \bar{\rightarrow} Pr], & \quad (q1) \\ \text{compute\_discounted\_price}(S, S\_price), S\_price \square 210. & \end{aligned}$$

For simplicity, the computation of the discounted price is performed by *compute\_discounted\_price*(S, S\_price).

The query is answered by first invoking the query planner, and subsequently executing the resulting plans. As discussed previously, query planning *unfolds* the query in terms of source calls while *propagating constraints* with the purpose of discarding the inconsistent plans before the actual source accesses. Note that constraint propagation is interleaved with query unfolding, so that inconsistencies are detected as early as possible.

In our example, query planning starts by unfolding the query (q1) to the following subgoals and constraint (as discussed previously, *G* stands for ‘goal’, while *H* stands for ‘holds’):

$$GS:\text{system}[\text{motherboard} \bar{\rightarrow} Mb[\text{brand} = \text{gigabyte}], \text{processor} \bar{\rightarrow} Pr] \quad (g1)$$

$$G\text{compute\_discounted\_price}(S, P) \quad (g2)$$

$$S\_price \leq 210 \quad (c1)$$

Then (g1) will be unfolded with (r7-B)<sup>5</sup> to:

$$GMb:\text{motherboard}[\text{brand} = \text{gigabyte}] \quad (g3)$$

$$GPr:\text{processor} \quad (g4)$$

Since motherboards can be acquired either from *flamingo* or from *oktal*, (g3) will be unfolded first with (r1-B) to the following (we first try to acquire the motherboard from *flamingo*):

$$GMb:\text{fl\_motherboard}[\text{brand} \bar{\rightarrow} \text{gigabyte}, m\_price \bar{\rightarrow} MP, \dots] \quad (g5)$$

This will propagate (with rule (gh)) *HMb:fl\_motherboard[brand  $\square$  gigabyte, m\_price  $\rightarrow$  MP, ...]* which in turn will activate (r1-F) and propagate

$$HMb:\text{motherboard}[m\_vend\bar{or} \bar{\rightarrow} \text{flamingo}, \text{brand} \bar{\rightarrow} \text{gigabyte}, \dots]. \quad (h1)$$

But now, (h1) will trigger the (ic2) integrity constraint, which fails because the required brand for the motherboard (*gigabyte*) is not included in the list of distributed brands of the vendor *flamingo*.

Alternatively, we could obtain a motherboard from *oktal*, i.e. unfold *GMb:motherboard* (g3) with (r2-B) to:

$$GMb:\text{okt\_motherboard}[\text{brand} \bar{\rightarrow} \text{gigabyte}, m\_price \bar{\rightarrow} MP, \dots] \quad (g5)$$

$$HMb:\text{okt\_motherboard}[\text{brand} \bar{\rightarrow} \text{gigabyte}, m\_price \bar{\rightarrow} MP, \dots] \quad (h2)$$

is propagated with rule(gh). In this case, (r2-F) will propagate *HMb:motherboard[m\_vend $\bar{or}$   $\rightarrow$  oktal, brand  $\rightarrow$  gigabyte, ...]*. (h3)

(h2) will trigger (ic2), which will be consistent because the brands known to be distributed by *oktal* are *gigabyte* and *msi*. (h3) will also trigger (ic3) and propagate the price constraint *Mb.m\_price  $\geq$  70* (c2).

Having selected a consistent source for the motherboard, we now consider selecting a processor. The goal (g2) *GPr:processor* will be first unfolded using (r3-B) to:

<sup>5</sup> (ri-B) is the goal reduction rule (*backward* version) of the rule (ri), while (ri-F) is the associated *forward* propagation rule.

$$GPr:fl\_processor[p\_price \rightarrow PP, \dots] \quad (g6)$$

and will propagate with (gh) and (r3-F):

$$HPr:fl\_processor[p\_price \rightarrow PP] \quad (h4)$$

which will trigger (ic5) and propagate the price constraint  $Pr.p\_price \geq 150$ . (c3)

Then the next goal,  $Gcompute\_discounted\_price(S, S\_Price)$  (g2), is unfolded and partially evaluated. For the simplicity of the presentation, we do not present here all the intermediate steps. Briefly, although at query planning time the prices of the components are not yet known, we can nevertheless propagate the constraints (c2) and (c3) and thus obtain a lower bound on  $S\_price$ .

More precisely, since we have different vendors for the two components ( $Pr.p\_vendor = flamingo$ ,  $Mb.m\_vendor = oktal$ ), any potential discounts will be applied separately to each component. But *flamingo* doesn't offer discounts and the discount threshold for *oktal* is too high (200) to be applicable. Therefore, the system price will be:  $S\_price = Pr.p\_price + Mb.m\_price \geq 220$  (due to (c3) and (c2)), which is inconsistent with the (c1) constraint  $S\_price \leq 210$ , so the combination of an *oktal* motherboard and a *flamingo* processor will be rejected.

Alternatively, we could acquire the processor from *oktal*, i.e. reduce  $GPr:processor$  (g1) with (r4-B), (r4-F) and (gh) to:

$$HPr:okt\_processor[p\_price \rightarrow PP, ..] \quad (h5)$$

As before, (h5) will trigger (ic6) and propagate the price constraint  $Pr.p\_price \geq 150$ . (c4)

But now the unfolding of  $compute\_discounted\_price(S, S\_price)$  (g2) will be different, as both components are acquired from the same vendor ( $Pr.p\_vendor = oktal$ ,  $Mb.m\_vendor = oktal$ ), so that the discount is applicable to both components. In fact, since the undiscounted price  $S\_price1 = Pr.p\_price + Mb.m\_price$  is above the discount threshold for *oktal* ( $S\_price1 \geq 20$  (due to (c4) and (c2)),  $oktal.discount\_threshold = 200$ ), the system price will be:  $S\_price = S\_price1 (1 - oktal.discount) = 0.9 \cdot S\_price$ . Thus, the lower bound on  $S\_price$  will be 198, which is consistent with the (c1) constraint. Query planning has therefore retained only the following combination of sources: *okt\_processor* and *okt\_motherboard*, while discarding the other three possible combinations at planning time, so that only *okt\_processor* and *okt\_motherboard* are actually queried.

## 5. Conclusions and Future Work

An exhaustive comparison with other information integration systems is impossible, due to the very large number of such systems as well as to the lack of space. Briefly, while database oriented approaches to integration (such as *multi-databases* and *federated databases*) use *fixed* global schemas and are appropriate only if the information sources and users do not change frequently, we deal with *dynamically evolving* schemas (especially if the LAV modeling approach is employed). On the other hand, more *procedural* intelligent information integration approaches, like TSIMMIS [11] and even some declarative systems like MedLan [1] or HERMES [21], use explicit query reformulation rules<sup>6</sup>, but *without the equivalent of our forward propagation rules* (which allow an early discovery and pruning of inconsistent plans before query execution). Our approach is closer to the more declarative systems like SIMS [2], Information Manifold [17] and Infomaster [6].

COIN [5] also uses a CLP framework (Eclipse) for abductive reasoning and CHRs for implementing integrity constraints. However, integrity constraints can be imposed in COIN *only on source predicates*. Thus, COIN domain knowledge reduces to Prolog reduction rules, which are used *only backwards* (during goal regression). The lack of forward propagation rules involving base predicates (and not just sources) makes the discovery of potential interactions between base predicates (and thus the full use of domain knowledge) impossible. Other related mediator-based Web integration systems are EMERAC [13] and Ariadne [15].

We are currently working towards extending our system to use F-logic as internal reasoning language.

<sup>6</sup> Such query templates correspond to our goal regression rules.



This is non-trivial as there are no Prolog environments allowing a combination F-logic and CHR. In this sense, our work is related to Xcerpt [4], an elegant declarative, rule-based query and transformation language for XML, for which a query planner is also currently under development. We also consider using WSDL [23] descriptions as source capabilities during query planning.

**Acknowledgements.** This work has been supported by the ROTEL project (CEEX contract 29/2005) of the Romanian research of Excellence Program. We are grateful to the members of the REVERSE Network of Excellence of the European Community (<http://www.reverse.net>), in which we are also involved.

## References

1. **AQUILINO, D., P. ASIRELLI, C. RENSO, F. TURINI:** MedLan: a Logic-based Mediator Language, IEI Technical Report B4-16, November 1997.
2. **ARENS, Y., C. A. KNOBLOCK, CHUN-NAN HSU:** Query Processing in the SIMS Information Mediator, Advanced Planning Technology, A. Tate (ed), AAAI Press, 1996.
3. **BADEA, L., D. TILIVEA:** Intelligent Information Integration as a Constraint Handling Problem, Proc. of the Fifth International Conference on Flexible Query Answering Systems (FQAS2002), October 27 - 29, 2002, Copenhagen, Springer Verlag, pp. 12-27.
4. **BERGER, S., F. BRY, S. SCHAFFERT, C. WIESER C.:** Xcerpt and visXcerpt: From Pattern-Based to Visual Querying of XML and Semistructured Data. Proceedings VLDB03, Sept. 2003.
5. **BRESSAN, S., C. H. GOH:** Answering queries in context. In Proceedings of the International Conference on Flexible Query Answering Systems, FQAS-98, Roskilde, 1998.
6. **DUSCHKA, O.M., M. R. GENESERETH:** Infomaster -An Information Integration Tool. Tool. Proc. International Workshop "Intelligent Information Integration", KI-97, Freiburg, 1997.
7. **DECKER, S., M. SINTEK:** Triple - an RDF query, inference, and transformation language. In Proc. of the 2002 International Semantic Web Conference (ISWC-2002).
8. **DECKER, S., D. BRICKLEY, J. SAARELA, J. ANGELE:** A Query and Inference Service for RDF. QL'98 – The Query Languages Workshop, World Wide Web Consortium, 1998.
9. **FENSEL, D., et al.:** On2broker: Semantic-based Access to Information Sources at the WWW, Proceedings of WebNet, 1999, pp. 366-371.
10. **FRUEWIRTH, T.:** Theory and Practice of Constraint Handling Rules, JLP 37:95-138, 1998.
11. **GARCIA-MOLINA H., et al.:** The TSIMMIS approach to mediation: Data models and Languages. In Journal of Intelligent Information Systems, 1997.
12. **KIFER, M., G. LAUSEN, J. WU:** Logical foundations of object-oriented and frame-based languages, Journal of the ACM, Volume 42 , Issue 4 (July 1995), pp. 741-843.
13. **KAMBHAMPATI, S., et al.:** Optimizing Recursive Information Gathering Plans in EMERAC. Journal of Intelligent Information Systems. Vol. 22, No. 2, March 2004, p. 119-153.
14. **KAKAS, A., R. KOWALSKI, F. TONI:** The role of abduction in logic programming, Handbook of logic in AI and LP 5, OUP 1998, pp. 235-324.
15. **KNOBLOCK, C. et al.:** The ARIADNE Approach to Web-Based Information Integration. In: International Journal of Cooperative Information Systems 10(1-2), 2001, pp. 145-169.

16. **LEVY, A.Y.:** Logic-Based Techniques, in Data Integration Logic Based Artificial Intelligence, Jack Minker (ed). Kluwer, 2000.
17. **LEVY, A.Y., A. RAJARAMAN, J. J. ORDILLE:** Querying Heterogeneous Information Sources Using Source. Proc. 22nd VLDB Conference, Bombay, India. 1996.
18. **LEVY, A.Y., M. C. ROUSSET:** Combining Horn Rules and Description Logics in CARIN. Artificial Intelligence Journal 104, September 1998.
19. **LUDASCHER, B., et al.:** Managing Semistructured Data with FLORID: A Deductive Object-oriented Perspective. Information Systems, 23 (8), 1998, pp. 589-613.
20. **MARCHIORI, M., J. SAARELA:** Query+Metadata+Logic=Metalog, <http://www.w3.org/TandS/QL/QL98/pp/metalog>
21. **SUBRAHMANIAN, V.S. et al.:** HERMES: A heterogeneous reasoning and mediator system. <http://www.cs.umd.edu/projects/hermes/overview/paper>
22. **WIEDERHOLD, G.:** Mediators in the architecture of future information systems, IEEE Comp. 25 (3) 1992, pp. 38-49.
23. \* \* \*: Web Services Description Language (WSDL) Version 2.0 <http://www.w3.org/TR/wsdl20/>.