

Massachusetts Institute of Technology
Engineering Systems Division

Working Paper Series

ESD-WP-2009-03

RECONCILIATION OF TEMPORAL SEMANTIC HETEROGENEITY
IN EVOLVING INFORMATION SYSTEMS

Hongwei Zhu¹ and Stuart Madnick²

¹College Business and Public Administration,
Old Dominion University
hzhu@odu.edu

²Sloan School of Management and Engineering Systems Division, MIT
smadnick@mit.edu

April 2009

Reconciliation of temporal semantic heterogeneity in evolving information systems

Hongwei Zhu* – Stuart E. Madnick**

**College Business and Public Administration, Old Dominion University
Norfolk, VA 23529, USA*

hzhu@odu.edu

***MIT Sloan School of Management, Room E53-321
Cambridge, MA 02139, USA*

smadnick@mit.edu

ABSTRACT. The change in meaning of data over time poses significant challenges for the use of that data. These challenges exist in the use of an individual data source and are further compounded with the integration of multiple sources. In this paper, we identify three types of temporal semantic heterogeneity. We propose a solution based on extensions to the Context Interchange framework, which has mechanisms for capturing semantics using ontology and temporal context. It also provides a mediation service that automatically reconciles semantic conflicts. We show the feasibility of this approach with a prototype that implements a subset of the proposed extensions.

KEY WORDS: temporal context, semantic heterogeneity, ontology, logic programming

1. Introduction

Effective management of the temporal aspects of data has become increasingly important in application domains ranging from day-to-day record keeping to counterterrorism efforts. It is often even required by law for organizations to store historical data and make sure it is accurate and easy to retrieve¹. While temporal databases can be used to manage the data, ensuring that the retrieved data is meaningful to the users is still an unsolved problem when data semantics changes over time.

As an example, suppose an arbitrage analyst in New York needs to compare Daimler-Chrysler's stock prices at New York and Frankfurt exchanges. He retrieved the data from Yahoo's historical database, see Figure 1. Two anomalies caught his eyes at a quick glance at the data. First, prices at two exchanges differ substantially; and second, the price at Frankfurt stock exchange dropped by almost 50% at the turn from 1998 to 1999!

Date	Open	High	Low	Close	Volume	Adj Close*
6-Jan-99	105.25	105.74	103.92	105.13	2,061,200	90.49
5-Jan-99	99.05	103.43	98.93	103.31	2,634,600	88.92
4-Jan-99	99.66	100.69	98.08	98.99	3,441,400	85.20
31-Dec-98	94.49	94.55	93.21	93.51	506,900	80.49
30-Dec-98	94.97	95.34	94.18	94.18	391,300	81.06
29-Dec-98	96.13	96.25	95.64	95.95	1,195,700	82.58
28-Dec-98	95.64	96.43	95.16	95.64	1,707,800	82.32
6-Jan-99	90.10	92.40	89.30	92.30	13,950,500	86.67
5-Jan-99	86.80	88.60	86.10	86.80	12,329,300	81.51
4-Jan-99	83.50	88.50	82.50	87.50	13,660,200	82.16
30-Dec-98	166.30	167.90	164.50	164.50	4,934,820	154.47
29-Dec-98	166.00	166.50	164.50	165.00	5,039,660	154.94
28-Dec-98	159.50	167.80	159.30	166.50	9,748,480	156.34

* Close price adjusted for dividends and splits.

Figure 1. Historical stock prices for Daimler-Chrysler. Top: New York Stock Exchange; Bottom: Frankfurt Stock Exchange

These anomalies result from unresolved semantic conflicts between the data sources. In this case, not only are the currencies for the stock prices different at the

¹ See Robert Sheier on "Regulated storage" in Computerworld, 37(46), November 17, 2003. Health Insurance Portability Act requires healthcare providers keep records till two years after death of patients; Sarbanes-Oxley Act requires auditing firms retain records of financial statements.

two exchanges, but the currency at Frankfurt also changed from German Marks to Euros at the beginning of 1999 – but that is not noted by the source. Once the data is normalized using this knowledge, it can be shown that there is neither significant arbitraging opportunity between the two exchanges nor an abrupt price plunge at Frankfurt for this stock. We call metadata knowledge such as the currency for price *context knowledge*, and the history of time-varying metadata *temporal context*.

To allow data receivers like the analyst to effectively use data from time-varying heterogeneous sources, we need to represent temporal context knowledge and incorporate it into data integration and query answering systems. Temporal database research has primarily focused on the management of temporal data in a constant and homogeneous context environment. Semantic data integration techniques developed so far are based on snapshot data models that ignore the time dimension.

The objective of this research is to fill this gap by developing techniques to effectively resolve temporal semantic conflicts between data sources and receivers. Specifically, we extend the Context Interchange (COIN) framework (Firat 2003, Goh *et al.* 1999) with temporal contextual knowledge representation and reasoning capabilities.

The rest of the paper is organized as follows. In section 2 we use examples to illustrate three types of temporal semantic heterogeneity. In Section 3 we give a brief review of related research. In section 4 we describe our solution approach. In section 5 we present some preliminary results. In the final section, we summarize and briefly discuss future research

2. Challenges of temporal data integration

2.1. A simple integration example

A temporal database is one that supports some aspect of time, not counting user-defined time such as birthday and hiring date (Jensen *et al.* 1998) . This rather informal definition is due to the fact that the temporal dimensions are often application specific, therefore it is either difficult or unnecessary to support all aspects of time. Nevertheless, most *temporal data* can be viewed as time-stamped propositions and represented as relational tuples with timestamps.

Table 1 gives an example of some time series data for a company. Intuitively, the example describes how the values of several attributes change over time. Each tuple represents a fact that can be viewed as a predicate with a timestamp argument and other non-temporal arguments. However, there are other unspecified metadata attributes, such as currency type and scale factor, that critically determine the truth value of each predicate. We call the specification of metadata attributes *context knowledge*. For metadata attributes whose value changes over time, a specification of their history is termed *temporal context*.

Table 1. *Company time series data*

Year	Num_Employee	Profit	Tax
...			
1999	5100	4.2	1.1
2000	12000	13000	2500
2001	25.3	20000	4800
2002	30.6	35.3	7.97
...			

Table 2 gives examples of the context knowledge in a simple information integration scenario involving the source in Table 1 and a receiver. The receiver context can be time-varying as well. Semantic conflicts arise because the source and the receiver have different contexts, which need to be reconciled for the receiver to meaningfully use the data. Imagine the complexity of scenarios that involve dozens of sources and receivers, each with time-varying heterogeneous contexts. We need effective technologies to manage this complexity.

Table 2. *Examples of temporal context*

	Source	Receiver
Currency	Francs(FRF), year \leq 2000 Euros, year \geq 2001	USD, always
Scale factor for profit and tax	1M, year \leq 1999 1K, 2000 \leq year \leq 2001 1M, year \geq 2002	1K, always
Scale factor for Num_Employee	1, year \leq 2001 1K, year \geq 2002	1K, always
Profit	Exclude tax, year \leq 2000 Include tax, year \geq 2001	Include tax, always

2.2. Temporal semantic heterogeneities

We see at least three categories of issues in the integration of temporal data.

Representational heterogeneity – the same relational attribute may be represented differently in the time span of a data source. In addition to *currency* changes for monetary concepts like *profit* and *tax*, there are also *scale factor* changes, as described in Table 2.

Ontological heterogeneity – the ontological concept represented by an attribute may change over time. In Table 2, profit on and before 2000 excludes taxes, afterwards it includes taxes. There are also cases where the entity referred to by an identifier changes over time. For example, stock symbol “C” on the New York

Stock Exchange (NYSE) used to refer to Chrysler but changed to refer to Citigroup on December 4, 1998 after Chrysler merged with Daimler-Benz. Similarly, country code “YUG” for Yugoslavia have different geographic boundaries at different times before, during, and after the recent Balkans war.

The derivation method or composition of complex concepts often changes over time. Many government published time series data sets often come with a “Change of Definitions” that explains changes to terminologies. For example, the national unemployment data may include undocumented immigrants in the workforce at one time and exclude them at another.

Heterogeneity in temporal entity – the abstraction and representation of time domain differs across systems and time. Although a temporal entity is just another data type, it has special properties and operations that warrant a category of its own. The example in Table 1 uses point representation for the timestamp attribute *year*. Another system may choose to use intervals, e.g., [1/1/1999, 12/31/1999] for the year 1999. Differences in calendar systems, time zones, and granularities present many challenges for integration of temporal data.

The semantics of the association between propositions described by the non-temporal attributes in a tuple and the temporal entity may differ across attributes. How the truth of a proposition over an interval is related to its truth over subintervals is described by the proposition’s *heredity* properties (Shoham 1987). Recognizing this property is useful for temporal query language design. For example, if in a bank account database the balance over an interval is known and the user queries the balance at a time within the interval, the query language should use the *liquidity* property of balance attribute to infer the result (Bettini *et al.* 2000). We observe that heredity is often attribute dependent and does not change over time or across data sources. Thus we need not consider heterogeneity of this property in the data integration setting.

In an effective integration framework, data receivers should not be burdened by these types of context heterogeneity; rather, there should be a system service to record contexts and reconcile context differences before delivering data to the receiver. Our temporal extension to the original COIN framework provides such a solution.

3. Review of Related Research

Related research can be found in the areas of temporal database, temporal reasoning, and data integration. Although each provides useful concepts and techniques, none address all the temporal semantic heterogeneity problems identified in this paper. The following brief review is not intended to summarize or criticize the findings in each area; rather, it is to identify the most relevant results and show what is missing from a temporal semantic data integration point of view.

Temporal databases. The time domain is often represented as time points with certain granularities. An interval is a set of contiguous points and can be represented as a pair of begin and end points. A time point may have a duration and thus is not an *instant* in time ontology (Hobbs 2002).

Over 40 temporal data models have been proposed (Özsoyoglu and Snodgrass 1995). Many of the models let the system manage timestamps, which effectively hide the timestamp attribute from the user. This approach is inconsistent with the relational theory (Date *et al.*, 2003). As commonly practiced, databases that store temporal data often have a schema with explicit timestamp attribute(s); standard SQL is used to retrieve data and temporal operations are selectively implemented in the application layer. Our framework targets the common situation where data sources have limited temporal support.

As in the case of conventional databases, temporal databases also fail to facilitate context knowledge management. As a result, context is often hard-coded into data transformations in data warehouses. This ad-hoc approach lacks flexibility and scalability.

Temporal reasoning. While a restricted set of temporal logics can be executed using logic programming, there seems to be a trend where temporal logics are transformed into temporal constraints to take advantage of the efficiency of constraint solvers. The framework provided in (Meiri 1996) combines qualitative and quantitative (metric) temporal relations over both time points and time intervals. These relationships can be considered as temporal constraints in constraint logic programming. Therefore, temporal reasoning can be treated as a constraint solving problem, to which a number of constraint solving techniques (Jaffar and Maher 1996) can be applied. We use a solver implemented using constraint handling rules (CHR) (Frühwirth 1998) as demonstrated in (Frühwirth 1994).

Temporal granularity research has developed logic- (Montanari 1996) and algebra-based (Bettini 2000) techniques for representing and reasoning about granularities and user-defined calendars. Conversions between granularities (Bettini *et al.* 2003) will be useful in dealing with heterogeneity in temporal entities.

Data integration. Approaches to achieving data integration largely fall into tight-coupling and loose-coupling categories depending on whether a global schema is used (Firat 2003, Goh *et al.* 1999). In these approaches, intended data semantics in sources and receivers are explicitly incorporated in either the view definitions or the user queries. The computation complexity (Lenzerini 2002) in rewriting user queries for the former approach and the user's burden of writing complex queries for the latter limit the flexibility and scalability of these approaches.

COIN (Firat 2003, Goh *et al.* 1999) is a middle ground approach that avoids these shortcomings by encapsulating data semantics into a context theory and maintaining accessibility of source schema by users. In COIN, a user issues queries as if all sources are in the user's context and a mediator is used to automatically rewrite the queries to resolve semantic differences.

Unfortunately, existing approaches, including COIN, use a static data model and ignore temporal context. Consequently, temporal concepts are missing in the ontologies used in these systems. In a seemingly relevant research of medical data integration (Nguyen et al. 1999, Shahar 1994), interpretation contexts are used to derive diagnostic conclusions from temporal data. For example, “having a fever” conclusion can be derived from a series of high temperatures in a regular interpretation context. But it cannot be derived in another interpretation context where doses of a certain drug indicate the use of a certain treatment protocol that has a side effect of causing high temperatures. Although interpretation contexts are constructed from temporal data, they do not vary over time once constructed.

This research focuses on the representation and reasoning of temporal context. Our framework incorporates context into the query evaluation process to automatically detect and reconcile temporal semantic conflicts. By combining the concepts and techniques from the three relevant research areas, we develop a scalable solution to temporal heterogeneity problems.

4. Temporal COIN approach

Below we provide a brief description of the COIN framework, followed by the descriptions of the extensions of temporal representation and reasoning to enable COIN to handle temporal semantic heterogeneity.

4.1. The COIN framework

The COIN framework consists of a deductive object-oriented data model, based on F-logic (Kiffer *et al.* 1995), for context knowledge representation and a mediation service module that detects and resolves semantic conflicts in user queries at run-time (see Figure 2).

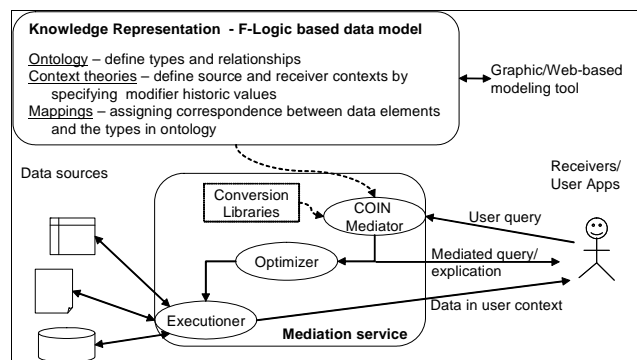


Figure 2. Architecture of COIN system

Knowledge representation in COIN consists of three components:

Ontology – to define a common semantic domain that consists of semantic data types and their relationships. A type corresponds to a concept in the problem domain and can be related to another in three ways: (1) as a subtype or super-type (e.g., *profit* is a subtype of *monetary value*); (2) as a named attribute (e.g., *temporal entity* such as year is a temporal attribute of *monetary value*), whose values are usually in the data sources; and (3) as a *modifier*, also known as a contextual attribute (Sciore *et al.*, 1994), whose value is specified in context axioms and can functionally determine the interpretation of instances of the type that has this modifier (e.g., *monetary value* type has a *scale factor* modifier). There is a distinguished type *basic* in the ontology that serves as the super type of all the other types and represents all primitive data types. Objects are instances of the semantic types;

Context theories – to specify the values of modifiers for each source or receiver and the conversions for transforming an attribute value in one context to another. The context of each source or receiver is uniquely identified with a context label. The value specification for modifiers can be a simple value assignment or a set of rules that specify how to obtain the value. Conceptually a context can be thought to be a set of $\langle \text{modifier}, \text{object} \rangle$ pairs, where *object* (i.e., the value of the modifier) is a singleton in most non-temporal cases; and

Semantic mappings – to establish correspondences between data elements in sources and the types in the ontology, e.g., *profit* attribute in the example of Table 1 corresponds to *profit* type, which is in turn of *monetary value* type by inheritance. The semantic mappings are also called *elevation rules*, or *elevations* for short.

The core component in the mediation service module is the COIN mediator implemented in abductive constraint logic programming (ACLP) (Kakas *et al.* 2000). It takes a user query and produces a set of mediated queries (*MQs*) that resolve semantic differences. This happens by first translating the user query into a Datalog query and using the encoded knowledge to derive the *MQs* that incorporate necessary conversions from source contexts to receiver context. The query optimizer and processor optimize the *MQs* using a simple cost model and the information on source capabilities, obtain the data, perform the conversions, and return final datasets to the user.

In addition to providing a rigorous theoretical foundation to our solution, ACLP has two other advantages: (1) abductive reasoning can produce explanations, which we use to generate *MQs* as the *intensional* answers to the original user query; and (2) the availability of efficient concurrent constraint solvers allows for good performance of the mediator.

The original COIN used a snapshot data model that did not allow temporal context representation; the mediator also lacked temporal reasoning capability.

4.2. Temporal extensions to representation

The extended framework admits temporal data sources, which are assumed to be relational with an explicit timestamp attribute in their schema. They accept SQL queries with usual comparison operators ($=$, $>$, $<$, *etc.*) on timestamp domain.

The ontology is augmented with temporal concepts as defined in the time ontology (Hobbs 2002). The most general one is *Temporal Entity*, which can be further specialized as *Instant* or *Interval*. Each element in the source schema is mapped to a corresponding semantic data type in the ontology by an *elevation axiom*. A timestamp can be elevated to Temporal Entity or a subtype. For types whose values are time dependent, we relate them to a temporal entity type via *temporal attribute*.

Definition The *temporal context* of a data source or a receiver is a specification of the history of all modifiers in the ontology. Conceptually, it can be considered to be a set of $\langle \text{modifier}, \text{history} \rangle$ pairs, where *history* is a set of $\langle \text{object}, \text{valid_interval} \rangle$ pairs.

Recall that in existing COIN, a context is simply a set of $\langle \text{modifier}, \text{object} \rangle$ pairs. The temporal extension allows us to represent the entire history of each modifier. If the value does not change over time, the *history* set is simply a singleton with the *valid_interval* covering entire time span of interest. We achieve backward compatibility by treating $\langle \text{modifier}, \text{object} \rangle$ as the shorthand for $\langle \text{modifier}, \{ \langle \text{object}, \text{entire_time_span} \rangle \} \rangle$.

Temporal entity type may also have modifiers, *e.g.*, granularity, time zone, *etc.*, to account for various contexts.

4.3. Temporal extensions to mediation

Given a user query expressed in the user's context, the mediator detects and reconciles semantic differences by examining involved modifiers and applying appropriate conversion functions if the values differ between any source and the receiver. With temporal extensions, modifiers are no longer singly valued. However, at any point in time, there is only one valid value for each attribute. The mediator needs to find the maximum time interval over which all involved modifiers are singly valued. Over this interval, an *MQ* can be generated as in the case of existing COIN; the interval appears in the *MQ* as additional constraints over the attribute of temporal entity type.

The mediator translates *history* set for modifiers into temporal constraints, which are posted into a constraint store concurrently solved by solvers in CHR. Through back tracking, all intervals over which modifiers are singly valued are found.

In our framework, contexts are declaratively defined using First Order Logic (FOL) rules. This can be done because of the equivalence of F-logic and FOL. The mediator is a general abductive reasoner. When new sources are added, we only

need to add context and semantic mapping rules for them. External functions can be called to convert between contexts. These features lend COIN great flexibility and scalability.

5. Prototype and preliminary results

We are able to solve a range of temporal heterogeneity problems exemplified in Section 2 by implementing a fraction of the suggested extensions.

5.1. Representation and mediation

Representation. Figure 3 shows the ontology for the example. Here, we use the most general concept *temporal entity*. Using elevation axioms, we create the mappings between attributes in the data source and the types in the ontology.

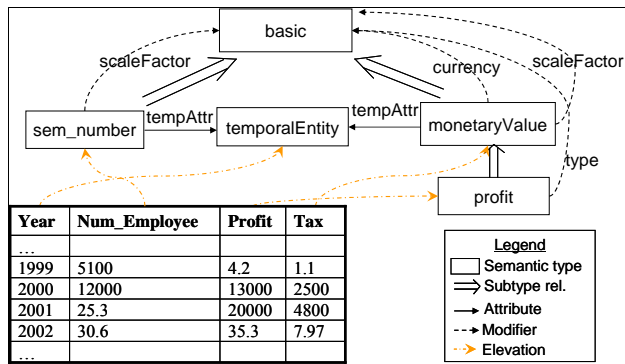


Figure 3. Example ontology and elevations

We model the time line as discrete and unbounded with both points and intervals as primitives. The past and future infinities are represented by constants *bottom* and *top*. We implement the \leq relation between points as a *tle* (i.e., temporal less than or equal to) constraint. The *contains* relation between an interval and a point is translated into *tle* constraints; the *overlaps* relation between intervals are also translated into *tle* constraints.

This simple model has sufficient expressive power to represent the temporal knowledge needed in Table 2. For example, internally we use the following Prolog statements to represent the source context for currency:

```

modifier(monetaryValue,O,currency,c_src,M):-
    containObj([bottom, 2000], O),
    cste(basic, M, c_src, "FRF").
modifier(monetaryValue,O,currency,c_src,M):-
    containObj([2001, top],O),
    cste(basic, M, c_src, "EUR").

```

The head of the statement reads: O of type *monetaryValue* has a modifier *currency*, whose value in source context c_src is M . Its body has two predicates. *containObject(I, O)* uses the *tempAttr* of O to obtain its temporal attribute T (which corresponds to *Year* attribute in the data source) of type *temporalEntity* and add constraint *contains(I, T)*. The helper predicate *cste* specifies the primitive value of M in c_src context. Thus, the history of each modifier is now a set of pairs $\langle V_i, I_i \rangle$, where $\bigcup I_i = [bottom, top]$.

For context that does not change over time, we could have used $[bottom, top]$ interval in *containObj* predicate. Since the translated constraints are always true, we will not include this predicate for this case.

Mediation. As described earlier, the mediation service needs to find the maximum interval over which all modifiers are singly valued. Figure 4 helps visualize this task by graphically representing the context knowledge in Table 2. For example, $[bottom, 1999]$ is such an interval where the source context can be described with a set of singly valued modifiers:

```

{<monetaryValue.currency, "FRF">,
 <monetaryValue.scale, "1000000">,
 <profit.type, "taxExcluded">,
 <sem_num.scale, "1">}.

```

	c_src: source context				c_target: receiver context			
	monetaryValue currency	profit scale	sem_num type	sem_num scale	monetaryValue currency	profit scale	sem_num type	sem_num scale
1999	FRF	1M	Excl Tax	1				
2000		1K			USD	1K	Incl Tax	1K
2001			Incl Tax					
2002	EUR	1M		1K				

Figure 4. Visualization of temporal contexts

Recall that we translate all temporal relations into *tle* constraint over points. Each modifier generates two *tle* constraints for the temporal variable. The above problem

is thus turned into to a problem of solving the constraints generated by all the modifiers, which is solved concurrently using a solver implemented in CHR.

Constraints over *bottom* and *top* can be removed using simplification rules so that these two literals do not appear in the list of abducted predicates. Constraints over other time points can be pair-wise simplified. We also implement *overlaps* to simplify *tle* constraints over four points at a time. These rules tighten the bounds of the temporal variable or signify a failure if inconsistencies are found.

Together with the rules that handle equality constraint, this point-based temporal constraint solver covers the 13 relations for temporal intervals in (Allen 1983). Relations *before*, *after*, *meets*, and *met_by* generate a failure, all the rest relations are subsumed into *overlaps*.

Through backtracking, the recursive algorithm finds all intervals over which modifiers are singly valued. Conversions are applied as in the case of existing COIN implementation. This simple temporal constraint-based extension transforms a temporal context problem into a set of non-temporal problems, thereby allows us to reuse the non-temporal implementation of the COIN mediator.

5.2. Preliminary results

These temporal extensions to COIN framework enable semantic interoperability for the integration example. The prototype can generate *MQs* that reconcile temporal context differences.

As an example, suppose a user in the receiver context wants to retrieve data from the company time series relation named *Financials* in Table 1 using the following SQL query:

```
Select Year, Num_Employee, Profit
From Financials;
```

and expects the returned data to be in his context. The query is translated into a well formed Datalog query in our prototype. The extended COIN mediator takes this query and the representation of the integration as input, and produces the following mediated query in Datalog (which COIN eventually converts to SQL):

```
%1. =<1999: currency,scale,type;scale
answer('V32', 'V31', 'V30') :-
  'V29' is 'V28' * 1000.0, 'V31' is 'V27' * 0.001,
  olsen("FRF", "USD", 'V26', 'V32'),
  'V28' is 'V25' * 'V26',
  financials('V32', 'V27', 'V25', 'V24'),
  'V32' =< 1999, 'V23' is 'V24' * 'V26',
  'V22' is 'V23' * 1000.0, 'V30' is 'V29' + 'V22'.
%2. 2000: currency and type;scale
answer(2000, 'V21', 'V20') :-
  'V21' is 'V19' * 0.001, 'V20' is 'V15' + 'V14',
```

```

financials(2000, 'V19', 'V18', 'V17'),
olsen("FRF", "USD", 'V16', 2000),
'V15' is 'V18' * 'V16', 'V14' is 'V17' * 'V16'.
%3. 2001: currency;scale
answer(2001, 'V13', 'V12') :-
'V13' is 'V11' * 0.001, 'V12' is 'V10' * 'V8',
financials(2001, 'V11', 'V10', 'V9'),
olsen("EUR", "USD", 'V8', 2001).
%4. >=2002: currency,scale;none
answer('V7', 'V6', 'V5') :-
olsen("EUR", "USD", 'V4', 'V7'),
financials('V7', 'V6', 'V3', 'V2'),
2002 =< 'V7', 'V1' is 'V3' * 'V4',
'V5' is 'V1' * 1000.0.

```

The mediated query has four subqueries, each resolves a set of semantic conflicts that exist in the time specified by the timestamp attribute. Note that *olsen* predicate corresponds to a currency conversion data source introduced by the conversion function for *currency* modifier. These subqueries resolve all the semantic conflicts in Table 2 or in Figure 4.

6. Discussion and future plan

We identified three types of semantic heterogeneity in the integration of temporal data. There is an ever increasing need to efficiently handle temporal heterogeneity as more historical data is used for auditing, forecasting, investigation, and many other purposes. We have described temporal extensions to the COIN framework. A prototype of the extensions shows that our approach is capable of solving temporal context problems. A demonstration of solving the problems of the historic stock prices example can be found in (Madnick and Zhu 2006). With its declarative knowledge representation and its capability of dynamically composing data conversions, this approach has been shown to have good flexibility and scalability (Zhu and Madnick 2004, Zhu and Madnick 2006).

Our future research aims to develop this approach in several aspects. Current representation of temporal context explicitly compares an interval with the temporal attribute of an object. The representation may be made cleaner by using an annotated temporal constraint logic (Frühwirth 1996). We need to investigate how this logic can be integrated with the ACLP based COIN mediator.

An important part of future research will be focused on the heterogeneity of temporal entities. We plan to add various modifiers to the temporal entity type in the ontology and use external functions to convert between contexts. If this is not expressive enough to represent the diversity of time, a richer time ontology may be necessary. We also need to incorporate metric temporal reasoning, which often involves computations of one or more calendars. We will investigate the feasibility of leveraging web services like those in (Bettini 2003). This is a challenging and

important research area because misunderstanding date and time can have serious consequences, as history has shown in an 1805 event (Madnick 1999) where the Austrian troops were forced to surrender largely because of the misunderstanding of a date in two different calendar systems.

Acknowledgements: The work reported herein has been supported, in part, by MITRE Corporation, MIT-MUST project, the Singapore-MIT Alliance, and Suruga Bank.

References

- Allen J. F. "Maintaining knowledge about temporal intervals", *Communications of the ACM*, vol. 26 no. 11, 1983, p. 832-843.
- Bettini C. "Web services for time granularity reasoning", TIME-ICTL'03, Los Alamitos, CA, 2003, p. 2-4.
- Bettini C., Jajodia S., Wang X. S., *Time granularities in databases, data mining, and temporal reasoning*, Springer, 2000.
- Date C. J., Darwen H., Lorentzos N. A., *Temporal data and the relational model*, Morgan Kaufmann Publishers, 2003.
- Firat A., Information integration using contextual knowledge and ontology merging, PhD Thesis, MIT, 2003.
- Frühwirth T., Temporal reasoning with constraint handling rules, ECRC-94-5, 1994.
- Frühwirth T., "Temporal annotated constraint logic programming", *Journal of Symbolic Computation*, vol. 22 no.5-6, 1996, p. 555-583.
- Frühwirth T., "Theory and practice of constraint handling rules", *Journal of Logic Programming*, vol. 37 no. 1, 1998, p. 95-138.
- Goh H., Bressan S., Madnick S., Siegel M., "Context interchange: new features and formalisms for the intelligent integration of information", *ACM TOIS*, vol. 17 no. 3, 1999, p. 270-293.
- Hobbs J. R., A DAML ontology of time, LREC, 2002.
- Jaffar J., Maher M. J., "Constraint logic programming: a survey", *Journal of Logic Programming*, vol. 19/20, 1999, p. 503-581.
- Jensen, C.S., Dyreson C.E., Böhlen M., Clifford J., Elmasri R., Gadia, S.K., Grandi F., Hayes P., Jajodia S., Käfer W., Kline N., Lorentzos N., Mitsopoulos Y., Montanari A., Nonen D., Peressi E., Pernici B., Roddick J.F., Sarda N.L., Scalas M.R., Segev A., Snodgrass R.T., Soo M.D., Tansel A., Tiberio P., Wiederhold G., "The consensus glossary of temporal database concepts-February 1998 version", LNCS 1399, 1998, p. 367-405.
- Kakas C., Michael A., Mourlas C., "ACLP: integrating abduction and constraint solving," *Journal of Logic Programming*, vol. 44 no. 1-3, 2000, p. 129-177.

- Kiffer M., Laussen G., Wu J., “Logic foundations of object-oriented and frame-based languages”, *Journal of the ACM*, vol. 42 no. 4, 1995, p. 741-843.
- Lenzerini M., “Data integration: a theoretical perspective” 21st ACM SIGMOD-SIGACT symposium on Principles of database systems, 2002, Madison, Wisconsin, p.233-246.
- Madnick S. E., “Metadata Jones and the Tower of Babel: the challenge of large-scale semantic heterogeneity”, IEEE Meta-Data Conference, April 6-7, 1999, p.1-13.
- Madnick S.E., Zhu H., “Improving data quality with effective use of data semantics”, *Data and Knowledge Engineering*, vol. 59 no. 2, 2006, p. 460-475.
- Meiri I., “Combining qualitative and quantitative constraints in temporal reasoning”, *Artificial Intelligence*, vol. 87, no. 1-2, 1996, p. 343-385.
- Montanari A., “Metric and layered temporal logic for time granularity”, PhD Thesis, University of Amsterdam, 1996.
- Nguyen H., Shahar Y., Tu S. W., Das A. K., Musen M. A., “Integration of temporal reasoning and temporal-data maintenance into a reusable database mediator to answer abstract, time-oriented queries: the Tzolkín system”, *Journal of Intelligent Information Systems*, vol. 13 no. 1-2, 1999, p. 121-145.
- Özsoyoglu G., Snodgrass R. T., “Temporal and real-time databases: a survey”, *IEEE Transactions on Knowledge and Data Engineering*, vol. 7 no. 4, 1995, p. 513-532.
- Sciore E. Siegel M., Rosenthal A., “Using semantic values to facilitate interoperability among heterogeneous information systems”, *ACM Transaction on Database Systems*, vol. 19 no. 2, 1994, p. 254-290.
- Shahar Y., A knowledge-based method for temporal abstraction of clinical data, Stanford University, 1994
- Shoham Y., “Temporal logics in AI: semantical and ontological considerations”, *Artificial Intelligence*, vol. 33 no. 1, 1987, p. 89-104.
- Zhu H., Madnick S., “Context interchange as a scalable solution to interoperating amongst heterogeneous dynamic services”, *3rd Workshop on eBusiness (WEB'04)*, December 11, 2004, Washington, D.C
- Zhu H., Madnick S., “A lightweight ontology approach to scalable interoperability”, *VLDB Workshop on Ontologies-based techniques or DataBases and Information Systems (ODBIS'06)*, September 11, 2006, Seoul, Korea