

# Annotating and Reasoning about Time and Events

**Jerry Hobbs**

USC/ISI  
4676 Admiralty Way  
Marina del Rey, CA 90292  
hobbs@isi.edu

**James Pustejovsky**

Dept. of Computer Science  
Brandeis University  
Waltham, MA 02454  
jamesp@cs.brandeis.edu

## Abstract

In this paper we discuss the relationship between TimeML, a rich specification language for event and temporal expressions in text, and the interpretation of these expressions in a temporal semantics. Specifically, we propose to demonstrate how a TimeML markup of text is interpreted within the DAML-Time Ontology and time framework of Hobbs (2002). We demonstrate the expressiveness of TimeML in English text, and discuss the relevance to content specification in the semantic web.

## Introduction

The Semantic Web will necessarily require specification of temporal information. For example, someone who does a Web search trying to find a place to buy a book needed before next Tuesday may or may not be able to use an online bookstore that promises delivery within five business days. Someone doing a genealogical search may want to specify that the birthdate of a person is between 15 and 45 years before a known marriage date. In response to this need, in connection with the DARPA-sponsored DAML program, we have been developing an ontology of temporal concepts, DAML-Time, that covers the basic topological temporal relations on instants and intervals, measures of duration, and the clock and calendar (Hobbs, 2002). This will enable the statement of and reasoning about the temporal aspects of the contents and capabilities of web resources and of the needs of users.

But most of the information on the Web is in natural language, and there is no chance that it will ever be marked up for semantic retrieval if that has to be done by hand. Natural language programs will have to process the contents of web pages to produce annotations. Remarkable progress has been made in the last decade in the use of statistical techniques for analyzing text. However, these techniques for the most part depend on having large amounts of annotated data, and annotations require an annotation scheme. It is for this reason that the ARDA-funded AQUAINT program has spon-

sored the development of an mark-up language TimeML for temporal information in texts.

These annotations are most relevant to the Semantic Web enterprise if the annotation scheme meshes well with temporal ontologies used in the Semantic Web. The aim of this paper is to define just such a mapping. In Section 2 of this paper we present TimeML. In Section 3 we present the relevant parts of DAML-Time. In Section 4 we discuss how TimeML is interpreted within DAML-Time in order to provide a temporal grounding for texts that will facilitate reasoning over events, their orderings, and their relative granularity.

## TimeML

### The Need for Temporal Annotation

The AQUAINT program is a multi-project effort to improve the performance of question answering systems over free text, such as that encountered on the Web. An important component to this effort is the access of information from text through content rather than keywords. Named entity recognition [refs] has moved the fields of information retrieval and information exploitation closer to access by content, by allowing some identification of names, locations, and products in texts. Beyond these metadata tags (ontological types), however, there is only a limited ability at marking up text for real content. One of the major problems that has not been solved is the recognition of events and their temporal anchorings. In this paper, we report on an AQUAINT project to create a specification language for event and temporal expressions in text.

Events in articles are naturally anchored in time within the narrative of a text. For this reason, temporally grounded events are the very foundation from which we reason about how the world changes. Without a robust ability to identify and extract events and their temporal anchoring from a text, the real “aboutness” of the article can be missed. Moreover, since entities and their properties change over time, a database of assertions about entities will be incomplete or incorrect if it does not capture how these properties are temporally updated. To this end, event recognition drives basic inferences from text.

For example, currently questions such as those shown below are not supported by question answering systems.

1. a. Is Gates currently CEO of Microsoft?
- b. When did Iraq finally pull out of Kuwait during the war in the 1990s?
- c. Did the Enron merger with Dynegy take place?

What characterizes these questions as beyond the scope of current systems is the following: they refer, respectively, to the temporal aspects of the properties of the entities being questioned, the relative ordering of events in the world, and events that are mentioned in news articles, but which have never occurred.

There has recently been a renewed interest in temporal and event-based reasoning in language and text, particularly as applied to information extraction and reasoning tasks (cf. Mani and Wilson, 2000, *ACL Workshop on Spatial and Temporal Reasoning*, 2001, Gaizauskas and Setzer, 2002). Several papers from the workshop point to promising directions for time representation and identification (cf. Filatova and Hovy, 2001, Schilder and Habel, 2001, Setzer, 2001). Many issues relating to temporal and event identification remain unresolved, however, and it these issues that TimeML was designed to address. Specifically, four basic problems in event-temporal identification are addressed:

- (a) Time stamping of events (identifying an event and anchoring it in time);
- (b) Ordering events with respect to one another (lexical versus discourse properties of ordering);
- (c) Reasoning with contextually underspecified temporal expressions (temporal functions such as *last week* and *two weeks before*);
- (d) Reasoning about the persistence of events (how long does an event or the outcome of an event last).

The specification language, TimeML, is designed to address these issues, in addition to handling basic tense and aspect features.

## The Basics of TimeML

Unlike most previous attempts at event and temporal specification, TimeML separates the representation of event and temporal expressions from the anchoring or ordering dependencies that may exist in a given text. There are four major data structures that are specified in TimeML (Ingria and Pustejovsky, 2002, Pustejovsky et al., 2002): **EVENT**, **TIMEX3**, **SIGNAL**, and **LINK**. These are described in some detail below. The features distinguishing TimeML from most previous attempts at event and time annotation are summarized below:

1. Extends the TIMEX2 annotation attributes;
2. Introduces **Temporal Functions** to allow intensionally specified expressions: *three years ago*, *last month*
3. Identifies signals determining interpretation of temporal expressions;

- (a) Temporal Prepositions: *for*, *during*, *on*, *at*;
  - (b) Temporal Connectives: *before*, *after*, *while*.
4. Identifies all classes of event expressions;
    - (a) Tensed verbs: *has left*, *was captured*, *will resign*;
    - (b) Stative adjectives: *sunken*, *stalled*, *on board*;
    - (c) Event nominals: *merger*, *Military Operation*, *Gulf War*;
  5. Creates dependencies between events and times:
    - (a) Anchoring: *John left on Monday*.
    - (b) Orderings: *The party happened after midnight*.
    - (c) Embedding: *John said Mary left*.

TimeML considers “events” (and the corresponding tag **<EVENT>**) a cover term for situations that *happen* or *occur*. Events can be punctual or last for a period of time. We also consider as events those predicates describing *states* or *circumstances* in which something obtains or holds true. Not all stative predicates are marked up, however, as only those states which participate in an opposition structure in a given text are marked up. Events are generally expressed by means of tensed or untensed verbs, nominalizations, adjectives, predicative clauses, or prepositional phrases. The specification of **EVENT** is shown below:

```
attributes ::= eid class tense aspect
eid ::= ID
{eid ::= EventID
EventID ::= e<integer>}
class ::= 'OCCURRENCE' | 'PERCEPTION' | 'REPORTING'
        | 'ASPECTUAL' | 'STATE' | 'I_STATE' |
        'I_ACTION' | 'MODAL'
tense ::= 'PAST' | 'PRESENT' | 'FUTURE' | 'NONE'
aspect ::= 'PROGRESSIVE' | 'PERFECTIVE' |
          'PERFECTIVE_PROGRESSIVE' | 'NONE'
```

Examples of each of these event types are given below:

1. **Occurrence**: *die*, *crash*, *build*, *merge*, *sell*
2. **State**: *on board*, *kidnapped*, *love*, ..
3. **Reporting**: *say*, *report*, *announce*,
4. **I-Action**: *attempt*, *try*, *promise*, *offer*
5. **I-State**: *believe*, *intend*, *want*
6. **Aspectual**: *begin*, *finish*, *stop*, *continue*.
7. **Perception**: *See*, *hear*, *watch*, *feel*.

The **TIMEX3** tag is used to mark up explicit temporal expressions, such as times, dates, durations, etc. It is modelled on both Setzer’s (2001) **TIMEX** tag, as well as the **TIDES** (Ferro, et al. (2002)) **TIMEX2** tag. There are types of **TIMEX3** expressions:

- (a) Fully Specified Temporal Expressions, *June 11, 1989*, *Summer, 2002*;
- (b) Underspecified Temporal Expressions, *Monday*, *Next month*, *Last year*, *Two days ago*;
- (c) Durations, *Three months*, *Two years*.

```
attributes ::= tid type [functionInDocument]
                [temporalFunction]
                (value | valueFromFunction) [mod]
                [anchorTimeID | anchorEventID]
```

```

tid ::= ID
{tid ::= TimeID}
TimeID ::= t<integer>
type ::= 'DATE' | 'TIME' | 'DURATION'
functionInDocument ::= 'CREATION_TIME' | 'EXPIRATION_TIME' |
  'MODIFICATION_TIME' | 'PUBLICATION_TIME' |
  'RELEASE_TIME' | 'RECEPTION_TIME' | 'NONE'
temporalFunction ::= 'true' | 'false'
{temporalFunction ::= boolean}
value ::= CDATA
{value ::= duration | dateTime | time | date |
  gYearMonth | gYear | gMonthDay |
  gDay | gMonth}
valueFromFunction ::= IDREF
{valueFromFunction ::= TemporalFunctionID}
TemporalFunctionID ::= tf<integer>
anchorTimeID ::= IDREF
{anchorTimeID ::= TimeID}
anchorEventID ::= IDREF
{anchorEventID ::= EventID}

```

The optional attribute, `functionInDocument`, indicates the function of the `TIMEX3` in providing a temporal anchor for other temporal expressions in the document. If this attribute is not explicitly supplied, the default value is "NONE". The non-empty values take their names from the temporal metadata tags in the Prism draft standard (available at [www.primstandard.org/](http://www.primstandard.org/)).

The treatment of temporal functions in TimeML allows any time-value dependent algorithms to delay the computation of the actual (ISO) value of the expression. The following informal paraphrase of some examples illustrates this point, where DCT is the Document Creation Time of the article.

1. *last week* = (`predecessor (week DCT)`): That is, we start with a temporal anchor, in this case, the DCT, coerce it to a week, than find the week preceding it.
2. *last Thursday* = (`thursday (predecessor (week DCT))`): Similar to the preceding expression, except that we pick out the day named 'thursday' in the predecessor week.
3. *the week before last* = (`predecessor (predecessor (week DCT))`): Also similar to the first expression, except that we go back two weeks.
4. *next week* = (`successor (week DCT)`): The dual of the first expression: we start with the same coercion, but go forward instead of back.

SIGNAL is used to annotate sections of text, typically function words, that indicate how temporal objects are to be related to each other. The material marked by SIGNAL constitutes several types of linguistic elements: indicators of temporal relations such as temporal prepositions (e.g. *on*, *during*) and other temporal connectives (e.g. *when*) and subordinators (e.g. *if*). The functionality of the SIGNAL tag was introduced by Setzer (2001). In TimeML it also marks polarity indicators such as *not*, *no*, *none*, etc., as well as indicators of temporal quantification such as *twice*, *three times*, and so forth. The specification for SIGNAL is given below:

```

attributes ::= sid
sid ::= ID
{sid ::= SignalID}
SignalID ::= s<integer>

```

To illustrate the application of these three tags, consider the example annotation shown below.

*John left 2 days before the attack.*

```

John
<EVENT eid="e1" class="OCCURRENCE" tense="PAST"
aspect="PERFECTIVE">
left
</EVENT>
<MAKEINSTANCE eiid="ei1" eventID="e1"/>
<TIMEX3 tid="t1" type="DURATION" value="P2D"
temporalFunction="false">
2 days
</TIMEX3>
<SIGNAL sid="s1">
before
</SIGNAL>
the
<EVENT eid="e2" class="OCCURRENCE" tense="NONE"
aspect="NONE">
attack
</EVENT>
<MAKEINSTANCE eiid="ei2" eventID="e2"/>

```

## LINKS

The set of LINK tags encode the various relations that exist between the temporal elements of a document. There are three types of link tags.

- **TLINK:**

A TLINK or Temporal Link represents the temporal relationship holding between events or between an event and a time, and establishes a link between the involved entities making explicit if they are: *simultaneous*, *before*, *after*, *immediately before*, *immediately after*, *including*, *holds*, *beginning*, and *ending*.

- **SLINK:**

An SLINK or Subordination Link is used for contexts introducing relations between two events, or an event and a signal. SLINKs are of one of the following sorts: *Modal*, *Factive*, *Counter-factive*, *Evidential*, *Negative evidential*, and *Negative*.

- **ALINK:**

An ALINK or Aspectual Link represents the relationship between an aspectual event and its argument event. Examples of the aspectual relations to be encoded are: *initiation*, *culmination*, *termination*, and *continuation*.

Below, we present the specification for each link relation.

- **TLINK:**

```

attributes ::= (eventInstanceID | timeID)
               [signalID] (relatedtoEvent
               | relatedtoTime) relType [magnitude]
eventInstanceID ::= ei<integer>

```

```

timeID ::= t<integer>
signalID ::= s<integer>
relatedToEvent ::= ei<integer>
relatedToTime ::= t<integer>
relType ::= 'BEFORE' | 'AFTER' | 'INCLUDES' |
            'IS_INCLUDED' | 'HOLDS' | 'SIMULTANEOUS' |
            'IAFTER' | 'IBEFORE' | 'IDENTITY' | 'BEGINS' |
            'ENDS' | 'BEGUN_BY' | 'ENDED_BY'
magnitude ::= t<integer>

```

To illustrate the function of this link, let us return to the sentence above,

*John left 2 days before the attack.*

now adding the annotation of the TLINK, which orders the two events mentioned in the sentence, with a magnitude denoted by the temporal expression.

```

<TLINK eventInstanceID="ei1" signalID="s1"
relatedToEvent="ei2"
relType="BEFORE" magnitude="t1"/>

```

#### SLINK:

```

attributes ::= [eventInstanceID]
              (subordinatedEvent |
               subordinatedEventInstance)
              [signalID] relType [polarity]
eventInstanceID ::= ei<integer>
subordinatedEvent ::= e<integer>
subordinatedEventInstance ::= ei<integer>
signalID ::= s<integer>
relType ::= 'MODAL' | 'NEGATIVE' | 'EVIDENTIAL' |
            'NEG_EVIDENTIAL' | 'FACTIVE' |
            'COUNTER_FACTIVE'

```

A modally subordinating predicate such as *want* is typed as introducing a SLINK, as shown below.

*Bill wants to teach on Monday.*

```

Bill
<EVENT eid="e1" class="I_STATE" tense="PRESENT"
aspect="NONE">
wants
</EVENT>
<MAKEINSTANCE eiid="ei1" eventID="e1"/>
<SLINK eventInstanceID="ei1" signalID="s1"
subordinatedEvent="e2" relType="MODAL"/>
<SIGNAL sid="s1">
to
</SIGNAL>
<EVENT eid="e2" class="OCCURRENCE" tense="NONE"
aspect="NONE">
teach
</EVENT>
<MAKEINSTANCE eiid="ei2" eventID="e2"/>
<SIGNAL sid="s2">
on
</SIGNAL>
<TIMEX3 tid="t1" type="DATE" temporalFunction="true"
value="XXXX-WXX-1">
Monday
</TIMEX3>
<TLINK eventInstanceID="ei2" relatedToTime="t1"
relType="IS_INCLUDED"/>

```

#### ALINK:

```

attributes ::= eventInstanceID [signalID]
              relatedToEvent relType
eventInstanceID ::= ei<integer>
signalID ::= s<integer>
eventID ::= e<integer>
relType ::= 'INITIATES' | 'CULMINATES' |
            'TERMINATES' | 'CONTINUES'

```

To illustrate the behavior of ALINKs, notice how the aspectual predicate *begin* is a separate event, independent of the modified event, where the “phase” is introduced in the relation within the ALINK.

*The boat began to sink.*

```

The boat
<EVENT eid="e1" class="ASPECTUAL" tense="PAST"
aspect="NONE">
began
</EVENT>
<MAKEINSTANCE eiid="ei1" eventID="e1"/>
<SIGNAL sid="s1">
to
</SIGNAL>
<EVENT eid="e2" class="OCCURRENCE" tense="NONE"
aspect="NONE">
sink
</EVENT>
<ALINK eventInstanceID="ei1" signalID="s1"
relatedToEvent="e2" relType="INITIATES"/>

```

We should point out that the ALINK captures aspectual phase information associated with an event, and is logically (temporally) distinct from the BEGIN relType in the TLINK above. For example, in a sentence *The boat began to sink when the torpedo hit it.*, the torpedo hitting the boat is TLINKed to the sinking event through the relType BEGIN, while the beginning of the sinking is ALINKed through the phase relType INITIATE.

In order to provide an interpretation of the TimeML specification of event and temporal expressions described above, we will adopt the DAML-Time ontology for time as a model.

## DAML-Time

DAML-Time is an ontology of temporal concepts, for describing the temporal content of Web pages and the temporal properties of Web services. Its development is being informed by temporal ontologies developed at a number of sites and is intended to capture the essential features of all of them and make them and their associated resources easily available to a large group of Web developers and users. In this paper we specify the ontology in predicate calculus; on the [daml.org](http://daml.org) Web site there is a DAML+OIL specification of those parts of DAML-Time that are easily expressible in DAML+OIL, and a complete specification, due to George Ferguson, in an encoding of logic in RDF/DAML developed by McDermott et al. (2001).

## Topological Temporal Relations

**Instants and Intervals** There are two subclasses of temporal-entity: *instant* and *interval*.

$$\begin{aligned} \text{instant}(t) &\supset \text{temporal-entity}(t)^1 \\ \text{interval}(T) &\supset \text{temporal-entity}(T) \end{aligned}$$

*start-of* and *end-of* are functions from temporal entities to instants.

$$\begin{aligned} \text{temporal-entity}(T) &\supset \text{instant}(\text{start-of}(T)) \\ \text{temporal-entity}(T) &\supset \text{instant}(\text{end-of}(T)) \end{aligned}$$

For convenience, we can say that the start and end of an instant is itself.

$$\begin{aligned} \text{instant}(t) &\supset \text{start-of}(t) = t \\ \text{instant}(t) &\supset \text{end-of}(t) = t \end{aligned}$$

*inside* is a relation between an instant and an interval.

$$\text{inside}(t, T) \supset \text{instant}(t) \wedge \text{interval}(T)$$

This concept of *inside* is not intended to include starts and ends of intervals.

*interval-between* is a relation among a temporal entity and two instants.

$$\begin{aligned} \text{interval-between}(T, t_1, t_2) \\ &\supset \text{temporal-entity}(T) \wedge \text{instant}(t_1) \\ &\quad \wedge \text{instant}(t_2) \end{aligned}$$

The two instants are the start and end points of the temporal entity.

$$\begin{aligned} \text{interval-between}(T, t_1, t_2) \\ \equiv \text{start-of}(T) = t_1 \wedge \text{end-of}(T) = t_2 \end{aligned}$$

The ontology is silent about whether the interval from  $t$  to  $t$ , if it exists, is identical to the instant  $t$ .

The ontology is silent about whether intervals *consist of* instants.

The ontology is silent about whether intervals are uniquely determined by their starts and ends.

We can define a proper interval as one whose start and end are not identical.

$$\begin{aligned} \text{proper-interval}(t) &\equiv \\ &\text{interval}(t) \wedge \text{start-of}(t) \neq \text{end-of}(t) \end{aligned}$$

The ontology is silent about whether there are any intervals that are not proper-intervals.

**Before** There is a *before* relation on temporal entities, which gives directionality to time. If temporal entity  $T_1$  is before temporal entity  $T_2$ , then the end of  $T_1$  is before the start of  $T_2$ . Thus, before can be considered to be basic to instants and derived for intervals.

$$\begin{aligned} \text{before}(T_1, T_2) \\ \equiv \text{before}(\text{end-of}(T_1), \text{start-of}(T_2)) \end{aligned}$$

The end of an interval is not before the start of the interval.

<sup>1</sup>A note on notation: Conjunction ( $\wedge$ ) takes precedence over implication ( $\supset$ ) and equivalence ( $\equiv$ ). Formulas are assumed to be universally quantified on the variables appearing in the antecedent of the highest-level implication.

$$\begin{aligned} \text{interval}(T) \\ \supset \text{before}(\text{end-of}(T), \text{start-of}(T)) \end{aligned}$$

The start of a proper interval is before the end of the interval.

$$\begin{aligned} \text{proper-interval}(T) \\ \supset \text{before}(\text{start-of}(T), \text{end-of}(T)) \end{aligned}$$

If one instant is before another, there is an interval between them.

$$\begin{aligned} \text{instant}(t_1) \wedge \text{instant}(t_2) \wedge \text{before}(t_1, t_2) \\ \supset (\exists T) \text{interval-between}(T, t_1, t_2) \end{aligned}$$

The ontology is silent about whether there is an interval from  $t$  to  $t$ .

If an instant is inside a proper interval, then the start of the interval is before the instant, which is before the end of the interval. The converse is true as well.

$$\begin{aligned} \text{instant}(t) \wedge \text{proper-interval}(T) \\ \supset [\text{inside}(t, T) \\ \equiv \text{before}(\text{start-of}(T), t) \\ \quad \wedge \text{before}(t, \text{end-of}(T))] \end{aligned}$$

Intervals are contiguous with respect to the *before* relation, in that an instant between two other instants inside an interval is inside the interval.

$$\begin{aligned} \text{before}(t_1, t_2) \wedge \text{before}(t_2, t_3) \\ \wedge \text{inside}(t_1, T) \wedge \text{inside}(t_3, T) \\ \supset \text{inside}(t_2, T) \end{aligned}$$

The *before* relation is anti-symmetric and transitive.

$$\begin{aligned} \text{before}(T_1, T_2) &\supset \neg \text{before}(T_2, T_1) \\ \text{before}(T_1, T_2) \wedge \text{before}(T_2, T_3) \\ &\supset \text{before}(T_1, T_3) \end{aligned}$$

The relation *after* is defined in terms of *before*.

$$\text{after}(T_1, T_2) \equiv \text{before}(T_2, T_1)$$

The ontology is silent about whether time is linearly ordered.

**Interval Relations** The relations between intervals defined in Allen's temporal interval calculus (Allen and Kautz, 1997) can be defined in a straightforward fashion in terms of *before* and identity on the start and end points. We illustrate this with the relations *int-meets* and *int-finishes*.

$$\begin{aligned} \text{interval}(T_1) \wedge \text{interval}(T_2) \\ \supset [\text{int-meets}(T_1, T_2) \\ \equiv \text{end-of}(T_1) = \text{start-of}(T_2)] \\ \text{interval}(T_1) \wedge \text{interval}(T_2) \\ \supset [\text{int-finishes}(T_1, T_2) \\ \equiv \text{before}(\text{start-of}(T_2), \text{start-of}(T_1)) \\ \quad \wedge \text{end-of}(T_1) = \text{end-of}(T_2)] \end{aligned}$$

**Linking Time and Events** The time ontology links to other things in the world through the predicates *at-time* and *during*. We assume that another ontology provides for the description of events—either a general ontology of event structure abstractly conceived, or specific, domain-dependent ontologies for specific domains.

The term “eventuality” will be used to cover events, states, processes, propositions, states of affairs, and anything else that can be located with respect to time. The possible natures of eventualities would be spelled out in the event ontologies.

The predicate *at-time* relates an eventuality to an instant, and is intended to say that the eventuality holds, obtains, or is taking place at that time.

$$at-time(e, t) \supset instant(t)$$

The predicate *during* relates an eventuality to an interval, and is intended to say that the eventuality holds, obtains, or is taking place during that interval.

$$during(e, T) \supset interval(T)$$

If an eventuality obtains during an interval, it obtains at every instant inside the interval.

$$during(e, T) \wedge inside(t, T) \supset at-time(e, t)$$

Whether a particular process is viewed as instantaneous or as occurring over an interval is a granularity decision that may vary according to the context of use, and is assumed to be provided by the event ontology.

The event ontology could extend temporal functions and predicates to apply to events in the obvious way, e.g.,

$$ev-start-of(e) = t \\ \equiv time-span-of(T, e) \wedge start-of(T) = t$$

This would not be part of the time ontology, but would be consistent with it.

## Measuring Durations

**Temporal Units** This development assumes ordinary arithmetic is available. We can consider temporal units to constitute a set of entities—call it TemporalUnits—and have a function *duration* mapping Intervals  $\times$  TemporalUnits into the Reals.

$$duration([5 : 14, 5 : 17], *Minute*) = 3$$

The arithmetic relations among the various units can be stated in axioms like

$$duration(T, *Second*) = 60 * duration(T, *Minute*)$$

The relation between days and months (and, to a lesser extent, years) is specified as part of the ontology of clock and calendar below. On their own, however, month and year are legitimate temporal units.

**Concatenation and Hath** The multiplicative relations above don’t tell the whole story of the relations among temporal units. Temporal units are *composed of* smaller temporal units. A larger temporal unit is a concatenation of smaller temporal units. We first defined a general relation of concatenation between an interval and a set of smaller intervals. Then a predicate *Hath* is defined that specifies the number of smaller unit intervals that concatenate to a larger interval.

Concatenation: A proper interval  $x$  is a concatenation of a set  $S$  of proper intervals if and only if  $S$  covers all of  $x$ , and all members of  $S$  are subintervals of  $x$

and are mutually disjoint. From this definition we can prove as theorems that there are elements in  $S$  that start and finish  $x$ ; and except for the first and last elements of  $S$ , every element of  $S$  has elements that precede and follow it.

*Hath*: The basic predicate for expressing the composition of larger intervals out of smaller temporal intervals of unit length is *Hath*, from statements like “30 days hath September” and “60 minutes hath an hour.” Its structure is

$$Hath(N, u, x)$$

meaning “ $N$  proper intervals of duration one unit  $u$  hath the proper interval  $x$ .” That is, if *Hath*( $N, u, x$ ) holds, then  $x$  is the concatenation of  $N$  unit intervals where the unit is  $u$ . For example, if  $x$  is some month of September then *Hath*(30, \*Day\*,  $x$ ) would be true.

*Hath* is defined as follows:

$$Hath(N, u, x) \equiv (\exists S)[card(S) = N \wedge \\ (\forall z)[member(z, S) \supset duration(z, u) = 1] \wedge \\ concatenation(x, S)]$$

That is,  $x$  is the concatenation of a set  $S$  of  $N$  proper intervals of duration one unit  $u$ .

We are now in a position to state the relations between successive temporal units, by means of axioms like the following:

$$duration(T, *Minute*) = 1 \supset Hath(60, *Second*, T)$$

The relations between months and days are dealt with in after the calendar has been characterized.

## Clock and Calendar Units

We take a day as a calendar interval to begin at and include midnight and go until but not include the next midnight. By contrast, a day as a duration is any interval that is 24 hours in length. The day as a duration was dealt with in the previous section. This section deals with the day as a calendar interval.

Including the beginning but not the end of a calendar interval in the interval may strike some as arbitrary. But we get a cleaner treatment if, for example, all times of the form 12:xx a.m., including 12:00 a.m. are part of the same hour and day, and all times of the form 10:15:xx, including 10:15:00, are part of the same minute.

For stating general properties about clock intervals, we use

$$clock-int(y, n, u, x)$$

This expression says that  $y$  is the  $n$ th clock interval of type  $u$  in  $x$ . For example, the proposition *clock-int*(10 : 03, 3, \*Minute\*, [10 : 00, 11 : 00]) holds. Here  $u$  is a member of the set of clock units, that is, one of \*Second\*, \*Minute\*, or \*Hour\*.

In addition, there is a calendar unit function with similar structure:

$$cal-int(y, n, u, x)$$

This says that  $y$  is the  $n$ th calendar interval of type  $u$  in  $x$ . For example, the proposition  $cal-int(12Mar2002, 12, *Day*, Mar2002)$  holds. Here  $u$  is one of the calendar units  $*Day*$ ,  $*Week*$ ,  $*Month*$ , and  $*Year*$ .

Each of the calendar intervals is that unit long; a calendar year is a year long.

$$cal-int(y, n, u, x) \supset duration(y, u) = 1$$

The distinction between clock and calendar intervals is because they differ in how they number their unit intervals. The first minute of an hour is labelled with 0; for example, the first minute of the hour [10:00,11:00] is 10:00. The first day of a month is labelled with 1; the first day of March is March 1. We number minutes for the number just completed; we number days for the day we are working on. Thus, if the larger unit has  $N$  smaller units, the argument  $n$  in  $clock-int$  runs from 0 to  $N-1$ , whereas in  $cal-int$   $n$  runs from 1 to  $N$ . To state properties true of both clock and calendar intervals, we can use the predicate  $cal-int$  and relate the two notions with the axiom

$$cal-int(y, n, u, x) \equiv clock-int(y, n-1, u, x)$$

DAML-Time includes a treatment of weeks and days of the week as well.

## Months and Years

The months have special names in English, as in

$$cal-int(y, 9, *Month*, x) \equiv September(y, x)$$

The number of days in a month have to be spelled out for individual months, by axioms of the following form:

$$\begin{aligned} September(m, y) \\ \supset (\exists S)Hath(S, 30, *Day*, m) \end{aligned}$$

The definition of a leap year is as follows:

$$\begin{aligned} (\forall z)[leap-year(y) \\ \equiv (\exists n, x)[year(y, n, (CE(z))) \\ \wedge [divides(400, n) \\ \vee [divides(4, n) \wedge \neg divides(100, n)]]]] \end{aligned}$$

Now the number of days in February can be specified.

$$\begin{aligned} February(m, y) \wedge leap-year(y) \\ \supset (\exists S)Hath(S, 29, *Day*, m) \\ February(m, y) \wedge \neg leap-year(y) \\ \supset (\exists S)Hath(S, 28, *Day*, m) \end{aligned}$$

Months can now be defined as a duration, namely, the duration from the  $n$ th day of one month to the  $n$  day of the next month.

To say that July 4 is a holiday in the United States one can write

$$\begin{aligned} (\forall d, m, y)[cal-int(d, 4, *Day* * m) \wedge July(m, y) \\ \supset holiday(d, USA)] \end{aligned}$$

## Time and Duration Stamps

Standard notation for times list the year, month, day, hour, minute, and second. It is useful to define a predication for this.

$$time-of(t, y, m, d, h, n, s, z)$$

For example, an instant  $t$  has the time

5:14:35pm PDT, October 4, 2002

if the following properties hold for  $t$ :

$$time-of(t, 2002, 10, 4, 17, 14, 35, *PDT*)$$

We can similarly define a predicate that specifies the duration of an interval in standard units.

$$duration-of(T, y, m, d, h, n, s)$$

It is straightforward to translate predications of  $time-of$  and  $duration-of$  into ISO 8601 format for dates, times, and durations.

## Interpreting TimeML in DAML-Time

The intended interpretation of the TimeML specifications can now be provided in terms of the DAML-Time ontology. This allows one to move back and forth between the two representations as required for linguistic analysis and reasoning.

There are some aspects of TimeML that relate not to a theory of time but a theory of the structure of events. Among these are aspect, such as “perfective”. We do not deal with these here.

The information in the SIGNAL tag is of importance in the automatic recovery of temporal relations, and therefore it is important that signals be annotated. But by the time annotation is complete, that information will have been incorporated into LINK tags, so it is not necessary to interpret signals with respect to the DAML-Time ontology.

The SLINK is about grammatical subordination, or predicate-argument structure, and the modal or cognitive relation between the dominating and subordinated events. Thus, it is not concerned with the ontology of time as such, although it is important in linking the subordinated possible event into the temporal structure of the whole text.

The TimeML attribute “functionInDocument” refers to a concept in a theory of documents and relates it to times. Essentially, it states an *at-time* relation between a document event, such as publication, and the time referred to by a temporal expression. It is required for the interpretation of deictic expressions such as “now” and “last week”. The treatment of deictic time will be addressed in a future release of DAML-Time.

The “tense” attribute of EVENT tags is expressed in terms of a *before* or equality relation, once the deictic anchor (“now”) is known.

The “value” in a TIMEX3 tag is a date, time, or duration in the ISO 8601 standard format, as extended by TIDES2. As pointed out above, there is a straightforward mapping between *time-of* predications and date and time formats, and between *duration-of* predications and duration formats. The “XX” notation in the ISO standard corresponds to existentially quantified variables in the logic.

The “relType” values of the TLINK tag correspond roughly to the interval relations in Allen’s interval calculus and can be defined either in terms of these or in a similar manner. For example, an IAFTER relation can be defined as the interval during which one event occurs meeting the interval during which the other occurs. IDENTITY is not a temporal relation but is expressed as equality involving reified events.

The “relType” values of the ALINK tag can be defined similarly. An event that INITIATES another, for example, occurs at the start of the interval during which the other occurs. The ALINK tag conveys something more than just temporal information. As mentioned above, the beginning of a sinking of a boat is not the same event as the event of a torpedo hitting the boat, even if they are simultaneous. This distinction, however, is not part of a theory of time but a theory of the structure of events. Similarly, the distinction between TERMINATES and CULMINATES is not a temporal one but something that would have to come from event theory.

The logical representation of the temporal information in the sentence

*John left 2 days before the attack.*

can now be read off the TimeML tags in the markup. Recall that the ordering between the events is encoded in a TLINK as:

```
<TLINK eventInstanceID="ei1" signalID="s1"
relatedToEvent="ei2" relType="BEFORE"
magnitude="t1"/>
```

Assume  $e_1$  is John’s leaving and  $e_2$  is the attack. If the TLINK relType is interpreted as *before*, and the magnitude attribute is associated with *duration*, then the DAML-Time representation of the temporal information from TimeML is as shown below.

$$\begin{aligned} &at-time(e_1, t_1) \wedge at-time(e_2, t_2) \\ &\wedge before(t_1, t_2) \wedge interval-between(T, t_1, t_2) \\ &\wedge duration(T, *Day*) = 2 \end{aligned}$$

## Relevance to Semantic Web

Significant efforts have been launched to annotate the temporal information in large textual corpora, resulting in the development of the mark-up language TimeML. Significant efforts have also taken place in the Semantic Web community for developing an ontology of time for expressing the temporal content of web sites and the temporal properties of web services, resulting in DAML-Time. There is much to be gained by linking these two efforts, and that is what we have begun to do in this paper.

**Acknowledgements** The authors would like to thank the other members of the TERQAS Working Group on TimeML for their contribution to the specification language presented here: Robert Ingria, Robert Gaizauskas, Graham Katz, Jose Castano, Andrea Setzer, Roser Sauri, Inderjeet Mani, Antonio Sanfilippo, Beth Sundheim, and Andy Latto.

This work was performed in support of the Northeast Regional Research Center (NRRC) which is sponsored by the Advanced Research and Development Activity in Information Technology (ARDA), a U.S. Government entity which sponsors and promotes research of import to the Intelligence Community which includes but is not limited to the CIA, DIA, NSA, NIMA, and NRO. It was also funded in part by the Defense Advanced Research Projects Agency as part of the DAML program under Air Force Research Laboratory contract F30602-00-C-0168.

## References

- Allen, J.F. (1984). Towards a general theory of action and time. *Artificial Intelligence* 23, pp. 123-154.
- Allen, James F., and Henry A. Kautz. 1985. “A Model of Naive Temporal Reasoning”, *Formal Theories of the Commonsense World*, ed. by Jerry R. Hobbs and Robert C. Moore, Ablex Publishing Corp., pp. 251-268.
- Allen, J.F. and P.J. Hayes (1989). Moments and points in an interval-based temporal logic. *Computational Intelligence* 5, pp. 225-238.
- Allen, J.F. and G. Ferguson (1997). Actions and events in interval temporal logic. In Oliveiro Stock (ed.), *Spatial and Temporal Reasoning*, Kluwer Academic Publishers, pp. 205-245.
- Ferro, Lisa, Inderjeet Mani, Beth Sundheim and George Wilson, 2001, TIDES Temporal Annotation Guidelines. Version 1.0.2 MITRE Technical Report, MTR 01W0000041.
- Filatova, Elena and Eduard Hovy, 2001. “Assigning Time-Stamps To Event-Clauses” in *Proceedings of the ACL-2001 Workshop on Temporal and Spatial Information Processing*, ACL-2001. Toulouse, France, 6-11 July. pp. 88-95.
- Gaizauskas, Robert and Andrea Setzer, 2002, editors, *Annotation Standards for Temporal Information in Natural Language*, LREC 2002.
- Ingria, Robert and James Pustejovsky, 2002. “TimeML Specification 1.0”, available at <http://time2002.org>.
- McDermott, Drew, Jonathan Borden, Mark Burstein, Doug Smith, and Richard Waldinger, 2001. “A Proposal for Encoding Logic in RDF/DAML”, available at <http://www.cs.yale.edu/homes/dvm/daml/>.
- Mani, Inderjeet and George Wilson. 2000, “Robust Temporal Processing of News”, *Proceedings of the ACL’2000 Conference*, 3-6 October 2000, Hong Kong.
- Pustejovsky, James, Robert Gaizauskas, Roser Sauri, Andrea Setzer, Robert Ingria, 2002. “Annotation Guideline to TimeML 1.0”, available at <http://time2002.org>.
- Schilder, Frank and Christopher Habel, 2001. “From Temporal Expressions To Temporal Information: Semantic Tagging Of News Messages” in *Proceedings of*



*the ACL-2001 Workshop on Temporal and Spatial Information Processing*, ACL-2001. Toulouse, France, 6-11 July. pp. 65-72.

Setzer, Andrea, 2001. "Temporal Information in Newswire Articles: an Annotation Scheme and Corpus Study", PhD dissertation, University of Sheffield.