Back and forth through time and events

(Extended Abstract)

Patrick Blackburn,[1] Claire Gardent[2] and Maarten de Rijke[3]

[1] Department of Philosophy, Rijksuniversiteit Utrecht, Heidelberglaan 8, 3584 CS Utrecht.
Email: patrick@phil.ruu.nl

[2] GRIL, Université de Clermont Ferrand, France; and Department of Computational Linguistics, Universiteit van Amsterdam, Spuistraat 134, 1012 VB Amsterdam.
Email: claire@mars.let.uva.nl

[3] CWI, P.O. Box 4079, 1009 AB Amsterdam.
Email: mdr@cwi.nl

Abstract. In this extended abstract *back-and-forth structures* are defined and applied to the semantics of natural language. Back-and-forth structures consist of an event structure and an interval structure communicating via a relational link; transitions in the one structure correspond to transitions in the other. Such structures enable us to view temporal constructions (such as tense, aspect, and temporal connectives) as methods of moving systematically between information sources. We illustrate this with a treatment of the English present perfect, and progressive aspect, that draws on ideas developed in Moens and Steedman (1988).

1 Introduction

Formal accounts of temporal constructions in natural language often disagree about the semantic ontology to be assumed — should it be point based, interval based or event based? We think that more adequate analyses of natural language will be obtained by *combining* ontologies, not choosing between them. We illustrate this by combining interval structures with (various forms of) event structures into what we call *back-and-forth structures* (BAFs). These consist of an interval structure and an event structure linked by a relation so that transitions in the one correspond to transitions in the other.

Such combined ontologies enable us to build our analyses round the following intuition: temporal constructions are means of systematically exploiting links between information sources. Consider the English present perfect. It is common to informally gloss this construction as 'a past tense of present relevance'. For example, '*John has gone to the store*' means that at some past time John went to the store and, moreover, that John's excursion is somehow of relevance to the present context. We see two important transitions here: a move backwards in time through an interval structure, and a move to an associated event in an event structure. The English present perfect coordinates these transitions, and BAFs enable us to model this.

Much of this abstract uses BAFs to explore the ideas of Moens and Steedman (1988); indeed, BAFs developed by thinking about the kind of machinery required to formalise their work. Moens and Steedman provide a wide ranging account of temporal semantics (topics considered include tense, temporal reference, aspect and adverbial modification) couched as a Winograd-style procedural semantics. Their work hinges on (at least) the following ideas: that non-temporal relations between events must be admitted if an adequate account is to be given of the semantics of '*when*' and various aspectual phenomena; that there are key event configurations (called 'nuclei') underlying the richness of event ontology; and that adverbial (and other forms of) modification are to be accounted for in terms of 'type coercion'. The Moens and Steedman account is attractive because while it is wide ranging, its explanations reduce to the interaction of a handful of intuitive ideas. Its weakness is that it is largely unformalised. We believe BAFs provide a setting in which substantial parts of their account can be made precise. BAFs can be seen as a way of modeling the insight that a systematic interplay between temporal and non-temporal relations is called for, and by progressively enriching the event structures they are built over one can model ever more of the Moens and Steedman system.

We proceed as follows. We first discuss the semantics of the English present perfect, indicating why the use of combined ontologies seems promising. We then introduce *simple BAFs.* These consist of interval structures combined with an extremely simple type of eventuality structure. Although such structures are too simple to cope with all the subtleties of natural language, their use permits the central idea underlying our proposal to be clearly presented. Following this, we (slightly) enrich the eventuality component to form *sorted BAFs.* This enables us to refine our discussion of the present perfect, and to provide an analysis of progressive aspect that does not run foul of the so-called imperfective paradox. We close the abstract by briefly discussing how we are extending this work, and noting other BAF-like proposals we have found in the literature.

2 The Present Perfect

While descriptive work on the English present perfect abounds, the construction has been notoriously resistant to formal analysis. In this section we discuss the problems the present perfect gives rise to, and argue that these indicate the need for combined ontologies.

It is often argued that the English present perfect is used to describe past events of present relevance. Perhaps the most well-known account of this intuition is that described in Reichenbach (1947), where a present perfect is analysed as describing a past event (the event temporally precedes the speech time) whose reference time coincides with the speech time. Reichenbach's reference point is meant to be the time talked about; or, in other words, the temporal perspective from which the described event is viewed. By insisting that reference and speech time coincide, the present perfect is analysed as relevant to the present. This contrasts with the simple past which is viewed as describing a past event whose reference time coincides with the event time rather than with the speech time.

Although Reichenbach's approach goes one step toward capturing the intuition underlying the use of the present perfect, two problems remain. First, what is the nature of reference times, and how are they determined? Second, the Reichenbachian account fails to account for many observations made in the literature concerning the restrictions governing the use of the present perfect. For instance, it does not explain why the sentence in (1) is infelicitous if uttered at a time occurring after the coffee has been cleaned.

(1) I have spilled my coffee.

Similarly, it does not account for the restrictions placed by verbal aspectual classes on the use of the present perfect, for example:

(2) a. ? The house has been empty (stative expression)
    b. ? I have worked in the garden (process expression)
    c. ? The star has twinkled (point expression).

Example (2a) shows that the present perfect is awkward in combination with stative expressions; (2b) and (2c) illustrate its awkwardness in combination with process expressions and point expressions, respectively.

As Moens and Steedman (1986) convincingly argue, these problems can be resolved if the internal structure of events is taken into account. Briefly, the idea is that an event (or *nucleus* in Moens and Steedman's terminology) is a tripartite structure consisting of a *preparatory phase*, a *culmination* and a *consequent state.* Given such a structure, the function of the present perfect is to situate the reference time in the consequent state of the core event being described (cf. Moens and Steedman (1986), p.20). Thus instead of the Reichenbach schema

```
       E              R, S
   ---|-------------|-----
```

Moens and Steedman describe the present perfect by means of the following diagram:

```
                E         R, S
  |----------|----------------
       PP    Cul          CS
```

Their account incorporates the central Reichenbachian intuition, while eliminating its problematic aspects:

- The reference point is given a (more) precise and more motivated location in time, namely within the time stretch of the consequent state.
- Example (1) is explained as follows. An obvious consequence of spilling one's coffee is that coffee is spilled. Under the Moens and Steedman theory, uttering a sentence in the present perfect indicates (i) that the reference time coincides with the speech time and (ii) that both these times are included

in the time stretch of the consequent state. Thus by uttering the present perfect (1), the speaker indicates that coffee is still spilled. Hence the oddity of (1) in a context where it isn't.

- The ill-formedness of the examples in (2) is explained by the fact that stative, process and point expressions are used to describe either states (i.e. unstructured entities) or these parts of the event structure which do not include the consequent state.[4] Since these expressions do not involve the notion of consequent state, they cannot be used in the present perfect whose semantics is defined in terms of this very notion.

The Moens and Steedman approach is intuitively appealing: how can it be made precise? We believe this can be done quite straightforwardly by combining ontologies.

Intuitively, their approach demands a mixture of ontologies: at the very least it seems to call for *temporal* structure, *eventuality* structure, and (crucially) a 'sensible fit' between these two ontologies. The 'past tense' component of the present perfect seems to require some notion of temporal structure; at the very least, this will involve some notion of temporal precedence. But this temporal structure does not suffice: in addition we need to invoke some notion of 'eventuality', and some sort of relation of 'relevance' between eventualities (for example, between the act of spilling the coffee, and the presence of the coffee on the floor). Intuitively this relevance relation isn't temporal; nonetheless, capturing the idea that we want an event of *present* relevance seems to presuppose that some sort of 'synchronisation' between the precedence relation on the temporal structure and the relevance relation on the eventuality structure is in force.

Actually, we will need even more structure than this. As examples (2a)–(2c) showed, the present perfect does not willingly combine with all verb types. We will need to work with a suitably fine-grained view of eventuality structure to capture these restrictions; in particular, by using eventuality structures *sorted* in a manner that reflects verbal aspectual classes we can model more of the Moens and Steedman account.

In the following two sections we will present simple formal models that capture these intuitions. We first present *simple* BAFs. These combine interval structures with a very simple notion of eventuality structure in a way that permits the intuition of 'present relevance' to be directly captured. (Or, to put it in the terminology of Moens and Steedman, they enable us to model the intuition that the present perfect works by locating the reference point in the run-time of consequent state induced by the eventuality being described.) We then refine this simple picture by enriching the eventuality structures used to make BAFs. This allows us to model the aspectual restrictions governing the use of the present perfect, and yields a simple solution to the imperfective paradox.

---

[4] These aspectual notions are discussed in more detail in section 4.

3 Simple BAFs

Simple BAFs consist of four components: an *interval structure*, an *eventuality structure*, and (most importantly) two *links* between them.

An *interval structure* **I** is a triple $\langle I, <, \sqsubseteq \rangle$ as defined in van Benthem (1991). Here I is a set of intervals, $<$ is the precedence relation, and $\sqsubseteq$ is the subinterval relation. We work with *linear*, *atomic* interval structures. That is, we assume that given any two intervals either one precedes the other or they overlap, and that our structures contain minimal, 'point-like' intervals.

An *eventuality* structure of signature $\mathcal{E}$ is (for the purposes of the present section) a triple $\mathbf{O} = \langle O, \texttt{GRiTo}, \{P_e\}_{e\in\mathcal{E}} \rangle$. Here O is a non-empty set, the set of *eventuality occurrences*; GRiTo is a binary relation on O; and all the P_e are unary relations on O. We assume $\mathcal{E} \neq \emptyset$. If e GRiTo e' then we say *e gives rise to e'*. The unary relations P_e can be thought of as 'eventualities' for example *runnings*, *jumpings* and *recitings of poems*.

Now the crucial step. A *back-and-forth structure* (BAF) of signature $\mathcal{E}$ is a quadruple $\langle \mathbf{O}, z, \mathcal{Z}, \mathbf{I} \rangle$, where **O** is an eventuality structure of signature $\mathcal{E}$, **I** is an interval structure, z is a function from O to I that returns the run-time or temporal extent of an eventuality and that preserves the relation GriTo: if e GriTo e' then $z(e) < z(e')$. That is, z is an order-preserving morphism from the eventuality structure to the interval structure; it is this morphism that synchronizes the two ontologies. $\mathcal{Z}$ is the relation with domain O and range I defined by $e\mathcal{Z}i$ iff $i \sqsubseteq z(e)$. That is, we assume that all eventualities are *downward persistent* to subintervals.

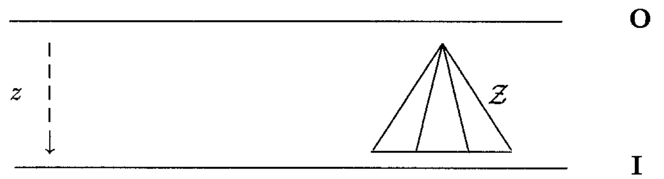

We now formulate a toy language for talking about BAFs: its vocabulary consists of all the items in $\mathcal{E}$, which we shall write as p, q, r, ...etc., and call *eventuality symbols*, and an operator PERF. If α is an eventuality symbol then PERF α is well formed (and nothing else is). Obviously it would be possible to add the Boolean operators and allow arbitrary embeddings of PERF; but while this leads to fairly interesting logical territory, it has little relevance to the semantics of natural language.

Now for the semantics. Let **B** $(= \langle \mathbf{O}, z, \mathcal{Z}, \mathbf{I} \rangle)$ be a BAF. Then, for all intervals i, and all eventuality symbols q, we define:

$$
\begin{aligned}
\mathbf{B}, i \models \textsc{Perf}\, q \text{ iff } & \exists i' \exists e' \exists e \big(i' < i \;\& \\
& i' = z(e') \;\& \\
& e' \in P_q \;\& \\
& e' \;\texttt{GRiTo}\; e \;\& \\
& e\mathcal{Z}i\big).
\end{aligned}
$$

Consider what this does. Suppose we have a sentence in the present perfect, say '*Fire has broken out on the oil rig*'. In our toy language this takes the form:

$$\text{PERF}(\textit{Fire breaks out on the oil rig}).$$

If we evaluate this at an interval i in $\mathbf{B}$, then we must 'complete a square' in a BAF back to the utterance interval i. That is, we move back in time to an interval i' which is the run-time for an event e'; this e' is an eventuality of the correct type (that is, e' is a breaking out of a fire) and moreover e' gives rise to an event e which is $\mathcal{Z}$ related to our utterance interval i. Intuitively, the eventuality of present relevance e would be the ongoing burning of the fire, that is the consequent state of the breaking out of the fire event. Roughly, this semantics relates to Reichenbach and Moens and Steedman's approaches as follows: i is the time of speech (S), i' is the event time (E) and e is the consequent state induced by the event being described, namely e'. The Reichenbachian constraint according to which speech and reference times coincide is replaced by the Moens and Steedman intuition that the time stretch of the consequent state includes the speech time. In this way, we capture the intuition of present relevance which characterises the English present perfect.

4 Sorted BAFs

Simple BAFs have the virtue of making clear the fundamental idea underlying our approach, but they are very crude. To encode the aspectual restrictions placed on the use of the present perfect, and to model further temporal constructions such as the progressive, we need to say more about the relation between time and aspect. This is the object of the present section. We will insist that the eventuality structures used to make BAFs embody the sortal distinctions (and additional relations) demanded by the various verb classes. We start by motivating these additions.

Eventualities

On the basis of the tenses, aspects and adverbials with which they occur, we classify eventualities into five types; our classification is similar to the one of Carlson (1981) and Moens & Steedman (1988). First we distinguish between indefinitely extending eventualities which we call *states*, and eventualities with defined beginnings and ends called *events*. Sentence (3) describes a state:

(3) Her hair is black.

Events are subdivided into *atomic* and *extended* events, depending on whether or not their runtimes are an atomic interval.

To motivate a further subdivision of the extended events, compare sentences (4) and (5) below.

3. $\forall e\,(\texttt{Culm}(e) \rightarrow \exists e'\,(\texttt{Culm\_Proc}(e')\ \&\ e'\,\texttt{Compl}\,e))$.

Now, a *sorted BAF* is a BAF $\mathbf{B} = \langle \mathbf{O}, z, \mathcal{Z}, \mathbf{I} \rangle$, where $\mathbf{O}$ is a sorted eventuality structure in which the following additional conditions are satisfied:

4. $\forall e\,(\texttt{Point}(e) \rightarrow z(e)$ is an atomic interval).
   $\forall e\,(\texttt{Culm}(e) \rightarrow z(e)$ is an atomic interval).
5. $\forall e\,(\texttt{Proc}(e) \rightarrow z(e)$ is an non atomic bounded interval).
   $\forall e\,(\texttt{Culm\_Proc}(e) \rightarrow z(e)$ is an non atomic bounded interval).
6. $\forall e\,(\texttt{State}(e) \rightarrow z(e)$ is an non atomic, non bounded interval).
7. $\forall e, i\,(i \sqsubseteq z(e) \leftrightarrow e\mathcal{Z}i)$.

Item 4 says that points and culminations are atomic events, item 5 that processes and culminating processes are non atomic bounded eventualities and item 6 that states are non atomic, unbounded eventualities; the seventh item ensures that eventualities are *downward persistent*. Note that BAFs *do* distinguish between points and culminations; only culminations can enter into the Culm relation. Similarly, the Compl relation differentiates between processes and culminating processes.

Present perfect and sentence aspect

As we observed in section 2, not all verbs may be naturally used with the present perfect. More specifically, the examples in (2) show that stative, process and point expressions are awkward in combination with the present perfect. Now consider the semantics we propose for this tense: we require that the event talked about gives-rise-to some other eventuality (the consequent state) whose time stretch includes the speech time. Since we also insist that the 'gives-rise-to' partial function is only defined on culminations, this means that no interpretation can be assigned to a natural language sentence which has tense present perfect and aspectual category anything other than a culmination.[5] In this way, we capture the intuition that only those expressions which evoke a consequent state may be used in the present perfect.

Progressive aspect and the imperfective paradox

We will now examine progressive aspect using sorted BAFs. Following Kamp and Reyle (1993), we assume that the function of the English progressive is to focus attention on the (culminating) process of some eventuality. This idea can be captured as follows. First, we enrich our toy language by adding the operators

[5] This is clearly too strong, for given sufficient contextual support such combinations may be naturally interpretable. In the full version of the paper these readings are captured by adding a relation to eventuality structures that explicitly codes this contextual dependence between events. This addition seems to reflect the intentions of Moens and Steedman (cf. their discussion of the *enablement* relation) and is also needed to cope with the semantics of *when*.

Typical examples are:

(a) be green, know
(b) recognize, complete a paper
(c) hiccup, twinkle
(d) build a house, write a thesis
(e) play the piano, sleep, waste time

To sum up: the aspectual category of a sentence determines the sort of eventuality being described. Process, state and point expressions refer to some unstructured entity whereby a stative expression describes some unstructured event stretching over an unbounded period of time, a process expression some unstructured event stretching over a bounded period of time and a point expression some unstructured atomic event. In contrast, culminating process and culmination expressions are used to talk about structured events, that is events consisting of a culmination process, a culmination and a consequent state. Furthermore, it has often been argued (see Kamp and Reyle (1993) and Moens and Steedman (1988)) that the function of the English grammatical aspectual markers (such as the perfective and the progressive aspect) is to indicate which parts of the event structure are being referred to. Roughly, a progressive refers to the culminating process of a structured event and a perfect to its consequent state. In what follows, we show how sorted BAFs allow us to capture these intuitions.

Sorting eventuality structures

We now want to formalize the above ideas by extending our earlier simple BAFs. First, a *sorted eventuality structure* is a tuple

$$\mathbf{O} = \langle \mathtt{Point}, \mathtt{Culm}, \mathtt{Proc}, \mathtt{Culm\_Proc}, \mathtt{State}; \mathtt{GRiTo}, \mathtt{Compl}; \{P_e\}_{e \in \mathcal{E}} \rangle,$$

where `Point`, `Culm`, `Proc`, `Culm_Proc` and `State` are mutually disjoint domains whose elements are used to interpret the various aspectual categories described above. `GRiTo` is a specialization of the 'gives-rise-to' relation defined in Section 3; here we insist that it only relates culminations to other eventualities. `Compl` is a binary relation between culminating processes and culminations. We think of the completion relation `Compl` as a partial function: if it is defined for an event e, it picks out a preferred or default consequence among all the consequences of e. (As not all events which have a natural culmination actually reach it, we can only have a *partial* function here.) Conversely, we assume that for every culmination there is a culminating process whose completion is this culmination. (We interpret the relation 'has-as-a-culminating process' using the converse of `Compl`.) More precisely, sorted eventuality structures should satisfy the following conditions:

1. `GRiTo` is a function whose domain is `Culm`.
2. `Compl` is a partial function whose domain is a subset of `Culm_Proc` and whose range is `Culm`.

(4) Bert was writing a thesis.

(5) Bert was sleeping.

The difference between sentences such as (4) and sentences such as (5) has been observed by numerous authors, and is often couched in terms of accomplishments and activities, cf. Vendler (1967). We express this distinction between (4) and (5) by saying that the event reported in (4) has a natural *culmination*, viz. the completion of the thesis; (5) has no such culmination. Processes that tend to have culminations in this sense are said to be *culminating*. Both the accomplishments of Vendler (1967) and the culminat*ed* processes of Moens & Steedman (1988) are composite events, consisting of a culminating process and a culmination; we feel it is more natural to split those composites and refer explicitly to the completion relation between culminating processes and their culminations.

Corresponding to the above distinction between processes and culminating processes, we divide atomic events into points and culminations. They differ in that culminations describe the culmination of a structured event (or nucleus) whereas points simply describe isolated atomic events; as a result a culmination may be associated with a culminating process and a consequent state whereas points cannot. To understand this division consider sentences (6) and (7) below.

(6) Bert completed his thesis.

(7) Bert hiccupped.

Sentence (6) reports a culmination; its culminating process is the writing of the thesis, its consequent state a state where the thesis is completed. Without further 'world knowledge' no natural culminating process or consequent state can be associated with the point event of (7).

Here, then, is a scheme of the eventualities we distinguish:

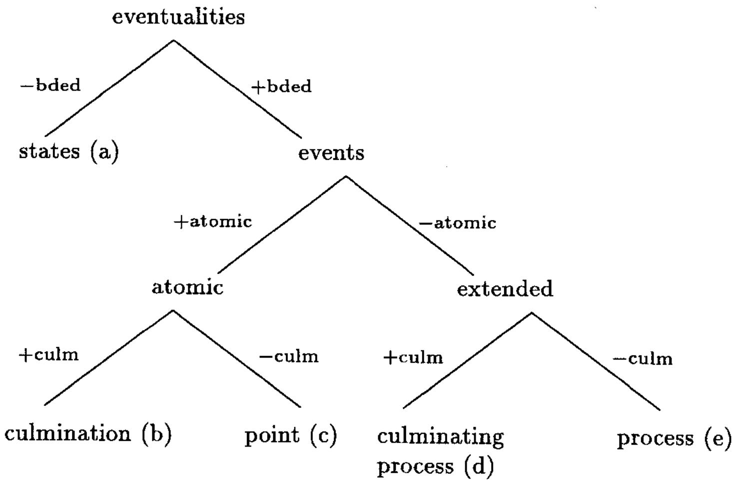

PAST and PROG, and allowing expressions of the form PAST q and PROG q and PASTPROG q to be well formed. As for the semantics, first, define $i \sqsubseteq^+ j$ to hold between two intervals i, j if the following is the case:

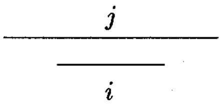

Let $\mathbf{B}$ $(= \langle \mathbf{O}, z, \mathcal{Z}, \mathbf{I} \rangle)$ be a sorted BAF. Then, for all intervals i, we define the relation $\mathbf{B}, i \models \phi$ as follows:

$\mathbf{B}, i \models \text{PROG}\, q$ iff $\exists e\, (e \in P_q \,\&\, (\texttt{Proc}(e) \vee \texttt{Culm\_Proc}(e)) \,\&\, i \sqsubseteq^+ z(e))$
$\mathbf{B}, i \models \text{PAST}\, q$ iff $\exists j, e\, (j < i \,\&\, e\mathcal{Z}j \,\&\, e \in P_q)$
$\mathbf{B}, i \models \text{PASTPROG}\, q$ iff $\exists j\, (j < i \,\&\, \mathbf{B}, j \models \text{PROG}\, q)$.

One of the merits of such a semantics for the progressive is that it yields a simple solution to the so-called 'imperfective-paradox'. Following Dowty (1979), this paradox has been discussed by numerous authors. Briefly, the paradox is this: how can we account for the meaning of a progressive sentence like (8) and (10) in such a way that (8) may be true without (9) ever becoming true, while on the other hand (10) would tautologically imply (11)?

(8) Bert was writing a thesis.

(9) Bert wrote a thesis.

(10) Bert was wasting valuable time and money.

(11) Bert wasted valuable time and money.

The key to a solution to the imperfective puzzle is the observation that there is an important difference between the pair of sentences (8), (9) and (10), (11): in asking whether (8) $\models$ (9) one asks whether a culminating process entails its culmination; in asking whether (10) $\models$ (11) the question is essentially whether processes are downward persistent. To be precise, *Bert's writing a thesis* is classified as a culminating process, and the culmination *Bert wrote a thesis* is its completion. According to our BAF account there is no contradiction in continuations of culminating processes that explicitly deny its culmination:

(12) Bert was writing a thesis, but he gave it up to join a heavy metal band.

Formally, in a sorted BAF failure of completion of a culminating process e is represented by the fact that the partial function `Compl` is not defined in e.

The above solves one half of the imperfective puzzle: (8) does not imply (9). How do we guarantee that (10) implies (11)? This is a simple consequence of clause 7 of the definition of a sorted BAF. Identifying *Bert's wasting* ... as a (non-culminating) process, we have for any sorted BAF $\mathbf{B}$, and any interval i in that BAF:

$$\mathbf{B}, i \models \text{PAST PROG}(Bert\ldots)$$
$$\text{iff } \exists j\,(j < i \;\&\; \mathbf{B}, j \models \text{PROG}(Bert\ldots))$$
$$\text{iff } \exists j, e\,(j < i \;\&\; e \in P_{Bert\ldots} \;\&\; \texttt{Proc}(e) \;\&\; j \sqsubseteq^{+} z(e)).$$

But this means that $j \sqsubseteq z(e)$, and hence $e\mathcal{Z}j$, and thus

$$\mathbf{B}, i \models \text{PAST}\,(Bert\ldots),$$

and (10) implies (11).

5 Conclusion

In this extended abstract we have sketched, in very simple terms, how combined ontologies can be used in the semantics of temporal constructions. To conclude we briefly discuss our ongoing work on richer, more realistic systems, and note other BAF-like proposals we have found in the literature.

Sorted BAFs incorporate some of the Moens and Steedman ideas, but a great deal remains to be done. For example, although the sorts and the `GRiTo` and `Compl` relations model something of the Moens and Steedman notion of subevent structure, they don't capture the important idea that this subevent structure is recursively formed out of entities called nuclei. A nuclei is essentially a little 'package' consisting of a culminating process, a culmination, and a consequent state. Sometimes one wants to look at the internal structure of such packages, and sometimes one wants to treat this package simply as a 'lump' which can be linked to other packages. We are currently working with what we term *nucleic* BAFs. These are BAFs in which the eventuality occurrences are recursively generated out of Moens and Steedman style nuclei. Using such structures makes it possible to give analyses of a number of phenomena: in particular, we have given a Moens and Steedman style analysis of adverbial modification, and moreover can account for the interaction of progressive and perfective aspect in a natural way. (This is a topic that Moens and Steedman do not consider.) We are working on the semantics of temporal connectives (such as '*when*' and '*until*') in the setting of nucleic BAFs. An important part of this work is to reconstruct in the (essentially static) BAF framework an analogue of the (essentially dynamic) notion of 'type coercion' used by Moens and Steedman.

But these are topics for the full version of the paper. What can be said at a more general level concerning the idea of using combined ontologies in the study of temporal semantics?[6] We find the approach appealing for a number of reasons. First, it is intuitive. Pre-theoretical talk is often couched in terms of a mixture of different sorts of entities and their interrelations. Rather than ignore these intuitions, it seems better to try and be precise about them. Second, it

[6] Actually, the idea of combining ontologies seems of importance in many other areas of applied logic as well; see Blackburn and de Rijke (1994) for further discussion.

seems to work. Formalisations couched in a single ontological setting tend to fare well with a handful of phenomena but can be extended only with difficulty. In contrast, we find the ease with which a wide range of phenomena can be modeled with BAFs striking. (We believe that most of the work of Moens and Steedman can be captured — and extended — in a manner that does no violence to its guiding intuitions.) Thirdly, the approach is, in a very useful sense of the word, conservative. It does not discard the work offered by point based, interval based or event based approaches: rather, it locates them in a richer setting. This retains what is good in earlier analyses, and lets the reasons for their shortcomings become clearly visible. To sum up, while BAFs as we have defined them here are only a crude approximation to the subtlety of temporal discourse, we feel that the underlying idea of combining ontologies will prove useful.

To close the abstract we briefly note some other multiple ontology or BAF-like approaches we are familiar with. First, Oversteegen (1989) analysed the semantics of various English and Dutch expressions in terms of certain moves between an 'objective' and a 'subjective' time flow. Although her structures differ from ours — the 'objective' flow is like an interval structure and the 'subjective' flow is a discrete time line — her approach has many ideas in common with ours. Tense, and perfective and progressive aspect are analysed in terms of a number of basic transition patterns between the structures. Her analysis of Dutch temporal constructions is quite detailed, and we think it would be interesting to formalise her discussion in terms of BAF-like structures.

Second a back-and-forth picture can be found in Seligman and ter Meulen (1992). This aspect of their work may not be immediately obvious, for most of their discussion is devoted to the construction of Dynamic Aspect Trees. Nonetheless, their idea of 'classifying interval frames' involves moving back-and-forth between two structures, and (we would argue) it is this that gives the needed flexibility to drive their dynamic system.

Lastly, our account seems to have affinities with Situation Semantics. This is clear if the Channel Theory initiated by Seligman (1990) is considered. In his terms we are using an interval structure to classify eventuality occurrences. Our treatment of the English present perfect essentially says that the peculiarities of the construction are due to the fact that it exploits this channel in a particularly strong way. More generally, Situation Semantics has long emphasized the importance of ontological diversity, and the way we evaluate formulas in BAFs could be regarded as an instance of their 'relational account' of meaning.

Acknowledgements. Patrick Blackburn and Maarten de Rijke would like to acknowledge the financial support of the Netherlands Organisation for Scientific Research (NWO), project NF 102/62-356 'Structural and Semantic Parallels in Natural Languages and Programming Languages'. Claire Gardent would like to acknowledge the financial support of the LRE 61-062 project 'Towards a Declarative Theory of Discourse'.

References

1. van Benthem, J.: 1991, *The Logic of Time*, sec. edn., Kluwer, Dordrecht.
2. Blackburn, P. and de Rijke, M.: 1994, Zooming in, zooming out. Unpublished manuscript.
3. Carlson, L.: 1981, Aspect and quantification, in Tedeschi, P.J. and Zaenen, A. (eds.): 1981, *Syntax and Semantics 14*, Academic Press, New York.
4. Kamp, H. and Reyle, U.: 1993, *From Discourse to Logic*, Kluwer, Dordrecht.
5. Moens, M, and Steedman, J.: 1986, Temporal information and natural language processing, Research paper EUCCS/RP-2, University of Edinburgh.
6. Moens, M, and Steedman, J.: 1988, Temporal Ontology and Temporal Reference, *Computational Linguistics*, **14**, pp. 15–28.
7. Oversteegen, L.: 1989, *Two track theory of time*, PhD thesis, University of Amsterdam.
8. Reichenbach, H.: 1947, *Elements of symbolic logic*, Free Press, New York.
9. Seligman, J.: 1990, *Perspectives: a Relativistic Approach to the Theory of Information*, PhD thesis, University of Edinburgh.
10. Seligman, J. and ter Meulen, A.: 1992, Dynamic Aspect Trees, in *Logic at Work*, CCSOM, Amsterdam.
11. Vendler, Z.: 1967, Verbs and times. Chapter 4 in *Linguistics and Philosophy*, Cornell U.P., Ithaca.