Temporal Prepositions and their Logic [*]

Ian Pratt-Hartmann

School of Computer Science
University of Manchester
Manchester M13 9PL, UK

Abstract

A fragment of English featuring temporal prepositions and the order-denoting adjectives first and last is defined by means of a context-free grammar. The phrase-structures which this grammar assigns to the sentences it recognizes are viewed as formulas of an interval temporal logic, whose satisfaction-conditions faithfully represent the meanings of the corresponding English sentences. It is shown that the satisfiability problem for this logic is NEXPTIME-complete. The computational complexity of determining logical relationships between English sentences featuring the temporal constructions in question is thus established.

Key words: natural language, temporal prepositions, interval temporal logic, computational complexity

1 Introduction

Consider the following sentences:

(1) An interrupt was received during every cycle
(2) The main process ran after the last cycle
(3) While the main process ran, an interrupt was received before loop 1 was executed for the first time.

These sentences speak of events and their temporal locations: of what happened and when. The principal devices they employ to encode this information are temporal prepositions and the adjectives first and last. The aim of this paper is to answer the question: What is the computational complexity of determining logical relationships between sentences encoding temporal information using such devices?

This question is of theoretical interest, because the events mentioned in (1)–(3)—cycles, executions of processes, receipts of interrupts—are *extended* in time; and temporal logics which deal with extended events—so-called *interval temporal logics*—typically exhibit high computational complexity. Given that the syntax of these logics has little affinity with that of temporal expressions in English, it is natural to ask whether the meanings of sentences such as (1)–(3) can be captured in a computationally manageable logic. The formal semantics of temporal constructions in English have been investigated by a succession of researchers (Dowty, 1979; Stump, 1985; Crouch and Pullman, 1993; Kamp and Reyle, 1993; Hwang and Schubert, 1994; Ogihara, 1996; ter Meulen, 1996). Yet in none of these accounts are the issues of expressive power and computational complexity to the fore. Indeed, many treatments of the semantics of temporal constructions in English represent sentence-meanings in a first-order language having variables which range over time-intervals and predicates which correspond to event-types and temporal order-relations—a logic which is easily shown to be undecidable. Given the recent surge of interest in logical fragments of limited computational complexity, this situation is unsatisfactory. There are evident practical and theoretical reasons for presenting the semantics of natural language constructions, where possible, using formal systems of limited expressive power.

The plan of this paper is as follows. Section 2 outlines the semantics of the English temporal constructions considered in this paper. Section 3 then uses a simple context-free grammar to define a fragment of English featuring these constructions; we call this fragment $\mathcal{TPE}$, an acronym for *temporal preposition English*. We show how the phrase-structures assigned to $\mathcal{TPE}$-sentences by this grammar can in fact be viewed as expressions in an interval temporal logic, which we call $\mathcal{TPL}$. Section 4 presents formal semantics for $\mathcal{TPL}$. Sections 5 and 6 provide matching upper and lower complexity-bounds for $\mathcal{TPL}$-satisfiability, showing that this problem is NEXPTIME-complete.

The following terminology and notation will be used throughout. We take a (time) *interval* to be a closed, bounded, convex (non-empty) subset of the real line. We denote the set of intervals by $\mathcal{I}$, and we use the (possibly decorated) letters I, J, ..., as variables ranging over $\mathcal{I}$. Observe that intervals may be punctual. If I and J denote the intervals $[a, b]$ and $[c, d]$, respectively, with $a, b, c, d \in \mathbb{R}$ and $a \leq c \leq d \leq b$, we let the terms $\text{init}(J, I)$ and $\text{fin}(J, I)$ denote the intervals $[a, c]$ and $[d, b]$, respectively. In other words, whenever $J \subseteq I$ is

true, we take $\mathrm{init}(J, I)$ to denote the initial segment of I up to the beginning of J, and $\mathrm{fin}(J, I)$ to denote the final segment of I from the end of J. More standardly, the symbol $\subset$ always denotes the strict subset relation, and $\subseteq$ the corresponding non-strict relation. Finally, we occasionally employ the definite quantifier $\iota x(\phi, \psi)$ with the standard (Russellian) semantics.

2   Semantics

In this section, we consider the semantics of the temporal constructions featured in the fragment of English defined below—principally, the temporal prepositions. Here, we follow modern usage and count temporal subordinating conjunctions as temporal prepositions taking clausal (rather than nominal) complements. We defer a formal specification of the fragment in question to Section 3, and the algorithmic derivation of sentence-meanings to Section 4.

2.1   Temporal preposition-phrases: basic semantics

Consider the following sentences:

(4) An interrupt was received
(5) An interrupt was received during every cycle
(6) An interrupt was received during every cycle until the main process ran
(7) After the initialization phase, an interrupt was received during every cycle until the main process ran.

Sentence (4) asserts that, within some contextually specified interval of interest, there is an interval over which an interrupt was received. Interpreting the unary predicate int-rec so that it is satisfied by all and only those time intervals over which an interrupt was received, we may thus represent the meaning of (4) by the formula

(8) $\exists J_0(\text{int-rec}(J_0) \wedge J_0 \subset I)$.

Notice that the temporal context to which the quantification in (4) is limited is represented by the free variable I in (8). That is: the meaning of (4) is a *temporal abstract*, receiving a truth-value (in an interpretation) only relative to a time interval. Viewing sentence meanings in this way greatly simplifies the semantics of temporal preposition-phrases.

Sentence (5) asserts that, within the given temporal context, every interval over which a cycle occurs includes some interval over which an interrupt was

3

received. Interpreting the unary predicate cyc so that it is satisfied by all and only those time intervals over which a cycle occurs, we may thus represent the meaning of (5) by the formula

(9) $\forall J_1(\mathrm{cyc}(J_1) \wedge J_1 \subset I \rightarrow \exists J_0(\mathrm{int\text{-}rec}(J_0) \wedge J_0 \subset J_1))$.

The normal type in (9) indicates the material contributed by the temporal preposition-phrase during every cycle, and the light type the material contributed by the sentence An interrupt was received, which it modifies. Observe that this material in light type is identical to the formula (8), except that the free temporal context variable has been bound by a quantifier introduced by the temporal preposition-phrase. On this view, the temporal preposition-phrase functions semantically as a modal operator, mapping one temporal abstract to another.

Sentences (6)–(7) can now be treated analogously. Making use of the notation introduced at the end of Section 1, and helping ourselves to a suitable signature of unary predicates of intervals, we may plausibly represent these sentences' truth-conditions as, respectively,

(10) $\quad \iota J_2(\mathrm{main}(J_2) \wedge J_2 \subset I,$
$\qquad \forall J_1(\ \mathrm{cyc}(J_1) \wedge J_1 \subset \mathrm{init}(J_2, I) \rightarrow \exists J_0(\mathrm{int\text{-}rec}(J_0) \wedge J_0 \subset J_1)))$

(11) $\quad \iota J_3(\mathrm{init\text{-}phase}(J_3) \wedge J_3 \subset I,$
$\qquad \iota J_2(\mathrm{main}(J_2) \wedge J_2 \subset \mathrm{fin}(J_3, I),$
$\qquad\quad \forall J_1(\mathrm{cyc}(J_1) \wedge J_1 \subset \mathrm{init}(J_2, \mathrm{fin}(J_3, I)) \rightarrow$
$\qquad\qquad \exists J_0(\mathrm{int\text{-}rec}(J_0) \wedge J_0 \subset J_1))))$.

We pass over the usual issues as to the faithfulness of the Russellian interpretation of definite quantification (either expressed or implied) in these sentences. Again, the normal type in (10) and (11) indicates the material contributed by the newly-added temporal preposition-phrases in (6) and (7) respectively, and the light type the material contributed by the sentences they modify. Again, this colouring scheme highlights the fact that the successive temporal preposition phrases function semantically as modal operators, binding the temporal context variables associated with the sentences they modify. This cascading quantification, typical of iterated temporal preposition phrases, was pointed out in Pratt and Francez (2001), and is discussed further in von Stechow (2002).

The fragment of temporal English considered here deals only with events, as opposed to states—that is, only with telic as opposed to atelic eventualities (Vendler 1967; see Steedman 1996 for an extended discussion). The thesis that all simple, event-reporting sentences are implicitly existentially quantified was proposed by Davidson (1967), and is defended in Parsons (1990). These authors take the quantification in question to be over events rather than time intervals; but this issue may be ignored for present purposes. A

4

recent collection of papers on this topic can be found in Higginbotham et al. (2000). One could doubtless quibble about whether the $\subset$ in (8)–(11) should be $\subseteq$; however, the operative concepts seem too vague for this issue to admit of resolution.

We drew attention above to the fact that the formulas (8)–(11) feature a free variable representing a temporal context. This naturally suggests an alternative representation using a propositional modal logic in which formulas are evaluated relative to time-intervals, and event-types are represented by propositional variables. Suppose, for example, such a logic features the modal operator $\langle D \rangle$, where $\langle D \rangle \phi$ is taken to be true at an interval of evaluation I if and only if, for some proper subinterval J of I, ϕ is true at J; and let $[D]$ be the modal dual of $\langle D \rangle$. Then the 1-place first-order formulas (8) and (9) can be equivalently—and more compactly—re-written as the propositional modal formulas

(12) $\langle D \rangle$int-rec
(13) $[D](\text{cyc} \rightarrow \langle D \rangle \text{int-rec})$.

It is obvious that, with the aid of appropriate modal operators, formulas (10) and (11) could be treated analogously.

Several such logics have in fact been proposed in the literature, of which the best-known are the systems usually referred to as CDT (Venema, 1991) and $\mathcal{HS}$ (Halpern and Shoham 1991; see also Venema 1990). The logic CDT is strictly less expressive than the first-order language employed in (8)–(11); and the logic $\mathcal{HS}$ is in turn strictly less expressive than CDT. Despite its æsthetic appeal, however, a reformulation along the lines of (12)–(13) yields no useful information on the computational complexity of the logic generated by temporal constructions in natural language. For Halpern and Shoham showed that $\mathcal{HS}$ is undecidable over all interesting temporal flows; and still very little is known about its decidable fragments. (For a discussion, see Goranko et al. 2004.) In fact, the most commonly encountered way to ensure decidability for modal interval temporal logics is to impose the restriction that the proposition-letters represent point-events. This move leads naturally to various well-known systems, for example, those of Paech (1988), Moszkowski (1985) and Bowman and Thompson (2003). While these logics are of considerable theoretical interest in their own right, they are of little use for representing the meanings of temporal constructions in natural language.

One striking characteristic of formulas (8)–(11) is the 'quasi-guarded' nature of the quantification they feature. Thus, for example, (8) existentially quantifies over intervals satisfying the predicate int-rec; likewise, (9) universally quantifies over intervals satisfying the predicate cyc; and so on. By contrast, the modal operator $\langle D \rangle$ suggested above (and its dual) quantify over all proper

subintervals of the current interval of evaluation without restriction; corresponding remarks apply to all the modal operators of CDT and $\mathcal{HS}$: they lack the 'quasi-guarded' character of formulas (8)–(11). It is precisely this feature which we shall exploit in our search for a computationally manageable logic to capture the meanings of temporal expressions in English.

2.2   Complications

It is impossible, within the space of a few pages, to do full justice to the complexities of the English constructions featured in this paper. Nevertheless, some elaboration of the foregoing account is required; we confine ourselves to those features of greatest relevance to the ensuing computational analysis. For a comprehensive guide to the grammar of English prepositions, see Huddleston and Pullum (2002, Ch. 7); for an account of the English temporal prepositions in particular, see e.g. Bennett (1975).

We begin with some remarks on the temporal preposition before. We take the sentence

(14) An interrupt was received before the main process ran

to be true in a temporal context I when there is a unique running of the main process during I, and an interrupt is received over some subinterval of I prior thereto. Ordinary usage is vague as to whether it is the beginning- or end-times of the events in question that are being compared. To resolve any uncertainly, we simply take (14) to require that some interrupt-event *finished* before the run of the main process *began*. We therefore propose to render the meaning of (14) by

(15) $\iota J_1(\mathrm{main}(J_1) \wedge J_1 \subset I, \ \exists J_0(\mathrm{int\text{-}rec}(J_0) \wedge J_0 \subset \mathrm{init}(J_1, I)))$.

Notice that these truth-conditions impose no limit on *how long* before the running of the main process the interrupt was received (except that imposed by the temporal context I). That is: before is here used in the sense of *some time before*. Sometimes, however, before is taken to mean *just before* or *shortly before* (The tablets are to be taken before dinner). This latter sense reflects the possibility of adding a time-measure as a specifier, as in the phrase five minutes before. In this paper, we ignore this latter sense of before entirely: incorporating it into our account would involve us in a discussion of either vagueness or the semantics of temporal measure-phrases, both of which we choose to avoid.

Actually, the previous paragraph is misleading in glossing the sense of before assumed here as *some time before*. For the existential quantification in the

6

meaning (15) of (14) is not provided by the **before**-phrase at all, but rather by the sentence **An interrupt was received** occurring in its scope; the **before**-phrase serves merely to specify a temporal context to which that quantification is restricted. In fact, there is no reason this quantification need be existential at all, thus:

(16) **An interrupt was received during every cycle before the main process ran.**

We take (16) to have the meaning (10); that is, we take it to be (truth-conditionally) synonymous with (6). Here again, the **before**-phrase in (16) serves merely to identify a temporal context to which the quantification in its scope is restricted; in particular, it provides no universal quantification of its own.

As for **before**, so for **until**: until-phrases serve only to create temporal contexts restricting the quantification provided by the sentences in their scope; but they do not provide that quantification. This is most apparent by considering the pair of sentences (5) and (6), where the universal quantification evidently arises from the determiner **every**. This treatment of **until** may surprise readers familiar with so-called until-operators in temporal logic, whose semantics do typically contribute universal quantification. Apparently, there is an association of **until** with universal quantification, at least in the minds of temporal logicians; and it is natural to ask how this apparent association can be reconciled with the view adopted here.

The answer is as follows. Sentence (5), which the **until**-phrase in sentence (6) modifies, is *downward monotonic*: if it is true over some interval I, then it is also true over all subintervals of I. (Downward monotonicity is, of course, characteristic of sentences which universally quantify over subintervals.) It transpires that **until**-phrases *require* a downward-monotonic scope, as witnessed by the anomalous

(17) ? **An interrupt was received until the main process ran**
(18) ? **An interrupt was received during some cycle until the main process ran.**

Thus, on our account, the universal quantification—or more accurately, downward monotonicity—is not provided by **until**; but the presence of **until** requires it to be provided by something else. **Before** imposes no such requirement, as we have seen. Thus, the difference between **before** (in the sense adopted here) and **until** lies not in their contribution to truth-conditions, but merely in the situations in which they can be used. Actually, downward monotonicity is not always sufficient for applicability of **until**-phrases (see e.g. Zucchi and White 2001). The exploration of this issue—and indeed of the myriad other differences between **before** and **until**—lies outside the scope of the present enquiry.

We note in passing that until, like before, also allows nominal complements. However, in the case of until, these complements must clearly denote an event or a time:

(19) An interrupt was received during every cycle until 5 o' clock/the first execution of the main process

(20) ? An interrupt was received during every cycle until the main process.

The preposition when creates another sort of difficulty. When serves primarily to indicate proximity between the events identified in its scope and complement, thus:

(21) An interrupt was received when the main process ran.

Sentences such as (21) in fact impose remarkably loose constraints on the temporal relation between the events in question, as various writers have noted. But whatever the final verdict on the nature of those constraints, we cannot usefully treat the associated vagueness in the present paper, and some further regimentation is necessary. To simplify matters, we treat (21) as synonymous with

(22) An interrupt was received *while* the main process ran,

and give it the semantics

(23) $\iota J_1(\text{main}(J_1) \wedge J_1 \subset I, \ \exists J_0(\text{int-rec}(J_0) \wedge J_0 \subset J_1))$.

Our excuse for doing so is simply that inclusion is an easier relation to work with than approximate collocation. Readers who find this expedient too brutal can simply omit when from our fragment.

We have so far discussed quantification in the *scope* of temporal prepositions; we now move to the issue of quantification in their *complements*. Nominal complements of temporal prepositions typically include determiners; and these determiners contribute quantification to the meanings of sentences containing them. This is evident, for example, with the occurrences of during *every* cycle in (5)–(7), which contribute the universal quantifiers in (9)–(11).

Clausal complements of temporal prepositions, by contrast, typically lack an overt quantifier; and the question therefore arises as to how the variables in these complements get quantified. The answer is that they are (almost always) *definitely* quantified—i.e. bound by an ι-operator. Thus, until the main process ran in (6) is interpreted as *until the unique time over which the main process ran*, as reflected by the ι-operator in (10). It may seem harsh to count (6) as false if there are two runs of the main process within the temporal context; it would perhaps be fairer to interpret the relevant until-phrase as picking out

the period before the *first* time over which the main process ran. But since this facility is available in our fragment anyway, as discussed in Section 2.3, the issue need not detain us.

The obvious exception to the rule that temporal prepositions interpret their clausal complements as definitely quantified is whenever. Thus, we take

(24) Whenever the main process ran, an interrupt was received

to have the truth-conditions

(25) $\forall J_1(\text{main}(J_1) \wedge J_1 \subset I \rightarrow \exists J_0(\text{int-rec}(J_0) \wedge J_0 \subset J_1))$.

That is: the variable contributed by the complement of the whenever-phrase is universally quantified. In the sequel, we shall assume that all quantification in clausal complements of temporal prepositions is definite, except in the case of whenever, where is it universal. Note that we are mimicking our earlier discussion of when in again taking the operative temporal relation here to be inclusion rather than approximate collocation. As before, this represents a certain deviation from ordinary usage; again, however, we cannot sensibly deal with vague truth-conditions here, and so we pass over the issue.

Some temporal prepositions have been conspicuous by their absence from the foregoing discussion. The temporal prepositions on and in, in phrases such as on Mondays or in January, are specific to certain categories of complements, but are otherwise equivalent to during. Since this detail clearly has no logical significance, we ignore these uses of in and on, and confine our attention to during. The preposition at, which in English is used in conjunction with clock-times (and some religious festivals) may also fall into this category, though there are further complications here concerning its inherent approximateness. The propositions for and in, in phrases such as for/in five minutes, take as complements temporal measure-phrases. These lie outside the scope of the logic considered here.

The preposition by, in its temporal sense, functions analogously to until, except that it prefers upward-monotonic sentences in its scope; moreover, like until, it dislikes complements which are not explicitly temporal, thus:

(26) An interrupt was received by 5 o'clock
(27) ? An interrupt was received by the first cycle.

(Note that (27) has a perfectly natural reading in which by is interpreted non-temporally.) In addition, by exhibits interesting interactions with aspect:

(28) The main process ran/had run/was running by 5 o'clock.

9

Finally, we observe that by occurs frequently in the construction by the time ... with a clausal complement, again with the same preference for qualifying upward-monotonic sentences. Dealing with the rather difficult behaviour of by in our fragment would complicate the grammar without adding anything of logical interest, and so we ignore it.

In some respects, the mirror-image of both until and by is since:

(29) An interrupt has been received since the main process ran
(30) An interrupt has been received during every cycle since the main process ran.

(When used in its temporal sense, since requires the sentence in its scope to have perfect aspect.) Unlike until and by, however, since resists embedding in contexts established by quantification, as we see by comparing

(31) During every cycle, an interrupt did not occur until the main process
     ran
(32) ? During every cycle, an interrupt has/had not occurred since the main
     process ran.

Because of these complications, we do not include since in our fragment. However, we do include after, which we take (again, ignoring some linguistic subtleties) to function as a mirror image of before. Given the inclusion of after, our omission of since does not affect the fragment's (truth-conditional) expressive power.

2.3   First and Last

Our fragment will also contain sentences such as

(33) An interrupt was received during the *first* cycle
(34) An interrupt was received before the main process ran for the *last* time.

Suppose that, in the relevant temporal context I, there is an unambiguously first cycle: that is, a cycle which begins and ends before all the others. Then (33) asserts that, if J is the interval over which this cycle occurs, then an interrupt was received over some subinterval of J. A corresponding account can of course be given for (34). Problems arise, however, when there is no unambiguously first cycle within I. Suppose, for example, cycles occur during intervals J_1, J_2, and nowhere else, in either of the following arrangements. (In such diagrams, left-to-right arrangement depicts temporal order; vertical arrangement has no significance.)

10

$$\overline{\hspace{4cm}}\,I \qquad\qquad \overline{\hspace{4cm}}\,I$$

$$\underline{\quad\text{cycle}\quad}\,J_1 \qquad\qquad \underline{\quad\text{cycle}\quad}\,J_2$$

$$\underline{\quad\quad\text{cycle}\quad\quad}\,J_2 \qquad\qquad \underline{\quad\quad\text{cycle}\quad\quad}\,J_1$$

It is unclear what the truth-value of (33) should be in such cases. Apparently, we need to legislate.

We take the mathematically simplest way out. Since we may assume that only *finitely* many events of any given type e occur within a given interval I, we proceed as follows. Let $\mathcal{J}$ be the collection of all proper subintervals of I over which an event of type e occurs, and assume $\mathcal{J}$ is nonempty. Since $\mathcal{J}$ is by hypothesis finite, we can select the (non-empty) subset $\mathcal{J}'$ whose elements have the (unique) earliest end-point. Now select the unique element $J \in \mathcal{J}'$ whose start-point is latest. Thus, J is the *smallest* of the *earliest-ending* proper subintervals of I over which an e-event occurs. In the sequel, then, we interpret the phrase the first e, within a temporal context I, to pick out this interval. (In the situations depicted above, these are the intervals marked J_1.) Similarly, we interpret the phrase the last e, within a temporal context I including at least one occurrence of e, to pick out the *smallest* of the *latest-beginning* proper subintervals of I over which an e-event occurs. To re-iterate, we are simply legislating here in the most convenient way in cases where native-speaker intuition returns no clear verdict; if readers prefer to say that the relevant sentences lack truth-values in such cases, then the results obtained below apply unproblematically. The only point at which we appeal to this legislation is in Lemma 3 of Section 5.

3 A Fragment of Temporal English

The task of this section is to define a fragment of temporal English. We do this by writing a context-free grammar to recognize its sentences. The grammar assigns phrase-structures to these sentences in the familiar way, and we shall see that, following some cosmetic re-arrangement, the phrase-structures in question can be regarded as formulas of the temporal logic $\mathcal{TPL}$ defined in Section 4.

3.1 Delineating the fragment

We begin with the simplest sentences in our fragment:

(35) An interrupt was received

11

Fig. 1. The structure of sentences (35)–(36)

(36) An interrupt was not received.

For present purposes, sentence (35) is taken as atomic: that is, we ignore its internal structure. Accordingly we treat such sentences as vocabulary items, of class S^0, and write the grammar rules:

$$S \to S^0 \qquad\qquad S^0 \to \text{an interrupt was received}/\text{int-rec}.$$

Moreover, the only property of sentence (36) which concerns us is its relation to (35): that is, we wish to ignore other aspects of its structure. Accordingly, we pretend that (36) is obtained by simply *prefixing* the word not to (35), and write the grammar rules

$$S \to \text{Neg}, S^0 \qquad\qquad \text{Neg} \to \text{not}/\neg.$$

This expedient removes needless clutter from our grammar, while affecting nothing of logical substance. (It is a simple exercise to restore the clutter.) Thus, our grammar assigns to (35) and (36) the phrase-structures shown in Fig. 1. These diagrams feature the symbols int-rec and $\neg$, as specified in the grammar rules. These symbols are simply mnemonics for the corresponding vocabulary items, which will be used later.

Temporal prepositions with nominal complements belong in our grammar to the category P_N, and occur in phrases such as

(37) during every cycle
(38) after the initialization phase
(39) before the first interrupt.

Nominal expressions such as cycle, initialization phase and interrupt are taken to be of (lexical) category N^0 and to denote event-types in the same way as items of category S^0. Again, we regard them as structureless:

$$N^0 \to \text{cycle}/\text{cyc} \quad N^0 \to \text{initialization phase}/\text{init} \quad N^0 \to \text{interrupt}/\text{int-rec}.$$

We allow these expressions to be optionally modified (once) by the order-specifying adjectives first and last, resulting in a phrase which in turn combines with a determiner to produce the complement of a temporal preposition.
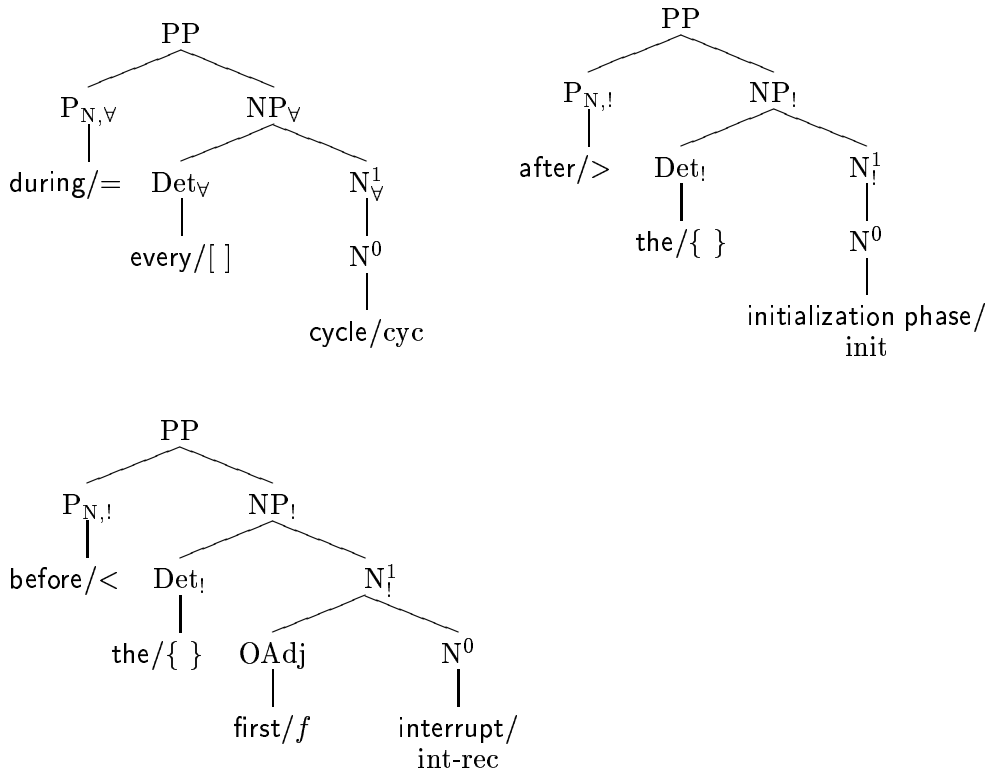
Tree 1 (PP):

PP
- $P_{N,\forall}$ — during/=
- $NP_\forall$
  - $Det_\forall$ — every/[ ]
  - $N^1_\forall$
    - N^0 — cycle/cyc

Tree 2 (PP):

PP
- $P_{N,!}$ — after/>
- $NP_!$
  - $Det_!$ — the/{ }
  - $N^1_!$
    - N^0 — initialization phase/init

Tree 3 (PP):

PP
- $P_{N,!}$ — before/<
- $NP_!$
  - $Det_!$ — the/{ }
  - $N^1_!$
    - OAdj — first/f
    - N^0 — interrupt/int-rec

Fig. 2. Structures of preposition-phrases with nominal complements

Accordingly, we write the grammar rules

$$PP \to P_{N,D},\ NP_D \qquad NP_D \to Det_D,\ N^1_D \qquad N^1_D \to N^0$$

$$N^1_! \to OAdj,\ N^0 \qquad OAdj \to first/f \qquad OAdj \to last/l$$

$$Det_\forall \to every/[\ ] \qquad Det_\exists \to some/\langle\ \rangle \qquad Det_! \to the/\{\ \}$$

$$P_{N,D} \to during/= \qquad P_{N,!} \to after/> \qquad P_{N,!} \to before/<,$$

where the variable subscript D in the above rules ranges over the set of tags $\{\forall, \exists, !\}$. Thus, our grammar assigns to (37)–(39) the respective phrase-structures shown in Fig. 2. As before, we have augmented terminal nodes with the corresponding mnemonics to the right of the obliques in the lexicon.

The tags $\{\forall, \exists, !\}$ simply indicate a subcategorization of NP, Det, N^1 and P_N. This subcategorization restricts the use of determiners in two ways. First, it requires that phrases involving first and last only ever combine with the definite article. This requirement reflects the observation that (outside university mathematics departments) locutions such as during a first interrupt and during every first interrupt are anomalous.

Our second restriction on the use of determiners requires that complements of the temporal prepositions until, before and after also incorporate the definite article. For until, this requirement serves to rule out some clearly anomalous sentences (it is the italicized *every* which causes the problem):

(40) ? An interrupt occurred during every cycle until *every* reset point.

For before and after, the requirement reflects our earlier decision to interpret before in the sense of *some time before*, rather than *shortly before*. To see why, note that common usage (again: professional mathematicians excepted) does not take the sentences

(41) An interrupt was received before every reset point
(42) An interrupt was received before the first reset point

to be equivalent in contexts where there is a unique first reset point, as our assumed sense of before would require. We conclude that the term before can only have the *shortly-before* sense in (41), and so we banish that sentence from our fragment. Admittedly, existentially quantified complements with these prepositions sound better, even with our chosen sense of before:

(43) An interrupt occurred before *some* reset point
(44) An interrupt occurred during every cycle until *some* reset point.

Indeed, such sentences could be admitted into our fragment without compromising the complexity-theoretic results derived below. However, banning sentences such as (41) while admitting those such as (43) would generate a logical fragment not fully closed under negation; and, while such fragments are unproblematic in principle, they tend to make for notational and conceptual clutter. For simplicity, therefore, we duck the issue, and simply decree that these temporal prepositions require complements with the definite article.

Temporal prepositions with clausal complements belong in our grammar to the category P_S, and occur in phrases such as

(45) before the main process ran
(46) whenever the main process ran
(47) while the main process ran for the last time.

Unmodified clausal complements are taken to be atomic, again of category S_0. Our grammar permits modification (once) of these clausal complements by the adverbials for the first/last time, analogous to the modification of nominal complements by the adjectives first/last. Accordingly, we write the grammar rules
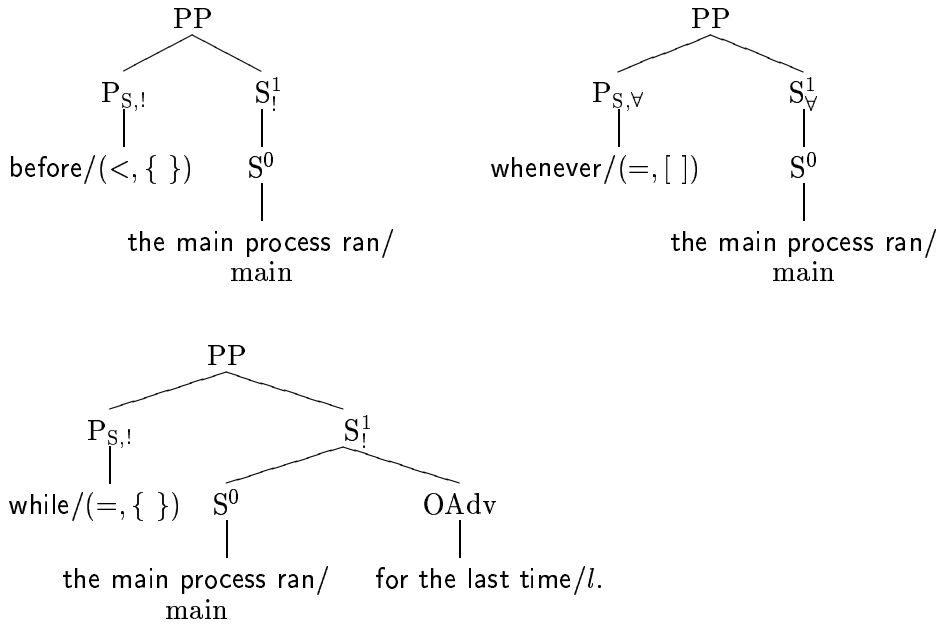
$$\text{PP}$$

$$\text{P}_{\text{S},!} \qquad \text{S}^1_!$$

before/$(<, \{\ \})$ S^0

the main process ran/
main

$$\text{PP}$$

$$\text{P}_{\text{S},\forall} \qquad \text{S}^1_\forall$$

whenever/$(=, [\ ])$ S^0

the main process ran/
main

$$\text{PP}$$

$$\text{P}_{\text{S},!} \qquad \text{S}^1_!$$

while/$(=, \{\ \})$ S^0 OAdv

the main process ran/   for the last time/l.
main

Fig. 3. Structures of preposition-phrases with sentential complements

$$\text{PP} \to \text{P}_{\text{S},D},\ \text{S}^1_D \qquad \text{P}_{\text{S},!} \to \text{while}/(=, \{\ \}) \qquad \text{OAdv} \to \text{for the first time}/f$$

$$\text{S}^1_! \to \text{S}^0,\ \text{OAdv} \qquad \text{P}_{\text{S},!} \to \text{before}/(<, \{\ \}) \qquad \text{OAdv} \to \text{for the last time}/l$$

$$\text{S}^1_D \to \text{S}^0 \qquad\qquad \text{P}_{\text{S},\forall} \to \text{whenever}/(=, [\ ]),$$

thus assigning to (45)–(47) the respective phrase-structures shown in Fig. 3. Recall that whenever is associated with universal, rather than definite, quantification of its complement. That is why the grammar rule for whenever incorporates the bracket-pair [ ], rather than { }, to the right of the oblique. The motivation for these mnemonics will be revealed in Section 4.

We allow that expressions of categories S^0 and N^0 may correspond to the *same* event-type, as indicated by the mnemonics in the lexicon, thus:

$$\text{S}^0 \to \text{the main process ran}/\text{main} \qquad \text{N}^0 \to \text{run of the main process}/\text{main}.$$

Since we want to finesse issues of subsentential and subnominal structure, we leave it to grammar-writers' common sense to spot such nominalizations where they occur. The task of providing a more complex grammar to automate this job is independent of the issues addressed here.

Finally, we have grammar rules to adjoin preposition-phrases to sentences and to handle sentence coordination using and and or. There are no surprises here:

$$\text{S} \to \text{S},\ \text{PP} \qquad \text{S} \to \text{S},\ \text{Conj},\ \text{S} \qquad \text{Conj} \to \text{and}/\wedge \qquad \text{Conj} \to \text{or}/\vee.$$

15

a)

S
|
S⁰
|
An interrupt was received/int-rec

b)

S
├── S
│   └── S⁰
│   └── An interrupt was received/int-rec
└── PP
     ├── $P_{N,\forall}$
     │   └── during/=
     └── $NP_\forall$
          ├── $Det_\forall$
          │   └── every/[ ]
          └── $N^1_\forall$
               └── N^0
                    └── cycle/cyc

c)

S
├── S
│   ├── S
│   │   └── S⁰
│   │   └── An interrupt was received/int-rec
│   └── PP
│   ├── $P_{N,\forall}$
│   │   └── during/=
│   └── $NP_\forall$
│   ├── $Det_\forall$
│   │   └── every/[ ]
│   └── $N^1_\forall$
│   └── N^0
│   └── cycle/cyc
└── PP
     ├── $P_{S,!}$
     │   └── until/({ }, <)
     └── $S^1_!$
          └── S⁰
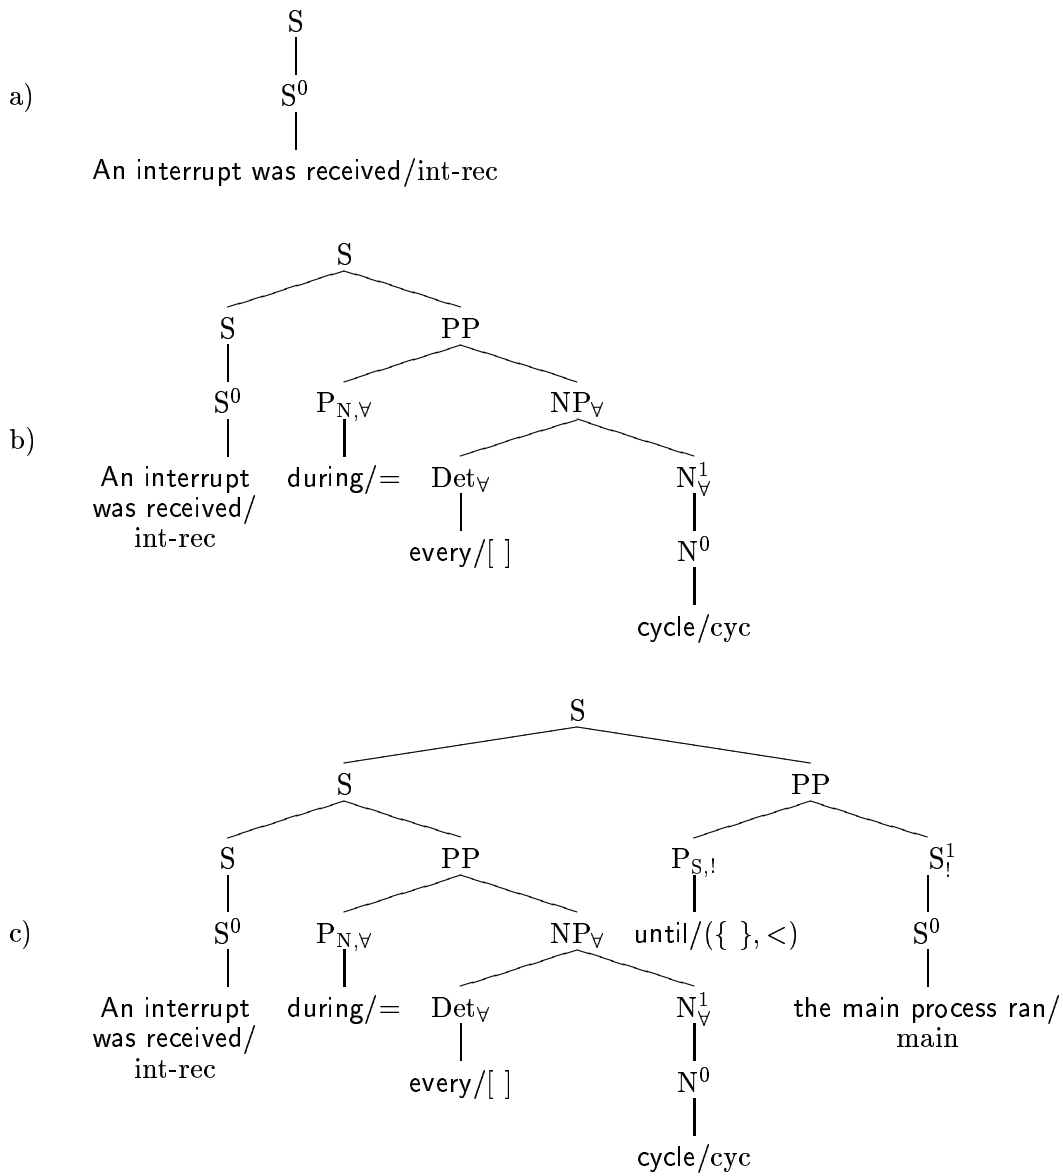   └── the main process ran/main

Fig. 4. Structures of sentences (4)–(6)

Fig. 4 shows the phrase-structures of sentences (4)–(6). Our grammar takes no account of *fronted* preposition-phrases, as illustrated, for example, by Sentence (7). It is obvious that this defect can easily be rectified. This completes our explanation of the fragment of English studied in this paper. We dub this fragment $\mathcal{TPE}$, an acronym for *temporal preposition English*; the full list of grammar rules is given in the Appendix.

In Section 4, we show how phrase-structures in $\mathcal{TPE}$ can be treated as formulas in a language for which a recursive semantics can be given in the style due to Tarski. Moreover, the satisfaction-conditions thus associated with $\mathcal{TPE}$-sentences convincingly systematize the meanings proposed for the various examples considered in Section 2. To facilitate the presentation, we first subject $\mathcal{TPE}$ phrase-structures to some minor geometrical re-arrangement, which we now proceed to describe. We have three base cases and three recursive cases to consider.

First base case: Any structure of the forms depicted in Fig. 1 will be re-written more compactly as follows:

$$
\begin{array}{c}
\text{S} \\
| \\
\text{S}^0 \\
| \\
e
\end{array}
\Rightarrow \langle e \rangle_{=} \top
\qquad
\begin{array}{c}
\text{S} \\
\diagup \quad \diagdown \\
\text{Neg} \qquad \text{S}^0 \\
| \qquad\quad | \\
\neg \qquad\quad e
\end{array}
\Rightarrow \neg \langle e \rangle_{=} \top .
$$

(Here and in the sequel, we have replaced all terminal nodes with the mnemonics to the right of the obliques: this simply unclutters the diagrams.)

**Second base case:** Any structure of category N^1 will be re-written more compactly as follows:

$$
\begin{array}{c}
\text{N}^1_D \\
| \\
\text{N}^0 \\
| \\
e
\end{array}
\Rightarrow e
\qquad
\begin{array}{c}
\text{N}^1_! \\
\diagup \quad \diagdown \\
\text{OAdj} \quad\; \text{N}^0 \\
| \qquad\quad | \\
f \qquad\quad e
\end{array}
\Rightarrow e^f
\qquad
\begin{array}{c}
\text{N}^1_! \\
\diagup \quad \diagdown \\
\text{OAdj} \quad\; \text{N}^0 \\
| \qquad\quad | \\
l \qquad\quad e
\end{array}
\Rightarrow e^l .
$$

Third base case: Any structure of category S^1 will be re-written analogously:

$$
\begin{array}{c}
\text{S}^1_D \\
| \\
\text{S}^0 \\
| \\
e
\end{array}
\Rightarrow e
\qquad
\begin{array}{c}
\text{S}^1_! \\
\diagup \quad \diagdown \\
\text{S}^0 \quad\; \text{OAdv} \\
| \qquad\quad | \\
e \qquad\quad f
\end{array}
\Rightarrow e^f
\qquad
\begin{array}{c}
\text{S}^1_! \\
\diagup \quad \diagdown \\
\text{S}^0 \quad\; \text{OAdv} \\
| \qquad\quad | \\
e \qquad\quad l
\end{array}
\Rightarrow e^l .
$$

First recursive case: Consider a structure of category S immediately dominating a structure Δ of category S and a PP with a nominal complement Γ. Assuming that we already know how to re-write Δ and Γ, such a structure

will be re-written more compactly as follows:

$$
\begin{array}{c}
\text{S} \\
\diagup\quad\diagdown \\
\text{S:}\Delta \qquad \text{PP} \\
\diagup\qquad\diagdown \\
P_{N,D}\qquad NP_D \\
\mid\qquad\diagup\quad\diagdown \\
\tau\quad Det_D\qquad N^1_D\text{:}\Gamma \\
\mid \\
\|\ \|
\end{array}
\qquad \Rightarrow \|\alpha\|_\tau \psi \qquad\qquad \text{if } \Gamma \Rightarrow \alpha \text{ and } \Delta \Rightarrow \psi,
$$

where $\|\ \|$ denotes any of the bracket-pairs $\langle\ \rangle$, $[\ ]$ or $\{\ \}$, and τ any of the symbols $<$, $>$ or $=$.

Second recursive case: Consider a structure of category S immediately dominating a structure Δ of category S and a PP with a clausal complement Γ. Assuming that we already know how to re-write Δ and Γ, such a structure will be re-written more compactly as follows:

$$
\begin{array}{c}
\text{S} \\
\diagup\quad\diagdown \\
\text{S:}\Delta \qquad \text{PP} \\
\diagup\qquad\diagdown \\
P_{S,D}\qquad S^1_D\text{:}\Gamma \\
\mid \\
(\tau, \|\ \|)
\end{array}
\qquad \Rightarrow \|\alpha\|_\tau \psi \qquad\qquad \text{if } \Gamma \Rightarrow \alpha \text{ and } \Delta \Rightarrow \psi,
$$

where $\|\ \|$ denotes either of the bracket-pairs $[\ ]$ or $\{\ \}$, and τ any of the symbols $<$, $>$ or $=$.

Third recursive case: Any structure of category S immediately dominating a node of category Conj will be re-written more compactly as an expression with major connectives $\wedge$ or $\vee$ in the obvious way. The details are routine and are left to the reader to spell out formally.

Consider, for example, the phrase-structures of the $\mathcal{TPE}$-sentence (4)–(6), as drawn in Fig. 4. Re-writing these phrase-structures yields the respective expressions

$$\langle\text{int-rec}\rangle_=\top, \qquad [\text{cyc}]_=\langle\text{int-rec}\rangle_=\top, \qquad \{\text{main}\}_<[\text{cyc}]_=\langle\text{int-rec}\rangle_=\top.$$

Apart from some unusual brackets and decorations, which will be explained later, the results of this re-arrangement look remarkably like formulas of propositional dynamic logic, with the event-classifying mnemonics occupying the

place of atomic programs. So they look; and so they are. We shall give a standard account of the semantics of these formulas along the lines of the usual semantics for propositional dynamic logic. We stress (though it is obvious) that no information has been created or destroyed in the above re-arrangement process: it is a simple graphical matter of replacing an unfamiliar arboreal typography with a familiar (and more compact) linear one. We could have stuck with trees if we had really wanted.

There is one further round of simplification before we proceed. We have demonstrated how PPs in the fragment $\mathcal{TPE}$ can be regarded (syntactically) as modal operators of the form $\|\alpha\|_\tau$, where α is an expression of one of the forms e, e^f or e^l, $\|\ \|$ is one of $\langle\ \rangle$, $[\ ]$ or $\{\ \}$, and τ is one of $=$, $<$ or $>$. However, our grammar imposes restrictions on the quantification in PP-complements ensuring that, if $\tau \in \{<, >\}$ or if α has one of the forms e^f, e^l, then $\|\ \|$ is $\{\ \}$. This cuts down the set of modal operators to the forms

$$\langle e \rangle_=, [e]_=, \{e\}_=, \{e\}_\tau, \{e^\sigma\}_=, \{e^\sigma\}_\tau,$$

where e corresponds to a vocabulary item (of category S^0 or N^0), $\tau \in \{<, >\}$ and $\sigma \in \{f, l\}$. Finally, to avoid clutter, we may take the $=$-subscripts as understood, e.g. writing $[e]$ instead of $[e]_=$. Thus, the final collection of operators is

$$\langle e \rangle, [e], \{e\}, \{e\}_\tau, \{e^\sigma\}, \{e^\sigma\}_\tau,$$

with e an event-atom, $\tau \in \{<, >\}$ and $\sigma \in \{f, l\}$.

Let us take stock. In Section 2, we proposed truth-conditions for a range of sentences involving temporal prepositions and the order-denoting adjectives first and last. By treating sentence-meanings as temporal abstracts, we showed how temporal preposition-phrases could be viewed (semantically) as modal operators. In this section, we have formalized the English fragment we are working with using a context-free grammar. We observed that the phrase-structures which this grammar associates with the sentences it recognizes can be re-arranged as formulas of a language whose syntax resembles propositional dynamic logic. Of course, the point of this re-arrangement is that the resulting formulas can be given a formal semantics which reproduce the truth-conditions proposed in Section 2. It is to that task we now turn.

4   The Temporal Logic

In the sequel, let E be a fixed infinite set. We refer to elements of E as *event-atoms*.

Definition 1 *Let e range over the set E of event-atoms. We define the cate-*

gories of event-relation α, $\mathcal{TPL}$-formula ϕ *and* $\mathcal{TPL}^+$-formula ψ *as follows:*

$\alpha := e \mid e^f \mid e^l$;

$\phi := \langle e \rangle \top \mid \neg \langle e \rangle \top \mid \langle e \rangle \phi \mid [e]\phi \mid \{\alpha\}\phi \mid \{\alpha\}_{>}\phi \mid \{\alpha\}_{<}\phi \mid \phi \wedge \phi' \mid \phi \vee \phi'$;

$\psi := \top \mid \neg \psi \mid \langle e \rangle \psi \mid [e]\psi \mid \{\alpha\}\psi \mid \{\alpha\}_{>}\psi \mid \{\alpha\}_{<}\psi \mid \psi \wedge \psi' \mid \psi \vee \psi'$.

It is easy to see that the syntax of $\mathcal{TPL}$ matches that of the English fragment $\mathcal{TPE}$ exactly: $\mathcal{TPL}$-formulas simply *are* phrase-structures in $\mathcal{TPE}$ and vice-versa. The language $\mathcal{TPL}^+$ is a slight extension of $\mathcal{TPL}$ in which negation is applied rather more freely. Of course, the real object of study is $\mathcal{TPL}$, not $\mathcal{TPL}^+$. The latter is introduced only for the purpose of simplifying the proofs of Section 5.

When dealing with $\mathcal{TPL}^+$, we avail ourselves of the Boolean connectives $\rightarrow$, $\leftrightarrow$ and $\perp$, understood as abbreviations in the usual way. Our first task is to give a formal semantics for $\mathcal{TPL}^+$, and show that, for the fragment $\mathcal{TPL}$, these semantics generate the satisfaction-conditions proposed in Section 2.

Recall from Section 1 that $\mathcal{I}$ denotes the set of *intervals*, where an interval is a closed, bounded, convex (non-empty) subset of $\mathbb{R}$. Recall also the partial functions $\mathrm{init}(J, I)$ and $\mathrm{fin}(J, I)$ defined on $\mathcal{I}$.

Definition 2 *A $\mathcal{TPL}^+$-interpretation $\mathcal{A}$ (henceforth: interpretation) is a finite subset of $\mathcal{I} \times E$. For any $J \in \mathcal{I}$, we write $\mathcal{A}(J)$ for $\{e \in E \mid \langle J, e \rangle \in \mathcal{A}\}$, and for any $e \in E$, we write $\mathcal{A}(e)$ for $\{J \in \mathcal{I} \mid \langle J, e \rangle \in \mathcal{A}\}$.*

Think of an entry $\langle J, e \rangle$ in an interpretation $\mathcal{A}$ as representing the occurrence of an event of type e over the interval J. The motivation for insisting that interpretations are *finite* sets is simply that we have in mind situations in which event-atoms denote everyday event-types instantiated in finite contexts.

We now turn to the interpretation of event-relations. Recalling our (rather artificial) stipulations about the meanings of the words first and last applied to event-types of which there *is* no unambiguously first or last instance, we adopt the following terminology.

Definition 3 *Let I be an interval and $J \subset I$, where J satisfies some property $\mathcal{P}$. We say that $J = [a, b]$ is the* minimal-first *subinterval of I satisfying $\mathcal{P}$ just in case for every $J' = [a', b'] \subset I$ satisfying $\mathcal{P}$, either (i) $b < b'$ or (ii) $b = b'$ and $a \geq a'$. Likewise, we say that $J = [a, b]$ is the* minimal-last *subinterval of I satisfying $\mathcal{P}$ just in case for every $J' = [a', b'] \subset I$ satisfying $\mathcal{P}$, either (i) $a > a'$ or (ii) $a = a'$ and $b \leq b'$.*

Definition 4 *Let α be an event-relation, $\mathcal{A}$ an interpretation, and $I, J \in \mathcal{I}$.*

We define $\mathcal{A} \models_{I,J} \alpha$ by cases as follows:

(1) $\mathcal{A} \models_{I,J} e$ iff $J \subset I$ and $e \in \mathcal{A}(J)$;
(2) $\mathcal{A} \models_{I,J} e^f$ iff $\mathcal{A} \models_{I,J} e$ and J is the minimal-first such interval;
(3) $\mathcal{A} \models_{I,J} e^l$ iff $\mathcal{A} \models_{I,J} e$ and J is the minimal-last such interval.

It is obvious that, since $\mathcal{A}$ is finite, if there exists any $J \subset I$ such that $\langle J, e \rangle \in \mathcal{A}$, then the minimal-first and minimal-last such J exist and are unique.

We are now ready to give the satisfaction-conditions for formulas in $\mathcal{TPL}^+$.

Definition 5 *Let ϕ be a formula, $\mathcal{A}$ an interpretation, and $I \in \mathcal{I}$. We define $\mathcal{A} \models_I \phi$ recursively as follows:*

(1) $\mathcal{A} \models_I \langle e \rangle \psi$ iff for some J, $\mathcal{A} \models_{I,J} e$ and $\mathcal{A} \models_J \psi$;
(2) $\mathcal{A} \models_I [e]\psi$ iff for all J, $\mathcal{A} \models_{I,J} e$ implies $\mathcal{A} \models_J \psi$;
(3) $\mathcal{A} \models_I \{\alpha\}\psi$ iff there is a unique $J \subset I$ such that $\mathcal{A} \models_{I,J} \alpha$, and for that J, $\mathcal{A} \models_J \psi$;
(4) $\mathcal{A} \models_I \{\alpha\}_{<}\psi$ iff there is a unique $J \subset I$ such that $\mathcal{A} \models_{I,J} \alpha$, and for that J, $\mathcal{A} \models_{\mathrm{init}(J,I)} \psi$;
(5) $\mathcal{A} \models_I \{\alpha\}_{>}\psi$ iff there is a unique $J \subset I$ such that $\mathcal{A} \models_{I,J} \alpha$, and for that J, $\mathcal{A} \models_{\mathrm{fin}(J,I)} \psi$;
(6) the usual rules for $\top$, $\wedge$, $\vee$ and $\neg$.

If $\mathcal{A} \models_I \phi$, we say that ϕ is true *at I in $\mathcal{A}$. For any set of formulas Φ, we write $\mathcal{A} \models_I \Phi$ if $\mathcal{A} \models_I \phi$ for all $\phi \in \Phi$. If, for all $\mathcal{A}$ and I, $\mathcal{A} \models_I \Phi$ implies $\mathcal{A} \models_I \phi'$, we say that Φ entails ϕ'; and if ϕ is the sole element in Φ, we say that ϕ entails ϕ'. If ϕ and ϕ' entail each other, we say they are* logically equivalent *and write $\phi \equiv \phi'$. A set of formulas Φ is said to be* satisfiable *if, for some $\mathcal{A}$ and I, $\mathcal{A} \models_I \Phi$.*

We remark that the condition in Definition 2 that interpretations are *finite* subsets of $\mathcal{I} \times E$ is significant. For example, the $\mathcal{TPL}$-formula $\langle e \rangle \top \wedge [e]\langle e \rangle \top$ is unsatisfiable.

Since any $\mathcal{TPL}$-formula ϕ is just the phrase-structure of a $\mathcal{TPE}$-sentence, the immediate question is whether the satisfaction-conditions assigned to ϕ in Definition 5 correctly reproduce the meanings proposed in Section 2.

A little thought shows that they do. For example, the grammar of Section 3 assigns to the sentences (4)–(6), which we repeat here for convenience as

(48) An interrupt was received
(49) An interrupt was received during every cycle
(50) An interrupt was received during every cycle until the main process ran,

the respective phrase-structures

(51) $\langle\text{int-rec}\rangle\top$
(52) $[\text{cyc}]\langle\text{int-rec}\rangle\top$
(53) $\{\text{main}\}_<[\text{cyc}]\langle\text{int-rec}\rangle\top.$

From Definition 5, we see that the satisfaction-conditions of these formulas correspond exactly to those of the respective first-order formulas

(54) $\exists J_0(\text{int-rec}(J_0) \wedge J_0 \subset I)$
(55) $\forall J_1(\text{cyc}(J_1) \wedge J_1 \subset I \to \exists J_0(\text{int-rec}(J_0) \wedge J_0 \subset J_1))$
(56) $\iota J_2(\text{main}(J_2) \wedge J_2 \subset I,$
$\qquad \forall J_1(\text{ cyc}(J_1) \wedge J_1 \subset \text{init}(J_2, I) \to \exists J_0(\text{int-rec}(J_0) \wedge J_0 \subset J_1))).$

But these are precisely the meanings proposed in (8)–(10). More generally, we took pains in Section 2 to show that temporal preposition phrases could be regarded, semantically, as modal operators, mapping one temporal abstract to another, and binding the free-variable in their arguments. The formal semantics of $\mathcal{TPL}$ reflect this observation. In particular, we see how the various components of such a modal operator are contributed by the temporal preposition and its complement. The appropriateness of the semantics for the modifiers $\cdot^f$ and $\cdot^l$ and the Boolean connectives should be self-evident.

This concludes the first part of the paper. We have defined an English fragment, $\mathcal{TPE}$, incorporating temporal prepositions and order-specifying adjectives. We have shown that sentences in this fragment can be regarded as formulas in an interval temporal logic $\mathcal{TPL}$, with satisfaction-conditions matching the meanings which speakers of English assign to them, modulo the various caveats and occasional stipulations mentioned in Section 2. In particular, the problems of determining the satisfiability of a set of sentences or the validity of an argument in $\mathcal{TPE}$ are identical to the corresponding problems in $\mathcal{TPL}$. In the second part of this paper, we proceed to determine the computational complexity of these problems.

5   Upper Complexity Bound

The aim of this section is to show that the satisfiability problem for $\mathcal{TPL}^+$ (and hence $\mathcal{TPL}$) is in NEXPTIME. This is achieved by establishing an exponential bound on the size of satisfying interpretations.

22

Fig. 5. Two interpretations making $\langle a \rangle \top \wedge \neg \{a\} \top$ true at any $I \supset I_1$

Lemma 1 *For all $e \in E$, $\phi \in \mathcal{TPL}^+$, $\tau \in \{<, >\}$ and $\sigma \in \{f, l\}$:*

$$\neg \langle e \rangle \phi \equiv [e] \neg \phi \qquad\qquad \neg [e] \phi \equiv \langle e \rangle \neg \phi$$

$$\neg \{e\} \phi \equiv \neg \{e\} \top \vee \{e\} \neg \phi \qquad \neg \{e^\sigma\} \phi \equiv [e] \bot \vee \{e^\sigma\} \neg \phi$$

$$\neg \{e\}_\tau \phi \equiv \neg \{e\} \top \vee \{e\}_\tau \neg \phi \qquad \neg \{e^\sigma\}_\tau \phi \equiv [e] \bot \vee \{e^\sigma\}_\tau \neg \phi.$$

PROOF. Trivial.

Lemma 2 *Every $\mathcal{TPL}^+$-formula is logically equivalent to one in which $\neg$ appears only in subformulas of the forms $\neg \{e\} \top$ and $\bot \ (= \neg \top)$.*

PROOF. The logical equivalences of Lemma 1, together with familiar propositional validities, allow negations to be moved successively inwards until the desired form is reached.

Definition 6 *Let $\mathcal{A} \neq \emptyset$ be an interpretation. The* depth *of $\mathcal{A}$ is the greatest m for which there exist $J_1 \supset \ldots \supset J_m$ with $\mathcal{A}(J_i) \neq \emptyset$ for all i $(1 \leq i \leq m)$. If $\mathcal{A}$ is empty, we take its depth to be 0.*

The next lemma shows that, in determining satisfiability of $\mathcal{TPL}^+$-formulas, we need never consider very deep interpretations. To illustrate the basic idea, let $I_1 \supset \cdots \supset I_4$ be a descending chain of intervals, and let $\mathcal{A}$ be the interpretation $\{\langle I_i, a \rangle | 1 \leq i \leq 4\}$, as shown in the left-hand diagram in Fig. 5. Evidently, for any $I \supset I_1$, $\mathcal{A} \models_I \langle a \rangle \top \wedge \neg \{a\} \top$. However, it is clear that we can remove the occurrence of a at I_1 (indeed, also at I_2) without compromising this fact. Thus, if $\mathcal{A}^*$ is the interpretation $\{\langle I_i, a \rangle | 2 \leq i \leq 4\}$ depicted in the right-hand diagram of Fig. 5, we still have, for any $I \supset I_1$, $\mathcal{A}^* \models_I \langle a \rangle \top \wedge \neg \{a\} \top$.

Lemma 3 *Let ϕ be a $\mathcal{TPL}^+$-formula, $\mathcal{A}$ an interpretation and I an interval such that $\mathcal{A} \models_I \phi$. Denote the number of symbols in ϕ by $|\phi|$. Then there exists an interpretation $\mathcal{A}^* \subseteq \mathcal{A}$ with depth at most $O(|\phi|^2)$ such that $\mathcal{A}^* \models_I \phi$.*

PROOF. We may assume that ϕ has the form guaranteed by Lemma 2, and further, that $\mathcal{A}$ involves no event-atoms not mentioned in ϕ. Let Φ be the set of subformulas of ϕ. For every event-atom e mentioned in ϕ and every interval J, define

$$L(J) = \{\psi \in \Phi \mid \mathcal{A} \models_J \psi\}$$
$$L_e(J) = L(J) \setminus \bigcup \{L(K) \mid K \subset J, K \in \mathcal{A}(e)\}.$$

Thus, $L_e(J)$ records which subformulas of ϕ are true at an interval J, ignoring those subformulas which are true at some proper subinterval of J satisfying e. Say that a pair $\langle J, e \rangle \in \mathcal{A}$ is *redundant* if $L_e(J) = \emptyset$ and there exist $K, K' \in \mathcal{A}(e)$ such that $K \subset K' \subset J$. Now set

$$\mathcal{A}^* = \mathcal{A} \setminus \{\langle J, e \rangle \mid \langle J, e \rangle \text{ is redundant}\}.$$

To illustrate, suppose for the moment that ϕ is $\langle a \rangle \top \wedge \neg \{a\} \top$ and $\mathcal{A}$ the interpretation depicted in the left-hand diagram of Fig. 5. It is routine to check that $L(I_1) = L(I_2)$, whence $L_a(I_1) = \emptyset$. On the other hand, $L_a(I_2)$, $L_a(I_3)$ and $L_a(I_4)$ are all non-empty, so that $\mathcal{A}^*$ is as depicted in the right-hand diagram of Fig. 5. As we observed, the reduction of $\mathcal{A}$ to $\mathcal{A}^*$ does not affect the truth-value of ϕ at any interval $I \supset I_1$.

Returning to the general case, it is obvious that, if $J \subset J'$ with $J, J' \in \mathcal{A}(e)$, then $L_e(J)$ and $L_e(J')$ are disjoint. It follows that the depth of $\mathcal{A}^*$ is bounded by $m(m' + 2)$, where m is the number of event-atoms occurring in ϕ and m' the number of subformulas of ϕ. It thus suffices to show that, for all I and all $\psi \in \Phi$, $\mathcal{A} \models_I \psi$ implies $\mathcal{A}^* \models_I \psi$.

We proceed by induction on the complexity of ψ. The base cases are of the forms $\psi = \top, \bot, \neg \{e\} \top$. The first two of these are trivial. For the case $\psi = \neg \{e\} \top$, suppose $\mathcal{A} \models_I \psi$. If there is no $J \subset I$ with $J \in \mathcal{A}(e)$, then since $\mathcal{A}^* \subseteq \mathcal{A}$, we certainly have $\mathcal{A}^* \models_I \psi$. Otherwise, there exist $J \subset I$ and $J' \subset I$ with $J \neq J'$ and $J, J' \in \mathcal{A}(e)$. If neither of the pairs $\langle J, e \rangle$ and $\langle J', e \rangle$ is redundant, then $J, J' \in \mathcal{A}^*(e)$. On the other hand, if $\langle J, e \rangle$ is redundant, there must exist $K \subset K' \subset J$ such that the pairs $\langle K, e \rangle$ and $\langle K', e \rangle$ are non-redundant elements of $\mathcal{A}$, whence $K, K' \in \mathcal{A}^*(e)$; and similarly if $\langle J', e \rangle$ is redundant. Either way, then, $\mathcal{A}^* \models_I \psi$.

The recursive cases are of the forms $\psi = [e]\theta, \langle e \rangle \theta, \{\alpha\}\theta, \{\alpha\}_\tau \theta$, where α is of the forms e, e^f or e^l, and $\tau \in \{<, >\}$. For the case $\psi = [e]\theta$, we need only observe that $\mathcal{A}^* \subseteq \mathcal{A}$. For the case $\psi = \langle e \rangle \theta$, suppose $\mathcal{A} \models_I \psi$. Then there exists $J \subset I$ such that $J \in \mathcal{A}(e)$ and $\mathcal{A} \models_J \theta$. By the finiteness of $\mathcal{A}$, choose such a J which is minimal under the order $\subset$, so that $J \in \mathcal{A}^*(e)$. By inductive hypothesis, $\mathcal{A}^* \models_J \theta$; hence $\mathcal{A}^* \models_I \psi$. For the case $\psi = \{e\}\theta$, suppose $\mathcal{A} \models_I \psi$. Then there exists a unique $J \subset I$ such that $J \in \mathcal{A}(e)$; and for

24

this J, $\mathcal{A} \models_J \theta$. In particular, there is no $K \subset J$ such that $K \in \mathcal{A}(e)$, whence $J \in \mathcal{A}^*(e)$. By inductive hypothesis and the fact that $\mathcal{A}^* \subseteq \mathcal{A}$, we then easily have $\mathcal{A}^* \models_I \psi$. The remaining cases are dealt with exactly as for $\psi = \{e\}\theta$, noting, in particular, that $\mathcal{A} \models_{I,J} e^f$ implies $\mathcal{A}^* \models_{I,J} e^f$ and $\mathcal{A} \models_{I,J} e^l$ implies $\mathcal{A}^* \models_{I,J} e^l$. (This is the point at which we rely on the rather artificial choice of semantics for e^f and e^l in Definition 4 in cases where there is no unambiguous first or last e-interval.)

Theorem 1 *Let ϕ be a formula of $\mathcal{TPL}^+$. If ϕ is satisfiable, then ϕ is satisfied in an interpretation of size bounded by $2^{p(|\phi|)}$, for some fixed polynomial p.*

PROOF. Suppose that $\mathcal{A} \models_{I_0} \phi$. We may assume that ϕ has the form guaranteed by Lemma 2; and by Lemma 3, we may assume that the depth of $\mathcal{A}$ is at most of order $|\phi|^2$. As before, let Φ be the set of subformulas of ϕ. Say that a formula $\chi \in \Phi$ is *basic* if the major connective of χ is neither $\wedge$ nor $\vee$. For any interval I and any $\psi \in \Phi$, denote by $S(\psi, I)$ the set of all maximal basic subformulas χ of ψ such that $\mathcal{A} \models_I \chi$. It is easy to see that, for any $\psi \in \Phi$ and $I \in \mathcal{I}$ with $\mathcal{A} \models_I \psi$, $S(\psi, I)$ entails ψ.

We now construct a sub-interpretation $\mathcal{A}^*$ of $\mathcal{A}$, starting with the interval I_0 and choosing witnesses, tableau-style, for formulas in Φ. More specifically, the procedure `tree`(ϕ, I_0) in Fig. 6 grows a labelled tree with nodes V, edges E, and the two labellings $\lambda : V \to \mathcal{I}$ and $L : V \to \mathbb{P}(\Phi)$; the interpretation $\mathcal{A}^*$ will then be extracted from $\mathcal{A}$ using this labelled tree. For $v \in V$, think of $\lambda(v)$ as the interval represented by v, and think of $L(v)$ as some collection of formulas which must all be true at this interval. The variable Q is simply a queue of nodes in V awaiting processing. Steps 1–5 ensure, roughly, that 'existential' formulas in Φ have witnesses as required; the embedded calls to `univ`(u) ensure that 'universal' formulas in Φ are not falsified by these witnesses. A straightforward check shows that the invariant $\mathcal{A} \models_{\lambda(v)} L(v)$ for all $v \in V$ is maintained by `tree`(ϕ, I_0). Note that the function λ is not required to be 1–1. Note also that the individual steps in `tree`(ϕ, I_0) need not be effective: all we require for the proof of the theorem is the existence of the interpretation $\mathcal{A}^*$ with the advertised properties.

We claim that `tree`(ϕ, I_0) terminates after finitely many iterations, and that, upon termination, the tree (V, E) satisfies the size bound of the Theorem. By inspection of Steps 1–5, whenever an edge (v, w) is added to E, we have $\lambda(w) \subset \lambda(v)$. Therefore, at any point in the execution of `tree`(ϕ, I_0), if the tree (V, E) contains a path $v_0 \to \cdots \to v_m$, then $\lambda(v_0) \supset \cdots \supset \lambda(v_m)$. Consider those values of i ($0 \leq i < m$) for which the call to `univ`(v_{i+1}) adds material to $L(v_{i+1})$. By inspection of `univ`, this can certainly happen only if, for at least one event-atom e', $e' \in \mathcal{A}(\lambda(v_{i+1}))$. Therefore, it can happen for at most

begin tree(ϕ, I_0)
  Choose some object v_0, and set
    $Q = V = \{v_0\}$; $\lambda(v_0) = I_0$; $L(v_0) = S(\phi, I_0)$; $E = \emptyset$.
  until $Q = \emptyset$ **do**
    Select $v \in Q$, set $I := \lambda(v)$, and set $Q := Q \setminus \{v\}$.
    for every $\psi \in L(v)$, **do**
      (1) If $\psi = \langle e \rangle \theta$, let J be such that $\mathcal{A} \models_{I,J} e$ and $\mathcal{A} \models_J \theta$. Select $w \notin V$
          and set $\lambda(w) := J$; $L(w) := S(\theta, J)$; $Q := Q \cup \{w\}$; $V := V \cup \{w\}$;
          $E := E \cup \{(v, w)\}$. Execute univ(w).
      (2) If $\psi = \{\alpha\}\theta$, let J be such that $\mathcal{A} \models_{I,J} \alpha$. Select $w \notin V$ and
          set $\lambda(w) := J$; $L(w) := S(\theta, J)$; $Q := Q \cup \{w\}$; $V := V \cup \{w\}$;
          $E := E \cup \{(v, w)\}$. Execute univ(w).
      (3) If $\psi = \{\alpha\}_{<}\theta$, let J be such that $\mathcal{A} \models_{I,J} \alpha$ and let $J' = \text{init}(J, I)$.
          Select $w, w' \notin V$ and set $\lambda(w) := J$; $\lambda(w') := J'$; $L(w) := \emptyset$;
          $L(w') := S(\theta, J')$; $Q := Q \cup \{w, w'\}$; $V := V \cup \{w, w'\}$; $E :=$
          $E \cup \{(v, w), (v, w')\}$. Execute univ$(w)$ and univ(w').
      (4) If ψ is $\{\alpha\}_{>}\theta$, proceed symmetrically.
      (5) If ψ is $\neg\{e\}\top$, and there exist $J \subset I$, $J' \subset I$ with $J \neq J'$ and $J, J' \in$
          $\mathcal{A}(e)$, choose any such J, J'. Select $w, w' \notin V$ and set $\lambda(w) := J$;
          $\lambda(w') := J'$; $L(w) := \emptyset$; $L(w') := \emptyset$; $Q := Q \cup \{w, w'\}$; $V := V \cup$
          $\{w, w'\}$; $E := E \cup \{(v, w), (v, w')\}$. Execute univ$(w)$ and univ(w').
    end for every
  end until
end tree

begin univ(u)
    for every formula $[e']\theta \in \Phi$ such that $\langle \lambda(u), e' \rangle \in \mathcal{A}$ and there
    exists $L \supset \lambda(u)$ with $\mathcal{A} \models_L [e']\theta$ **do**
        Set $L(u) := L(u) \cup S(\theta, \lambda(u))$.
    end for every
end univ

Fig. 6. Construction of small interpretations in $\mathcal{TPL}^+$

D different values of i, where D is the depth of $\mathcal{A}$. Moreover, any call to
univ(v_{i+1}) adds at most $|\phi|^2$ symbols to $L(v_{i+1})$; and if the call to univ(v_{i+1})
adds no material to $L(v_{i+1})$, then $L(v_{i+1})$ contains strictly fewer symbols than
$L(v_i)$. Since D is at most of oder $|\phi|^2$, the length of the path $v_0 \to \cdots \to v_m$
is therefore at most of order $|\phi|^4$. The bound on the eventual size of V then
follows from the fact that the out-degree of any node in V is bounded by $2|\phi|$.

Now let $\mathcal{A}^* = \{\langle J, e \rangle \in \mathcal{A} | \text{for some } v \in V, J = \lambda(v)\}$. Evidently, $|\mathcal{A}^*|$ satisfies
the size bound of the Theorem; it thus suffices to show that $\mathcal{A}^* \models_{I_0} \phi$. In fact,
we show by structural induction that, for any node $v \in V$ and any formula
ψ, $\psi \in L(v)$ implies $\mathcal{A}^* \models_{\lambda(v)} \psi$. Denote $\lambda(v)$ by I. (Hence $\mathcal{A} \models_I L(v)$.) The

base cases are of the forms $\psi = \top, \bot, \neg\{e\}\top$. The case $\psi = \top$ is trivial. For the case $\psi = \bot$, the fact that $\mathcal{A} \models_I L(v)$ ensures that $\psi \notin L(v)$. For the case $\psi = \neg\{e\}\top$, if $\psi \in L(v)$, $\mathcal{A} \models_I L(v)$ ensures that either (i) there is no $J \subset I$ such that $J \in \mathcal{A}(e)$ or (ii) there exist $J \subset I$, $J' \subset I$ with $J \neq J'$ such that $J, J' \in \mathcal{A}(e)$. In the former case, since $\mathcal{A}^* \subseteq \mathcal{A}$, then $\mathcal{A}^* \models_I \psi$. In the latter case, Step 5 of $\texttt{tree}(\phi, I_0)$ ensures that, for some such J, J', we have $w, w' \in V$ with $\lambda(w) = J$ and $\lambda(w') = J'$; hence $J, J' \in \mathcal{A}^*(e)$ and $\mathcal{A}^* \models_I \psi$. The inductive cases are almost as straightforward:

(1) Suppose ψ is $\langle e \rangle \theta$. If $\psi \in L(v)$, then, by Step 1 of $\texttt{tree}(\phi, I_0)$, there exists $w \in V$ and $J \subset I$ such that $\lambda(w) = J$, $S(\theta, J) \subseteq L(w)$, $\langle J, e \rangle \in \mathcal{A}$, and $\mathcal{A} \models_J \theta$. By inductive hypothesis, $\mathcal{A}^* \models_J S(\theta, J)$, and since $\mathcal{A} \models_J \theta$, $S(\theta, J)$ entails θ, whence $\mathcal{A}^* \models_J \theta$. By construction, $\langle J, e \rangle \in \mathcal{A}^*$. Hence, $\psi \in L(v)$ implies $\mathcal{A}^* \models_I \psi$.

(2) Suppose ψ is $[e]\theta$. If $\psi \in L(v)$, then $\mathcal{A} \models_I \psi$. Consider any $J \subset I$ with $J \in \mathcal{A}^*(e)$. Certainly, then, $J \in \mathcal{A}(e)$; hence $\mathcal{A} \models_J \theta$, so that $S(\theta, J)$ entails θ. Moreover, by the construction of $\mathcal{A}^*$ there exists $w \in V$ with $\lambda(w) = J$, in which case the call to $\texttt{univ}(w)$ ensures that $S(\theta, J) \subseteq L(w)$. By inductive hypothesis, $\mathcal{A}^* \models_J S(\theta, J)$, whence $\mathcal{A}^* \models_J \theta$. Hence, $\psi \in L(v)$ implies $\mathcal{A}^* \models_I \psi$.

(3) The remaining cases are handled similarly to Case 1, or are trivial.

Corollary 1 *The satisfiability problem for $\mathcal{TPL}^+$ is in NEXPTIME.*

PROOF. Let ϕ be a formula of $\mathcal{TPL}^+$, and let d be the maximum depth of nesting of modal operators in ϕ. By Lemma 3, if ϕ is satisfiable, then it is satisfiable in an interpretation whose size is bounded by some fixed exponential function of $|\phi|$. Guess such an interpretation $\mathcal{A}$ and an interval I. Let $\mathcal{J}_0$ be the set of intervals mentioned in $\mathcal{A}$ together with I, and for any $i \geq 0$, let $\mathcal{J}_{i+1}$ be $\mathcal{J}_i$ together with all intervals expressible as $\text{init}(J_0, J)$ or $\text{fin}(J_0, J)$, where $J_0 \in \mathcal{J}_0$ and $J \in \mathcal{J}_i$. Now, for all i $(0 \leq i \leq d)$ and all $J \in \mathcal{J}_{d-i}$, guess which subformulas of ϕ having modal depth i are true at J in $\mathcal{A}$. It is then straightforward to check, in time bounded by some fixed exponential function of $|\phi|$, that these guesses are correct, and thence to determine whether $\mathcal{A} \models_I \phi$.

The proofs of Lemma 3 and Theorem 1 thus make essential use of the 'quasi-guarded' nature of $\mathcal{TPL}^+$, which we observed in Section 2.1, together with the assumption that only finitely many events occur in a bounded time-interval. Note that the construction employed in the proof of Theorem 1 does not, as formulated there, constitute a tableau decision procedure for $\mathcal{TPL}^+$, because the steps are not necessarily effective. We remark that a (non-terminating) tableau procedure has been devised for the interval temporal logic CDT, interpreted over branching-time structures (Goranko et al., 2003). It is not immedi-

ately clear whether such an approach could be adapted to yield a terminating procedure for $\mathcal{TPL}^+$, interpreted over a linear time flow, and incorporating the assumption that only finitely many events can occur over a bounded time-interval. However, the results of the next section indicate that any such tableau method is likely to require extensive backtracking.

6 Lower Complexity Bound

In this section, we show that the satisfiability problem for $\mathcal{TPL}$ (and hence $\mathcal{TPL}^+$) is NEXPTIME-hard. Denote by $\mathbb{N}_n$ the natural numbers less than n. Define an *exponential tiling problem* to be a triple (C, H, V), where $C = \{c_0, \ldots, c_{M-1}\}$ is a set and H and V are binary relations over C. We call the elements of C *colours*, and we call H and V the *horizontal constraints* and the *vertical constraints*, respectively. An *instance* of (C, H, V) is a list $c_0', \ldots, c_{n-1}'$ of elements of C (repetitions allowed). Such an instance is *positive* if there exists a function $\tau : \mathbb{N}_{2^n} \times \mathbb{N}_{2^n} \to C$ such that: (i) $\tau(i, 0) = c_i'$ for all i ($0 \leq i \leq n - 1$); (ii) $\langle \tau(i, j), \tau(i+1, j) \rangle \in H$ for all i, j ($0 \leq i < 2^n - 1, 0 \leq j \leq 2^n - 1$); (iii) $\langle \tau(i, j), \tau(i, j+1) \rangle \in V$ for all i, j ($0 \leq i \leq 2^n - 1, 0 \leq j < 2^n - 1$); and (iv) $\tau(0, 2^n - 1) = c_0$. We refer to τ as a *tiling*. Intuitively, the elements of C represent colours of unit square tiles which must be arranged so as to fill a grid of $2^n \times 2^n$ squares, with the top left-hand square required to have the colour c_0. The constraints H (respectively, V) list which colours are allowed to go to the right of (respectively, above) which others. The problem instance $c_0', \ldots, c_{n-1}'$ lists the colours of the first n tiles in the bottom row. For a discussion of exponential tiling problems, see Börger et al. (1997, Sec. 6.1.1).

To show that a problem $\mathcal{P}$ is NEXPTIME-hard, it suffices to show that, for any exponential tiling problem (C, H, V), any instance of (C, H, V) may be encoded, in polynomial time, as an instance of $\mathcal{P}$. We now proceed to do this where $\mathcal{P}$ is $\mathcal{TPL}$-satisfiability. The main technical challenge is to encode, using a succinct formula of $\mathcal{TPL}$, the information that there are *exactly* 2^{2n} pairwise disjoint intervals satisfying some event-atom t within a given interval I^*. We begin by tackling this problem; the remainder of the reduction is routine.

6.1 Fixing a large number of tiles

Let $m \geq 2$ and let a_0, $a_1^0, \ldots, a_{m+1}^0$, $a_1^1, \ldots, a_{m+1}^1$ and z be pairwise distinct event-atoms. To simplify the notation, we write a_0 alternatively as a_0^0 or a_0^1. The event-atom z will always function as a harmless 'dummy'; it occurs in subformulas $\langle z \rangle \top$ whose only purpose is to ensure that we remain inside the temporal logic $\mathcal{TPL}$, rather than the more general $\mathcal{TPL}^+$. The following ter-

$$\begin{array}{c}
\underline{\qquad\qquad\qquad\qquad a_0 \qquad\qquad\qquad\qquad} \\
\underline{\quad a_1^0 \quad} \qquad\qquad \underline{\quad a_1^1 \quad} \\
\cdots \qquad \cdots \qquad\qquad \cdots \qquad \cdots \\
\underline{a_{m-1}^0 \quad a_{m-1}^1} \quad \cdots \quad \underline{a_{m-1}^0 \quad a_{m-1}^1} \\
\underline{a_m^0}\ \underline{a_m^1}\ \underline{a_m^0}\ \underline{a_m^1} \quad \cdots \quad \underline{a_m^0}\ \underline{a_m^1}\ \underline{a_m^0}\ \underline{a_m^1}
\end{array}$$

Fig. 7. Arrangement of i-witnesses ($0 \leq i \leq m$)

minology will be used to aid readability. Where an interpretation $\mathcal{A}$ is clear from context, we say that an interval I *satisfies* an event-atom e if $\langle I, e \rangle \in \mathcal{A}$; alternatively, we say that I is an *e-interval*.

Define ψ_1 to be the conjunction of the following formulas, where $0 \leq i \leq m$ and $0 \leq h \leq 1$:

(57) $\{a_0\}\langle z\rangle\top, \quad [a_i^h]\{a_{i+1}^0\}_{>}\langle a_{i+1}^1\rangle\top, \quad [a_i^h]\{a_{i+1}^1\}\langle z\rangle\top.$

Let $\mathcal{A}$ be an interpretation and I^* an interval such that $\mathcal{A} \models_{I^*} \psi_1$. For all i ($0 \leq i \leq m$), define an *i-witness* inductively as follows:

(1) I is a 0-witness if and only if I is the unique proper subinterval of I^* satisfying a_0.
(2) J is an $(i+1)$-witness if and only if there exists an i-witness I such that J is either the unique proper subinterval of I satisfying a_{i+1}^0 or the unique proper subinterval of I satisfying a_{i+1}^1.

Given that $\mathcal{A} \models_{I^*} \psi_1$, each i-witness I properly includes exactly one interval J satisfying a_{i+1}^0 and exactly one interval J' satisfying a_{i+1}^1, with J preceding J'. Thus, there are exactly 2^i i-witnesses for all i ($1 \leq i \leq m$); moreover, these are pairwise disjoint and alternate between intervals satisfying a_i^0 and a_i^1, as depicted in Fig. 7. Note however that, in general, the i-witnesses will be a subset of the subintervals of I^* satisfying a_i^0 or a_i^1 in $\mathcal{A}$.

The formula ψ_1 thus provides a succinct way of guaranteeing that *at least* 2^{m-1} proper subintervals of I^* satisfy a_m^0 in $\mathcal{A}$—viz, every other m-witness. A much greater challenge is to write a succinct collection of formulas ensuring that no *other* proper subintervals of I^* satisfy a_m^0. This task occupies the remainder of Section 6.1.

Let $b_1, \ldots, b_m, p_0^0, \ldots, p_{m-1}^0$ and $p_0^1, \ldots, p_{m-1}^1$ be new event-atoms (i.e. pairwise distinct and distinct from z, a_0 and the a_i^h). Intuitively, the event-atoms b_i will be used to prevent 'additional' a_i^0-events and a_i^1-events slipping in between successive i-witnesses; the event-atoms p_j^0 and p_j^1 will function as 'nails', holding the whole rickety structure together. Let ψ_2 be the conjunction of the following formulas, where $0 \leq i < m$, $0 \leq h \leq 1$ and $0 \leq h' \leq 1$:

$$\overline{\quad\overset{\textstyle a_i^{h'}}{\rule{0pt}{0pt}}\quad}\,I$$
$$\overline{\quad\overset{\textstyle b_{i+1}}{\rule{0pt}{0pt}}\quad}\,L$$
$$\overline{\overset{a_{i+1}^0}{\rule{0pt}{0pt}}}\,J \qquad \overline{\overset{a_{i+1}^1}{\rule{0pt}{0pt}}}\,J'$$
$$\overline{\overset{p_i^0}{\rule{0pt}{0pt}}}\qquad \overline{\overset{p_i^1}{\rule{0pt}{0pt}}}$$
$$\overline{\overset{a_{i+2}^1}{\rule{0pt}{0pt}}}K \qquad \overline{\overset{a_{i+2}^0}{\rule{0pt}{0pt}}}K'$$

(a)

$$\overline{\quad\overset{\textstyle b_i}{\rule{0pt}{0pt}}\quad}\,I$$
$$\overline{\quad\overset{\textstyle b_{i+1}}{\rule{0pt}{0pt}}\quad}\,L$$
$$\overline{\overset{a_{i+1}^1}{\rule{0pt}{0pt}}}\,J \qquad \overline{\overset{a_{i+1}^0}{\rule{0pt}{0pt}}}\,J'$$
$$\overline{\overset{q_i^0}{\rule{0pt}{0pt}}}\qquad \overline{\overset{q_i^1}{\rule{0pt}{0pt}}}$$
$$\overline{\overset{a_{i+2}^1}{\rule{0pt}{0pt}}}K \qquad \overline{\overset{a_{i+2}^0}{\rule{0pt}{0pt}}}K'$$

(b)

Fig. 8. Representative arrangement of intervals: (a) under each $a_i^{h'}$-interval, and (b) under each b_i-interval

(58)
$$[a_i^{h'}]\{b_{i+1}\}\langle z\rangle\top, \quad [a_i^{h'}]\{p_i^h\}\langle z\rangle\top, \quad [a_{i+1}^h]\langle p_i^h\rangle\top,$$
$$[b_{i+1}]\langle p_i^h\rangle\top, \qquad\qquad [p_i^h]\langle a_{i+2}^{1-h}\rangle\top.$$

Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_2$. Formula ψ_1 guarantees that, for all i ($0 \le i < m$), any subinterval $I \subset I^*$ satisfying either a_i^0 or a_i^1 properly includes a unique J satisfying a_{i+1}^0 and a unique J' satisfying a_{i+1}^1, with J preceding J'. Moreover, there exists a unique $K \subset J$ satisfying a_{i+2}^1, and a unique $K' \subset J'$ satisfying a_{i+2}^0. In addition, ψ_2 guarantees that I also properly includes a unique L satisfying b_{i+1}; moreover, the event-atoms p_i^0 and p_i^1, which are satisfied uniquely within I, effectively 'nail' the L, J, J', K and K' together so that $L \cap J \supset K$ and $L \cap J' \supset K'$. A representative situation conforming to these constraints is depicted in Fig. 8(a).

Let $q_1^0, \ldots, q_{m-1}^0$ and $q_1^1, \ldots, q_{m-1}^1$ be new event-atoms, and let ψ_3 be the conjunction of the following formulas, where $1 \le i < m$ and $0 \le h \le 1$:

(59)
$$[b_i]\{a_{i+1}^1\}_{>}\langle a_{i+1}^0\rangle\top, \quad [b_i]\{a_{i+1}^0\}\langle z\rangle\top, \quad [b_i]\{b_{i+1}\}\langle z\rangle\top,$$
$$[b_i]\{q_i^h\}\langle z\rangle\top, \qquad\qquad [b_{i+1}]\langle q_i^h\rangle\top, \qquad [a_{i+1}^h]\langle q_i^{1-h}\rangle\top, \quad [q_i^h]\langle a_{i+2}^{1-h}\rangle\top.$$

Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_3$. Looking at the first row of (59), ψ_3 guarantees that, for all i ($1 \le i < m$), any subinterval $I \subset I^*$ satisfying b_i properly includes a unique J satisfying a_{i+1}^1 and a unique J' satisfying a_{i+1}^0 (with J preceding J'), as well as a unique L satisfying b_{i+1}. Further, ψ_1 guarantees that there is a unique $K \subset J$ satisfying a_{i+2}^1, and a unique $K' \subset J'$ satisfying a_{i+2}^0. Looking now at the second row of (59), the event-atoms q_i^0 and q_i^1, which are satisfied uniquely within I, effectively 'nail' the L, J, J', K and K' together so that $L \cap J \supset K$ and $L \cap J' \supset K'$. A representative situation conforming to these constraints is depicted in Fig. 8(b).

These observations help us establish:

Claim 1 *Let $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_2 \wedge \psi_3$, and let K, K' be consecutive $(i+1)$-witnesses,*

with $0 \le i < m$. Then there exists an interval $L \subset I^*$ properly including both K and K', such that L satisfies one of a_i^0, a_i^1 or b_i.

PROOF. We proceed by induction on i. If $i = 0$, the result is trivial, because the only 1-witnesses are by definition properly included in the 0-witness.

For the inductive case, suppose the statement of the Lemma holds with $0 \le i < m - 1$; we show the same statement holds with i replaced by $i + 1$. Let K, K' be consecutive $(i+2)$-witnesses, then; without loss of generality, we can suppose that K precedes K'. Each $(i + 2)$-witness is by definition properly included in a unique $(i + 1)$-witness; so let J be the $(i + 1)$-witness such that $K \subset J$ and J' be the $(i + 1)$-witness such that $K' \subset J'$. Since K and K' are consecutive, J and J' are identical or consecutive. In the former case, we may put $L = J = J'$, and L satisfies either a_{i+1}^0 or a_{i+1}^1 as required by the Lemma. So assume the latter. By inductive hypothesis, then, J and J' are properly included within an interval $I \subset I^*$ such that I satisfies a_i^0, a_i^1, or b_i. Moreover, since K and K' are consecutive but not included in a common $(i+1)$-witness, K satisfies a_{i+2}^1 and K' satisfies a_{i+2}^0.

If I satisfies $a_i^{h'}$ ($0 \le h' \le 1$), then ψ_1 guarantees that I properly includes exactly one interval satisfying a_{i+1}^0 and exactly one interval satisfying a_{i+1}^1, with the former preceding the latter; these must be, respectively, J and J', therefore. Again by ψ_1, J properly includes exactly one interval satisfying a_{i+2}^1 and J' exactly one interval satisfying a_{i+2}^0; these must be, respectively, K and K', therefore. Thus, we have the arrangement of Fig. 8(a). In particular, we have seen that ψ_2 guarantees the existence of an interval L satisfying b_{i+1} and properly including both K and K', as required by the Lemma.

If I satisfies b_i, then ψ_3 guarantees that I properly includes exactly one interval satisfying a_{i+1}^1 and exactly one interval satisfying a_{i+1}^0, with the former preceding the latter; these must be, respectively, J and J', therefore. By ψ_1, J properly includes exactly one interval satisfying a_{i+2}^1 and J' exactly one interval satisfying a_{i+2}^0; these must be, respectively, K and K', therefore. Thus, we have the arrangement of Fig. 8(b). In particular, we have seen that ψ_3 guarantees the existence of an interval L satisfying b_{i+1} and properly including both K and K', as required by the Lemma.

Claim 2 *Let $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_2 \wedge \psi_3$. If K and K' are consecutive i-witnesses (in that order), with $1 \le i \le m$, then no subinterval $H \subset I^*$ satisfying either a_i^0 or a_i^1 can begin after K begins and end before K' ends.*

31

PROOF. Suppose for contradiction that such an H exists. By Claim 1, we have some $L \subset I^*$ satisfying one of a_{i-1}^0, a_{i-1}^1 or b_{i-1}, with $L \supset K$ and $L \supset K'$. Thus, $L \supset H$. But ψ_1 and ψ_3 contain conjuncts requiring L to properly include exactly one interval satisfying a_i^0 and exactly one interval satisfying a_i^1. Contradiction.

Let $l_1, \ldots, l_m, l_1', \ldots, l_{m+1}'$ be new event-atoms, and let ψ_4^l be the conjunction of the following collection of formulas, where i $(1 \le i \le m)$:

$$(60) \qquad \{a_0\}\langle l_1'\rangle\top, \qquad \{l_i'\}\{a_i^0\}\langle l_{i+1}'\rangle\top,$$
$$\{l_i'\}\{a_i^0\}_>\langle l_i\rangle\top, \quad \{l_i\}_<\{a_i^0\}\langle z\rangle\top, \quad \{l_i\}_<\neg\langle a_i^1\rangle\top.$$

Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_4^l$ and $1 \le i \le m$. For all i $(1 \le i \le m)$, let J_i be the first-occurring i-witness, and let L_i' be the unique proper subinterval of I^* satisfying l_i'. Then the conjuncts in the first row of (60) enforce the arrangement

$$L_1' \supset J_1 \supset L_2' \supset J_2 \supset \cdots \supset L_m' \supset J_m.$$

Further, for all i $(1 \le i \le m)$, let L_i be the unique proper subinterval of I^* satisfying l_i. Then the conjuncts in the second row of (60) ensure that J_i ends before L_i begins, and, moreover, J_i is the only subinterval of I^* satisfying either a_i^0 or a_i^1 which ends before L_i begins. In particular, no subinterval of I^* satisfying either a_i^0 or a_i^1 ends before J_i ends.

Symmetrically, let $r_1, \ldots, r_m, r_1', \ldots, r_{m+1}'$ be new event-atoms, and let ψ_4^r be the conjunction of the following collection of formulas, where i $(1 \le i \le m)$:

$$\{a_0\}\langle r_1'\rangle\top, \qquad \{r_i'\}\{a_i^1\}\langle r_{i+1}'\rangle\top,$$
$$\{r_i'\}\{a_i^1\}_<\langle r_i\rangle\top, \quad \{r_i\}_>\{a_i^1\}\langle z\rangle\top, \quad \{r_i\}_>\neg\langle a_i^0\rangle\top.$$

Let ψ_4 be $\psi_4^l \wedge \psi_4^r$. We thus have:

Claim 3 *If $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_4$ and $1 \le i \le m$, then no subinterval of I^* satisfying either a_i^0 or a_i^1 can end before the first i-witness ends or begin after the last i-witness begins.*

We are now ready to achieve the main ask of Section 6.1. Fix $n > 0$. Set $m = 2n + 1$, let $\psi_1, \ldots, \psi_4$ be as above, and let ψ_5 be the conjunction of the following formulas, where $1 \le i \le m$, $0 \le h \le 1$ and $0 \le h' \le 1$:

$$(61) \qquad\qquad\qquad [a_i^h]\neg\langle a_i^{h'}\rangle\top.$$

Claim 4 *Let $\mathcal{A} \models_{I^*} \psi_1 \wedge \cdots \wedge \psi_5$. Then, for all i ($0 \leq i \leq m$), there exist exactly 2^i proper subintervals of I^* satisfying either a_i^0 or a_i^1. These intervals are arranged as in Fig. 7. Hence, there are exactly 2^{2n} proper subintervals of I^* satisfying a_m^0.*

PROOF. Suppose $0 \leq i \leq m$. Certainly, there are exactly 2^i i-witnesses. It suffices to show that no other proper subinterval of I^* satisfies a_i^0 or a_i^1. Suppose, for contradiction, $J \subset I^*$ and J satisfies a_i^h, but J is not an i-witness. By ψ_5, J neither properly includes nor is properly included in any i-witness. Hence, the following possibilities are exhaustive: (i) J ends before the first i-witness ends; (ii) J begins after one i-witness begins and ends before the next one ends; and (iii) J begins after the last i-witness begins. But Claims 2 and 3 rule out all these possibilities. Hence, all proper subintervals of I^* satisfying a_i^0 or a_i^1 are i-witnesses.

As a final trick, we show how the 2^{2n} a_m^0-intervals identified in Claim 4 can be consecutively numbered. Let $d_1^0, \ldots, d_{m-1}^0$, $d_1^1, \ldots, d_{m-1}^1$ be new event-atoms. Think of d_i^h ($1 \leq i \leq m-1$, $0 \leq h \leq 1$) as stating that the ith digit in a certain $(m-1)$-digit binary numeral is h, where the first digit is the most significant and the $(m-1)$th the least significant. Let ψ_6 be the conjunction of the following formulas, where $1 \leq i < m$ and $0 \leq h \leq 1$:

(62) $\qquad [a_i^h][a_m^0]\langle d_i^h\rangle\top, \quad [a_m^0](\neg\langle d_i^0\rangle\top \vee \neg\langle d_i^1\rangle\top).$

Claim 5 *Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \psi_6$, and consider the 2^{2n} m-witnesses which satisfy a_m^0. Let these intervals be numbered in order of temporal precedence as $J_0, \ldots, J_{2^{2n}-1}$. For all k ($0 \leq k < 2^{2n}$), and all i ($1 \leq i \leq 2n$) denote the ith digit in the $2n$-digit binary numeral for k (counting the most significant as the first) by $k[i]$. Then we have:*

$$
k[i] = \begin{cases} 1 \text{ iff } \mathcal{A} \models_{J_k} \langle d_i^1\rangle\top \\ 0 \text{ iff } \mathcal{A} \models_{J_k} \langle d_i^0\rangle\top. \end{cases}
$$

PROOF. By formula ψ_6 and inspection of Fig. 7.

Let us refer to the 2^{2n} a_m^0-intervals identified in Claim 4 as *tiles*, and let us write a_m^0 more suggestively as t. We continue to denote the tiles in order of temporal precedence as $J_0, \ldots, J_{2^{2n}-1}$, and we say that J_k ($0 \leq k < 2^{2n}$) *has index* k. If J is any tile, denote its index by k_J. In that case, Claim 5

Fig. 9. Arrangement of g_i^0- and g_i^1-intervals within an i-witness ($1 \leq i \leq m$)

lets us read $\mathcal{A} \models_J \langle d_i^h \rangle \top$ as 'saying' that the ith digit in the $2n$-digit binary representation of k_J is h.

6.2 Organizing the tiles into a grid

Group the 2^{2n} tiles into 2^n blocks, each containing 2^n consecutive tiles. Regarding each block as a row gives us a $2^n \times 2^n$ grid. If J and J' are tiles, then J' lies immediately above J in this grid in case $k_{J'} = k_J + 2^n$; similarly, J' lies immediately to the right of J in the grid in case $k_{J'} = k_J + 1$ and the last n bits of k_J are not all 1s. We now write formulas ensuring that, for all tiles J, J' such that $k_{J'} = k_J + 2^n$, we can identify an interval L such that J is the first tile included in L and J' is the last.

Continuing to write m for $2n + 1$, let $g_1^0, \ldots, g_m^0$, $g_1^1, \ldots, g_m^1$, be new event-atoms, and let ψ_7 be the conjunction of the following formulas, where $0 \leq i < m$ and $0 \leq h \leq 1$:

(63) $[a_i^h]\{g_{i+1}^0\}_{>}\langle a_{i+1}^0 \rangle \top, \quad [a_i^h]\{g_{i+1}^1\}_{<}\langle a_{i+1}^0 \rangle \top, \quad [a_i^h]\{g_{i+1}^1\}_{>}\langle a_{i+1}^1 \rangle \top.$

Fig. 9 illustrates how the g_{i+1}^0- and g_{i+1}^1- intervals are arranged under an i-witness if $\mathcal{A} \models_I \psi_1 \wedge \psi_7$. It helps to think of the g_i^h-intervals as 'short' intervals separating consecutive i-witnesses.

Now let $f_0, f_1^0, \ldots, f_{2n}^0, f_1^1, \ldots, f_{2n}^1$ be new event-atoms, write f^0 alternatively as f_0^0 or f_0^1, and let ψ_8 be the conjunction of the following formulas, where $0 \leq i < 2n$, $0 \leq h \leq 1$ and $0 \leq h' \leq 1$:

(64) $\langle f_0 \rangle \top, \qquad\qquad [f_{2n}^h]\langle a_{2n}^h \rangle \top, \qquad\qquad [f_{2n}^h]\{(a_{2n}^h)^f\}_{<}\neg\langle a_{2n+1}^{h'} \rangle \top,$

$\quad [f_i^h]\{(a_i^h)^f\}_{<}\neg\langle a_{i+1}^{h'} \rangle \top, \quad [f_i^h]\{f_{i+1}^{h'}\}_{<}\neg\langle g_{i+1}^{h'} \rangle \top$

To motivate this construction, it helps to imagine the f_i^h-intervals guaranteed by ψ_8 as distributed similarly to the corresponding a_i^h-intervals in Fig. 7, except that the end-point of every f_i^h-interval is shifted right by a 'large' fixed amount—specifically, an amount equal to the time occupied by 2^n consecutive tiles. Fig. 10 illustrates how f_i^h- and $f_{i+1}^{h'}$-intervals are arranged in such an interpretation.

Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \cdots \wedge \psi_8$. Ignoring for the moment all intervals which

34

Fig. 10. Arrangement of f_i^h- and $f_{i+1}^{h'}$-intervals

are not proper subintervals of I^*, any f_i^h-interval $(1 \leq i \leq 2n, 0 \leq h \leq 1)$ properly includes an a_i^h-interval; and any f_i^h-interval $(1 \leq i < 2n, 0 \leq h \leq 1)$ properly includes a unique f_{i+1}^0-interval and a unique f_{i+1}^1-interval. Now let k be an integer with $0 \leq k < 2n$, and denote the $2n$ digits of k by $k[i]$ $(1 \leq i \leq 2n)$ as in Claim 5. Then we can form a chain of intervals $L_1 \supset \cdots \supset L_{2n}$ such that, for all i $(1 \leq i \leq 2n)$, L_i is an $f_i^{k[i]}$-interval. Moreover, for all i $(1 \leq i \leq 2n)$, L_i properly includes some $a_i^{k[i]}$-interval; so let K_i be the *first* such interval. We claim that $K_1 \supset \cdots \supset K_{2n}$. To see this, suppose $1 \leq i < 2n$, and write h for $k[i]$ and h' for $k[i+1]$. Let K be the unique $a_{i+1}^{h'}$-interval properly included in K_i. From $[f_i^h]\{(a_i^h)^f\}_< \neg \langle a_{i+1}^{h'} \rangle \top$, L_i cannot include any $a_{i+1}^{h'}$-interval which finishes before the start of K_i. By Claim 4 and inspection of Fig. 7, we see that K is therefore the first $a_{i+1}^{h'}$-interval properly included in L_i. From $[f_i^h]\{f_{i+1}^{h'}\}_< \neg \langle g_{i+1}^{h'} \rangle \top$, L_{i+1} starts before the start of K, and, since it is an $f_{i+1}^{h'}$-interval, properly includes at least one $a_{i+1}^{h'}$-interval. It follows that K is the first $a_{i+1}^{h'}$-interval properly included in L_{i+1}: in other words, $K = K_{i+1}$. Thus, $K_i \supset K_{i+1}$ as required. Hence we have:

Claim 6 *Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \cdots \wedge \psi_8$ and $0 \leq k < 2^{2n}$. Then there exists $L \subset I$ such that L is either an f_{2n}^0-interval or an f_{2n}^1-interval, and the first tile properly included in L is J_k.*

PROOF. Consider the chain $K_1 \supset \cdots \supset K_{2n}$ constructed above. The first tile properly included in L is properly included in K_{2n}. The result follows from Claim 5.

In the sequel, we use v to denote either f_{2n}^0 or f_{2n}^1 indifferently. Thus, if $\mathcal{A} \models_{I^*} \psi_1 \wedge \cdots \wedge \psi_8$, then there are at least 2^{2n} v-intervals properly included in I^*—one 'starting with' each of the 2^{2n} tiles. We now proceed to ensure that, if L is a v-interval starting with tile J_k, where $0 \leq k < 2^{2n} - 2^n$, then L includes exactly $2^n + 1$ consecutive tiles. (See Fig. 11.) To aid readability, we occasionally employ $\mathcal{TPL}^+$-formulas in the sequel; their conversion into logically equivalent $\mathcal{TPL}$-formulas is completely routine.

As a preliminary, let let $d^-, d_1^*, \ldots, d_{2n}^*$ be new event-atoms, and ψ_9 be the

35

$$
\begin{array}{c}
\overline{\qquad\qquad v \qquad\qquad} \\[2pt]
\overline{\qquad\qquad\quad v \qquad\qquad\quad} \\
\cdots \\
\cdots
\end{array}
$$

$$
\frac{t}{0} \quad \frac{t}{1} \cdots \frac{t}{2^n} \quad \frac{t}{2^n+1} \cdots \frac{t}{2^{2n}-2^n-2} \quad \frac{t}{2^{2n}-2^n-1} \cdots \frac{t}{2^{2n}-2} \; \frac{t}{2^{2n}-1}
$$

Fig. 11. Arrangement of event-atoms indicating vertical neighbourhood in the grid conjunction of the following formulas, where $1 \leq i \leq n$:

$$
(65) \quad [t]\left(\langle d^- \rangle \top \leftrightarrow \bigvee_{1 \leq j \leq n} \langle d_j^0 \rangle \top \right), \quad [t]\left(\langle d_i^* \rangle \top \leftrightarrow \left(\langle d_i^0 \rangle \top \wedge \bigwedge_{i < j \leq n} \langle d_j^1 \rangle \top \right) \right).
$$

The purpose of ψ_9 is to enable us to simulate addition of 2^n to binary numerals representing integers less than $2^{2n} - 2^n$. Suppose $\mathcal{A} \models_{I^*} \psi_1 \wedge \cdots \wedge \psi_9$. Then it is routine to check that: (i) for any tile J, $\mathcal{A} \models_J \langle d^- \rangle \top$ if and only if k_J is in the range $0 \leq k_J < 2^{2n} - 2^n$; (ii) for any tile J with $0 \leq k_J < 2^{2n} - 2^n$, $\mathcal{A} \models_J \langle d_i^* \rangle \top$ if and only if i is the least integer such that the jth digit in the $2n$-digit binary representation of k_J is 1 for all j in the range $i < j \leq n$. With this interpretation in mind, let ψ_{10} be the conjunction of the following formulas, where $1 \leq i \leq n$:

$$
\begin{array}{c}
\bigwedge_{1 \leq i \leq n} [v](\{t^f\}\langle d_i^* \rangle \top \rightarrow \{t^l\}(\bigwedge_{i < j \leq n} \langle d_j^0 \rangle \top \wedge \langle d_i^1 \rangle \top)), \\
(66) \quad \bigwedge_{1 \leq i \leq n} \bigwedge_{1 \leq j < i} [v](\{t^f\}\langle d_i^* \rangle \top \rightarrow (\{t^f\}\langle d_j^1 \rangle \top \leftrightarrow \{t^l\}\langle d_j^1 \rangle \top)), \\
\bigwedge_{n < j \leq 2n} [v](\{t^f\}\langle d^- \rangle \top \rightarrow (\{t^f\}\langle d_j^1 \rangle \top \leftrightarrow \{t^l\}\langle d_j^1 \rangle \top)).
\end{array}
$$

If $\mathcal{A} \models_{I^*} \psi_1 \wedge \cdots \wedge \psi_{10}$, then we can read ψ_{10} as stating that, for every subinterval $L \subset I^*$ satisfying v, if the first tile included in L has index less than $2^{2n} - 2^n$, then the indices of the first and last tiles included in L differ by precisely 2^n. Pictorially, we have the arrangement of v-satisfying intervals shown in Fig. 11. The corresponding formulas $\psi_{11}, \ldots, \psi_{14}$ required to establish a suitable arrangement of event-types h encoding horizontal neighbourhood are analogous and need not be spelled out here.

6.3 Encoding Tiling Problems

We are now ready to prove the main result of this section.

Theorem 2 *The satisfiability problem for $\mathcal{TPL}$ is NEXPTIME-hard.*

PROOF. Let (C, H, V) be any exponential tiling problem and $c'_0, \ldots, c'_{n-1}$ an instance of size n. Construct the formulas $\psi_1, \ldots, \psi_{14}$ as above. If $C = \{c_0, \ldots, c_{M-1}\}$, take the c_j ($0 \le j < M$) to be event-atoms, and let ψ_T be the conjunction of the following two formulas:

$$[t] \bigvee_{0 \le j < M} \langle c_j \rangle \top, \quad [t] \bigwedge_{0 \le j < j' < M} (\neg \langle c_j \rangle \top \vee \neg \langle c_{j'} \rangle \top).$$

Given a tile J, we regard the satisfaction of an event-atom c_j by a proper subinterval of J as indicating that the tile J is coloured by c_j. The formula ψ_T simply states that each tile has exactly one colour chosen from C.

Let ψ_H be the conjunction of the following formulas, where $(c_i, c_j) \notin H$:

$$[h](\{t^f\} \neg \langle c_i \rangle \top \vee \{t^l\} \neg \langle c_j \rangle \top).$$

Let ψ_V be the conjunction of the following formulas, where $(c_i, c_j) \notin V$:

$$[v](\{t^f\} \neg \langle c_i \rangle \top \vee \{t^l\} \neg \langle c_j \rangle \top).$$

The motivation for ψ_H and ψ_V should be obvious. Finally, we encode the fact that the initial tile J_0 is required to have colour c'_0 using the formula

$$\langle t \rangle (\langle d_0^0 \rangle \top \wedge \cdots \wedge \langle d_{2n}^0 \rangle \top \wedge \langle c'_0 \rangle \top),$$

and similarly for the other tiles which are required to have a particular colour. Denote the conjunction of all these formulas by ψ_I. From the above constructions, it is routine to verify that the instance $c'_0, \ldots, c'_{n-1}$ of (C, H, V) is positive if and only if

$$\psi_1 \wedge \ldots \wedge \psi_{14} \wedge \psi_T \wedge \psi_H \wedge \psi_V \wedge \psi_I$$

is satisfiable. This completes the reduction.

Corollary 2 *The satisfiability problems for $\mathcal{TPL}$ and $\mathcal{TPL}^+$ are NEXPTIME-complete.*

Actually, a glance at the proof of Theorem 2 reveals that it shows a little more. The satisfiability problem for the fragment of $\mathcal{TPL}$ in which the modal depth of formulas is limited to 3 is still NEXPTIME-hard, since only formulas from this fragment were used to encode instances of tiling problems. The fact that only very simple $\mathcal{TPL}$-formulas figure in this proof is crucial for the argument of Section 6.4.

We have now established that satisfiability in $\mathcal{TPL}$ is NEXPTIME-complete. Since $\mathcal{TPL}$ closely matches our English fragment $\mathcal{TPE}$, we are close to answering the question with which we began: What is the computational complexity of determining logical relationships between sentences employing the temporal constructions featured in sentences such as (1)–(3)? However, one small matter remains. It should be obvious that the grammar of $\mathcal{TPE}$, restricted as it is, accepts strings whose status as English sentences—on syntactic, semantic or pragmatic grounds—is doubtful. What we require, then, is an assurance that no linguistically motivated tightening of our fragment $\mathcal{TPE}$ could affect the above complexity result.

At first sight, this seems an impossible demand, since we cannot know in advance what refinements might be made to our English grammar. However, it turns out that the proof strategy employed above yields an easy solution. Obviously, eliminating marginal or awkward sentences from $\mathcal{TPE}$ can only cause the fragment to contract, and so cannot increase the computational complexity of its satisfiability problem. The only possibility we must guard against is that such a contraction might invalidate the NEXPTIME-hardness result. And this is where the details of the proof of that result come to our rescue. For that proof depends on the encoding of tiling problems by the formulas ψ_1–ψ_{14}, ψ_T, ψ_H, ψ_V and ψ_I. All we need do then is to examine these formulas one by one and check that they can be generated, using the grammar presented above, by good, idiomatic English sentences. If so, we know that any linguistically motivated restrictions on $\mathcal{TPL}$ will still include these sentences, and will assign them the advertised satisfaction-conditions. Thus, the NEXPTIME-completeness result will still apply to any linguistically motivated tightening of the grammar.

The formulas ψ_1–ψ_3, given in (57)–(59), consist of conjuncts of the forms

$$\{a_0\}\langle z\rangle\top, \ \ [a_{i+1}^h]\langle p_i^h\rangle\top, \ \ [a_i^h]\{a_{i+1}^0\}_{>}\langle a_{i+1}^1\rangle\top, \ \ [a_i^h]\{a_{i+1}^1\}\langle z\rangle\top.$$

But these formulas express the meanings of the unobjectionable $\mathcal{TPE}$-sentences

(67) During the occurrence of a_0, z occurred

(68) During every occurrence of a_{i+1}^h, p_i^h occurred

(69) During every occurrence of a_i^h, a_{i+1}^1 occurred after the occurrence of a_{i+1}^0

(70) During every occurrence of a_i^h, z occurred during the occurrence of a_{i+1}^1.

For added naturalness, we have fronted one preposition-phrase in each of

these sentences; this facility could easily be incorporated into our grammar, of course.

Formula ψ_4^l, given in (60), additionally involves conjuncts of the forms

$$\{l_i'\}\{a_i^0\}\langle l_{i+1}'\rangle\top, \quad \{l_i'\}\{a_i^0\}_>\langle l_i\rangle\top, \quad \{l_i\}_<\neg\langle a_i^1\rangle\top, \quad \{l_i\}_<\{a_i^0\}\langle z\rangle\top;$$

these formulas express the meanings of the unobjectionable $\mathcal{TPE}$-sentences

(71) During the occurrence of l_i', l_{i+1}' occurred while a_i^0 occurred
(72) During the occurrence of l_i', l_i occurred after a_i^0 occurred
(73) a_i^1 did not occur before l_i occurred
(74) Before l_i occurred, z occurred during the occurrence of a_i^0.

Formulas ψ_5 and ψ_6, given in (61)–(62), additionally involve conjuncts of the forms
$$[a_i^h]\neg\langle a_i^{h'}\rangle\top, \quad [a_i^h][a_m^0]\langle d_i^h\rangle\top, \quad [a_m^0](\neg\langle d_i^0\rangle\top \vee \neg\langle d_i^1\rangle\top);$$
these are generated by the unobjectionable $\mathcal{TPE}$-sentences

(75) During every occurrence of a_i^h, $a_i^{h'}$ did not occur
(76) During every occurrence of a_i^h, d_i^h occurred during every occurrence of a_m^0
(77) During every occurrence of a_m^0, either d_i^0 did not occur or d_i^1 did not occur.

For added naturalness, we have helped ourselves to the word either, which could be easily incorporated into our grammar.

Formula ψ_7, given in (63), presents no new difficulties. Formula ψ_8, given in (64), additionally involves conjuncts of the forms

$$\langle f_0\rangle\top, \quad [f_i^h]\{(a_i^h)^f\}_<\neg\langle a_{i+1}^{h'}\rangle\top, \quad [f_i^h]\{f_{i+1}^{h'}\}_<\neg\langle g_{i+1}^{h'}\rangle\top;$$

these are generated by the unobjectionable $\mathcal{TPE}$-sentences

(78) f_0 occurred
(79) During every occurrence of f_i^h, $a_{i+1}^{h'}$ did not occur before the first occurrence of a_i^h
(80) During every occurrence of f_i^h, $g_{i+1}^{h'}$ did not occur before the occurrence of $f_{i+1}^{h'}$.

In the presence of the preceding formulas, formula ψ_9, given in (65), can be equivalently expressed as a conjunction of formulas of the forms

$$[t]\left(\neg\langle e_0\rangle\top \vee \langle e_1\rangle\top \vee \ldots \vee \langle e_l\rangle\top\right), \quad [t]\left(\langle e_0\rangle\top \vee \langle e_1\rangle\top \vee \ldots \vee \langle e_l\rangle\top\right),$$

39

for various collections of event-atoms $e_0, \ldots, e_l$; these correspond to the $\mathcal{TPE}$-sentences

(81) During every occurrence of t either e_0 did not occur or e_1 occurred or
     ... or e_l occurred
(82) During every occurrence of t either e_0 occurred or ... or e_l occurred.

These sentences are certainly grammatical. Admittedly, huge disjunctions might be said not to belong to English as she is spoken; however, it is a simple matter to convert the relevant formulas equisatisfiably into formulas where disjunctions involve no more than three disjuncts, thus avoiding even this degree of unnaturalness. Formulas ψ_{10}–ψ_{14}, ψ_T, ψ_H, ψ_V and ψ_I present no new difficulties. We conclude that no linguistically motivated tightening of our fragment $\mathcal{TPE}$ could change the above complexity result. Determining the satisfiability of sets of sentences featuring the temporal constructions studied in this paper is indeed NEXPTIME-complete.

7 Conclusion

In this paper, we defined the fragment of temporal English $\mathcal{TPE}$, together with a matching interval temporal logic $\mathcal{TPL}$. The satisfiability problem for $\mathcal{TPL}$ was shown to be complete for the complexity class NEXPTIME. In view of the intimate connection between $\mathcal{TPE}$ and $\mathcal{TPL}$, we take this result to indicate the complexity of performing logical deductions in the fragment of temporal English in question, and thus to give a rough measure of the expressive resources which the grammatical constructions it features—primarily, temporal prepositions—put at speakers' disposal. By the standards of most interval temporal logics, $\mathcal{TPL}$ has low complexity. In the search for logics of limited expressive power, fragments owing their salience to the syntax of natural language are a good place to look.

We endeavoured throughout to be faithful to the facts of English usage while retaining a reasonably perspicuous formal system, amenable to mathematical analysis. These two aims are to some extent antagonistic, of course. Natural languages are products of human biology and human civilization, and as such do not always admit of a comfortable mathematical description. Thus, even the simple fragment of English considered here skirts many delicate issues of syntax, and includes sentences about whose exact semantics even native speakers are uncertain. In this situation, we have occasionally had to legislate, sometimes in whatever way is mathematically most convenient. Nevertheless, while faithfulness to the linguistic data is a virtue, it is all too easy, in pursuit of this virtue, to lose sight of the remarkable logical regularity of the constructions studied here; and it is this regularity that has been the focus of

our attention. To what extent this analysis can be usefully extended to cover other temporal constructions in English (and other natural languages), and what effects such extensions will have on the complexity of satisfiability in the accompanying logic, remain open.

Appendix: The grammar rules for $\mathcal{TPE}$

Syntax   Open-class lexicon

$S \rightarrow S, PP$ $NP_D \rightarrow Det_D, N_D^1$ $S^0 \rightarrow$ an interrupt was received/int-rec
$S \rightarrow S, Conj, S$ $N_D^1 \rightarrow N^0$ $S^0 \rightarrow$ the main process ran/main
$S \rightarrow S^0$ $N_!^1 \rightarrow OAdj, N^0$ $\ldots$
$S \rightarrow Neg, S^0$ $PP \rightarrow P_{N,D}, NP_D$ $N^0 \rightarrow$ cycle/cyc
$S_D^1 \rightarrow S^0$ $PP \rightarrow P_{S,D}, S_D^1$ $N^0 \rightarrow$ run of the main process/main
$S_!^1 \rightarrow S^0, OAdv$ $\ldots$

Closed-class lexicon

$Det_\forall \rightarrow$ every/[ ] $OAdj \rightarrow$ first/f $OAdv \rightarrow$ for the first time/f
$Det_\exists \rightarrow$ some/$\langle\,\rangle$ $OAdj \rightarrow$ last/l $OAdv \rightarrow$ for the last time/l
$Det_! \rightarrow$ the/{ } $P_{N,D} \rightarrow$ during/$=$ $P_{S,!} \rightarrow$ while/$(=, \{\ \})$
$Neg \rightarrow$ not/$\neg$ $P_{S,!} \rightarrow$ when/$(=, \{\ \})$
$Conj \rightarrow$ and/$\wedge$ $P_{N,!} \rightarrow$ until/$<$ $P_{S,\forall} \rightarrow$ whenever/$(=, [\ ])$
$Conj \rightarrow$ or/$\vee$ $P_{N,!} \rightarrow$ before/$<$ $P_{S,!} \rightarrow$ until/$(<, \{\ \})$
    $P_{N,!} \rightarrow$ after/$>$ $P_{S,!} \rightarrow$ before/$(<, \{\ \})$
        $P_{S,!} \rightarrow$ after/$(>, \{\ \})$.

References

Bennett, D. C., 1975. Spatial and Temporal Uses of English Prepositions. Vol. 17 of Longman Linguistics Library. Longman, London.

Börger, E., Grädel, E., Gurevich, Y., 1997. The Classical Decision Problem. Perspectives in Mathematical Logic. Springer-Verlag, Berlin.

Bowman, H., Thompson, S., 2003. A decision procedure and complete axiomatization of finite interval temporal logic with projection. Journal of Logic and Computation 13 (2), 195–239.

Crouch, R., Pullman, S., 1993. Time and modality in a natural language planning system. Artificial Intelligence 63, 265–304.

Davidson, D., 1967. The logical form of action sentences. In: Rescher, N. (Ed.),

The Logic of Decision and Action. University of Pittsburgh Press, Pittsburgh, pp. 81–95.

Dowty, D., 1979. Word Meaning and Montague Grammar. D. Reidel, Dordrecht.

Goranko, V., Montanari, A., Sciavicco, G., 2003. A general tableau method for propositional interval temporal logics. In: Mayer, M. C., Pirri, F. (Eds.), Tableaux 2003. Vol. 2796 of Lecture Notes in Artificial Intelligence. Springer, Berlin, pp. 102–117.

Goranko, V., Montanari, A., Sciavicco, G., 2004. A road map of interval temporal logics and duration calculi. Journal of Applied Non-classical Logics 14 (1–2), 9–54.

Halpern, J. Y., Shoham, Y., Oct. 1991. A propositional modal logic of time intervals. Journal of the ACM 38 (4), 935–962.

Higginbotham, J., Pianesi, F., Varzi, A. C. (Eds.), 2000. Speaking of Events. Oxford University Press, New York.

Huddleston, R., Pullum, G., 2002. The Cambridge Grammar of the English Language. Cambridge University Press, Cambridge.

Hwang, C. H., Schubert, L. K., 1994. Interpreting tense, aspect and time adverbials. In: Gabbay, D. M., Ohlbach, H.-J. (Eds.), Proc. first international conference on temporal logic. Vol. 827 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, pp. 238–264.

Kamp, H., Reyle, U., 1993. From Discourse to Logic: Introduction to Modeltheoretic Semantics of Natural Language, Formal Logic and Discourse Representation Theory. Kluwer Academic Publishers, London, Boston, Dordrecht, studies in Linguistics and Philosophy, Volume 42.

Moszkowski, B., 1985. A temporal logic for multilevel reasoning about hardware. Computer 18, 10–19.

Ogihara, T., 1996. Tense, Attitudes and Scope. No. 58 in Studies in Linguistics and Philosophy. Kluwer, Dordrecht.

Paech, B., 1988. Gentzen-systems for propositional temporal logics. In: Börger, E., Kleine Büning, H., Richter, M. (Eds.), Proceedings of the 2nd Workshop on Computer Science Logic. Vol. 385 of Lecture Notes in Computer Science. Springer-Verlag, Berlin, pp. 240–253.

Parsons, T., 1990. Events in the Semantics of English: a Study in Subatomic Semantics. The MIT Press, Cambridge, Massachusetts.

Pratt, I., Francez, N., 2001. Temporal prepositions and temporal generalized quantifiers. Linguistics and Philosophy 24, 187–222.

Steedman, M., 1996. Temporality. In: van Benthem, J., ter Meulen, A. (Eds.), Handbook of Logic and Language. Elsevier, pp. 895–935.

Stump, G., 1985. The Semantic Variability of Absolute Constructions. Kluwer, Dordrecht.

ter Meulen, A. G., 1996. Representing Time in Natural Language. MIT Press, Cambridge, MA.

Vendler, Z., 1967. Linguistics in Philosophy. Cornell University Press, Ithaca.

Venema, Y., Fall 1990. Expressiveness and completeness of an interval tense

logic. Notre Dame Journal of Formal Logic 31 (4), 529–547.

Venema, Y., 1991. A modal logic for chopping intervals. Journal of Logic and Computation 1 (4), 453–476.

von Stechow, A., 2002. Temporal preposition-phrases with quantifiers: Some additions to Pratt and Francez (2001). Linguistics and Philosophy 25, 755–800.

Zucchi, S., White, M., 2001. Twigs, sequences and the temporal constitution of predicates. Linguistics and Philosophy 24, 223–270.