Chapter 10

TEMPORAL LOGIC

Yde Venema

1   Introduction

Time must be the most paradoxical concept our minds have to deal with.
To quote from the *Confessions* of St. Augustine

> What then, is time? If no one asks me, I know; but if I wish to
> explain it to someone who should ask me, I do not know.

One of time's most puzzling aspects concerns its ontological status: on the
one hand it is a subjective and relative notion, based on our conscious ex-
perience of successive events; yet on the other hand, our civilization and
technology are based on the understanding that something like objective,
absolute Time exists. Some philosophers have taken this paradox so far as
to conclude that time is unreal; others, accepting the existence of absolute
time, have engaged in heated debates regarding its structure, be it linear or
circular, bounded or unbounded, dense or discrete.

But even if we leave these metaphysical issues aside, it is obvious that
time plays such a fundamental role in our thinking that there is a clear need
for *precise reasoning* about it, such as we see in Physics, formal Linguistics,
Computer Science, and Artificial Intelligence. While these enterprises are
not necessarily concerned with the *same* concept of time, they all could go
under the heading of Temporal Logic. Often, however, a more restricted,
technical definition is used in which temporal logic—or tense logic—is a
branch of *modal* logic, an approach that began about forty years ago with
the work of Arthur Prior. We will largely confine ourselves to this modal
perspective here, though, as we shall see, this still includes a great variety
of systems.

In the next section we discuss some of the most well-known mathematical modellings of time. These are the structures the formal languages of temporal logic are designed to talk about. The main part of this paper, Section 3, is devoted to a fairly detailed exposition of Prior's basic tense logic; the aim of this is not only to present this particular system, but perhaps even more to introduce the kinds of questions that temporal logicians tend to ask. In the Sections 4 and 5 we describe some extensions and alternatives to this base system. In Section 6 we sketch some developments that have taken place over the last ten years or so. Finally, in the Epilogue we try to answer the question what Temporal Logic *is*.

2   Flows of Time

Before we can start a discussion of various logics of time, it helps to look at some standard mathematical models of time. When asked to think of time in an abstract way, many people will form a picture of a line — only the simplest of the many spatial metaphors that people use for temporal concepts! The mathematics of this picture is given by a set of time points, together with an ordering relation and perhaps a metric measuring the distance between two points. Later on, we will discuss some objections and alternatives to this point-based paradigm. For now, let us formally represent time as a *frame*; that is, a structure $\mathcal{T} = (T, <)$ such that $<$ is a binary relation on T, called the *precedence relation*. Elements of T are called *time points*; if a pair (s, t) belongs to $<$ we say that s is *earlier than* t. In the remainder of this section we will discuss a number of more or less intuitive conditions that have been imposed on such structures in order to make them useful as models of time. (We will frequently use first and second order logic for describing these properties; the *first order frame language* we use will have only one dyadic predicate symbol, which is denoted by R and interpreted as $<$.)

Obviously, many frames will not qualify as intuitively acceptable representations of time. At a minimum one should require that $<$ be irreflexive and transitive. A frame satisfying these conditions will be called a *flow of time*. Flows of time are known from mathematics as strict partial orders, and in accordance with this we will use familiar notation like $s > t$ for 's is later than t' or $s \leq t$ for 'either $s = t$ or $s < t$'. For a point t, the set $\{s \in T | t < s\}$ will be called the *future of t*; the *past of t* is defined likewise. (In the sequel, we will omit definitions pertaining to the past if they mirror an obvious counterpart for the future.)

Standard candidates are given by the familiar orderings of well-known number sets: $\mathcal{N} = (\mathbb{N}, <)$ (the natural numbers), $\mathcal{Z} = (\mathbb{Z}, <)$ (the integer numbers), $\mathcal{Q} = (\mathbb{Q}, <)$ (the rational numbers) and $\mathcal{R} = (\mathbb{R}, <)$ (the real numbers). Less familiar examples are the *binary tree* $\mathcal{B} = (\mathbb{B}, \prec)$ (where $\mathbb{B}$ is the set of sequences of 0s and 1s, while $s \prec t$ holds if s is an initial segment of t), and four-dimensional *Minkowski spacetime* $\mathcal{S} = (\mathbb{R}^4, \lhd)$; here we put $(x_0, x_1, x_2, t) \lhd (x_0', x_1', x_2', t')$ if not only the temporal component of the first point is smaller than that of the second one ($t < t'$), but also the spatial distance between the two points should enable one to *reach* the one point from the other without having to travel faster than the speed of light.

Observe that our definition excludes circular time: if there were a series of time points $s_1 < s_2 < \ldots s_n < s_1$ then by transitivity we would have $s_1 < s_1$ which is not possible since we assumed our flow of time to be irreflexive. Since it is not the *logician*'s task to choose between different ontologies, why not allow circular time? After all, many civilizations have regarded time as being essentially cyclic in nature. Also, practical applications of circular time are easily conceivable, such as the construction of rotas. Our only reason is simply that circular time has received very little attention in the logical literature.

On the other hand, the reader may have missed one condition in the definition of a flow of time, namely linearity. A strict partial ordering is called *linear* if any two distinct points are related; expressed in first order logic, the structure is to satisfy the sentence $\forall xy\, (Rxy \vee x = y \vee Ryx)$. This perspective on time is dominant in science and, probably for that reason, has become the standard in most people's minds; in particular, all of our number examples are linear flows of time. Nevertheless, so-called *branching-time* structures such as $\mathcal{B}$ and $\mathcal{S}$ have received a lot of attention in the literature on temporal logic. A structure is called *branching to the future* if there is some point having two unrelated points in its future, and *not branching to the future*, if on the contrary, the future of each point is linearly ordered. A flow of time is *not branching* if it is neither branching to the future nor branching to the past; note that this condition differs from linearity in that it does not exclude 'parallel' time lines. In the literature one often encounters the condition that flows of time are allowed to branch to the future, but not to the past; this condition reflects the idea that at any moment, the past is determined while the future is not. As we will see later, the logic of branching time ties up with the logic of necessity and possibility, that is, with alethic modal logic [See Chapter 7]. In the sequel of this section we will confine ourselves to linear time, but this is not to say that the concepts to be defined would not make sense outside of this context.

Questions concerning the boundedness of time have occupied philosophers, theologians and physicists for centuries, but for the logician this is generally not the most interesting issue. Let us just mention the definitions pertaining to the future: $\mathcal{T}$ *has a first point* if it satisfies $\exists x \forall y\, (Rxy \lor x = y)$, while it is called *right-serial* if each point has a non-empty future.

A more fundamental choice, it seems, is that between *denseness* and *discreteness* of a flow of time. A linear ordering $\mathcal{T}$ is *dense* if between any two distinct points we can find a third point; formally: $\forall xy\, (Rxy \to \exists z\, (Rxz \land Rzy))$. Because this way of representing time is very convenient for modeling the notion of movement, we have become quite familiar with dense flows of time such as the orderings of the rational or the real numbers. However, for computer scientists and economists time has a very different flavor in that it is supposed to proceed in discrete *steps*; that is, with each non-final point we associate a *next* point or *immediate successor*. In first order logic: $\forall xy\, (Rxy \to \exists z\, (Rxz \land \neg \exists u\, (Rxu \land Ruz)))$. Standard examples of discrete flows of time are given by the natural or the integer numbers.

Density is often confused with *continuity*. Suppose that we cut the set of rational numbers into a left and a right half, of numbers smaller and bigger than $\sqrt{2}$, respectively. Such a cut, without a proper point on either edge, is called a *gap*, and a flow of time is called *continuous* if it has no gaps. $\mathcal{Q}$ thus forms the standard counterexample, whereas $\mathcal{R}$ and $\mathcal{Z}$ are continuous.

Unlike the properties discussed before, continuity is essentially a second order notion, its definition necessarily involving a quantification over *sets* of time points. There are many other interesting second order conditions that one may impose on temporal structures. For example, one might argue that abstract time exists independently of the events 'filling it', and that therefore, the structure of time should be 'the same everywhere'. One way of making this precise is to demand that a flow of time is *homogeneous*: for any two points s and t of $\mathcal{T}$, there should be an automorphism of T (that is, a bijection f from T onto T satisfying $x < y$ iff $f(x) < f(y)$ for all x and y in T) mapping s to t. Another second order property that we will meet further on is that of having *finite intervals*; this means that there can be at most finitely many points between any two points. Observe that this condition implies discreteness of both $<$ and $>$.

3   Basic Tense Logic

In this section we will show temporal logic at work. That is, we introduce Prior's basic system of temporal logic, and discuss some of the fundamental logical questions pertaining to it.

Syntax and Semantics

In order to define the syntax and semantics of temporal logic, we should first note that temporal logic is an extension of classical propositional logic. Recall that classically, propositional formulas are interpreted as truth values (either 1 for 'true' or 0 for 'false'); this truth value is inductively determined by a *valuation*: a function mapping propositional variables to truth values. Once we know the valuation, the truth value of any formula is fixed. Now what to do with the fact that the truth value of statements like 'it is raining' or 'I am wearing an umbrella' will *change* from time to time? For instance, it may be raining today but sunny tomorrow; or, I may be wearing my umbrella now but fold it in some time after the rain stops.

The first basic idea underlying temporal logic is to address this issue by making valuations time-dependent; more precisely, one associates a separate valuation with each point of a given flow of time. Formally, let $\mathcal{T} = (T, <)$ be a flow of time; a *valuation on* $\mathcal{T}$ is a map $\pi : (T \rightarrow (\Phi \rightarrow \{0, 1\}))$; here Φ denotes the set of propositional variables. A *model* is a pair $\mathcal{M} = (\mathcal{T}, \pi)$ consisting of a flow of time and a valuation.

Observe that with this definition we can already interpret classical formulas in each point of a model, in a standard way. For instance, we will say that the formula $p \wedge \neg q$ is true at a time point t precisely if $\pi(t)(p) = 1$ and $\pi(t)(q) = 0$. The spice of temporal logic, however, lies in its second basic idea, namely to use *new*, non-classical connectives to relate the truth of formulas in possibly *distinct* time points. In this section, we discuss two such operators: F and P. These names are mnemonics for 'future' and 'past' respectively: the intended meaning of the formula '$F\varphi$' is 'at some time in the future, φ is the case', while '$P\varphi$' is to be read as 'at some time in the past, φ holds'.

Formally, we define the set $\mathcal{L}_t$ of *Priorean formulas* as the smallest set containing the propositional variables that is closed under constructing new formulas using the boolean connectives $\neg$ and $\wedge$, and the temporal operators G and H. For technical reasons, we take F and P to be *defined* operators in our set-up; $F\varphi$ abbreviates $\neg G\neg\varphi$ and $P\varphi$ abbreviates $\neg H\neg\varphi$. $G\varphi$ and $H\varphi$ are read as 'henceforth, φ' and 'hitherto, φ', respectively. As further

abbreviations we use $\bot$, $\top$, $\wedge$, $\rightarrow$ and $\leftrightarrow$ in their usual meaning. We will also frequently refer to the *mirror image* of a formula; this is simply the formula one obtains by simultaneously replacing all Hs with Gs and vice versa.

Bringing the previous observations together, we can give the following inductive definition of the notion of *truth* of a formula φ at a time point t in a model $\mathcal{M} = (T, <, \pi)$:

(1)
$$
\begin{array}{lll}
\mathcal{M}, t \Vdash q & \text{if} & \pi(t)(q) = 1, \\
\mathcal{M}, t \Vdash \neg\varphi & \text{if} & \text{not } \mathcal{M}, t \Vdash \varphi, \\
\mathcal{M}, t \Vdash \varphi \wedge \psi & \text{if} & \mathcal{M}, t \Vdash \varphi \text{ and } \mathcal{M}, t \Vdash \psi, \\
\mathcal{M}, t \Vdash G\varphi & \text{if} & \mathcal{M}, s \Vdash \varphi \text{ for all } s \text{ with } t < s, \\
\mathcal{M}, t \Vdash H\varphi & \text{if} & \mathcal{M}, s \Vdash \varphi \text{ for all } s \text{ with } t > s.
\end{array}
$$

If $\mathcal{M}, t \Vdash \varphi$ we say that φ *holds* or *is true at* t.

As an example, consider the ordering $\mathcal{N}$ of the natural numbers; let τ be the valuation making q true at all numbers bigger than $1\,000$, and r at all even numbers. With this valuation, it is easy to see that the formula FGq holds at the point 0. For, the formula Gq holds at those points of which the future is a subset of the set of 'q-points', and this is the case for any number bigger than 999. But from $\mathcal{M}, 1\,000 \Vdash Gq$ and $0 < 1\,000$ it follows that $\mathcal{M}, 0 \Vdash FGq$. It is likewise easy to see that the formula FGr does *not* hold at 0, or indeed, at any point in this model; the formula GFr on the other hand holds throughout $\mathcal{M}$.

Finally, observe that from the *technical* point of view, this system is very similar to the systems defined in Chapter 7 on Modal Logic: G and H are very much like the operator L of alethic modal logic. The difference is that in Priorean temporal logic we are dealing with *two* modal operators instead of one. One might then expect that we would interpret this language in structures with *two* accessibility relations, say, R_F and R_P. And in fact we may adopt a perspective in which we see $<$ and $>$ as these two distinct accessibility relations; however, it is a crucial aspect of temporal logic that these two accessibility relations are each other's converse. The main distinction between alethic modal logic and temporal logic is thus one of aim: temporal logic starts with structures (flows of time), for which one is trying to find good modal description languages; whereas in alethic modal logic it has more often been the other way around.

Validity and Definability

Temporal logicians are generally not so much interested in the truth or falsity of formulas in specific models, but rather in those formulas that remain true throughout the flow of time even if we change the valuation. It is felt that such formulas provide essential information concerning the structure of the underlying flow of time. Formally, we say that a formula φ is *valid* on a flow of time $\mathcal{T}$, notation: $\mathcal{T} \Vdash \varphi$, if for every valuation π on $\mathcal{T}$, and every point of $\mathcal{T}$, we have $(\mathcal{T}, \pi), t \Vdash \varphi$. A formula is valid in a class of flows of time if it is valid on each member of the class. The notion of satisfiability is defined dually: we say that a formula φ is *satisfiable* in a flow of time (a class of flows of time) if its negation is not valid on the flow of time (in the class of flows of time, respectively).

As an example, we show that the formula $Fq \to FFq$ is valid on the class of dense linear orderings. Assume that $\mathcal{T}$ is a dense linear flow of time; in order to show that $Fq \to FFq$ holds on it, consider an arbitrary valuation π on $\mathcal{T}$, and an arbitrary point t in $\mathcal{T}$ such that $(\mathcal{T}, \pi), t \Vdash Fq$. By the truth definition, there is a later point s where q holds. But by density, there must be some point u between t and s; from $u < s$ we derive that Fq holds at u; but then from $t < u$ we may infer that FFq holds at t; since t and π were arbitrary, this suffices to show that $\mathcal{T} \Vdash Fq \to FFq$.

On the other hand, it is easy to see that the formula $Fq \to FFq$ is not valid on the ordering of the integers. For, take the points 0 and 1 and consider the valuation π that makes q true *only* at 1; then obviously, Fq is true at 0; but since there is no integer number between 0 and 1, the formula FFq cannot be true at 0. This shows that indeed $\mathcal{Z} \not\Vdash Fq \to FFq$. We can in fact generalize this argument to show that the formula $Fq \to FFq$ can be falsified on *every* non-dense frame. For, any non-dense frame must contain two points $s < t$ without intermediate points; so the valuation making q true only at t will make the formula $Fq \to FFq$ false at s. Hence, the formula $Fq \to FFq$ is very informative; it is a reliable witness of the density of a flow of time.

In general, we say that a Priorean formula φ *defines* a class C of flows of time within a class K if for every flow of time $\mathcal{T}$ in K, $\mathcal{T} \Vdash \varphi$ iff $\mathcal{T}$ belongs to C. If C is given as the class of frames satisfying some first order property α, we also say that φ *corresponds* to α (within K). For instance, we have just seen that the formula $Fq \to FFq$ corresponds to density.

Not every property of flows of time is definable; for instance, we can prove that there is no Priorean formula that defines the class of branching flows of time. On the other hand, there is a formula defining the flows of time

7

that are *not* branching; for, the formula $PFq \rightarrow (Pq \lor q \lor Fq)$ corresponds to non-branchingness to the future. Hence, the conjunction of this formula and its mirror image defines the flows of time that are not branching.

Especially if we confine ourselves to linear orderings, many interesting properties of flows of time *can* be defined in the Priorean language. The following table lists a number of such correspondences holding for linear flows of time; here $\diamond\varphi$ abbreviates $P\varphi \lor \varphi \lor F\varphi$, and $\square\varphi \equiv H\varphi \land \varphi \land G\varphi$. For future reference, we have given names to the modal formulas.

having a first point	$H\bot \lor PH\bot$	(A1)
left-seriality	$P\top$	(A2)
having a final point	$G\bot \lor FG\bot$	(A3)
right-seriality	$F\top$	(A4)
discreteness	$(F\top \land q \land Hq) \rightarrow FHq$	(A5)
density	$Fq \rightarrow FFq$	(A6)
continuity	$(Fq \land \diamond\neg q \land \square(q \rightarrow Hq)) \rightarrow$	
	$\rightarrow \diamond((q \land G\neg q) \lor (\neg q \land Hq))$	(A7)
having finite intervals	$G(Gq \rightarrow q) \rightarrow (FGq \rightarrow Gq) \land$	
	$\land H(Hq \rightarrow q) \rightarrow (PHq \rightarrow Hq)$	(A8)

Finally, since Priorean formulas may be interpreted on *all* frames (also ones that are not strictly partial orders), the question naturally arises whether the class of flows of time itself is definable. Since, analogous to the case of ordinary modal logic, transitivity may be defined by the formula $Gp \rightarrow GGp$, this boils down to the problem of finding a correspondent for irreflexivity (within the class of transitive frames). Unfortunately, there is *no* such formula.

Axiomatics

As we mentioned already, temporal logic starts with flows of time; but obviously, this does not diminish the interest in finding complete calculi for various classes of flows of time. Obviously, there are close connections with the axiomatics of alethic modal logic as discussed in Chapter 7. In particular, analogous to $\mathbf{K}$, there is a *minimal tense logic* for the Priorean language as well; it is called $\mathbf{K}_t$ and defined as the smallest class of Priorean formulas that is closed under the following axioms and derivation rules:

(CT)	all classical propositional tautologies	
$(DB_)$	$G(q \rightarrow r) \rightarrow (Gq \rightarrow Gr)$	
	$H(q \rightarrow r) \rightarrow (Hq \rightarrow Hr)$	(Distribution)
(CV)	$q \rightarrow GPq$	
	$q \rightarrow HFq$	(Converse)
(4)	$Gq \rightarrow GGq$	(Transitivity)
(US)	if φ is a theorem, then so is $\varphi[\psi/q]$	(Uniform Substitution)
(MP)	if φ and $\varphi \rightarrow \psi$ are theorems, then so is ψ	(Modus Ponens)
(TG)	if φ is a theorem, then so are $G\varphi$ and $H\varphi$	(Temporal Generalization)

Here $\varphi[\psi/q]$ denotes the result of substituting the formula ψ for the propositional variable q, uniformly throughout φ.

Most of these axioms and all of these rules are, perhaps under different names, familiar from ordinary modal logic. The exception is the Converse axiom (CV); as we will see, this axiom is needed to ensure that the accessibility relations for the operators G and H are each other's converse. In the chapter on Modal Logic it is discussed in detail that the formula (4) reflects the transitivity of the intended accessibility relation of a modal operator; thus, our constraints on flows of time explain the presence of (4) as an axiom. Recall that in the previous section we already saw that the property of being irreflexive is not definable in the Priorean language; now we see that irreflexivity does not even yield any extra validities. (This is not the rule in modal logics: frame conditions that are not definable in the modal language may nevertheless *imply* the validity of modal formulas.)

Theorem 3.1 *The logic K_t is sound and complete with respect to the class of all flows of time.*

For lack of space, we omit the proof of Theorem 3.1. Instead, we concentrate on completeness for the class of linear flows of time. Let **Lin** be the extension of $\mathbf{K}_t$ with the axiom (NB), which is the conjunction of the axiom $PFq \rightarrow (Pq \vee q \vee Fq)$ (defining non-branching to the future) and its mirror image $FPq \rightarrow (Fq \vee q \vee Pq)$.

Theorem 3.2 *The logic $\mathbf{Lin}$ is sound and complete with respect to the class of linear flows of time.*

PROOF. Just as the completeness proofs of Chapter 7, our proof method will make use of canonical models. Hence, let W^c be the set of maximal **Lin**-consistent sets of formulas (for unexplained terminology we refer to the modal completeness proof), and define the relation R^c on W^c by $R^c wv$ iff

$\varphi \in v$ for all $G\varphi \in w$. The structure $\mathcal{F} = (W^c, R^c)$ is called the *canonical frame*; on it, we define the *canonical valuation* π^c so that $\pi^c(q)(w) = 1$ iff $p \in w$.

Our first aim is to prove a Truth Lemma for this model, stating that for all Priorean formulas and every point w of the *canonical model* $\mathcal{M}^c = (\mathcal{F}^c, \pi^c)$ we have that 'truth coincides with membership':

(2)
$$\mathcal{M}^c, w \Vdash \varphi \text{ iff } \varphi \in w.$$

Analogous to the proof of Theorem 4 in Chapter 7, (2) is proved by formula induction. There is only one minor problem, caused by the fact that we now have two modal operators, and only one accessibility relation. This is precisely where the Converse axioms come in: they enable us to show that the canonical accessibility relation does not only work well for G but also for H. For, we can prove (details are left to the reader) that $R^c wv$ iff $\varphi \in w$ for all $H\varphi \in v$.

Now it follows easily from (2) that every **Lin**-consistent set of formulas is satisfiable in the canonical model, but unlike the case of modal logics like **S4** we are not finished here. We need to satisfy our **Lin**-consistent set of formulas in a linear flow of time. Now it is easy to verify that the canonical accessibility relation is transitive (use the axiom (4), as in the modal completeness proofs); it is not very difficult to show that R^c is not branching (but we leave the details of this proof to the reader — use the axiom (NB)); but it is impossible to prove that R^c is a linear ordering, because in general this will not be true! The main problem is that nothing guarantees irreflexivity of canonical accessibility relation. The difficult part of the proof consists in showing that we can *transform* the canonical frame into a strict linear order, while truth of formulas is preserved.

Let us agree to call a frame $\mathcal{F} = (W, R)$ a *pseudo-line* if R is transitive and strongly connected (that is, satisfying $\forall xy(Rxy \lor x = y \lor Ryx)$). Now given any maximal **Lin**-consistent set Σ, we may restrict ourselves to the part of the canonical frame that is connected (via R^c) to Σ and still prove the analogue of the Truth Lemma (2). It thus follows that every consistent formula is satisfiable in a pseudo-line. But then the missing link in the proof of the completeness theorem for **Lin** is the following claim.

(3) If φ is satisfiable on a pseudo-line, then also on a linear flow of time.

In order to prove claims like (3), several methods of 'frame surgery' have been developed; in order to give the reader an idea of such techniques, we briefly sketch the *bulldozing* method here. Assume that φ is satisfiable in

the model $\mathcal{M} = (\mathcal{F}, \pi)$ based on the pseudo-line $\mathcal{F} = (W, R)$. The first observation is that $\mathcal{F}$ may be represented as a linear ordering $\prec$ of so-called *clusters* which are special subsets of W. Each point s of W belongs to a unique cluster C_s which is either *degenerate* (consisting of a single irreflexive point) or *proper* (if R is universal on it). The relation $\prec$ is defined such that $C_s \prec C_t$ if and only if $C_s \neq C_t$ and Rst.

The key idea is now to 'bulldoze' each proper cluster into a special linear ordering $\mathcal{L}_C$ and to replace each C with $\mathcal{L}_C$. Obviously, replacing each proper cluster with a linearly ordered model yields a linear order; but is φ still satisfiable in the new model? To understand the positive answer to this question, note that any proper cluster introduces a infinity of information recurrence in both the forward and backward directions: we can follow paths within C, moving either forwards and backwards along R, for as long as we please. Thus, when we replace a cluster C with a linear ordering, we must ensure that the linear ordering duplicates all the information in C infinitely often, and in both directions. Bulldozing does precisely this, in the most straightforward way possible. For instance, suppose that the cluster C has three elements only: s_0, s_1 and s_2, with associated classical valuations σ_0, σ_1 and σ_2. Then $\mathcal{L}_C$ is given as the model $(\mathcal{Z}, \pi_C)$; here π_C is given by $\pi_C(z) = \sigma_{z \bmod 3}$; that is, $\mathcal{L}_C$ consists of an unbounded (in both directions) series of points with associated classical valuations σ_0, σ_1 and σ_2; as in $\cdots \sigma_1 \sigma_2 \sigma_0 \sigma_1 \sigma_2 \sigma_0 \cdots$.

There is thus an obvious relation linking points in the new, transformed model to points in the old one; using this we may prove that φ is indeed satisfiable in the new model. This finishes the proof sketch of (3).   QED

Turning to the axiomatics of specific structures, let us define the following logics:

Lin.N:	**Lin** +A1 +A4 +A8
Lin.Z:	**Lin** +A2 +A4 +A8
Lin.Q:	**Lin** +A2 +A4 +A6
Lin.R:	**Lin** +A2 +A4 +A6 +A7

For these logics we have the following result.

Theorem 3.3 *The logics* **Lin.N**, **Lin.Z**, **Lin.Q** *and* **Lin.R** *are sound and complete axiomatizations of the set of validities of the flows of time* $\mathcal{N}$, $\mathcal{Z}$, $\mathcal{Q}$ *and* $\mathcal{R}$, *respectively.*

We may conclude that temporal logicians have been rather successful in axiomatizing the standard flows of times and the most natural classes of

flows of time. Nevertheless, it would be wrong to conclude that conversely, (axiomatically defined) tense logics are always characterized by a class of flows of time. As in modal logic, incompleteness is the rule; in fact, the very first example of an incomplete (poly-)modal logic was found in tense logic.

Decidability and Complexity

The completeness theorems that we mentioned in the previous subsection are all very nice, but of course, if one wants to do actual reasoning in one of these logics, further properties are required. Minimally, one wants the logic to be decidable; that is, the existence is required of a terminating algorithm separating the logic's theorems from its non-theorems. Fortunately, all the complete logics defined in the previous subsection have this property. We mention only the following results explicitly.

Theorem 3.4 *The Priorean tense logics of the classes of all flows of time, and of all linear flows of time, are decidable.*

This follows from the fact that these logics are finitely axiomatizable and have the *finite model property*. The latter may be proved through the method of *filtrations* or the method of minimal models of Chapter 7, with allowance for complexities analogous to the proof of completeness for **Lin**.

For practical purposes decidability is not enough, however; one would like to have an efficient calculus. A more fine-grained analysis is needed to reveal the *computational complexity* of temporal logics. There is not enough space to go into details here; we only mention the result that the satisfiability problem for linear time is in NP. To be more precise, one can devise a non-deterministic Turing machine algorithm that correctly tells whether a Priorean formula φ is satisfiable in a linear frame or not, while coming up with this answer within $f(\varphi)$ computation steps. Here f is a linear function that grows at the same rate as the length of the formula φ.

4 Extending the Language

Since and Until

Our discussion in the previous section was based on the basic temporal language having G and H as its only primitive operators. For many applications, however, this language is too poor in expressivity, and several extensions with new operators have been suggested. The most important of these are the dyadic operators S and U introduced by H. Kamp; their

intended meaning is respectively 'since' and 'until', as in the sentences 'Ever since the roof caved in, it's been wet in the house' and 'Until we get the roof fixed, it will be damp in the house'. Let $\mathcal{L}_{su}$ denote the extension of the Priorean language with these two new connectives, the formal truth definition of which is given as follows.

$$
\begin{aligned}
\mathcal{M}, t \Vdash U\varphi\psi \quad &\text{if} \quad \mathcal{M}, s \Vdash \varphi \text{ for some } s \text{ such that } t < s \\
&\qquad\text{and } \mathcal{M}, u \Vdash \psi \text{ for all } u \text{ with } t < u < s, \\
\mathcal{M}, t \Vdash S\varphi\psi \quad &\text{if} \quad \mathcal{M}, s \Vdash \varphi \text{ for some } s \text{ such that } s < t \\
&\qquad\text{and } \mathcal{M}, u \Vdash \psi \text{ for all } u \text{ with } s < u < t.
\end{aligned}
$$

It is interesting to observe that the 'old' operators can be expressed in this new language, for instance $F\varphi$ may be seen to abbreviate $U\varphi\top$. But conversely, the new operators really add expressive power to the language; we can prove that they cannot be defined in terms of the old.

Another interesting temporal operator is the so-called *nexttime* or *tomorrow* operator X; the formula $X\varphi$ holds at a time point t if φ holds at the next moment in time (if there is such a next moment). Obviously, such an operator only makes sense in a discrete flow of time, as, for instance, in computer science, where one wants to talk about the next state of a process. However, adding this new connective to $\mathcal{L}_{su}$ would not add any expressive power, since $X\varphi$ can already be defined as an abbreviation for $U\varphi\bot$.

This raises the question whether perhaps *every* temporal operator can be defined in this apparently expressive language $\mathcal{L}_{su}$. The answer to this question is positive; that is, we can prove some sort of *functional completeness* result for $\mathcal{L}_{su}$.

Theorem 4.1 (Kamp) *Over the class of linear, continuous orderings, every temporal operator can be defined in $\mathcal{L}_{su}$.*

where by 'temporal operator' we mean any operator whose truth definition is expressible in first order logic. The restriction in the theorem to certain flows of time is essential. In particular, once we drop the condition of linearity, the results tend to be negative; for instance, over the class of all flows of time we cannot find a *finite* expressively complete set of operators.

Finally, for the language $\mathcal{L}_{su}$ one can ask the same kind of questions as for $\mathcal{L}_t$; and indeed, several results have been proved concerning definability, axiomatizability and decidability. In general these results are positive, but for lack of space we cannot go into detail here.

13

Branching Time Languages

As mentioned above, allowing flows of time that branch to the future means that we can no longer assume that the past determines everything that is going to happen. But if our formalism has to take into account that there are many *different* courses of events possible, it seems appropriate to pay somewhat more attention to the truth definition of our future operator F. For, the intuitive meaning of $F\varphi$, namely 'it will be the case that φ', is now more ambiguous than in linear flows of time. Recall that the interpretation of $F\varphi$ that we may calculate from the truth definition (1) yields 'φ holds at some future moment of some possible course of events'. But it does not seem to be unreasonable to assume that 'it will be the case that φ' expresses the speaker's conviction that φ will be the case, in the *actual* course of events, or perhaps *no matter what* course of events. These two interpretations give rise to respectively the Ockhamist and Peircean schools in branching time logic.

   In order to compare these two approaches, assume our flows of time to be *trees*, that is, connected strict partial orders that do not branch to the past. (Connectedness forbids for instance parallel time lines.) A *branch* of a tree $\mathcal{T} = (T, <)$ is a maximal linearly ordered subset of T; the intuitive idea is that each branch through t represents a possible course of events (for a point t and a branch b, we say that t *lies on* b or that b *goes through* t if t belongs to b). In this way, we can imagine a *possible future* of t as the set of all *later* points on some fixed branch b through t; since $\mathcal{T}$ is a tree, each point will have a unique past.

   Now Peircean branching time logic interprets the proposition 'it will be the case that φ' in the second way indicated above, namely that φ is bound to happen in every possible future. To make this more precise, define the Peircean tense language as the extension of the Priorean one with the future operator $F_\square$; this operator has a second order definition, involving a quantification over all branches through the actual time point:

$\mathcal{M}, t \Vdash F_\square \varphi$ if on each branch through t there is some $s > t$ with $\mathcal{M}, s \Vdash \varphi$. (4)

   In the Ockhamist approach on the other hand, it is *meaningless* to ask about the truth value of formulas of the form $F\varphi$ or $G\varphi$ at a time point t, unless we have specified which of the possible futures of t we have in mind. In order to be able to express that something that will be the case no matter what form the future will take, Ockhamists extend the language with an alethic modal operator $\square$. Ockhamist tense logic is thus an interesting combination of modal and tense logic; perhaps the easiest way to work out

the idea formally, is to require that in Ockhamist semantics the truth value of *any* formula is evaluated at a *pair* consisting of a time point and a branch through this point (representing the actual course of events). We thus arrive at the following truth definition.

$$(5) \quad \begin{array}{lll} \mathcal{M}, t, b \Vdash q & \text{if} & \pi(t)(q) = 1, \\ \mathcal{M}, t, b \Vdash \neg\varphi & \text{if} & \text{not } \mathcal{M}, t, b \Vdash \varphi \\ \mathcal{M}, t, b \Vdash \varphi \wedge \psi & \text{if} & \mathcal{M}, t, b \Vdash \varphi \text{ and } \mathcal{M}, t, b \Vdash \psi, \\ \mathcal{M}, t, b \Vdash G\varphi & \text{if} & \mathcal{M}, s, b \Vdash \varphi \text{ for all } s \text{ on } b \text{ with } t < s, \\ \mathcal{M}, t, b \Vdash H\varphi & \text{if} & \mathcal{M}, s, b \Vdash \varphi \text{ for all } s \text{ on } b \text{ with } t > s, \\ \mathcal{M}, t, b \Vdash \Box\varphi & \text{if} & \mathcal{M}, t, c \Vdash \varphi \text{ for all branches } c \text{ through } t. \end{array}$$

It is interesting to note that the Peircean language can be seen as a *fragment* of the Ockhamist one; consider the inductively defined translation $(\cdot)^o$ mapping Peircean formulas to Ockhamist ones. The only non-trivial clause of this map concerns the future operators: $(F_\Box\varphi)^o = \Box F\varphi^o$ and $(G\varphi)^o = \Box G\varphi^o$. It is straightforward to prove that for all tree models $\mathcal{M}$, all points t in $\mathcal{M}$ and all branches b through t, we have that

$$\mathcal{M}, t \Vdash \varphi \text{ iff } \mathcal{M}, t, b \Vdash \varphi^o.$$

Many results are known concerning Peircean and Ockhamist logic; for instance, axiomatizations have been found for the Peircean logic of the class of all trees. This logic is also known to be decidable, as is its Ockhamist alternative. It is an outstanding open problem to find an explicit axiomatization for the Ockhamist tree logic.

Finally, it is obvious that one can extend these branching time logics even further, for instance with the Since and Until operators defined earlier. The 'future fragment' of such systems is closely related to so-called *computational tree logics* that have been developed within theoretical computer science for the purpose of reasoning about paths through labeled transition systems, which in their turn form perhaps the simplest mathematical models of the notion of computation. It is interesting to note that the Peircean and the Ockhamist approaches in philosophical logic find (much more technically inspired) counterparts in the development of the computational tree logics: CTL and CTL*, respectively.

5 Time Periods

So far we have applied a point-based paradigm to represent time. Nevertheless, it seems that in every field where temporal logics are used or studied,

at a certain moment systems are designed in which *periods* are the central entities, or at least, play a more prominent role.

Motivations   The point-based perspective has never been without philosophical objections. For instance, Zeno's paradox of the flying arrow, which, it is argued, cannot change position at a isolated moment of time and thus cannot move at all, makes it clear that there is something problematic concerning the representation of time as a series of durationless moments if we want to describe the concept of movement. Some temporal predicates seem simply not to apply to time points. Suppose that p is a proposition formalizing the statement that Zeno's arrow moves. Obviously, the flying of an arrow is an activity that is extended in time; hence, one might argue that it is pointless to evaluate the truth of p at moments of time. It thus seems that we at least need the existence of time periods for the evaluation of certain expressions.

Apart from such semantic considerations, it is clear that time points are not the kind of objects that we can directly perceive. Due to years of exposure to the scientific view on time we may not always realize this, but if we want to base reality on our direct experience, then time points will come out as highly abstract and complex artifacts. Thus, it has been argued, it is a dubious enterprise to take points as having primitive ontological status; periods form a far sounder base. This second argument has been taken up, with a more practical twist, within Artificial Intelligence. Here the idea has been advocated that period-based representations of time are simpler and more natural in formalizing common sense reasoning than the standard scientific models. (Obviously, this argument may be pushed further, questioning the Newtonian perspective in which absolute Time exists regardless of anything happening in it. Such objections may lead to *event-based* ontologies which due to lack of space we cannot discuss here.)

Finally, in our discussion until now we have assumed that there is a clear and intuitive distinction between points and periods. This is questionable as well, however; one can quite convincingly argue that there is a notion of *granularity* involved here. A good example can be taken from Computer Science, where the addition of two numbers may be taken as an atomic, durationless action of a high-level programming language, whereas it may be implemented in terms of many operations on the lower level of the machine language.

Time in periods It is important to observe that the need for a more prominent role of periods does not necessarily commit one to model time in structures in which periods are *primitive* entities; they might as well be *derived* objects.

Indeed, one could well start from a flow of time $\mathcal{T} = (T, <)$ as described earlier, and then consider the question how to represent chunks of time within such a structure. For instance, periods could be defined as *convex* sets: subsets C of T that are uninterrupted in the sense that whenever s and a later point t belong to C, then so does any point between s and t. A set-theoretically slightly simpler option is to only consider (closed) intervals; in this approach, the period $\{u \in T \mid s \leq u \leq t\}$ can simply be represented as the *pair* $[s, t]$. Observe that this approach has the advantage that properties of periods can be expressed by binary predicates in the first order frame language, whereas for convex sets we have to use some kind of higher-order logic.

If one opts for periods as primitive entities, the simplest mathematical modeling will involve structures consisting of a set P of periods equipped with a collection of natural relations on P. But in contrast to the point-based approach where the temporal precedence relation is *the* candidate for such a relation, we now have many options. For instance, since one is obviously still interested in temporal precedence, the relation $\prec$, with $p \prec q$ holding if the entire period p precedes the entire period q, is a natural candidate, but so is the inclusion relation $\sqsubset$, with $p \sqsubset q$ holding if p is a proper part of q. And in fact, one widespread period-based modeling of time is that in structures of the form $\mathcal{P} = (P, \prec, \sqsubset)$. But $\prec$ and $\sqsubset$ are not the only candidates. If we are interested in relations that are close to our common sense experience, then the relation of one period overlapping with another is quite relevant as well. And we are not confined to binary relations at all: we may need a unary predicate informing us whether a period is of zero duration (and hence, point-like), whereas there are also interesting ternary relation such as the relation C holding of a triple p, q, r if p can be 'chopped' into the two pieces q and r. Of course, just like in the point-based case, one needs to impose conditions on period structures to make them useful as models of time. For instance, in a structure of the kind $\mathcal{P} = (P, \prec, \sqsubset)$ one will want $\prec$ and $\sqsubset$ to be strict partial orders that are related by conditions like $\forall xyz \, (x \sqsubset y \prec z \rightarrow x \prec z)$ and others.

The reader may have realized how hard it is to gather one's intuitions and make a complete list of such conditions without taking resort to talking about points after all. The concept of a point in time has obviously been very useful in our thinking about time. Hence, even if periods are to be

taken as the primitive entities of one's ontology, it is at least interesting, if not a test for the viability of the proposal, to see whether one can *construct* point-based flows of time from period structures. Various ways have been worked out for this purpose. Perhaps the simplest method is to take as points those periods that have zero duration — of course, this only works if such entities are around and we have access to this information (for instance, through a zero-duration predicate as we mentioned above). But even if our period structure does not have atomic periods, there are ways to extract a point structure from it, for instance, by defining a point to be any maximal set of mutually overlapping pairs of periods. Finally, once there are ways to construct point structures from period structures and vice versa, the obvious question is to see how such constructions interact. This line of research has been taken up with great mathematical sophistication, in a number of cases even leading to interesting categorical dualities.

Interval-based temporal logic   Just as in the case for point-based temporal logics, we may choose a class of period structures, design a formal language to talk about it, and study the resulting temporal logic.

For instance, suppose that we are working with intervals in point-based flows of time, as described above. Taking the modal approach, we find ourselves in a multi-dimensional setting; that is, we want to evaluate formulas at *pairs* of points representing the beginning and the end point of the interval, respectively. Typical modal operators are $\langle D \rangle$ and $\circ$ with truth tables given by

$$\mathcal{M}, [s,t] \Vdash \langle D \rangle \varphi \quad \text{if} \quad \mathcal{M}, [u,v] \Vdash \varphi \text{ for some } t, u \text{ with } s \leq u \leq v \leq t,$$
$$\mathcal{M}, [s,t] \Vdash \varphi \circ \psi \quad \text{if} \quad \mathcal{M}, [s,u] \Vdash \varphi \text{ and } \mathcal{M}, [u,t] \Vdash \psi \text{ for some } u \text{ with } s \leq u \leq t.$$

In words, $\langle D \rangle \varphi$ holds at an interval if φ holds at some interval *during* it, while $\varphi \circ \psi$ holds at an interval if can be chopped into a φ- and a ψ-part. In period terms, one would say that $\sqsubseteq$ and C are the accessibility relations of $\langle D \rangle$ and $\circ$, respectively.

For such modal systems, one may investigate meta-logical properties like completeness and decidability. The general picture here is that one has a price to pay for the increase in expressivity: complete axiomatizations are scarce and hard to find, and undecidability is the rule rather than the exception. On a technical level, the modal logic of time periods thus seems to be more complex (and hence, more intriguing) than point logics over the same flows of time, but the *kinds* of questions that are asked do not differ much.

Hence, let us finish this section mentioning some issues that are of specific interest to period logics. To start with, period logics differ from point logics in the sense that in many cases it is natural to correlate the interpretation of atomic propositions. A condition that one often encounters is that of *homogeneity* requiring that an atomic proposition holds at a period if and only if it holds at each of its parts. It is obvious that such a condition only has intuitive appeal for the propositions corresponding to the event categories of states and activities. And even in the latter case, one may raise objections to the 'only if' part of this condition: I can truthfully say that I have been *walking* through town for hours when in fact, I have paused a couple of times to take a coffee.

Now suppose that we are implementing this condition on some interval structure $\mathcal{I}(\mathcal{T})$ induced by the flow of time $\mathcal{T}$ by demanding that for each propositional variable p and each point-based valuation π we have

(6) $(\mathcal{I}(\mathcal{T}), \pi), [s, t] \Vdash p$ iff $\pi(u)(p) = 1$ for all u with $s \leq u \leq t$.

Observe that thus we have effectively reduced period predicates to point predicates. Such a reduction would have considerable computational advantages, something that can easily be explained by taking a first order perspective. It is obvious that the particular proposal (6) is rather naive: Zeno's moving arrow will lead us into trouble. But perhaps there are more inventive modellings in which formulas can be evaluated at periods, while the valuations remain point-based?

In any case, regardless of the technical advantages of reducing period predicates to point predicates, it is clear that there is a rather general philosophical issue at stake here, namely the problem of which *kinds* of predicates apply to periods and points, respectively, and how these are correlated. This issue is in fact a matter of ongoing, and at times heated, debate.

6   Temporal Logic Now

As we mentioned before, temporal logic has become a vast and active research area with applications in many disciplines. In this section we will briefly sketch some of these recent developments. Since not all of the work mentioned here is covered by the monographs mentioned at the end of the next section, we provide references to the literature.

Richer ontological structures   One common trend in temporal logic is to study logics of richer ontological structures since it is obvious that for se-

rious real-world applications the kind of temporal logics that we have been describing until now are far too simple. For example, one shortcoming of standard temporal logics is that they only deal with qualitative timing properties, whence they are inadequate for applications such as reasoning about real-time behaviour of software. In order to overcome this deficiency, people have designed logics for describing two-sorted structures consisting of a linear flow of time connected with some metric domain. Such approaches can be found both in the point-based and in the period-based paradigm, cf. [16] and [10], respectively. Another example of a multiple sorted ontology we have already met in the semantics of Ockhamist branching time logic, where branches appeared as a second kind of entities, next to points. One might vary on this 'standard' Ockhamist logic by admitting only some instead of all branches, perhaps a collection satisfying some addition constraints [24]. Applying this idea of using multiple sorted temporal ontologies to the discussion of the previous section, one can envisage structures in which points, periods and events co-exist, linked by suitable relations [18]. One possibility for such a link involves the notion of granularity: atomic objects might suddenly turn out to be divisible when approached at a different level. This obviously ties up with our way of classifying periods of time (months, weeks, days); modal logics for such layered structures are described in [15].

Temporal logic at work Turning temporal logics into actual working systems has created a number of interesting problems and challenges. For instance, one of the most fundamental contributions that Artificial Intelligence has made to the field of temporal logic, is that of identifying the *frame problem*. This is the problem of formalizing the properties of an application area that are unaffected by the performance of some action without explicitly summing up *all* such properties. This problem appears to be independent of the particular formalism employed, and has to be faced by anyone wishing to give a formal account of reasoning about change [19]. The computer science literature on modal logics of time has yielded an interesting perspective on the modal truth relation $(\mathcal{M}, t \Vdash \varphi)$ between a model $\mathcal{M}$, which is supposed to be finite, and a formula φ; in this perspective φ represents some property of a program and $\mathcal{M}$ some implementation of the program. For obvious reasons then, a considerable amount of effort has been devoted to finding fast *model checking* algorithms deciding whether a given formula holds in a given finite model [21]. As a last example we mention the *dynamic turn* which research in the semantics of natural language has taken. In this way of thinking, the meaning of a formula does not lie so much in its truth

condition; linguistic expressions are rather like programs that update the information state of some agent. For instance, in Discourse Representation Theory [12] temporal expressions in natural language are used to extend and refine temporal representations of the discourse; these representations on their turn are syntactic items themselves that can be interpreted in standard models.

Temporal logic in context   There is an increasing tendency to study modal formalisms not as isolated systems but in connection with other branches of logic, as in *Correspondence Theory* which relates modal logic to first and second order logic. For instance, the use of game-theoretic methods has deepened our understanding of the relative expressive power of modal logics of time: in particular, variants of Ehrenfeucht-Fraïssé games have provided an interesting perspective on expressive completeness results such as Theorem 4.1 [11, 22]. Recent approaches to decidability questions concerning modal and temporal logics use insights from algebraic logic and automata theory. This has lead to the identification of a variety of decidable fragments of first order logic, each of which is obtained from atomic formulas using all boolean connectives but allowing only a specific, guarded pattern of quantification [1]. As a last example we mention the emergence of so-called hybrid languages which aim to boost the expressive power of modal languages by adding some features from first order logic like 'names' (special variables that are to be true at a single state), over which quantification is allowed [9, 4].

7   Epilogue

What then, is temporal logic?

In the narrowest sense temporal logic comprises the design and study of specific systems for representing and reasoning about time, such as Prior's tense logic. These enterprises may have both an applied and a theoretical side, the former consisting of designing a system (that is, making choices in the fields of ontology, syntax and semantics), formalizing temporal phenomena in it, and then putting it to work (perhaps through implementing it). On the theoretical side, one aims at proving formal properties of the system, such as completeness or decidability.

On a slightly wider scale, temporal logicians may thus provide a supply of general tools and techniques for answering questions pertaining to specific systems. As an example we mention the method of filtration which is a quite

general method of proving decidability of a temporal logic, and the canonical model method which is very useful in proving completeness results.

A more ambitious aim for temporal logicians is to come up with frameworks for comparing and connecting different modellings of time. This aim can be realized both at a technical and at a philosophical level. As an example of the first, think of the game-theoretic analysis of the expressive power of modal languages, or of the duality between point and period-based representations of time, respectively. On a philosophical level, a thorough classification of event types and of the correlation between predicates pertaining to points and to periods, respectively, would be an extremely useful tool in any discussion on formal representations of temporal phenomena.

Since all of this is relevant for each of the disciplines where formal reasoning about time is needed, Temporal Logic forms a prime example of the growing role of Logic as a source and channel of ideas and techniques applicable in related disciplines. Ultimately, one would hope that temporal logic can provide a unifying perspective on our sometimes confusing thoughts about this highly puzzling thing we call time.

Suggested Further Reading   In this chapter we have only scratched the surface of Temporal Logic. The following monographs, each surveying part of the field of temporal logic, would form a good start for a bibliography. Concerning the philosophy of time I do not believe there is one standard reference, but WHITROW [23] is a very comprehensive study of the concept of time, while LE POIDEVIN & MACBEATH [13] brings together some seminal articles on the subject. ØHSTRØM & HASLE [17] gives a good treatment of philosophical aspects of temporal logic from a historical perspective. In GOLDBLATT [8] the reader finds a concise and very accessible treatment of the most important modal logics of time; GABBAY ET ALII [5] is a more extensive mathematical treatment. MANNA & PNUELI [14] is a classic on applications of temporal logic in computer science; GABBAY ET ALII [6] gives a good overview of the applications of temporal logic in artificial intelligence. There seems to be no monograph on the treatment in formal linguistics of temporal aspects of natural language, but STEEDMAN [20] surveys the field well. VAN BENTHEM [3] is a stimulating blend of much of the above. Finally, for an overview of recent developments in temporal logic the reader is referred to the proceedings of the first two conferences devoted solely to temporal logic, ICTL'94 [7] and ICTL'97 [2].

References

[1] H. Andréka, J. van Benthem, and I. Németi. Modal logics and bounded fragments of predicate logic. *Journal of Philosophical Logic*, 27:217–274, 1998.

[2] H. Barringer, M. Fisher, D. Gabbay, and G. Gough, editors. *Proceedings of the Second International Conference on Temporal Logic – ICTL'97*, to appear.

[3] J. van Benthem. *The Logic of Time.* Kluwer, Dordrecht, second edition, 1991.

[4] P. Blackburn and M. Tzakova. Hybrid languages and temporal logic. *Logic Journal of the IGPL*, to appear.

[5] D.M. Gabbay, I. Hodkinson, and M. Reynolds. *Temporal Logic: Mathematical Foundations and Computational Aspects.* Oxford Logic Guides. Oxford University Press, 1994.

[6] D.M. Gabbay, C.J. Hogger, and J.A. Robinson, editors. *Handbook of Logic in AI and Logic Programming*, volume 4 (Epistemic and Temporal Reasoning). Oxford University Press, Oxford, 1995.

[7] D.M. Gabbay and H.J. Ohlbach, editors. *Temporal Logic. Proceedings of the First International Conference, ICTL'94*, volume 827 of *Lecture Notes in Computer Science*, Berlin, 1994. Springer.

[8] R. Goldblatt. *Logics of Time and Computation.* CSLI Lecture Notes. Center for the Study of Language and Information, Stanford University, second edition, 1987.

[9] V. Goranko. Temporal logic with reference pointers. In Gabbay and Ohlbach [7], pages 133–148.

[10] M. Hansen and Zhou Chaochen. Duration Calculus: logical foundations. *Formal Aspects of Computing*, 9:283–330, 1997.

[11] N. Immerman and D. Kozen. Definability with bounded number of bound variables. In *Proceedings of the Conference on Logic in Computer Science (LICS'87)*, Washington, 1987. Computer Society Press.

[12] H. Kamp and U. Reyle. *From Discourse to Logic*. Kluwer Academic Press, Dordrecht, 1993.

[13] R. Le Poidevin and M. MacBeath, editors. *The Philosophy of Time*. Oxford University Press, Oxford, 1993.

[14] Z. Manna and A. Pnueli. *The Temporal Logic of Reactive and Concurrent Systems*. Springer Verlag, 1991.

[15] A. Montanari. *Metric and Layered Temporal Logic for Time Granularity*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, 1996.

[16] A. Montanari and M. de Rijke. Two-sorted metric temporal logic. *Theoretical Computer Science*, 183:187–214, 1997.

[17] P. Øhrstrøm and P. Hasle. *Temporal Logic: From Ancient Ideas to Artificial Intelligence*. Studies in Linguistics and Philosophy. Kluwer Academic Publishers, Dordrecht, 1995.

[18] C. Gardent P. Blackburn and M. de Rijke. Back and forth through time and events. In Gabbay and Ohlbach [7], pages 225–237.

[19] E. Sandewall and Y. Shoham. Non-monotonic temporal reasoning. In Gabbay et al. [6], pages 439–498.

[20] M. Steedman. Temporality. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 895–938. Elsevier Scientific Publishers, Amsterdam, 1997.

[21] C. Stirling. Bisimulation, modal logic, and model checking games. *Logic Journal of the IGPL*, to appear.

[22] Y. Venema. Expressiveness and completeness of an interval tense logic. *Notre Dame Journal of Formal Logic*, 31:529–547, 1990.

[23] G.J. Whitrow. *The Natural Philosophy of Time*. Clarendon Press, Oxford, 1980.

[24] A. Zanardo. Branching-time logic with quantification over branches. *Journal of Symbolic Logic*, 61:1–39, 1996.