Temporal Logic

Yde Venema
Institute of Logic, Language and Information
University of Amsterdam
Plantage Muidergracht 24
1081 TV Amsterdam
The Netherlands
E-mail: yde@wins.uva.nl

To appear in:

Lou Goble (ed.), *Blackwell Guide to Philosophical Logic*, Blackwell Publishers.

1 Introduction

Time must be the most paradoxical concept our minds have to deal with. To quote from the *Confessions* of St. Augustine[1]:

> 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 experience 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 for the time being, it is obvious that time plays such a fundamental role in our thinking that there is a clear need for *precise reasoning* about it. Naturally then, the problem of representing temporal knowledge and reasoning has been addressed in many academic disciplines besides philosophy as well. In *physics*, time and change have always been central notions, although in mathematical modelings time has usually not received a higher ontological status than that of being one of many dimensions. This is quite different in formal *linguistics*, where a lot of work has gone into modeling temporal phenomena — and necessarily so, if only by the observation that every single sentence we utter is put in some temporal perspective, through *tense*. *Computer scientists* are thinking of the behaviour of programs in terms of the step by step evolution of an artificial system; thus, specifying the behaviour of programs or reasoning about their correctness naturally involves a temporal component. More in particular, the field of *artificial intelligence* is permeated with temporal concerns; problems related to planning, such as efficiently maintaining temporal knowledge, are at the core agenda of this discipline.

While these enterprises are not necessarily concerned with the *same* concept of time, they all could go under the heading of Temporal Logic. However, often a more restricted, technical definition is used in which temporal logic is a branch of *modal* logic. We will largely confine ourselves to this modal perspective here, mainly for reasons of limitations on the length of this Chapter.

The modal approach towards logics of time was initiated about forty years ago by Arthur Prior. His aims were to develop temporal logics that are close to natural language and could be used for investigating philosophical problems related to time, such as determinism. The alternative name 'tense logic' that one often encounters for temporal logic stems from Prior: the names 'F' and 'P' of his basic modal operators were chosen to reflect their intended correspondence to the future and past tenses of natural languages. Prior's work has been a source of inspiration for modal logicians; the core of our knowledge of metalogical properties of modal formalisms for reasoning about time stems from the sixties and seventies, through the work of van Benthem, Burgess, Gabbay, Kamp, Segerberg and others. In this period the modal approach was also very influential in the formal semantics of temporal phenomena in natural language, but this dominance has decreased from the eighties onwards. Since then,

[1]In the original Latin: *Quid est ergo tempus? Si nemo ex me quaerat, scio; si quaerenti explicare velim, nescio.* (Book XI, Chapter XIV).

theoretical computer science has become the main area for applying modal logics of time; pioneering work in this area was done by Amir Pnueli. This change of setting also influenced the agenda of temporal logicians: next to expressiveness, completeness and decidability issues also computational aspects of logics are investigated.

Nowadays, Temporal Logic is a vast and active research area with applications in many areas and even a biannual conference entirely devoted to it. Even in the narrower framework of *modal* logics of time, there is such a bewildering variety of systems that a short introductory chapter like this cannot do justice to the whole field. Nevertheless, in this tangle there are some common threads to be identified, of which we mention two. First, most of the systems are based on a rather classical conceptual approach. That is, the starting point of the generic temporal logic is the selection of some class of mathematical structures that are to represent time; then, a formal language is designed to talk about these structures, one starts to formalize temporal phenomena in this language and perhaps some calculus is constructed for it. Second, there is a standard repertoire of questions that are asked concerning these systems, regarding the expressiveness of the language, the completeness and decidability of the calculus, and, depending on the degree of applicability one has in mind for the calculus, the existence of an efficient proof theory or model checking algorithm.

Hence, rather than vainly trying to give a complete survey, we aim at providing the reader with a flavour of what this field is like. We start with a section in which we discuss some of the most well-known mathematical modellings of time. 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 introduce the readers to this particular system, but perhaps even more to acquaint them with the kind of questions that temporal logicians tend to ask. In the sections 4 and 5 we present 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*; this section also contains a short list of monographs surveying the field of temporal logic.

2 Flows of time

Before we can start a discussion of various logics of time, it seems a good idea to acquaint ourselves with some standard mathematical models of time. Here already there is a wide choice of approaches, even if we confine ourselves to those in which time is represented as a *structure* consisting of a set of temporal objects connected by relations and/or operations.

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. It goes without saying that the notion of a 'time point' is a highly abstract one, and indeed it has been argued that it is far more natural and closer to common sense to model time in an ontology having time *periods* or *events* as the basic entities. We will take up this discussion in section 5, whereas quantitative aspects of representing time will be touched upon only very briefly (in section 6).

Hence, let us for the moment accept a point-based paradigm in which time is represented 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 of these frames will not qualify as intuitively acceptable representations of time. Minimal conditions that one will impose on a frame is that $<$ be irreflexive (no time point is earlier than itself) and transitive (if s is earlier than t and t is earlier than u, then s is also earlier than u). 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 would 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), or 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. Now in our opinion it is not the logician's task to choose between different ontologies, so why not allow circular time? After all, the linear representation of time is a product of the modern Weltanschauung — most civilizations have regarded time as being essentially cyclic in nature. And also, practical applications of circular time such as the construction of rotas, are well conceivable. Hence, our only reason is simply that circular time has hardly received any 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.

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 these are generally not the most interesting issues. Let us just mention some definitions: $\mathcal{T}$ *has a first point* if $\mathcal{T} \models \exists x \forall y\,(Rxy \vee x = y)$, it *has a final point* if it satisfies $\exists x \forall y\,(Ryx \vee x = y)$. On the other hand, $\mathcal{T}$ is *right-serial* if each point has a non-empty future, and *left-serial* if each point has a non-empty past.

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: $\forall xy\,(Rxy \to \exists z\,(Rxz \wedge Rzy))$. 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: $\forall xy\,(Rxy \to \exists z\,(Rxz \wedge \neg \exists u\,(Rxu \wedge Ruz)))$. Standard examples of discrete flows of time are given by the natural or the integer numbers.

Density is often confused with *continuity*; a temporal flow of time $\mathcal{T}$ is *continuous* if every way of cutting up T into an earlier and a later part will yield at least one point on the edge of the cutting. The familiar ordering of the rational numbers is a typical counterexample: consider the sets $\{q \in \mathbb{Q} \mid q < \sqrt{2}\}$ and $\{q \in \mathbb{Q} \mid q > \sqrt{2}\}$. Formally, define a *cut* of the flow of time $\mathcal{T} = (T, <)$ to be a non-empty, proper subset $X \subset T$ such that each point in X comes before each point not in X; a *gap* is a cut X such that neither X has a final point nor does its complement have a first one. Continuity is then defined as the absence of gaps. The standard example of a continuous flow of time is given by the real numbers, but note that for instance also the integer numbers are continuously ordered. Observe that unlike the properties discussed before, continuity has been expressed using a second order quantification; this is essentially so: one can prove that continuity cannot be expressed in first order logic (for instance through showing that the structures $\mathcal{Q}$ and $\mathcal{R}$ are elementary equivalent).

There are many other interesting second order conditions that one may impose on temporal structures. To motivate one of them, 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 $>$.

This list is by no means exhaustive; but since there are no further properties that play a role in this chapter, we leave it at this.

3 Basic tense logic

In this section, which forms the heart of the chapter on Temporal Logic, 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 $\pi : (T \to (\Phi \to \{0,1\}))$; here Φ denotes the set of propositional variables. A *model* is 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 $\pi(t)(q) = 0$. However, the spice of temporal logic 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', cf. Figure 1.

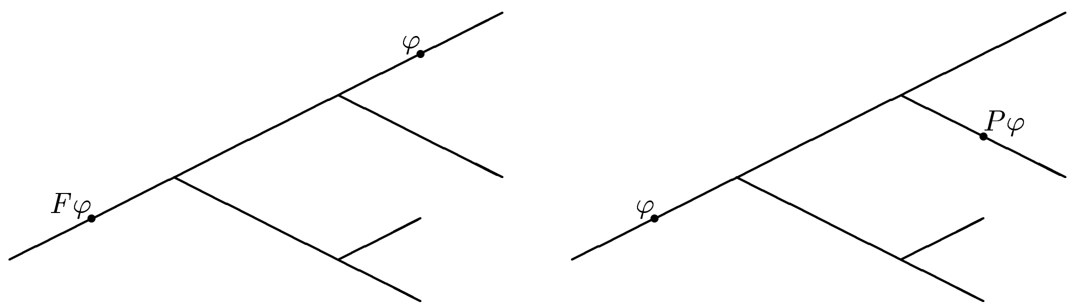

Figure 1: The operators F and P.

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 $\bot$, $\top$, $\wedge$, $\to$ 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 $\mathcal{M} = (T, <, \pi)$:

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

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 the Chapter 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, analogous to the case of Epistemic Logic (cf. Chapter **??**), 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.

Correspondence with first order logic

It is not only the aim of temporal logicians to study systems of temporal logic in isolation. For instance, in the *correspondence theory* of temporal logic one compares the expressive strength of various temporal languages with that of first order logic. The key observation enabling such comparisons is that any temporal model $\mathcal{M} = (T, <, \pi)$ can also be seen as a structure for the so-called *corresponding first order model language* $\mathcal{L}_m$. This language is an extension of the first order frame language with a monadic predicate A_q for each propositional variable q of the Priorean language (we write $\mathcal{L}_m(\Phi)$ if we want to make this set of variables explicit). For, given a model $\mathcal{M} = (T, <, \pi)$, let $\mathcal{M}^*$ be the first order $\mathcal{L}_m$-structure $(T, <, V(q))_{q \in \Phi}$; that is, each monadic predicate symbol q is interpreted as the set $V(q)$. Thus it is possible to compare the modal perspective of a formula being true *at* a point, to the first order view point of a formula being true *of* a time point.

Now consider the following *standard translation* mapping temporal formulas to formulas in $\mathcal{L}^1_m$ (consisting of those fomulas in $\mathcal{L}_m$ that have one free variable x and thus express properties of time points):

$$
\begin{array}{lrcl}
(2) \quad & ST_x(q) & = & A_q x \\
& ST_x(\neg\varphi) & = & \neg ST_x(\varphi) \\
& ST_x(\varphi \wedge \psi) & = & ST_x(\varphi) \wedge ST_x(\psi) \\
& ST_x(F\varphi) & = & \exists y\,(Rxy \wedge ST_x(\varphi)[y/x]) \\
& ST_x(P\varphi) & = & \exists y\,(Ryx \wedge ST_x(\varphi)[y/x]).
\end{array}
$$

Here $ST_x(\varphi)[y/x]$ denotes the result of substituting the variable y for x in the first order formula $ST_x(\varphi)$.

It is not difficult to see that the Priorean temporal language can be identified with a *fragment* of this first order language — this is the meaning of the following Proposition, the *proof* of which is by a straightforward induction on φ.

Proposition 3.1 *For any temporal model $\mathcal{M} = (T, <, \pi)$, any time point $t \in T$ and any Priorean formula φ,*

$$\mathcal{M}, t \Vdash \varphi \text{ iff } \mathcal{M}^* \models ST_x(\varphi)[x \mapsto t].$$

But if Proposition 3.1 informs us that every Priorean formula can equivalently be seen as a first order formula in one free variable, the question naturally arises whether the converse holds as well. Is perhaps *every* $\mathcal{L}_m$-formula in one free variable equivalent to the standard translation of a Priorean one? The answer to this question is negative and there is an elegant way to arrive at this answer. Our vehicle will be the notion of a temporal bisimulation, defined as follows.

Let $\mathcal{M} = (T, <, \pi)$ and $\mathcal{M}' = (T', <', \pi')$ be two models. A relation $Z \subseteq T \times T'$ is called a *temporal bisimulation* between $\mathcal{M}$ and $\mathcal{M}'$ if it satisfies the following conditions for all s, t in T and all s', t' in T':

- (*facts*) for any propositional variable q, $\pi(s)(q) = \pi(s')(q)$,
- (*F-forth*) if $s < t$ and sZs' then there is some $t' \in T'$ such that $s' <' t'$ and tZt',
- (*P-forth*) if $s > t$ and sZs' then there is some $t' \in T'$ such that $s' >' t'$ and tZt',
- (*F-back*) if $s' <' t'$ and sZs' then there is some $t \in T$ such that $s < t$ and tZt',
- (*P-back*) if $s' >' t'$ and sZs' then there is some $t \in T$ such that $s > t$ and tZt'.

We say that the points s of $\mathcal{M}$ and s' of $\mathcal{M}'$ are *bisimilar* if there is some bisimulation between $\mathcal{M}$ and $\mathcal{M}'$ linking s and s'. The intuitive idea underlying this notion is that every step from some point in either of the structures, be it towards the future or towards the past, has a suitable simulation from any bisimilar point in the other model.

As an example, consider the valuation τ' on the binary tree $\mathcal{B}$ that makes the variable q true at all words of odd length, and r at all words of at least 1,000 symbols. We leave it to the reader to verify that the relation L linking each word s in $\mathcal{B}$ to its *length* $|s|$, is a temporal bisimulation between the models $(\mathcal{B}, \tau')$ and $(\mathcal{N}, \tau)$ given right after (1).

The importance of temporal bisimulations is that the Priorean language is blind for the differences between bisimilar points; that is, two bisimilar points will satisfy exactly the same Priorean formulas.

Proposition 3.2 *Let Z be a bisimulation between the temporal models $\mathcal{M}$ and $\mathcal{M}'$. Then for any Priorean formula φ we have*

$$(3) \qquad \mathcal{M}, s \Vdash \varphi \text{ iff } \mathcal{M}', s' \Vdash \varphi, \quad \text{for all } s \text{ and } s' \text{ such that } sZs'.$$

Proof. We prove (3) by formula induction on φ. For atomic φ (3) is just the basic clause of the definition of a bisimulation. The case where φ is of the form $\neg\psi$ or $\psi \wedge \chi$ is rather easy and left to the reader. The interesting cases are when the main connective of φ is a temporal operator; we only treat the case where φ is of the form $G\psi$.

Consider two bisimilar points s and s', and suppose that $\mathcal{M}, s \Vdash G\psi$. In order to show that this implies the truth of $G\psi$ at s', take an arbitrary point t' in the future of s'. By clause (*F-back*) there is a point t in T such that $s < t$ and tZt'. By the truth definition, $s < t$ implies that $\mathcal{M}, t \Vdash \psi$; hence, it follows by the inductive hypothesis that $\mathcal{M}', t' \Vdash \psi$. Since t' was arbitrary, this immediately shows that $\mathcal{M}', s' \Vdash G\psi$. The proof that conversely, $\mathcal{M}', s' \Vdash G\psi$ only if $\mathcal{M}, s \Vdash G\psi$ follows by symmetry (using (*F-forth*) instead of (*F-back*)). QED

As we will see in section 4, this result can be used to prove that new temporal operators cannot be expressed in the Priorean language.

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 $(\mathcal{T}, V), t \Vdash Fq$. 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 adapt an earlier example of a temporal bisimulation to prove that there is no Priorean formula that defines the class of branching flows of time. For, suppose that φ were such a formula; since the structure $\mathcal{N}$ is not branching, φ would not be valid on $\mathcal{N}$; in other words, we could find a valuation π and a number n such that $(\mathcal{N}, \pi), n \not\Vdash \varphi$. But now consider the valuation π' on the binary tree $\mathcal{B}$ that makes a propositional variable q true at a string s if (and only if) π makes q true at the length $|s|$ of s. As before, we see that the function mapping any sequence of 0s and 1s to its length, provides a temporal bisimulation. Hence, it follows by Proposition 3.1 that the formula φ is not valid on $\mathcal{B}$; but this contradicts the fact that the formula φ *must* be valid on $\mathcal{B}$, since $\mathcal{B}$ is branching!

On the other hand, there is a formula defining the flows of time that are *not* branching; for, the formula $PFq \to (Pq \vee q \vee Fq)$ 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. The reader who is puzzled by the fact that a property may be definable while its negation is not, should realize that the law of the excluded middle does not apply to validity of modal formulas in flows of time. There is nothing mysterious about this, it is simply due to the fact that the concept of validity involves a quantification over valuations.

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 $P\varphi \vee \varphi \vee F\varphi$, and $\Box\varphi \equiv H\varphi \wedge \varphi \wedge G\varphi$. For future reference, we have given names to the modal formulas.

having a first point	$H\bot \vee PH\bot$	(A1)
left-seriality	$P\top$	(A2)
having a final point	$G\bot \vee FG\bot$	(A3)
right-seriality	$F\top$	(A4)
discreteness	$(F\top \wedge q \wedge Hq) \to FHq$	(A5)
density	$Fq \to FFq$	(A6)
continuity	$(Fq \wedge \Diamond\neg q \wedge \Box(q \to Hq)) \to$ $\to \Diamond((q \wedge G\neg q) \vee (\neg q \wedge Hq))$	(A7)
having finite intervals	$G(Gq \to q) \to (FGq \to Gq)$	(A8)

The last two examples show that the Priorean language is expressive enough to define essentially *second order* frame properties. This is of course due to the just mentioned fact that the definition of a formula being valid in a flow of time involves a quantification over all valuations on the flow of time. In fact, it is rather miraculous that so many Priorean formulas do have such a simple first order correspondent!

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 $Gp \to GGp$, this boils down to the problem of finding a correspondent for irreflexivity (within the class of transitive frames). Unfortunately, there is *no* such formula, as a simple argument (in the style of the one given for the undefinability of branchingness) could reveal.

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. The completeness theory of temporal logic is the topic of this subsection; obviously, there are close connections with the axiomatics of alethic modal logic as discussed in Chapter **??**. 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	
$(DB_)$	$G(q \to r) \to (Gq \to Gr)$	
	$H(q \to r) \to (Hq \to Hr)$	(Distribution)
(CV)	$q \to GPq$	
	$q \to HFq$	(Converse)
(4)	$Gq \to GGq$	(Transitivity)
(US)	if φ is a theorem, then so is $\varphi[\psi/q]$	(Uniform Substitution)
(MP)	if φ and $\varphi \to \psi$ 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 others 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 3.3 *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 3.3. 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 $PFq \to (Pq \vee q \vee Fq)$ (defining non-branching to the future) and its mirror image $FPq \to (Fq \vee q \vee Pq)$.

Theorem 3.4 *The logic K_t is sound and complete with respect to the class of linear flows of time.*

PROOF. Just as the completeness proofs of Chapter **??**, 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':

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

Analogous to the proof of Theorem 4 in Chapter **??**, (4) is proved by a 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 (4) 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 $\forall xy(Rxy \vee x = y \vee Ryx)$). 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 (4). 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.

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

In order to prove claims like (5), 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 between R and $\prec$ is such that Rst iff $C_s \prec C_t$.

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, 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 $\ldots \sigma_1\sigma_2\sigma_0\sigma_1\sigma_2\sigma_0 \ldots$

There is thus an obvious relation linking points in the new, transformed model to points in the old one; the fact that this relation is a temporal bisimulation ensures that φ is indeed satisfiable in the new model. This finishes the proof sketch of (5). 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 3.5 *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 the tense logic **KTh**. This logic is defined as the extension of $\mathbf{K}_t$ with the axioms (NB), (A4), $H(Hq \to q) \to Hq$ and $FGq \to GFq$. It can be proved that there is no single frame on which all of these axioms are valid simultaneously; hence, if **KTh** were the set of validities of some frame class C, then $\perp$ would also be valid on this class. But, using a more general kind of semantics, one can show that **KTh** is consistent; meaning that $\perp$ is *not* derivable in it. Hence, this logic cannot be characterized by any class of frames.

Decidability

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

Proof. We only give a proof sketch for the case of linear time. It is our aim to show that a formula is satisfiable on a linear flow of time iff it is satisfiable on a *finite* pseudo-line. From this, decidability follows easily: we start with effectively making a list of all models based on a finite pseudo-line; with the help of this list, we recursively enumerate the set of all formulas of which the negation is satisfiable on a finite pseudo-line. By the claim above, these are precisely the formulas that are not valid on the class of linear flows of time. But also, we can use Theorem 3.4 to give a recursive enumeration of the set of formulas that *are* valid on the class of linear flows of time. But if we can recursively enumerate both the theorems and the non-theorems, decidability follows by Post's Theorem.

Hence, it is left to prove that

(6) if φ is satisfiable on a linear flow of time, then also on a finite pseudo-line.

The proof of (6) is by a so-called *filtration* argument. Suppose that φ holds at the point t_0 in the model $\mathcal{M} = (\mathcal{T}, \pi)$ which is based on the linear ordering $\mathcal{T}$. We say that two points s and

t of T are *φ-equivalent*, notation: $s \sim_\varphi t$, if they agree on the truth of all subformulas of φ. This is clearly an equivalence relation, and since φ has only finitely many subformulas, there are only finitely many equivalence classes. The key idea underlying filtration is that we will take these equivalence classes as the basis of a *new*, pseudo-linear model $\mathcal{M}^f$, to be called the *filtration* of $\mathcal{M}$. It usually requires some creativity to define the accessibility relation on such a filtrated model.

For a formal definition, let S be the set of equivalence classes of T under $\sim_\varphi$; the equivalence class of a point s is denoted as $[s]$. The relation R^f on S is defined by $R^f[s][t]$ iff for all subformulas $G\psi$, $H\chi$ of φ: $\mathcal{M}, s \Vdash G\psi$ only if $\mathcal{M}, t \Vdash G\psi \wedge \psi$ and $\mathcal{M}, t \Vdash H\psi$ only if $\mathcal{M}, s \Vdash H\psi \wedge \psi$. Observe that this relation is well-defined (that is, does not depend on the particular representants of the equivalence classes) by our definition of $\sim_\varphi$. Finally, the filtrated valuation π^f on this frame is given by $\pi^f([s])(q) = \pi(s)(q)$, for each variable q occurring in φ.

We leave it to the reader to prove that R^f is indeed transitive and strongly connected. What makes this idea work is that for all subformulas ψ of φ and all points s of $\mathcal{M}$, truth of ψ in $\mathcal{M}$ at s coincides with truth of ψ in $\mathcal{M}^f$ at $[s]$. For, it follows from this claim and our assumption, that φ is satisfiable in $\mathcal{M}^f$ at $[t_0]$. The claim itself is proved by induction on ψ; the atomic case follows by the definition of π^f. We leave the boolean case of the induction step to the reader, and only treat the case where ψ is of the form $G\chi$ (the $H\chi$ case can be proved similarly). Observe that if ψ is a subformula of φ, then so is χ, whence we may apply the inductive hypothesis to χ. First assume that $\mathcal{M}, s \Vdash G\chi$. In order to prove that $\mathcal{M}^f, [s] \Vdash G\chi$, take an arbitrary equivalence class, say $[t]$, such that $R^f[s][t]$; it suffices to show that $\mathcal{M}^f, [t] \Vdash \chi$. It follows by the definition of R^f that $\mathcal{M}, t \Vdash \chi$, so by the inductive hypothesis, we find that indeed, $\mathcal{M}^f, [t] \Vdash \chi$.

For the other direction, assume that $\mathcal{M}^f, [s] \Vdash G\chi$, and let t be an arbitrary point in the future of s. It is straightforward to verify that $s < t$ implies $R^f[s][t]$; thus we obtain that $\mathcal{M}^f, [t] \Vdash \chi$. Hence, by the inductive hypothesis we find that $\mathcal{M}, t \Vdash \chi$. Since t was arbitrary, this gives $\mathcal{M}, s \Vdash G\chi$. QED

For practical purposes decidability is not enough; 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, given a Priorean formula φ of length, tells you before performing $f(n)$ steps, whether this formula is satisfiable in a linear frame or not; here f is in fact a *linear* function.

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 operators. For many application 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.

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

In a picture (assuming linear time):

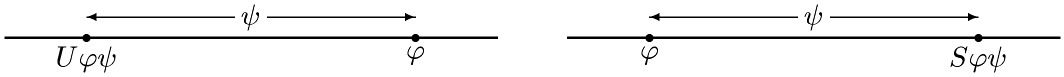

Figure 2: The operators U and S.

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 this by applying temporal bisimulations. Consider the valuations π and π' on the flow of time $\mathcal{R}$ of the real numbers; π makes q true at all non-zero integers, and in addition, π' makes q true in all points of the form $\frac{1}{z}$ with z a non-zero integer. Then according to π, q is true for an uninterrupted period of time after 0: $(\mathcal{R}, \pi), 0 \Vdash U\top q$; but this is not the case for π'. On the other hand, we can define a temporal bisimulation between these two models that links 0 (in the one model) to itself (in the other model); but then it follows from Proposition 3.1 that the Priorean language cannot distinguish the one 0 from the other. The definition of this bisimulation is very simple: simply link any 'q-point' in the one model to any 'q-point' in the other, and similarly with the 'non-q-points'.

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. For instance, in computer science, where one wants to talk about the next state of a process, this operator plays an important role. 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 $\mathcal{L}_{su}$. To be more specific, we need to say what we understand with the notion of a temporal operator: our only constraint will be that the truth definition of a temporal operator is expressible in first order logic. An equivalent way of expressing this is to say that for any temporal operator we can add a first order clause to the definition of the standard translation given in (2); for instance, the clause for U itself is:

$$ST_x(U\varphi\psi) \;=\; \exists y\, (ST_x(\varphi)[y/x] \wedge \forall z\, (x < z < y \rightarrow ST_x(\psi)[z/x])),$$

and a similar definition applies to S.

Theorem 4.1 (Kamp) *Over the class of linear, continuous, flows of time, every formula $\alpha(x)$ in the first order model language is equivalent to the standard translation of an $\mathcal{L}_{su}$-formula. Hence, over the class of linear, continuous orderings, every temporal operator can be defined in $\mathcal{L}_{su}$.*

The restriction to linear, continuous flows of time is essential; if we drop the requirement of continuity, there are temporal operators that can *not* be defined with the help of S and U. Nevertheless, two new operators (with a rather complex truth definition involving gaps) may be added to the language in order to obtain a new expressively complete formalism. Once we drop the condition of linearity however, the results tend to be negative; for instance, over the class of all flows of time we can not 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.

Branching Time Languages

As we have already mentioned, allowing flows of time that branch to the future we can no longer assume that the past already 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 approaches 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 $F_\Box$; this operator has a second order definition, involving a quantification over all branches through the actual time point:

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

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 $\Box$. 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.

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

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

As we already mentioned in section 2, there are contenders to the point-based paradigm in representing time: 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. In this section we briefly discuss such alternative perspectives on time.

Motivations It seems that 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.

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.

Events It is obvious that the second argument against point-based models of time may be pushed even further, questioning the Newtonian perspective in which absolute Time exists regardless of anything happening in it. For, if one is advocating a framework which is ontologically based on our everyday experience, then periods of time are just as abstract artifacts as points. If time does not exist without change, perhaps one should take *events* are the central objects?

Not that this notion is without problems. In particular, although we may have an intuitive idea of what events *are*, there seems to be no consensus whatsoever when it comes to a precise characterization of this notion. For instance, it is obvious that events have a composite nature: the generic event will not only occupy time but also space, and on top of that, involve the manifestation of one or more agents. But are events determined by these components? Is the event of me buying a car from you the same as the event of you selling this car to me? It will come as no surprise that the problem of how to *identify* events is highly controversial.

Nevertheless, it is obvious that events have to be taken into serious account when linguists or computer scientists start to formalize our talking or reasoning about change. An interesting notion stemming from the formal semantics of natural language is that of *aspect*, usually described as the intrinsic temporal constitution of an event. An influential categorization of events, due to Vendler, is that in the following four classes: achievements ('arriving'), activities ('walking') accomplishments ('building a house') and states ('knowing the answer'). Clearly, this classification involves criteria concerning the duration of an event, but also facets like telicity (goal-directedness) that are not purely temporal.

Regrettably, we have no space for a proper discussion of event-based temporal logic.

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 other than in 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 alternatives. 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

$$\begin{array}{lll} \mathcal{M},[s,t] \Vdash \langle D\rangle\varphi & \text{if} & \mathcal{M},[u,v] \Vdash \varphi \text{ for some } t,u \text{ with } s \le u \le v \le t, \\ \mathcal{M},[s,t] \Vdash \varphi \circ \psi & \text{if} & \mathcal{M},[s,u] \Vdash \varphi \text{ and } \mathcal{M},[u,t] \Vdash \psi \text{ for some } u \text{ with } s \le u \le t. \end{array}$$

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) that 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

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

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 (as in the part on correspondence theory in section 3): with the proposed reduction, the corresponding first order model language could do with unary predicates, and hence be decidable for many classes of linear flows of time. It is obvious that the particular proposal (9) 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 research area with applications in many disciplines. This field is quite active, and while many researchers see it branching off in various and diverging directions, one can still identify an huge common trunk, which comprises more wood than what we could describe in the previous sections. And also, one may discover that many of the branches have a common growth pattern. In this section we will try to 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 It is obvious that for serious real-world applications the kind of temporal logics that we have been describing until now are far too simple; one common trend in temporal logic is to study logics of richer ontological structures. To mention an 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 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 Over the last decades, 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 formalisms but in connection with other branches of logic such as algebraic or first 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. 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. Such enterprises 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/or periods is or 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.

Further reading In this chapter we have only scratched the surface of Temporal Logic. The following collection of 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].

Acknowledgements The research of the author has been made possible by a fellowship of the Royal Netherlands Academy of Arts and Sciences. Personal thanks are due to Johan van Benthem for helpful comments on a draft version of this chapter, and to Marco Aiello for a scrutinous reading of the manuscript.

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.