The Knowledge Engineering Review, Vol. 4: 2, 1989, 141–162

A review of temporal logics

DEREK LONG
Department of Computer Science, University College London, Gower Street, London, UK

Abstract

A series of temporal reasoning tasks are identified which motivate the consideration and application of temporal logics in artificial intelligence. There follows a discussion of the broad issues involved in modelling time and constructing a temporal logic. The paper then presents a detailed review of the major approaches to temporal logics: first-order logic approaches, modal temporal logics and reified temporal logics. The review considers the most significant exemplars within the various approaches, including logics due to Russell, Hayes and McCarthy, Prior, McDermott, Allen, Kowalski and Sergot. The logics are compared and contrasted, particularly in their treatments of change and action, the roles they seek to fulfil and the underlying models of time on which they rest. The paper concludes with a brief consideration of the problem of granularity—a problem of considerable significance in temporal reasoning, which has yet to be satisfactorily treated in a temporal logic.

1 Temporal Reasoning: The importance of time in problem solving

Time, despite its elusive nature, is pervasive throughout the human reasoning occurring in everyday situations. Probably the most common use for temporal reasoning is in understanding natural language, where one of the first hurdles to be crossed is that of disentangling the temporal references implicit in tensed verbs. However, time plays a major role in reasoning about physical systems, both in predicting and in explaining behaviours. It can be seen, then, that reasoning about and with time is a task with wide application in many problems in the domain of artificial intelligence.

The most primitive role for temporal reasoning is in maintaining a temporal database. In order to support a system able to reason about data that includes temporal information, it is necessary for that data to be kept consistent and accessible and that any immediate consequences of what is known be added to the database. Temporal database management can be seen as an extended version of the usual problem of database management—Dean's time map manager is an approach to this task which follows that line (Dean and McDermott, 1987). The event calculus, which is reviewed in section 4.4.3, is also, at least in part, a temporal database manager, which follows the conventional approach of a logic programming language database (Kowalski and Sergot, 1986). Temporal database management underlies all other temporal reasoning tasks, so it can be assumed that its needs must be addressed when considering the more sophisticated forms of temporal reasoning that follow.

Shoham (1988) claims that temporal reasoning can be divided according to four main roles:

- prediction,
- planning,
- explanation and
- learning about the physical behaviour of the world.

It is clear that all four of these tasks are valuable components in the armoury of human reasoning and would be equally valuable tools with which to equip AI systems.

Prediction requires that a starting state be given, together with some knowledge about the way that the world evolves (at least some of which is likely to be causal knowledge), and from these some estimation of the future states of the world is made. More specific forms of prediction involve considering specific properties or states and predicting their behaviour or occurrence, without concern for the rest of the world, or attempting to predict the state of the world at some specific time. Simulation techniques can be seen as a particular approach to prediction.

In his discussion of these reasoning tasks, Shoham points out that planning and explanation are closely related: in planning one decides on a current state and a goal and searches for a sequence of events which can be invoked to bring about the goal, while in explanation one also determines a goal event (the event to be explained) and a starting state (a state known to have existed previously) and then searches for a series of events which could have led from the start to the goal. However, planning and explanation *are* different. The events which are of interest in planning are generally ones which the planner can influence and, ideally, do so with minimal effort or cost. In explanation there is no such constraint on the events to be considered. Further differences arise as a consequence of the different purposes of planning and explanation: the purpose of planning is to achieve the final identified state with the least cost; the purpose of explanation can be to enhance one's own or someone else's understanding of a system to enable future prediction of similar behaviour, or to diagnose a current state in order to understand its significance and so on. Thus, explanation covers a broad spectrum of tasks. Finally, an important difference between planning and explanation is found in the difference between those events for which plans are constructed and those events for which explanations are constructed. Typically events that have to be planned require the expenditure of effort directed towards achieving the plan, while the same is not necessarily true of events which must be explained. For example, one might attempt to explain why a wet match fails to light when struck, but it does not appear to have much meaning to *plan* that a wet match should not light when struck. From this account it can be seen that explanation and planning are temporal reasoning tasks with different structures.

There is a difference between planning and prediction: in the latter only a starting state is known, and the progress of the world is considered without the possible effects of intervention by the temporal reasoner, while in the former start and goal states are given and the progress of the world is guided by actions of the reasoner in an attempt to bring about the goal state. Of course, it is also possible to see planning as the prediction of the effects of certain actions on the world. However, it is more than that, since the actions are chosen in the light of their expected and intended effects. The problem of prediction is one that is associated with the idea of *possible worlds*—the collection of all those states of affairs that might be reached from the current world by some particular series of actions or events. It is useful to distinguish the possible worlds which arise out of a lack of knowledge (that is, insufficient is known about the present and the past to enable the future to be predicted accurately—the inaccuracy being reflected by the series of possible worlds), from those which arise out of non-deterministic behaviour. In the latter case one or more agents might have choices which can be made non-deterministically (at least, within the model of the world assumed as a basis for the logic). The choices lead to different possible worlds. The relationship between these differently generated possible worlds and the *actual* future is an interesting one: possible worlds generated out of ignorance might be modelled by making *uncertain* statements about the future, while the non-deterministic actions might be modelled using statements about hypothetical futures, based on assumptions about the courses of action.

The problem of learning new physics is certainly a very hard one, which Shoham argues must rest on the preceding reasoning tasks. In addition, one of the underlying reasoning tasks involved in learning new physics is that of performing induction. Induction is a reasoning technique which frequently relies on temporal behaviour—generalizations about the properties of objects or systems implicitly generalize across different instances through time.

Linguistic philosophers and researchers in natural language understanding would certainly add to Shoham's list the use of temporal reasoning in understanding language. Not only tense, but also more subtle problems, rely on reasoning about time. For example, consider the following sentences:

- I read a book about evolution two million years ago.
- I read a book about evolution two days ago.

To disambiguate these sentences one must understand typical time scales of various activities and reason about their relationship to specific periods of time.

Further consider:

- Yesterday John smoked twenty cigarettes.
- Yesterday John bought twenty cigarettes.

There is an implicit reference in the first of these sentences to twenty separate periods of time, while the second (in its most natural interpretation) refers to only one period. To be sure, both of these examples rely on aspects of natural understanding which have application far beyond temporal reasoning (handling typical or default values, implicit reference and natural interpretation), but nevertheless, there is an important temporal component in each of the examples which cannot be ignored. A valuable effect of considering the use of time in natural language is that it forces a treatment of time in a way other than simply as another variable, which is the approach of mathematical sciences: the distinction between temporal quantities and other quantitative values, such as temperature, distance, mass and so on, is highlighted by the rich vocabulary that surrounds the use of time in natural language ("always", "sometimes", "never", "now", "while" and so on). In equational models of physical systems time is not generally afforded a privileged role in comparison with other physical quantities, although once the equations and their solutions are interpreted, the particular nature of time reasserts itself.

An interesting role for temporal reasoning is to be found in attempting to understand the behaviour of concurrent computing processes and parallel hardware architectures (Moszkowski, 1985). Prediction and planning are both employed in this application. One aspect of this application which makes it of special interest is that it can be important to model and reason about the possible simultaneity of actions performed by different processes.

Finally, an important use of temporal reasoning is in rational reconstruction of world states. That is, given observations of a current world state, rather than predicting its future evolution, attempting to reconstruct its past history. This is intimately linked to explanation, although it is not the same thing. For example, an attempt to describe the events leading up to an air crash might well rely on explanations of certain observations and might yield an explanation of its occurrence, but the reconstruction of the events preceding the crash do not themselves constitute an explanation of the crash.

To summarize, temporal reasoning plays a role in at least seven tasks:

- temporal database management
- prediction
- planning
- explanation
- learning new physics
- natural language understanding
- historical reconstruction.

2 The case for temporal logics

Having identified a role for temporal reasoning, some means of supporting it must be found. Here we stumble into the familiar debate about the role of logics in reasoning. There are many arguments for wishing to use logics as a support for the description, understanding and recreation of reasoning which are nicely summarised by Shoham (Shoham, 1988), following Drew McDermott (1978), as a way of maintaining "mental hygiene". That is, logics require a precision and formality which encourage care and depth of understanding in the expression of ideas. In addition, they can pave the way towards a computational model for the concepts they are used to express, based, for example, on logic programming languages. Finally, we might hope that since logic has become the *de facto* common currency of the expression of ideas in AI, the integration of approaches to largely orthogonal problems, such as temporal reasoning, the treatment of uncertainty, reasoning about actions and spatial relations, should be simplified.

Nevertheless, it must not be thought that we are proposing that logic is the only way to approach problems in AI in general, or in temporal reasoning in particular. An extremely important support for temporal reasoning, allowing accurate prediction of the behaviour of certain physical systems, is the thoroughly tried-and-tested method of mathematical modelling and equation-solving. This certainly should not be underestimated nor neglected when facing problems in temporal reasoning. Indeed, to retain any justification for pursuing a study of logics, it is necessary to keep the goals of the study, in its applications to the problems of AI, clearly in mind.

3 Modelling time

In order to develop a temporal logic, it is necessary to establish some notion of what structure of time the logic is intended to reflect. Although this question can be considered as a fundamental philosophical question about the absolute nature of time, in the context of temporal reasoning it is more fruitful to consider it as a pragmatic question, to be answered in the light of the expected applications of the logic. Thus, the various options considered as responses to the following issues need not be viewed as contenders for a position of absolute authority, but rather as possibilities to be considered in the light of a particular application.

Amongst the issues which must be considered is the question of whether time is considered discrete or continuous. There are many applications in which time can be naturally and conveniently considered to be discrete—in reasoning about computation, for example, time can be modelled as clock-ticks. It is also possible, when dealing with a finite set of properties, to consider time as if it were discrete, by considering the smallest (indivisible) unit to be the largest period over which none of the properties changes.

The main argument in favour of continuous models for time is that it corresponds to both the usual intuitive structure for time and also the usual model of time adopted in classical physics. In order to model continuous change it appears inevitable that there should be continuous time. A more fundamental question is whether continuity actually affects the behaviour of a system from the point of view of reasoning about its interaction with the reasoner and with the rest of the world. What often seems more important than the continuity itself is the fact, inherited through the denseness of continuous models, that between any two points it is possible to place a third point. This property ensures that whenever two events are separated in time, it is always possible to find a time between the events at which to interpose some further event known to occur after one of the events and before the other. Frequently continuous change is only of interest because it grants this property. For example, the continuously filling sink considered in (Shanahan, 1988) allows that if any point on the side of the sink is marked, then there will be a point in time at which the water level reaches the mark. In any finite computation, however, it will only be possible to identify and reason about a finite number of such marks, so only a finite number of points of time need to be added to the interval between the tap being turned on and the sink overflowing. So, it can be seen that the reasoner need only deal with discrete units of time in any finite piece of reasoning.

A related issue is discussed in section 5, that of *granularity* of time. Even if time is considered to have an underlying continuous structure, it appears, at least in natural language, that time carries a natural discretisation with respect to any given event or action type, corresponding to an appropriate *grain size* for that event or action.

The important question of the ordering of time must also be addressed. Time might be taken to be a single time-line, or a branching structure (forward or backward or both), parallel or even circular. A branching structure is one in which the temporal primitives (intervals, instants, states or events, or whatever they might be) are only partially ordered. The primitives then fall into classes which are totally ordered—chains of primitives—forming a possible direction for the evalution of the world. Typically the events within separate chains can be mutually exclusive. For example, the event of it raining in London at a certain time might fall in one chain, while the event of it being dry in London at a time which is the same distance (in time) from the present as the former event falls in a different chain. Thus, the two events represent mutually exclusive possible futures—one in which it rains over London and the other where it does not. Notice that the time that the events refer to cannot be the same, since that would lead to a statement being true and false at one and the same time, but the relationship between the times of the two events and the present moment can be the same, thus capturing the ideas that the two events are both possible future events for some particular time in the future.

If a single continuous time line is adopted, it is necessary to consider further whether it is well-ordered, so that it is possible to define a "next moment" with respect to any given time. Well-ordering is not a property of time modelled after the real line, so that in this case there is not a way to give meaning to a "next moment" operator. However, in natural language "at the next moment" is a common phrase and it would be useful to adopt a model in which it can be given an intuitively appealing meaning. Linear time corresponds to the classical physical model of time; branching time offers an attractive way to handle possible worlds, uncertainty about the past or the future and the effects of alternative actions when planning. Branching time does not capture the fact that of all possible futures or pasts there is precisely one *actual* future and past, while all the others will always remain hypothetical. Parallel time lines are a way of modelling separate and asynchronous processes, but are only interesting once some way of introducing interaction between time lines is constructed. This model might prove useful in developing logics for reasoning about parallel computation and concurrent processes. Finally, circular time might seem strongly counter-intuitive, but it has a place in modelling the behaviour of repetitive, cyclical processes. Many manufacturing systems might be modelled in this way and reasoned about usefully.

A further issue is the boundedness of time: should time be considered infinite in either or both directions? In most applications this question is irrelevant, since only a finite, though perhaps arbitrarily large, span of time is ever required. However, there are time-critical applications in which the sub-world of the application could be usefully considered to be finite and bounded. For example, real-time planning of satellite-probe operations requires that decisions about allocation of resources to experiments be performed taking into account that the descent of the probe is finite and that the time-span for experimentation is critically short.

Finally, there is the issue of whether time *instants* exist, or only intervals. Again, this issue is strongly linked to that of granularity considered in section 5, since instants might be considered to be the objects of smallest grain size. It might be argued that points could always be modelled as infinitesimally short intervals, but this need not be the case—Allen, for example, explicitly denies access to points in time within his interval calculus. Similarly, modelling intervals by taking their end-points in a point-based model can lead to problems: the annoying question of whether end-points are in the interval or not must be addressed, seemingly without any satisfactory solution. If the end-points are in an interval then adjacent intervals have end-points in common, which when adjacent intervals correspond to states of truth and falsity of some property, can lead to stituations in which a property is both true and false at an instant. Similarly, if end-points are excluded, there will be points at which the truth or falsity of a property will be undefined. Finally, the solution in which intervals are taken to be half-open, so that they sit conveniently next to one another, seems arbitrary and unsatisfactory.

Not all of these points are addressed explicitly in the review of temporal logics that follows, but each logic carries with it an implicit answer to each of these questions. What must be borne in mind is that the answers to these questions can affect the applications to which the logic is best suited and ought, therefore, be amongst the first issues considered when treating particular problems in temporal reasoning. Further more detailed discusion of the topologies of time can be found in (Newton-Smith, 1980).

4 A review of temporal logics

Existing temporal logics fall into several different groups. The first is best described as the first-order logic approach, led initially by Russell (Russell, 1903). Broadly, this approach involves introducing an additional argument into first-order predicates to represent a time stamp, but otherwise using first-order predicate calculus, frequently with the minor complication that it is sorted, enabling quantification specifically over time variables. The second approach is that of model logics, in which a series of new connectives are introduced into standard first-order predicate calculus, enabling statements about the future and the past to be formulated. Thirdly there is the reified logic approach. Here the ontology of the logic is complicated by the addition of new primitives representing not only temporal quantities, but also, for example, events, processes and properties.

We now proceed to consider these approaches in greater detail. In doing so we will seek to identify certain important features of each approach which are relevant to the tasks listed above. In particular, the way that the logics allow treatment of change, the assumptions that are made about the nature of time and the expressive power with respect to intuitively acceptable temporal concepts.

4.1 The first-order approach

The first-order approach is in many ways the most intuitive way to construct a temporal extension to first-order predicate calculus. The simplest way to perform the construction is by adding an extra argument to each predicate which is used for time values. Formulae now refer to the truth of predicates when applied to particular terms *at a particular time.* To state that there is a time when it is raining over London we write:

$$\exists t.\ \mathrm{rain\_over}(\mathrm{London}, t).$$

The logic is usually *sorted* so that the variable t is a time variable and the quantification is only considered over time values (thus, we need not be concerned with the truth or falsity of rain_over(London,London), which is the nonsensical statement "it is raining over London at time London").

A suggested modification to this approach (Quine, 1965) is to consider objects not as single entities progressing through time, but as a history, or progression, of entities forming a time-line. Instead of adding the time argument as an additional place in the predicate, using this approach requires that the time which is to be linked to an assertion is actually added to each object in order to isolate a "time-slice" of the history of the object. For example, the above formula becomes:

$$\exists t.\mathrm{rain\_over}(\mathrm{London\text{-}at\text{-}}t),$$

where London-at-t is the slice of the entire progression of London through time, taken at just the time t. One effect of this formulation, although very similar to Russell's approach, is to emphasize the Platonistic aspect of the world of classical logic—the objects in the world are forever present, so that the histories of objects are infinite in both directions, if one allows a model of time which is infinite in both directions. This means that even objects which appear to have come into existence at some definite time must still be modelled as if they had always existed, if not always with the same physical aspect. So, for example, it has to be possible to give meaning to statements about London-at-t, where t is, for example, some time several millions of years ago. In the world of abstract and often (ideally) unchanging concepts, such as mathematics, this Platonistic view does not clash quite so strongly with intuition as it does in this attempt to model the temporal progression of a changing world.

Russell's temporal logic does not, in the first instance, require any assumption to be made about the topology of time, other than that it is possible to refer to points in time and the same basic assumption shared by all the first-order approaches reviewed in this section: that time can be treated as discrete, with changes modelled as discrete steps. The questions of continuity, ordering and boundedness are all addressed in the selection of axioms to treat the relations between different points in time. This allows a wide flexibility in modelling time. For example, in developing an ordering on time, using the predicate later($t1$,$t2$), which is true precisely when $t1$ is later than $t2$, it is possible to make the ordering a linear ordering, a well-ordering or only a partial ordering (granting branching future, past or both), by choosing from amongst axioms such as:

$\forall t1 . \forall t2 . \forall t3 . \text{later}(t1,t2)\ \&\ \text{later}(t2,t3)$ ɑ $\text{later}(t1,t3)$

enforcing transitivity;

$\forall t1 . \forall t2 . \text{later}(t1,t2) \text{ or } \text{later}(t2,t1)$

enforcing comparability (that is, a total ordering)

and so on.

McCarthy's and Hayes' Situation Calculus (McCarthy and Hayes, 1969) also falls into this group of first-order approaches. The calculus differs very slightly from the logics of Russell and Quine in that McCarthy and Hayes do not choose the time variables to refer to points in time, but to *situations*, or "states of the world".

The ways that Russell's logic and the Situation Calculus allow change to be treated are also slightly different. Russell considers change to be "the difference, in respect of truth or falsehood between a proposition concerning an entity and a time, T, and the proposition concerning the same entity and a time, T′". McCarthy and Hayes, whose primary interest is in planning, treat change through functions from situations to situations, representing actions. So, the action of painting a blue house red would be expressed as follows:

colour(house1,blue,*s*) → colour(house1,red,paint(house1,red,*s*)),

where paint(house1,red,*s*) is a new situation created out of the old situation, *s*, by painting house1 red. This treatment allows McCarthy and Hayes to consider the construction of plans as the development of a series of situation-changing actions which proceed from the initial situation to one in which the desired final properties hold of the system.

Both of these approaches suffer from deficiencies, which are apparent in the assumptions that are made about the nature of time and change. A problem that is common to both approaches is that change must be considered to be discrete, not continuous. In the applications that McCarthy and Hayes envisage for the Situation Calculus this is not a serious problem, since all the actions available are discrete (or may be considered discrete without losing any of their interesting properties), such as moving blocks around in the typical "block-world". This view of change is not satisfactory, however, when what is required is the ability to express statements about continuous change, such as "my car is slowly rusting away", or "the temperature in the fermentation vessel is steadily rising".

A further disadvantage to the Situation Calculus lies in the use of functions to describe change: since change is expressed through the application of an action-inducing function to a given situation, it is not possible to express simultaneous action unless further functions are introduced which combine the effects of two (or more) separate actions. This is an entirely unsatisfactory technique, suggesting as it does, for example, that there is a significant difference between the action of crossing a road and the action of crossing a road while a pigeon flies overhead, both requiring entirely different functions to change the state:

∀*s*. by_road(John,side1,*s*) →
by_road(John,opposite(side1),cross_road(John,side1,*s*))

∀*s*. by_road(John,side1,*s*) →
by_road(John,opposite(side1),cross_road_with_pigeon(John,side1,*s*))

In many uses of temporal reasoning—for example in reasoning about concurrent computing processes—simultaneity of action is important, rendering the treatment of action in the Situation Calculus at best expensive and clumsy, at worst useless.

A final problem, which is particularly difficult to face with the Situation Calculus approach, but is actually a problem which must be considered with all approaches to temporal reasoning, is that of persistence. It is intended that each situation of the Situation Calculus contains a complete description of the world (at least that part of the world which is relevant to the system in question). When this description is very large, as is inevitable for complex systems, it becomes extremely important to be able to infer that changes in state which are brought about by certain actions do not change the state of components of the world that are untouched by the action. For example, changing the brightness control on a television does not affect the volume (at least, on a fully functional television), nor does moving a block-world affect this position of blocks not touching the first block before during or after the movement. This problem is another aspect of the so-called "frame problem" in which a huge, possibly infinite, number of conditions must be recorded to unambiguously specify the behaviour and condition of a system. Although this problem is insidious and finds its way into most aspects of AI, it is particularly acute with the Situation Calculus approach, since a situation cannot even be defined, strictly, if the world is large and complex, and this makes the use of situations as the fundamental temporal primitive cumbersome at best.

Perhaps the most important contribution that these logics have made to temporal reasoning is to show that it is both possible and potentially useful to add a temporal dimension to a logical formalism. All of the approaches considered in this section are somewhat dated in comparison with those reviewed in the sections that follow, which demonstrate a closer regard for the problems they seek to address.

4.2 The modal logics approach

Modal logics were first considered in order to express possibility and necessity. The potential of both of these operators is only fully realized when they are used with an evolving model. The semantic model usually used to give meaning to them is that of Kripke's possible worlds (Kripke, 1963) (or some variation of it). The model postulates a series of "worlds", with an accessibility relation which expresses the possible transitions between these worlds. If the accessibility relation is treated as a temporal ordering of the worlds then time is smoothly introduced into the model. Possibility and necessity can then be reinterpreted in a temporal framework, where they have a natural meaning associated with prediction and future possibility.

In order to fully exploit the concept of time within modal logics, the following connectives, due to Prior (1955), are commonly used:

F p meaning "p will be true at some time in the future"
P p meaning "p was true at some time in the past"
G p meaning "p is always going to be true in the future"
H p meaning "p has always been true in the past".

Modal temporal logics are also referred to as "tense logics" and their practitioners as "tensers", as opposed to the "detensers" who advocate other approaches. It is easy to understand why this is the case if one considers a statement in the first-order approach say, such as the following:

happy(Bill,t),

which states that Bill is happy at time t. Notice that the same statement is made regardless of whether t is in the future or the past, or is even the present. Thus, the statement requires and implies no tense. It is not possible to make an identical statement in modal logics, since not every time point is explicitly labelled. That is, with the absolutist approach of the first-order logics, time exists independently of the events which occur within it and every point in time can be identified by name; in the relativist approach of the modal logics, time does not necessarily exist independently of the events which occur within it, but, instead, is defined by those events, so that the only points in time which are identifiable and are named are those at which some event occurs, and the name of the point in time can be taken from the event. For example, the time at which the event occurs of the minute hand of a watch arriving at the twelve and the hour hand at the two can be given the name "two o'clock". This example serves to illustrate that a relativist approach allows a straightforward construction of a common reference time, similar to the time structure of the absolutist approach. The difference then lies in the ontological precedence of the structures of events and time. With the relativist approach statements are relativised to the present and to other events, thus:

F happy(Bill)

would mean that Bill will be happy (in the future, relative to the time now, or the event of our considering the formula, now). Similarly, we could state that Bill was happy, or that Bill is happy. Each statement is relativised to the present and is therefore tensed. Later, after the time at which Bill was to be happy has passed, we might wish to state:

P (**F** happy(Bill))

or "it was true in the past that Bill would be happy". In a situation in which different agents might have different (subjective) views of the flow of time, the absolutist approach can become more clumsy, requiring that a separate time structure be constructed for each agent.

This ability to express tense in modal logics makes the approach a stronger one for treatment of natural language. The work of Reichenbach (1947) is particularly interesting in this context, in which it is proposed that tenses of natural language can be all understood using just three times for each tense: the utterance time (the time at which the statement is made), the reference time and and the event time. For example, "John should have gone" can be understood using the utterance time (at which the statement is made), the reference time (at which it will be true that "John has gone") and the event time (the time at which John actually leaves), where the reference time lies ahead of both the event time and the utterance time. Reichenbach shows that all tenses can be considered to imply an ordering on these three times. One conclusion which can be drawn from this work is that natural language does not require the arbitrarily complex structures of modal temporal connectives which modal logics allow (by nesting modal connectives to an arbitrary depth) in order to achieve a rich and apparently unconstrained expressive power.

Without explicit names for points in time, modal temporal logics are less useful for reasoning about systems in which precise times are required, such as attempting to reason about concurrent programs, although Barrington et al. (1984) have shown how this can be done if one uses a discrete model of time, with an additional modal connective, next, allowing reference to be made to the truth of a formula at "the next moment in time".

The treatment of change in modal temporal logics has been considered, in a similar style to Russell's treatment, as a difference in truth values between statements made about the same object, and the same properties, but at different times. Von Wright (1965, 1966) suggested the use of the T connective:

$p\mathrm{T}q$ meaning "p is true now and q is going to be true at some later time".

Although this does not in itself suggest change, if p and q are inconsistent then the truth of the formula indicates some change in the world between the time at which p is true and the later time at which q is true. It has been pointed out that in order to give meaning to the idea that p is true at some time and then q is true at a later time, there must be a concept of an interval of time during which first p and then q is true. Galton (1984) suggest a different approach, turning attention away from the interval over which the propositions are true towards the transition from p to $\tilde{}p$—an instantaneous event. An occurrence of the event can be denoted by (**P** p & $\tilde{}p$). This formulation only requires that p has been true at some time in the past and that $\tilde{}p$ holds at the present. The transition itself might have occurred at any time (or indeed many times) in the past. To ensure that the transition is pinpointed in the present we could use the formula: ((**H** p) & $\tilde{}p$). That is, "p has always been true, and $\tilde{}p$ is true at the present".

This treatment of change still suffers from a deficiency noted already in the first-order approach, which is that change is modelled as a discrete transition from one truth value to another—continuous change does not fit happily into this model.

It is instructive to compare the expressive power of modal logics with that of the first-order logics.

A translation between the modal connectives and formulae of the classical first-order logic is proposed as follows:

$$\mathbf{P}\ p = \exists t.\ \text{later}(\text{now},t)\ \&\ p(t)$$
$$\mathbf{F}\ p = \exists t.\ \text{later}(t,\text{now})\ \&\ p(t)$$
$$\mathbf{H}\ p = \forall t.\ \text{later}(\text{now},t) \rightarrow p(t)$$
$$\mathbf{G}\ p = \forall t.\ \text{later}(t,\text{now}) \rightarrow p(t).$$

The predicate later($t1$,$t2$) is true precisely when $t1$ is later than $t2$. These translations all correspond with an intuitive interpretation of the connectives, at least on the face of it. However, the contrast between the Platonistic world of the first-order approaches and the more Intuitionistic world of the modal logics is highlighted when this translation is applied to a formula such as:

$$\mathbf{F}\ \exists x.\ p(x) = \exists t.\ \text{later}(t,\text{now})\ \&\ (\exists x.\ p(x,t))$$

which becomes: $\exists x.\ \exists t.\ \text{later}(t,\text{now})\ \&\ p(x,t)$

or: $\exists x.\ \mathbf{F}\ p(x)$.

This translation might seem reasonable, but there is a subtle difference between the two modal formulae: the first asserts the future existence of an object satisfying p, while the second asserts the *current* existence of an object which will later satisfy p. Thus, in the former there is an implicit assertion of the "birth" of an object. Of course, no specific semantics has been given here for the connectives of modal temporal logics, and it would be perfectly possible to take a semantics which upheld the translation (that is, a semantics of a Platonistic world). However, in much of the "common-sense" reasoning in everyday use, to speak of an object which is to be created in the future is very different to speaking of an object which already exists.

Pursuing the comparison a little, it is interesting to note that axioms of modal logic, such as ($\mathbf{F}\ \mathbf{F}\ p \rightarrow \mathbf{F}\ p$), which asserts that anything true at a time beyond some future point is simply true in the future, can require that additional axioms be added to the first-order logic in order to ensure their truth in that framework. In this case, for example, an additional axiom forcing the transitivity of the predicate "later" is required. In the opposite direction, there are axioms which might be employed in the first-order logics, such as ($\forall t.\ \tilde{}\text{later}(t,t)$) which states that no time is later than itself, that simply cannot be expressed in the modal framework. This is because axioms like this are specifically concerned with the underlying machinery of the absolute time-line, which is not considered a part of modal temporal logics, so it is not accessible from within them.

Work in temporal modal logics is very active, inspired particularly by its possible applications in natural language understanding. Galton (1984, 1987) is specifically concerned with application of modal logics in this area, while further work in modal temporal logics includes that of Gabbay (1986). Gabbay's interest in modal temporal logics includes the issues of implementation, as well as their applications to natural language understanding.

4.3 Further possible approaches to the treatment of actions

Related to change is the problem of action—the means of bringing about change in the world. One of the most influential pieces of work in this area is that of Davidson (1967), in which all actions are described through the formulation of events. For example, consider the following sentences:

- Tom hit the ball with a bat.
- Tom hit the ball yesterday.

As Kenny points out (Kenny, 1963), it is not satisfactory to propose (in a first-order style) a predicate hit(x,y,z) to express the first sentence, where the predicate is true when x hit y with z, since this obscures the relationship with the second sentence. More attractive might be to postulate a predicate hit(x,y,z,t) which is true when x hit y with z at time t. The sentences then become:

$$\exists t.\ \mathrm{hit}(\mathrm{Tom},\mathrm{ball1},\mathrm{bat1},t)$$

$$\exists z.\ \exists t.\ (\mathrm{hit}(\mathrm{Tom},\mathrm{ball1},z,t)\ \&\ \mathrm{during}(\mathrm{yesterday},t))$$

In this formulation the relationship between the two sentences is clarified. For a modal formulation one does not wish to refer to time in this way and the sentences become:

$$\mathbf{P}\ \mathrm{hit}(\mathrm{Tom},\mathrm{ball1},\mathrm{bat1})$$

$$\exists z.\ \mathbf{Y}\ \mathrm{hit}(\mathrm{Tom},\mathrm{ball1},z),$$

where **Y** is a modal connective meaning "Yesterday".

Davidson's approach is somewhat different to these: he claims that both sentences make a similar statement initially—that an event occurred (that of Tom hitting the ball) and that the adverbial qualifiers indicate properties which can be associated with the event. Thus:

$$\exists e.\ (\mathrm{hit}(\mathrm{Tom},\mathrm{ball1},e)\ \&\ \mathrm{with}(\mathrm{bat1},e))$$

$$\exists e.\ (\mathrm{hit}(\mathrm{Tom},\mathrm{ball1},e)\ \&\ \mathrm{during}(\mathrm{yesterday},e)).$$

Note that the predicate "during" here is slightly different to the predicate used previously. In this instance it requires an *event* as a second argument, where previously the second argument was a *time*.

In this way Davidson creates the intuitively appealing concept of events, which are then given attributes, such as times, places, actors and so on. This approach is another in the style of the first-order approaches, with time names being events. However, the concept of events corresponds closely to that of the event calculus approach reviewed in section 4.4.3.

A problem that must be faced in any general treatment of change, time and action is that of causality. This is, of course, a significant philosophical problem that has given, and continues to give, a major source of debate. Although all of the logics offer some way of treating causality, this review is not primarily concerned with this problem, which is not to belittle it. Shoham (1988) treats causality in some detail and he considers the comments of philosophers, including Russell (1913), Mackey (1974), Lewis (1973) and Suppes (1970). Suffice to say that treatment of causality is, perhaps even more than other issues in temporal modelling, guided strongly by pragmatic considerations inherited from the applications. Causal relationships can be modelled between events (the approach of the event calculus), facts, processes, states or some grouping of these objects. There is probably less argument about the ontological status of causality than there is about the basic temporal primitives.

4.4 The reified logics approach

Reified logics arise when expressions which might naturally be considered as atoms (that is, expressions with truth values) are required to be objects of study—that is, terms. In order to meet this need, the "predicates" are modelled as functions and can then become the arguments of predicates in a first-order logic. In order to attribute some concept of truth or falsity to these "predicates" a predicate is used, such as "True". So, colour(house1,red), a statement asserting that the predicate colour holds of arguments house1 and red (and presumably carrying the interpretation that house1 is red in colour) would become True(colour(house1,red)), in which colour(house1,red) is only a term.

The advantage in this is that it becomes possible to discuss statements which would naturally be considered true or false and add further information to that simple truth or falsity. For example, uncertainty might be modelled through the use of a predicate "Possibly_true", so that, if the degree of certainty in the knowledge that house1 is red is c, the fact can be expressed as:

Possibly_true(colour(house1,red),c)

In order to preserve a distinction between those terms which are intended to behave as if they have truth values and those which are more typical terms (such as roof(house1) in the example True(colour(house1),black))), it is necessary to employ a sorted-first-order logic, in which terms which are intended to have truth values associated with them are *proposition-types*.

It is typical within a reified logic, as will be seen in the examples discussed below, to introduce a much richer ontology and type-structure than is required simply to distinguish proposition-types. Events, processes and properties are all candidates for individual types and are all used in one or more of the logics discussed below.

The important contributions in this area are to be found in the work of Allen (1984), McDermott (1982) and Kowalski and Sergot (1986). Lee *et al.* (1985) also consider the problems of temporal database management in a style not dissimilar to that of Kowalski's and Sergot's event calculus. Moszkowski (1986) has developed the interval temporal logic, which is a modal logic with a semantics based on intervals. It enjoys the significant advantage of having been implemented as a programming language, *Tempura*, in Lisp by Moszkowski and in C by Hale (1987). This logic has been applied particularly in the area of reasoning about computer hardware, digital circuits and networks. More recently, Shoham (1988) has proposed a temporal logic which starts from a reified foundation, but includes modal operators and has a semantics to support non-monotonicity. The development of a semantics is something which Shoham regards as particularly important and he has argued (Shoham, 1985) that despite the popularity of reified approaches, none has been given a clear and explicitly defined semantics—something he sets out to rectify.

4.4.1 McDermott's temporal logic

McDermott is concerned specifically with the problem of reasoning with and about time and temporal primitives in the context of planning. McDermott's intention is to provide a versatile "common-sense" model for temporal reasoning, with a similar flavour to Hayes' naive physics (Hayes, 1978).

McDermott, as McCarthy and Hayes, adopts states as the primitive temporal element. However, he makes several crucial decisions: time is a continuum, with an infinite number of states between any pair of distinct states, and the future is branching—that is, there are many possible futures branching forward in time from the present. Each single branch, consisting of a connected series of states isomorphic to the real line, is called a chronicle.

McDermott is at pains to point out that his treatment is not equivalent to that of McCarthy and Hayes. The main difference that he claims is in the treatment of actions and change. McDermott argues that change is often a continuous process and that the continuous nature is fundamental to many processes of change. Thus, the treatment that the Situation Calculus offers, in which change is instantaneous (the transition from one state to the next) is inadequate. McDermott also argues that events brought about by actions are not always best characterised by a transition between states, or fact changes, as McCarthy and Hayes do; he cites as examples such things as running around a track three times, eating a gourmet meal and so on, where the changes brought about by the action are, at best, only partly relevant. Thus, McDermott distinguishes between facts (propositions held to be true at some indicated time) and events.

McDermott's solution to the problem of characterising events, particularly those which bring about minor or irrelevant state changes, is to identify with the event the intervals over which one occurrence of the event takes place. Intervals are denoted by the states marking their end points—the question of whether the end points are included in the intervals or not is raised, with a rather inconclusive argument favouring closed intervals. He points out, quite rightly, that if statements can be made about points of time, then at least some intervals are closed (the point intervals). He also claims that "it doesn't seem very important for most events whether they include two extra instants or not"—the two instants being the end points. It would appear to be a common feature to treatments of time based on intervals taken from the real line (or some structure isomorphic to the real line) that the question of whether intervals should be closed or open is an awkward one and there is no intuitively reasonable solution (as noted in section 3). The problem becomes most acute when intervals sharing an end point are occupied by incompatible states or events. For example, when a light is switched from on to off, is the light on or off at the instant of switching?

A second problem associated with this treatment of events (or facts) as sets of intervals is that there is a danger of an "event" which consists of an infinite number of intervals within a finitely bound super-interval. An axiom is employed to prevent this, which McDermott attributes to Ernie Davis. This has the effect of ensuring that no state can change infinitely often during a finite interval.

A final problem with this characterization is a more fundamental and also a more philosophical one: in order to define an event in this way one must begin by knowing all the intervals over which it occurs, in every chronicle. This last point is important, since it is only by considering the event over every chronicle that one is able to discard irrelevant features of an event. For example, the event of John going home is characterized by all those intervals in which John starts away from home and ends up at home. To ensure that such extraneous details as the route he took home are not bound into the definition of the event there must be intervals in the characterization in different chronicles in which John's routes are different. In order to precisely characterize the event that is desired, it is necessary to find sufficient chronicles to abstract out not only the route, but the weather, the colour of John's clothes and so on. In fact, it would appear that the only way to ensure that the intervals chosen actually do characterize the required event is to have the event defined in advance! Thus, this construction of events does not allow an internal characterization, which neither corresponds with intuition nor offers any explanation of how events might be constructed initially.

In order to treat continuous change, McDermott enforces an axiom which states that when a quantity is changing continuously, then during any interval for which the values of the quantity at the end points are known, the quantity must take all intermediate values. This certainly captures a more intuitively appealing notion of change and the consequence of action than that offered by the preceding logics. McDermott shows how this property of continuity can be used in reasoning about the behaviour of a tank filling with water. He shows that the axioms can be used to prove that if the water rises to the level of the overflow it will pass every level in between and then spill out through the overflow. He goes on to show how the logic can be used to produce a plan to stop the spilling water.

Having developed a logic to treat action, change and temporal relationships, McDermott is most interested in its applications to planning. He defines a plan to be a sequence of actions, but includes amongst his actions a set of more *passive* actions, such as observing, promising and maintaining, and also negative actions such as preventing and avoiding. Treating actions such as these requires some notion of expectation, which McDermott is able to treat quite readily using his branching time, so that, for example, prevention corresponds to a situation in which in all worlds where the preventative action is not taken, the event that is to be prevented necessarily occurs.

In his treatment of planning, McDermott notes the importance of *choice* in planning—choice of actions. Once again, this highlights the difference between planning and explanation noted in section 1.

McDermott also shows himself to be concerned with implementation issues. One interesting point that arises from his speculation on implementation is the point considered in section 3: using a branching model for time does not distinguish the real time-line from possible time-lines. McDermott admits that any implemented system would have to be allowed such a distinction "for its own sanity". A major implementation problem when taking the possible worlds approach is how to maintain all of the possible worlds. In practice, of course, only those worlds which are relevant to the problem in hand are maintained (or even constructed). The problem then shifts to that of knowing precisely which of the possible worlds are those which are to be considered relevant. Furthermore, using McDermott's characterization of events if (as is inevitable) only a finite set of chronicles are actually kept, there is no opportunity for abstracting new events from the knowledge of the world (since they are characterized by a set of intervals drawn from all possible chronicles). This impossibility emphasises the unsatisfactory nature of McDermott's definition of events.

McDermott's contribution to temporal logics is a significant one. It emphasizes the pragmatic approach to the needs of temporal reasoning in the context of specific tasks—in this case planning.

Although an earlier work than the other reified logics reviewed here, it nevertheless maintains a position of some authority amongst temporal logics. The logic attempts, with some success, to address several far-reaching and central problems of temporal reasoning—causality, persistence (and the frame problem), continuous change, planning and uncertainty about the future—and it is this which accounts for its deserved prominence in the field.

4.4.2 Allen's interval calculus

Allen introduced his "General Theory of Action and Time" (Allen, 1984) in order to provide a framework for the treatment of two problems: firstly, the specific problem of providing an adequately expressive formalism for treating natural language and, secondly, the more general problem of reasoning about actions for planning. Allen argues that the latter is a necessary prerequisite for the former, in any case, so that the whole theory will provide several forms of support in attempting to process natural language.

Allen's temporal logic offers a very different approach to modelling time to that of McDermott and, indeed, to the general "flavour" of all of the preceding approaches. Instead of adopting time points (or states which are associated with time points) he takes intervals as the primitive temporal quantity. Between pairs of these intervals there can hold any one of thirteen (mutually exclusive) relations, including relations describing coincidence of end points (at either end, or at both ends), intersection and inclusion. Not only does Allen take intervals as his primitive temporal quantity, but he also specifically excludes time points in claiming that *any* quantity of time must be subdivisible. This ruling eliminates the possibility of instantaneous events from Allen's treatment. It seems strange that a theory intended to support the expressive power of natural language does not support the expression of instantaneous events, particularly instants such as "now" or the commencement or termination of intervals. Allen's reasoning is that *all* quantities of time denotable in natural language are subdivisible, so must be intervals. He does not suggest how this argument faces the apparently instantaneous moment, "now", but perhaps the answer lies in considering all temporal references to carry an implicit "grain size" (see section 5), including the moment "now", so that subdivision corresponds to taking a finer grain size.

Allen then introduces a fairly rich ontology of temporal primitives: events, properties and processes. Events occur over maximal intervals—that is, if an event, e, occurs on some interval, I, expressed by $\mathrm{OCCUR}(e,I)$, then e does not occur on any subinterval of I. This is intuitively acceptable, corresponding to events such as turning a light on or off and other "indivisible" events. The concept of indivisibility is actually a slippery one and we shall return to it briefly in section 5. Properties are similar to facts in McDermott's treatment, and Allen introduces the predicate HOLDS to deal with them: $\mathrm{HOLDS}(p,t)$ is true precisely when the property p holds over the interval t. Allen demands that, in contrast to events, properties hold on all subintervals of any interval on which they hold. In fact, he goes further than this and demands that, when a property holds on an interval, then every subinterval contains a further subinterval on which the property holds:

$$\mathrm{HOLDS}(p,T) \text{ iff } \forall t.\mathrm{IN}(t,T) \rightarrow (\exists s.\mathrm{IN}(s,t) \;\&\; \mathrm{HOLDS}(p,s)),$$

where $\mathrm{IN}(t,T)$ is true precisely when t is a proper sub-interval of T. This axiom leads to a strange consequence: every interval over which a property holds and which contains one subinterval must contain an infinite hierarchy of nested subintervals. To see this, suppose s(0) is a subinterval of T, over which the property p holds. Then the axiom asserts the existence of s(1), a subinterval of s(0), on which p holds. Since IN is transitive, s(1) is a subinterval of T and, thus, the axiom grants a further subinterval, within s(1), s(2) say, on which p holds. The axiom can be applied repeatedly in this way to generate the infinite sequence of subintervals, s(0),s(1),.... The predicate IN excludes equality and therefore this hierarchy is one of strictly decreasing intervals. It is not made clear whether intervals have an associated (non-negative, real) size as one might expect, which would ensure that this kind of hierarchy converged on some limiting interval.

Properties are intended, by Allen, to be such things as "owning a car", which hold on every subinterval of an interval on which they hold. In contrast, Allen's third set of primitives, processes, is an unpleasant compromise between events and properties. If a process is OCCURRING over an interval, t, then it is OCCURRING on some subinterval of t. That is:

$$\text{OCCURRING}\ (p,T) \rightarrow \exists t.IN(t,T)\ \&\ \text{OCCURRING}(p,t).$$

This definition leads to a similar infinite hierarchy of nested subintervals being posited inside any interval on which process is occurring to that asserted to exist inside any interval on which a property holds. Allen intends the definition to treat processes such as "writing a book", in which the process may be said to be occurring even when the literal activity is not. Unfortunately it is very weak, since it allows processes to be occurring over every superinterval of an interval on which the process is occurring. Thus, for example, it is consistent to assert that John wrote a book over the last 2,000 years, provided John has been writing a book for some subinterval of that period. This does not appeal to intuition about such processes, where it would seem reasonable to expect the process to be occurring over "some reasonable number" of subintervals with a "reasonably large" union, in proportion to the size of the original interval. Allen considers this problem and notes dissatisfaction with the axiom he is forced to employ. He claims that some idea of *grain size* is required, so that a process must be occurring in all subintervals which are larger than the appropriate grain size, in order to be occurring on an interval. This idea is further considered in section 5.

An important difference between McDermott's and Allen's treatments of time is in the decision about the structure of the temporal model, particularly in respect of the future. McDermott, as has already been noted, employs a branching model of the future in order to reflect the notions of possibility and necessity. Allen argues that reasoning about the future is no different to reasoning about the past—both require hypothetical reasoning. Thus, time is modelled through a single timeline, representing the "actual" past and "actual" future. All hypotheses must be reasoned about explicitly as hypotheses. The most plausible of these hypotheses might be considered to represent the best current approximation to the "actual" past and future. Given the problems in implementation of branching models of time foreseen by McDermott, it seems that Allen's argument is supplemented by pragmatic considerations.

Allen's interest in the two temporal reasoning problems of problem-solving and natural language processing leads him to continue the development of his logic in two important directions: a treatment of causation and a treatment of intension and belief. Causation is expressed through two structures: a predicate, ECAUSE(e1,t1,t2) which is true precisely when the occurrence of event e1 at time t1 causes the occurrence of event e2 at time t2, and the function, ACAUSE(a,e1) which occurs precisely when an *agent*, a, causes the occurrence e1. There is a strange assymetry in the treatment of these two forms of causation, taking one form to be predicate and the other a function; Allen offers no reason for this, but it would seem that the decision was influenced by the second direction Allen pursues, intension and belief. Allen is concerned with the relationships between the actions and intentions of agents, so that the ability to make assertions about causal relationships between the actions of agents and their subsequent effects is extremely valuable. Since Allen does not pursue any attempt to model complex physical systems, he does not find it necessary to form expressions about causal relationships between events in the physical world. It seems likely that such a treatment would require that ECAUSES become a function, in order that assertions could be formed about it. The remainder of Allen's treatment of intension and belief is of some interest for the fact that it indicates the important relationship between planning and intention, in contrast to explanation (as discussed in section 1). However, it falls outside the scope of this review and therefore will not be further pursued here.

In Allen and Koomen (1983) a system is described in which Allen's interval calculus is used as a foundation for a planning system—planning for problem-solving. The idea, in outline, is that an initial state and a goal state are each given and a plan is constructed by "simulation" of the evolution of the system, modified where the planner discovers "causal gaps" by the addition of actions which generate an appropriate effect. Thus, the planner requires a complete causal model of the world in which it is to reason if it is to successfully plan its actions. There is not (nor is there intended to be) a way of generating new components for the causal model within Allen's system. This places a great onus on the database constructor, to ensure that the world model is causally complete in the areas where the system is expected to reason about plans.

In addition to the planning application, work has also been conducted, at the University of Rochester, on using Allen's calculus for dialogue understanding, which is, in fact, one of the primary motivations for Allen's work. Indeed, even the planning, given its emphasis on the aspects of intention and belief, can be seen as directed squarely at the problems associated with natural language understanding.

4.4.3 Kowalski and Sergot: The event calculus

Kowalski and Sergot (1986) introduce a temporal logic based on events. This offers a significantly different approach to temporal logic to the proposals outlined above. In the first place, it offers a relativist approach, in which time is considered an emergent phenomenon associated with the ordering of *events*, which are considered to be the primitive temporal objects. Where the situation calculus deals with global states, the event calculus is concerned with local events. The first goal in the development of the event calculus is to provide a framework for managing temporal databases. Such a database can be seen as a necessary precursor to planning, prediction, explanation and the other tasks identified as part of the domain of temporal reasoning. With this purpose in mind, the event calculus is strongly associated with a treatment of default reasoning and an ability to revise information, conclusions and predictions, non-monotonically. As has already been noted, treatments of these problems can be viewed as orthogonal to the problems of temporal reasoning, at least in the first instance. However, Kowalski and Sergot choose to treat the event calculus in a Horn-clause logic, augmented with negation-as-failure, leading to a PROLOG-like implementation. As they point out, employing a similar implementation style is not precluded by adopting one of the other approaches discussed (indeed Gabbay does something similar for a form of modal temporal logic in Gabbay, 1987), but this style does impose a certain "flavour" on the presentation and treatment.

Events, in the event calculus, are considered to be structureless "points" in time, where "points" is used here only to convey the lack of internal structure. Events are not prevented from having duration and work is in progress on ways to consider events at different levels of granularity. Events start and end periods of time, during which states are maintained. Events are considered to be after the time periods that they end and before the time periods that they start, not contained within either of these periods. It is important to note that the idea of states having a duration (between starting and terminating events) introduces a contrasting picture of states to that of McCarthy, say, where states are considered to be instantaneous "snapshots" of the universe. The two types of states are distinct, although both have intuitive appeal. For example, it coincides with common usage to say that the state of a light is either on or off (which corresponds to states with duration), but equally it is valid to speak of a state of a machine during a computation, which is generally considered to be an instantaneous state.

The emphasis on events and the time periods that are started or ended by events is held, to some extent, in common with Allen's approach, although Allen's decision to make time and, more particularly, intervals of time the fundamental primitive leads to a significant onotological difference between the logics. However, in common with Allen, Kowalski and Sergot reject branching time in favour of the single time line, also noting that time can be treated symmetrically in the past and in the future.

Although Kowalski and Sergot take a relativist approach to time, they do introduce an explicit time line to act as a "common currency" for comparing events. To associate events with times (where it is possible—it is not always so), the predicate Time is introduced:

Time(e1,t1) iff event e1 occurs at time t1.

To associate time periods with the events which start and end them, there are two predicates, Start and End:

Start(p,e1) iff event e1 starts the period p;

End(p,e1) iff event e1 ends the period p.

Times are assumed to have a total ordering which must coincide with the ordering of those events which are directly linked with times.

The event calculus also includes a Holds predicate, in a similar vein to Allen's interval calculus, in order to assert the fact that a property holds at some time (interval or period), but with two differences. The first is that Holds is only a one place predicate in the event calculus, requiring only a time period as an argument. The second difference is that Holds is true of a period only when it is a maximal period for which the associated property holds. Thus, Holds(*p*) (where *p* is a period) is true precisely when the property *associated with p* is true throughout the time period *p*. A predicate Rel(*P*,*p*) is used to associate a property, *P*, with a time period *p*. This formulation has been modified in more recent presentations of the event calculus, where a collection of different holds-style relations are given, but adopting a two-place notation rather than the one-place notation described above.

The predicates Initiates(*e*,*p*) and Terminates(*e*,*p*) are used to link events with the properties that they initiate or terminate. These properties are precisely the properties which hold over the period started or ended by the events. The negation-by-failure that the event calculus employs ensures that persistence of properties is the default assumption. In order to identify the end of a period (during which a property holds), it is necessary to discover an event which explicitly terminates the associated property and which occurs at some time following the start of the period. Another way in which a period can be ended is by the occurrence of an event which initiates a property incompatible with that associated with the original period. Incompatibility of properties is expressed through a predicate which must be defined for specific pairs of incompatible properties as they arise.

Events are further enriched by associating with them actions, subjects and objects, as well as the times at which they occur, properties they initiate or terminate and periods they start or end.

The event calculus is developed primarily as a means of supporting a temporal database, so is not immediately concerned with the problem of modelling change. Change is assumed to be associated with events and is reflected in the discrete transition from one state to another, punctuated by an intervening event. The event calculus has been extended to allow treatment of continuous processes (Shanahan, 1988). He does this principally by the introduction of a new predicate, Trajectory(*q*,*t1*,*p*,*t2*) which describes the evolution of a property, *p*, by giving its precise value at time *t2*, when the *second-order* property, *q*, was initiated at time *t1*. A property is considered to be second-order if when it holds a further property is continuously changing. For example, consider the following Horn clause:

Trajectory(moving(train1),*t1*, distance(*d2*,train1),*t2*) ←
  Holds_at(distance(d*1*,train*1*),*t1*, speed(*s*,train1), $d2 = s^*(t2) - t1 + d1$,

where the second-order property is that train1 is moving and its distance travelled, *d2*, by a time *t2*, is given by the property distance(*d2*,train1).

An extension is also a topic of research to offer a way to treat granularity, by allowing events to be decomposed into sub-events. The larger event might be considered as a *process* when perceived as a collection of smaller events. For example, the event of a plane landing could be seen as a process of landing when it is considered as a sequence of events such at losing altitude, gaining the correct attitude, setting throttles, undercarriage and so on. These extensions serve to demonstrate one of the most impressive strengths of the event calculus, which is its versatility. The approach to temporal reasoning through events, which can be enriched with more detail as desired, is certainly an attractive one, and makes the event calculus a powerful vehicle for tackling temporal reasoning in a wide array of domains, from natural language understanding to reasoning about physical systems.

5 Granularity in time

References, implicit and explicit, to time in natural language frequently if not, in fact, without exception, carry an associated "scale of accuracy". For example, consider the following sentences:

- The train leaves at 1407 hours.
- The party starts at 8.

As natural language expressions these sentences both imply a degree of precision in the temporal references. So, it would be reasonable to expect that the train actually leaves within a minute or so of the stated time, while the party might start anything up to an hour or more after the proposed time. These degrees of precision can be referred to as the *granularity* of time references associated with the events mentioned. It is this concept of grain size which enables the disambiguation of the two sentences about reading a book on evolution quoted in section 1. The grain associated with reading books is of the order of hours to days, while evolution has an associated grain of the order of hundreds of thousands to millions of years. So, grain size is the scale of units sensible for the events and states under consideration. Unfortunately, a definition such as this, though intuitively appealing, is of little value in attempting to set up any formal treatment of granularity.

The issue of grain is one that is generally ignored in treatment of temporal logics (Allen mentions the problem in connection with his definition of processes, but offers no solution). Absolutist approaches to modelling time require that some explicit tag be given to temporal primitives in the time structure (such as the points on a time line). The form that such tags might take is usually ignored in favour of the assumption that the tags are ordered (either partially or totally, as appropriate). This assumption is not unjustified, but it does not aid in applying the logics to specific problems. For example, suppose Allen's interval calculus is to be used to express a statement such as "last Tuesday John sold his car". This is certainly not a property, since John did not sell his car at all subintervals of last Tuesday. Intuitively it would appear to be an event, but there is a significant problem in identifying the precise interval over which the sale took place, which is necessary in order to define the event. Did the sale begin at the point when a customer approached John, or when the customer agreed to buy the car, or when money changed hands? An event of this sort, with "fuzzy" boundaries sits well in the event calculus, where the event of the sale can be considered a primitive which terminates John's ownership of the car and starts the customer's ownership of the same car. Within Allen's framework, however, it seems that the only possible conclusion is that the sale must be a process, in which case it is reasonable to state that the process occurred on last Tuesday, despite it not having occurred over all of last Tuesday, but only over some subinterval of that time. Unfortunately, Allen's axioms then force the existence of a subinterval of that time, during which the sale process is occurring, and a subinterval of that subinterval and so on, creating an infinite sequence of properly nested subintervals over which the sale process is occurring. This sequence of intervals must converge to some limiting interval which is precisely the interval at which the event of the sale can be considered to take place. Somehow the intuitively "correct" view of the borders of the sale event being "fuzzy" is lost and replaced with a picture of a sharply defined event. It seems that the granularity naturally associated with the sale of a car is such that the degree of precision with which the borders of the event can be specified is actually only to within some degree—perhaps a few minutes, or even half-an-hour. In order to discuss the event in a temporal logic it would provide an intuitively appealing framework if the time of the event could indeed be specified with that associated degree of accuracy, that is, with the appropriate unit of grain size. Allen's calculus could then be amended so that intervals could only be subdivided down to the appropriate level of granularity.

In practice, examples in which the temporal logics are used do take all the temporal references to be appropriately scaled (of appropriate grain size) for the problem being treated. However, no formalization of the idea appears, nor any means of treating problems in which the actions and events are interacting, but carry widely differing grain sizes.

To some extent the problem is one that can be ignored when taking a relativist view of time, since the events which form the foundation of the time structure carry their own inherent granularity. However, even in this case it is not a problem that can be completely ignored, since events which are considered primitive when viewed with one degree of granularity are compound when viewed with another. For example, the event of landing an aircraft can be seen as a single "point-like" event when viewed in the context of the problem of air-traffic control, but must be seen as a compound series of subevents when viewed in the context of the problem of actually bringing a safe landing about. This example serves to further emphasize how inappropriate it would be to propose as a solution adopting the smallest degree of granularity as the common denominator for expressing all the events in a problem. In that case, reasoning about air-traffic control would become impossibly complex, requiring that all the events involved in actual control of an aircraft be explicitly mentioned in order to model the larger structures, such as the landing itself. It is equally inappropriate to offer the artificial solution of taking an arbitrary smallest-grain event, such as the moment of touchdown of the aircraft wheels on the tarmac', and identifying that with the compound event of the landing.

Advantages that might be had from a formalization of the idea of granularity could include a discretization of the time structure, allowing a well-ordering and computational efficiency, while still maintaining the possibility of a behaviour rather like a dense time-line. Whenever a time unit is required between two existing units, but in a gap which is smaller than the existing grain size, a new unit could be introduced by switching grain size to a smaller unit. This corresponds to "opening out an event or interval", and is an attractive idea. Continuous change could still be modelled by allowing all continuous changes to be opened out arbitrarily many times without requiring a change in the basic process considered. For example, the process of erosion of a rock is a continuous one, which might have a grain size of the order of centuries. In order to consider the formation of a particular striation on its surface it might be appropriate to open up the grain size and consider units of decades. The fact that the erosion process is continuous would allow the same process—the erosion—to be applied with meaning to units of decades. In a similar vein, the erosion process could be treated as a meaningful one through ever decreasing grain sizes. In contrast, the process of landing a plane is not continuous, and when it is opened up and individual pilot actions considered, it becomes meaningless to treat the process of landing during each of those shorter periods of action.

Other thoughts on granularity, which is an important if underdeveloped area of research, can be found in, for example (Hobbs, 1985).

6 Summary

The development of temporal logics, which began with the attempt to demonstrate that time can be introduced into classical logics and reasoned about in a meaningful and useful way, has diversified into the study of several different classes of logic and members of those classes. The study is motivated mainly by a desire to treat one or more of the problems in temporal reasoning, outlined in section 1, particularly the tasks of planning, natural language understanding, temporal database management and, to a lesser extent, prediction, explanation and historical reconstruction. This aim has influenced the directions of development, as the logics have been tailored in a pragmatic manner, to suit the problems they address. This has led to a series of different temporal models and widely different formalisms in which to express temporal relationships. Although all of the approaches can claim their merits, it is fair to say that none of them represent a universally and uniformly good approach to all of the problems. Thus, as with most areas in the field of artificial intelligence, there is still a great deal of room for further work.

References and Bibliography

The list of references and other relevant material has been split into three separately ordered sections. The first of these includes a list of valuable review material. The second list is the set of references (other than those appearing in the first list) that are made in the text. The third list contains a more general bibliography of material which is connected, directly or indirectly, to temporal logics.

Van Benthem's book is a relatively recent and extremely thorough treatment of temporal logics. Galton's article has provided a considerable source of inspiration for the material in this review and is a valuable summary of the main work in temporal logics. It appears in a collection of papers edited by Galton, all of which are valuable sources in the area of temporal logics. There is, in particular, an interesting opportunity to contrast three implemented temporal logic languages: Hale introduces Tempura, a language based on the interval logic of Moszkowski (1983), Sadri discusses the event calculus (a Prolog-based implementation of the temporal logic) and Gabbay develops a modal temporal logic implementation around a Prolog core. McArthur's book and that of Rescher and Urquhart both cover the model theory for modal logics in detail. The work of Kripke (1963) is perhaps the most important contribution to models for temporal (and Intuitionistic) logics: Rescher and Urquhart provide an accessible review of this and compare and contrast it with alternative models. The book by Rescher and Urquhart, along with the earlier work by Prior, is considered a classic text in the area. Turner's book is a more recent work covering a wider range of logics used in artificial intelligence, including some temporal logics.

van Benthem, JFAK, 1983. *The logic of time*, Dordrecht: D. Reidel

Galton, AP, 1987. "Temporal logic and computer science: An overview, *Temporal logics and their applications*, Galton, AP ed, pp 1–50, New York: Academic Press

McArthur, RP, 1976. *Tense logic*, Dorcrecht: D. Riegel

Prior, AN, 1967. *Past, present and future*, Oxford: Clarendon Press

Rescher, N and Urquhart, A, 1971. *Temporal logic*, New York: Springer-Verlag

Turner, R, 1984. *Logics for artificial intelligence*, Chichester: Ellis-Horwood

The next list includes all the references that are made from the text. The works by Allen, McCarthy and Hayes, McDermott and Kowalski and Sergot are particularly important papers in the field of temporal logics. Hayes' Naive Physics Manifesto (Hayes, 1978) contains useful comments on the relationship between logics and reasoning with intuitive models of the environment. It also provides further motivation for studies of temporal logics and, indeed, logics that support other constructs of the physicist's models of the world. The list contains several philosophical treatments of causation (Suppes, Lewis and Mackey, for example). This review has only skirted the issue of causation—an issue that is extremely important in many envisaged applications of temporal logics. Modelling causation, reasoning with it, and, more fundamentally, its relationship to time, are all essential problems to consider in the development and application of temporal logics. Shoham (1988) gives a useful treatment of causation, and (Allen, 1984) also contains some interesting ideas on modelling causation, particularly considering intentions of agents. The works of Prior and Quine are both worthy of particular attention—Quine's Elementary Logic is another classic work in the field of logics.

Allen, JF, 1984. "Towards a general theory of action and time", *Artificial Intelligence* **23**, 123–154

Barringer, H, Kuiper, R and Pnueli, A, 1984. "Now you may compose temporal logic specifications" *Proc. 16th ACM Symp. on Theory of Computing*, pp 51–63

Davidson, D, 1967. "The logical form of action sentences" *The logic of decision and action*, Recher, N, ed, pp 81–95, Pittsburgh: University of Pittsburgh Press

Dean, TL and McDermott, D, 1987. "Temporal data base management" *Artificial Intelligence* **32** pp 1–55

Eshgi, K, 1988. "Abductive planning with event calculus", Tech. Report, Dept. of Comp. Sci., Imperial College, London

Gabbay, D, 1986. "Executable temporal logic for interactive systems, Technical Report, Imperial College, London

Gabbay, D, 1987. "Modal and temporal logic programming" *Temporal logics and their applications*, Galton, A, ed, pp 197–239, New York: Academic Press

Galton, AP, 1984. *The logic of aspect*, Oxford: Oxford Clarendon Press

Galton, AP, 1986. "A critical examination of J. F. Allen's theory of time and action", Report 6.86, Centre for Theoretical Comp. Sci., Univ. of Leeds, England

Galton, AP, 1987. "The logic of occurrence" *Temporal logics and their applications*, Galton, AP, ed, pp 169–175, New York: Academic Press
Hale, R, 1987. "Temporal logic programming" *Temporal logics and their applications*, Galton, AP, ed, pp 91–119, New York: Academic Press
Hobbs, J, 1985. "Granularity" *IJCAI* **9**
Kenny, A, 1963. *Action, emotion and will*, London: Routledge and Kegan Paul
Kowalski, RA and Sergot, MJ, 1986. "A logic-based calculus of events" *New Generation Computing* **4** 67–95
Kripke, S, 1963. "Semantical considerations on modal logic" *Acta Philosophica Fennica* **16** 83–94
Lee, RM, Coelho, H and Cotta, JC, 1985. "Temporal inferencing on administrative databases" *Information Systems* **10** 197–206
Lewis, D, 1973. "Causation" *J. of Philosophy* **70** 556–567
Mackey, JL, 1974. *The cement of the universe: A study of causation, Oxford: Oxford University Press*
McCarthy, J and Hayes, PJ, 1969. "Some philosophical problems form the standpoint of artificial intelligence" *Machine intelligence*, Meltzer, B and Michie, D, eds, Vol. 4, Edinburgh: Edinburgh University Press
McDermott, DV, 1978. "Tarskian semantics or, no notation without denotation" *Cog. Sci.* **2**(3)
McDermott, DV, 1982. "A temporal logic for reasoning about processes and plans" *Cog. Sci.* **6** 101–155
Moszkowski, B, 1986. *Executing temporal logic programs*, Cambridge: Cambridge University Press
Newton-Smith, WH, 1980. *The structure of time*, London: Routledge and Kegan Paul
Prior, AN, 1955. "Diodoran modalities" *Philosophical Quarterly* **5** 205–213
Quine, WV, 1965. *Elementary logic*, New York: Harper and Row
Reichenbach, H, 1947. *Elements of symbolic logic*, New York: Macmillan
Russell, B, 1903. *Principles of mathematics*, London: George and Unwin
Russell, B, 1913. "On the notion of cause" *Proc. of the Aristotelian Society* **13** 1–26
Shanahan, M, 1988. "A single logical framework for prediction problems", Tech. Report, Dept. of Comp., Imperial College, London
Shoham, Y, 1985. "Ten requirements for a theory of change" *New Generating Computing* **3** 467–77
Shoham, Y, 1988. *Reasoning about change*, Massachusetts: MIT Press
Suppes, P, 1970. *A probabilistic theory of causation*, Amsterdam: North Holland
von Wright, GH, 1965. "And next" *Acta Philosophica Fennica* **18** 293–304
von Wright, GH, 1966. "And then" *Commentationes Physico-Mathematicae of the Finnish Society of Sciences* **32**(7)

The bibliography that follows includes a wide range of relevant material, covering issues in modelling time—the work by Allen, Hamblin, Shoham and Tichy in particular, causality (Kim, Piaget and Garcia) and other issues. A significant proportion of this work is devoted to tense logics (modal temporal logics) and is strongly linked with linguistic and also philosophical issues in temporal reasoning. Further implementations of temporal logic languages which are of interest are Tokio, presented in Aoyagi et al. and Fujita et al., the further work of Gabbay, Kahn and Gorry, and also Moszowski's work. Pnueli and Manna have produced some useful work in the application of temporal logics to reasoning about program correctness. Although the work by Newton-Smith (1980), listed amongst the references, is the most comprehensive treatment of topologies for temporal structures, the work by Hamblin considering intervals and instants is a detailed consideration of the specific question about the nature of temporal primitives, worthy of attention.

Allen, JF, 1981. "An interval based representation of temporal knowledge" *Proc. IJCAI* **7** 221–226
Allen, JF, 1983. "Maintaining knowledge about temporal intervals" *CACM* **26**(11) 832–843
Allen, JF and Koomen, JA, 1983. "Planning using a temporal world model" *Proc. IJCAI* **8** 741–747
Allen, JF and Hayes, PJ, 1985. "A common sense theory of time" *Proc. IJCAI* **9** 528–531
Aoyagi, T, Fujita, M and Mota-Oka, T, 1985. "Temporal logic programming language Tokio: programming in Tokio" *Logic Programming '85*, vol. 221, pp 128–137, LNCS, New York: Springer-Verlag
Ben-Ari, M, Pnueli, A and Manna, Z, 1981. "The temporal logic of branching time" *Proc. 8th ACM Symp. on Principles of Progrmming Languages*, pp 164–176
Bruce, BC, 1972. "A model for temporal references and its application in a question answering program" *Artificial Intelligence* **3** 1–25
Burgess, JP, 1982. "Axioms for tense logic ii: time periods", *Notre Dame J. of Formal Logic*, **23**(4) 375–383
Clarke, EM, Emerson, EA and Sistla, AP, 1986. "Automatic verification of finite-state concurrent systems using temporal logic specifications" *ACM Trans. on Programming Languages and Systems* **8** 244–263
Cresswell, MJ, 1977. "Interval semantics and logical words" *On the logical analysis of tense and aspect*, Rohrer, C, ed., Tubingen: Gunter Narr

Dean, TL, 1986. "Temporal imagery: An approach to reasoning about time for planning and problem solving" PhD Thesis, Yale University

Doyle, J, 1979. "A truth maintenance system", *Artificial Intelligence* **12** 231–272

Emerson, EA and Halpern, JY, 1983. "'Sometimes' and 'not never' revisited: on branching vs. linear time" *Proc. 10th ACM Symp. on Principles of Programming Languages*, pp 127–140

Farinas del Cerro, L, 1985. "Resolution modal logics" *Logics and Models of Concurrent Programs*, Apt, K, ed., pp. 27–56, New York: Springer-Verlag

Findlay, JN, 1941. "Time: a treatment of some puzzles" *Austalasian J. of Philosophy* **19** 216–235

Fujita, M et al., 1986. "Tokio: logic programming language based on temporal logic and its compilation to Prolog" *Proc. 3rd Int. Conf. on Logic Prog.* Vol. 225, pp 695–709, LNCS, New York: Springer-Verlag

Gabbay, D et al., 1980. "On the temporal analysis of fairness" *Proc. 7th ACM Symp on Principles of Programming Languages*, pp 163–173

Ginsberg, ML, 1986. "Counterfactuals" *Artificial Intelligence* **30**(1) 35–81

Ginsberg, ML and Smith, DE, 1986. "Reasoning about action I: A possible worlds approach", Tech. Report KSL-86-37, Stanford Knowledge Systems Lab

Hanks, S and McDermott, D, 1985. "Temporal reasoning and default logics", CS Research Report, Vol. 430, Yale University

Hanks, S and McDermott, D, 1986. "Default reasoning, nonmonotonic logics and the frame problem" *Proc. 5th Nat. Conf. on AI, AAAI*, Vol. 1, pp 328–333

Hamblin, CL, 1971. "Instants and intervals" **24** 127–134

Halpern, J, Manna, Z and Moszkowski, B, 1983. "A high level semantics based on interval logics" *Proc. ICALP* 278–291

Halpern, J and Shoham, Y, 1986. "A proposition modal logic of time intervals" *Proc. Symp. on Logic in CS*, Boston, Massachusetts: IEEE

Hayes, P, 1978. "The naive physics manifesto" *Expert systems in the microelectronic age*, Michie, D, ed., Edinburgh: Edinburgh University Press

Humberstone, IL, 1979. "Interval semantics for tense logic" *J. Philosophical Logic* **8** 171–196

Kahn, K and Gorry, GA, 1977. "Mechanizing temporal knowledge" *Artificial Intelligence* **9** 87–108

Kamp, JAW, 1968. "Tense logic and the theory of linear order", PhD Thesis, University of California

Kim, J, 1971. "Causes and events: Mackey on causation" *J. of Philosophy* **68** 426–441

Kowalski, RA, 1986. "Database updates in the event calculus" DoC, 86/12, Dept. of Comp., Imperial College, London

Lamport, L, 1980. "'Sometimes' is sometimes 'not never'': on the temporal logic of programs" *Proc. 7th ACM Symp. on Principles of Programming Languages*, pp 174–185

Massey, G, 1969. "Tense logic! Why bother?" *Nous* **3** 17–32

Moens, M and Steedman, M, 1987. "Temporal ontology in natural languages" *Proceedings of the 25th ACL*, pp 1–7

Moens, M and Steedman, M, 1989. "Temporal ontology and temporal reference" *Computational Linguistics* **14**

Moszkowski, B, 1983. "Reasoning about digital circuits" PhD thesis, Stanford University

Moszkowski, B, 1985. "A temporal logic for multi-level reasoning about hardware" *Computer* **18** 10–19

Piaget, J and Garcia, R, 1974. "Understanding causality" W. W. Morton

Pnueli, A, 1977. "The temporal logic of programs" *Proc. 18th IEEE Symp. on Foundations of Comp. Sci.* pp. 46–67

Reichgelt, H, 1987. "Semantics for reified temporal logic", Hallam, J and Mellish, C, eds, *Advances in artificial intelligence, Proceedings of the 1987 AISB Conference*, pp 49–61, Chichester: Wiley and Sons

Richards, B, 1982. "Tense, aspect and adverbials" *Linguistics and Philosophy* **5** 59–107

Robinson, JA, 1965. "A machine-oriented logic based on the resolution principle" *J. of ACM* **12** 23–41

Roper, P, 1980. "Intervals and tenses" *J. of Philosophical Logic* **9**

Sadri, F, 1987. "Three recent approaches to temporal reasoning" *Temporal logics and their applications*, Galton, A, ed., pp 121–167, New York: Academic Press

Schwartz, RL, Melliar-Smith, PM and Vogt, FH, 1983. "An interval logic for higher-level temporal reasoning", SRI Int., CS Lab.

Shoham, Y, 1986. "Chronological ignorance: Time, nonmonotonocity and necessity" *Proc. AAAI*, Philadelphia, PA

Shoham, Y, 1987. "Temporal logics in AI: Semantical and ontological considerations" *Artificial Intelligence* **33** 89–104

Taylor, B, 1985. *Modes of occurrence*, Oxford: Basil Blackwell

Tichy, P, 1985. "Do we need interval semantics?" *Linguistics and Philosophy* **8** 263–282

Williams, BC, 1986. "Doing time: Putting qualitative reasoning on firmer ground" *Proc. AAAI*

Wolper, P, 1981. 'Temporal logics can be more expressive" *Proc. 22nd IEEE Symp. on Foundations of Comp. Sci.*, pp 340–348