11

TEMPORAL LOGIC

Ian Hodkinson and Mark Reynolds

1 INTRODUCTION

Time has always been with us, though few of us have enough of it. The nature of time itself is a conundrum that we nowadays leave to physicists. But we have always had to find our way through time, plan our activities, and cope with the uncertain future. This can be, indeed, has to be done without a deep scientific knowledge of what makes time tick.

We use language and its rich tense structure to express and reason about events in time. This of course throws up linguistic and philosophical conundrums of its own. With the rise in the 20th century of formal logical languages, it became natural to try to express temporal concepts and arguments in formal terms, and so it was that Arthur Prior from the 1950s came to develop *tense logics.* These were modal logics, with box-modalities H and G for 'always in the past' and 'always in the future', motivated by tenses in natural language. The advent of Kripke semantics in the 1960s gave the enterprise a boost, because a Kripke frame is so naturally seen as a set of time points endowed with an 'earlier-later' relation.

Temporal logic today is a large, busy subject with stakeholders from many disciplines. Philosophers and linguists have continued to make major contributions to it. Since Pnueli's pioneering 1977 paper [147], several branches of computer science and related fields — such as databases, specification and verification, synthesis of programs, temporal planning, temporal knowledge representation — have had a huge influence, and the use of temporal logic in some of these areas has developed to the point of commercial application. There is even some contact with physics, but so far this has been limited.

Temporal logic is in a way a branch of applied modal logic, but modal logicians may be disconcerted by what they find here. Temporal logic has always focused on handling time, it has developed whatever methods it found useful for this end, and not all of them are modal in a narrow sense. Connectives such as Until and Since, again mimicking the natural language constructs, go beyond boxes and diamonds and are of great importance in the subject. Indeed, completely general first-order-definable connectives are used as well. Bearing in mind the evaluation and reference points of natural language, it is natural that many-dimensional evaluation has long been of importance in temporal logic, whereas it only recently attracted great interest in modal logic proper. The focus in temporal logic is on a fairly narrow range of Kripke frames — nearly always irreflexive and transitive, and typically linear orders or trees, though relativistic and circular time are sometimes considered. The natural numbers are the dominant model of linear time, though dense and continuous and indeed arbitrary linear orders have found their way in (and in this chapter we are happy to consider them). Sometimes the pressures of time have led to a style of evaluation of formulas that seems non-modal at first sight (see Section 3.7). A very influential strand of work, started by Kamp in 1968, compares the expressive power of modal and first-order languages on the model (rather than frame) level. Rather than be content with limited but well-behaved modal expressiveness, the thrust of the work created temporal languages as strong as classical first-order logic and even monadic second-order logic. Perhaps because the proofs rely heavily on the assumption that time is linear or even natural number-like, not much similar research in modal logic has been done. Classical logic is not just a benchmark for the expressiveness of 'real' temporal logics: using first-order logic for handling time is itself a respectable tradition. In temporal logic there is an unusual (for modal logic) use of methods from

classical mathematical logic and combinatorial techniques such as automata. Problems such as model-checking (considered in Chapter 17) are all-important in temporal logic but do not appear much in modal logic.

Nonetheless, from Prior's work onwards, modal ideas have been prominent in temporal logic. Its most basic syntax and Kripke semantics are (multi-)modal; one often comes across modal techniques such as canonicity and Sahlqvist's theorem, filtration and non-standard inference rules; and problems of axiomatisation, decidability, and complexity are ubiquitous in modal and temporal logic. Sophisticated results on modal logics above K4 have been transferred to temporal logic. Chapters 2, 3, 4, 9, and 12 are very relevant to temporal logic. In Chapter 17, the reader will find a concentrated discussion of modal and temporal logic in computer science. In the current chapter, we will examine some topics in temporal logic that are considered both in computer science and in other fields. As we have not the space to provide a rigorous development from scratch, the chapter is intended more as a gateway to the subject. It is mostly a survey-style commentary on some important strands, with directions to the literature for those wishing to find out more. Our priority is range rather than depth, but we cannot be comprehensive. A chapter of definitions would be indigestible, so we have tried to include some of the arguments, but space limitations have meant that their level of detail veers wildly from a few words to (occasionally) something approaching a full proof. Readers may of course skip details if they so desire.

We start out in Section 2 with a basic round-up of the semantic options for handling time. In Section 3 we cover some of the logics (syntax and evaluation) that can be used. Bearing in mind the remarks above, it will be no surprise that we do not confine ourselves to modal-style logics: first- and second-order logics, and others, find their way in, and our lack of consideration of mu-calculi is only because chapter 12 is devoted to them. Cn Section 4 we compare the expressivity of classical and modal-style logics. Kamp's famous 1968 expressive completeness theorem makes its appearance here. In Section 5 we discuss temporal reasoning, mainly avoiding automata (see Chapter 17 for them) but covering Hilbert systems, tableaux, resolution, filtration and the finite model property, and other methods.

A word about first-order temporal logic. This is a complex issue. There is a confusing variety of ways to add first-order logic to a temporal system, and undecidability results obtained in the 1960s, accompanied by later expressive incompleteness results, also cast their shade over the development of this part of the subject. But at the time of writing, there is something of a resurgence of interest in it from the database and reasoning communities. We will discuss the rudiments of first-order temporal logic in Sections refchapter11:sec2–3, and also some of the recent results on expressive completeness and decidability in Sections refchapter11:sec4–5. Chapter 9 is also relevant of course.

Temporal logic, then, is a branch of applied logic that brings to bear a gamut of powerful methods from many fields to study time and temporal phenomena. It is not wholly modal, but rests on a modal base — it is a meeting ground for concepts from modal logic, classical first-order logic, and higher-order logic. It has found very successful application in computing, and embodies seminal contributions from philosophy and linguistics as well. We hope our chapter, and other chapters here, will serve as a guide for the reader wishing to discover more about this intensely active field.