Hybrid Languages and Temporal Logic

Patrick Blackburn
Miroslava Tzakova

MPI–I–98–2-006 February 1998

Authors' Addresses

Patrick Blackburn
Computerlinguistik, Universität des Saarlandes
66041 Saarbrücken, Germany
patrick@coli.uni-sb.de

Miroslava Tzakova
Max-Planck-Institut für Informatik
Im Stadtwald, 66123 Saarbrücken, Germany
tzakova@mpi-sb.mpg.de

Publication Notes

This paper has been submitted for publication elsewhere and will be copyrighted, if accepted.

1 Introduction

Arthur Prior proposed using modal languages for temporal reasoning more than 40 years ago, and since then the approach has become widespread in a variety of disciplines. Over this period, a wide range of (often very powerful) modalities has been used to reason about time. This is unsurprising. After all, different choices of temporal ontology (such as instants, intervals, and events) are relevant for different purposes, and (depending on the application) considerable expressive power may be needed to cope with the way information can be distributed across such structures. But inventing new modalities is not the only way of boosting modal expressivity. There is a largely overlooked alternative called *hybridisation*, and this paper explores its relevance for temporal logic.[1]

Hybridisation is best introduced by example. Consider the following sentence from the language we call $\mathrm{ML} + \forall$:

$$\forall x(x \rightarrow \neg\Diamond x).$$

The x in this expression is a *state variable*, and all its occurrences are bound by the *binder* $\forall$. Syntactically, state variables are *formulas*: after all, the expression $x \rightarrow \neg\Diamond x$ is built using $\rightarrow$, $\neg$ and $\Diamond$ in the same way that $p \rightarrow \neg\Diamond p$ is. Semantically, however, state variables are best thought of as *terms*. Our semantics will stipulate

[1] The literature on hybrid languages consists of a handful of papers published over the last thirty years by researchers with very different interests. Confining ourselves to the main line of development, the idea can be traced back Prior (1967), and the posthumously published Prior and Fine (1977) contains some of Prior's unfinished papers on the subject together with an appendix by Kit Fine. Prior's concerns were largely philosophical; technical development seems to have started with Bull (1970). Bull investigated a very strong hybrid temporal language, and introduced a number of key ideas such as the $\forall$ binder, the use of the universal modality A, and direct quantification over paths. In addition, he initiated the algebraic study of such systems. The paper never attracted the attention it deserved; in fact, apart from citations in the hybrid literature, the only mention we know of is the following passage from Burgess's survey of tense logic:

> *Other hybrids of a different sort — not easy to describe briefly — are treated in an interesting paper of Bull [1970].* (Burgess (1984, page 128)).

(This is probably the first use of 'hybrid' in connection with such languages.) The idea was independently invented by the Sofia School as a spin-off of their investigation of modal logic with names. The best guide to the Bulgarian tradition is the persuasively argued Passy and Tinchev (1991), drafts of which were in circulation as early as 1985. The discussion of hybridisation is confined to one section, and deals with Propositional Dynamic Logic enriched with both $\forall$ and the universal modality. Recent papers on the subject include Goranko (1994) (probably the first published account of hybrid languages containing the $\Downarrow^0$ binder), Blackburn and Seligman (1995), and Seligman (1997) (which investigates hybrid natural deduction and sequent calculi for applications in Situation Theory).

that state variables are satisfied at exactly one state in any model. In effect, state variables act as names; they 'label' the unique state they are true at.

The use of 'formulas as terms' gives hybrid languages their unique flavour: they are formalisms which blend the operator based perspective of modal logic with the classical idea of explicitly binding variables to states. Unsurprisingly, this combination offers increased expressive power. The above sentence, for example, is true at any irreflexive state in any model, and false at all reflexive ones. No ordinary modal formula has this property.

Now, the language $\mathrm{ML} + \forall$ is not the only hybrid language, and for many purposes it is not the most natural one. Arthur Prior laid great stress on the *locality* of modal logic. It should be intuitively clear (we shall be precise later) that $\forall$ is not local; as our notation suggests, $\forall$ quantifies across *all* states. So, if we want a local hybrid language, $\mathrm{ML} + \forall$ is not a suitable choice.

But what are the alternatives? To the best of our knowledge only one has been previously considered, namely the binder we here call $\Downarrow^0$. Now, $\Downarrow^0$ does something simple and natural: it binds a variable to the current state. Unfortunately, while $\mathrm{ML}+\Downarrow^0$ is a local language, it has two drawbacks. First, it is not expressive enough for many applications (for example, we shall show that it is not strong enough to define the *Until* operator). Second, and in stark contrast to $\mathrm{ML} + \forall$ which has an elegant axiomatisation, the only known axiomatisation for $\mathrm{ML} + \Downarrow^0$ makes use of rather complex rules of proof (both axiomatisations are introduced and proved complete in Blackburn and Tzakova (1997)).

To overcome these difficulties we here introduce a second local binder called $\Downarrow^1$. The $\Downarrow^1$ binder is a universal quantifier over *accessible* states and, as we shall show, $\mathrm{ML} + \Downarrow^0 + \Downarrow^1$ is a well-behaved temporal language in two important respects: it is strong enough to define *Until* and its minimal temporal logic (that is, the set of formulas valid on strict partial orders) has a straightforward axiomatisation. Admittedly, $\Downarrow^1$ does come at a price. While the minimal temporal logic of $\mathrm{ML}+\Downarrow^0$ is decidable (see Blackburn and Seligman 1996), we shall show that the minimal temporal logic of $\mathrm{ML} + \Downarrow^0 + \Downarrow^1$ is not.

Much of the paper is devoted to exploring $\mathrm{ML} + \Downarrow^0 + \Downarrow^1$, and the reader who works through the details will emerge with a solid grasp of how hybrid languages work and what they have to offer. So, to round off the paper, we broaden our discussion by posing the following question: why should we rest content with binding variables to *states*? In particular, from the point of view of temporal logic it would be nice to be able to bind variables to *paths*. Accordingly, we introduce a binder $\Downarrow^\pi$ which allows us to do this, and briefly examine the (interesting but difficult) territory to which this extension leads.

But we are jumping too far ahead. There is much to be done before we can usefully discuss $\Downarrow^\pi$, so let us call a halt to our introductory remarks, and start developing the idea of hybridisation systematically.

2 The standard modal language

One of the simplest languages for temporal reasoning is the propositional modal language that contains just two modalities: an operator $\Box$ (read as: *at all future states*) together with its dual operator $\Diamond$ (read as: *at some future state*). For most of this paper we will be working with various hybrid extension of this simple language (which we will call ML). The purpose of the present section is to fix notation and terminology, to remind the reader of various standard concepts (in particular, *generated submodels* and *bisimulations*), and to present a wish-list of properties for hybrid temporal languages.

Given a (countable) set of *propositional symbols* PROP $= \{p, q, r, \ldots\}$ the well-formed formulas of ML are defined as follows:

$$\text{WFF} := p \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box\varphi.$$

Other Boolean operators ($\vee$, $\rightarrow$, $\leftrightarrow$, $\bot$, $\top$, and so on) are defined in the usual way, and we define $\Diamond\varphi$ to be $\neg\Box\neg\varphi$.

ML is interpreted on *models*. A model $\mathcal{M}$ is a triple (S, R, V) such that S is a non-empty set of *states*, R is a binary relation on S (the *temporal precedence relation*), and $V : \text{PROP} \longrightarrow Pow(S)$, is the *valuation*, which tells us which states (if any) each propositional symbol is true at. Two remarks are in order. We call the elements of S 'states' because this seems a reasonably neutral term that covers a wide range of temporal entities. The reader is free to think of 'states' as instants of time, as intervals, or as time-stamped events; for the purposes of the present paper, the choice between these interpretations is largely irrelevant. What *is* important, however, is the choice of binary relation R. As our interest is temporal logic, we will only be interested in *strictly partially ordered models*; that is, those models where R is both transitive and irreflexive.

The satisfaction definition for ML is defined as follows. Let $\mathcal{M} = (S, R, V)$ and $s \in S$. Then:

$$\begin{array}{lll}
\mathcal{M}, s \models p & \text{iff} & s \in V(p), \text{ where } p \in \text{PROP} \\
\mathcal{M}, s \models \neg\varphi & \text{iff} & \mathcal{M}, s \not\models \varphi \\
\mathcal{M}, s \models \varphi \wedge \psi & \text{iff} & \mathcal{M}, s \models \varphi \ \& \ \mathcal{M}, s \models \psi \\
\mathcal{M}, s \models \Box\varphi & \text{iff} & \forall s'(sRs' \ \Rightarrow \ \mathcal{M}, s' \models \varphi).
\end{array}$$

If $\mathcal{M}, s \models \varphi$ we say that φ is *satisfied* in $\mathcal{M}$ at s. If φ is true at all states in a model $\mathcal{M}$ we say it is *valid* in $\mathcal{M}$ and write $\mathcal{M} \models \varphi$. In what follows, we will be particularly interested in those formulas that are valid on all strictly partially ordered model; we regard this set of formulas as the *minimal temporal logic*.

Perhaps the key intuition to note about the satisfaction definition is its locality: formulas are evaluated *inside* models at some particular state (called the *current* *state*), and the function of the $\Box$ and $\Diamond$ operators is to scan the states accessible from the current state via the precedence relation R. This locality intuition is arguably *the* central intuition underlying the modal approaches to temporal logic; it is certainly the intuition which prompted Arthur Prior to pioneer the "modal logic of time" (which he called *tense logic*). As he observed, we are situated *inside* the temporal flow, and our temporal intuitions and language reflect this internal perspective. Accordingly, it seemed to him that modal approaches to temporal reasoning were likely to be the most revealing.[2]

The locality of ML has an obvious mathematical consequence: satisfaction of ML formulas is *preserved under the formation of generated submodels*. Given a model $\mathcal{M} = (S, R, V)$ and a state s of S, the submodel of $\mathcal{M}$ generated by s is the smallest submodel $\mathcal{M}^s$ of $\mathcal{M}$ that contains s and is R-closed. (That is, the submodel generated by s contains just those states of $\mathcal{M}$ that are accessible from s by a finite number of transitions along R. Note that in strict partial orders, every state in the generated submodel distinct from s is accessible from s in one step.) It follows by an easy induction that for all formulas φ:

$$\mathcal{M}, s \models \varphi \text{ iff } \mathcal{M}^s, s \models \varphi.$$

In the work that follows, we shall use preservation under generated submodels as one of our criterion for judging hybrid languages. We are interested in developing *local* hybrid languages, and will reject those which lead to a loss of the preservation result for generated submodel.

Now for a key question: does ML have the expressivity needed for temporal reasoning? There is no absolute answer: it depends on the application. For some applications, ML will often be too *strong*. For example, if one is interested in pursuing the line of work initiated by Hennessy and Milner, in which modal languages are used to characterise various types of bisimulation invariance, it may be necessary to work with sublanguages of ML containing no propositional symbols (wffs would built using the constant $\perp$) or to shed some Boolean expressivity. But for some other applications, ML is too *weak*. For a start, as has already been mentioned, no formula of ML is capable of distinguishing irreflexive from reflexive states in all models; this means that one of our fundamental constraints on temporal precedence simply isn't reflected in this language. Moreover, recall the definition of the *Until* operator:

$$\mathcal{M}, s \models \mathit{Until}(\varphi, \psi) \text{ iff } \exists s'(sRs' \;\&\; \mathcal{M}, s' \models \varphi \;\&\; \forall t(sRt \;\&\; tRs' \;\Rightarrow\; \mathcal{M}, t \models \psi))$$

This is an extremely natural local operator (note that formulas built using *Until* are preserved under the formation of generated submodels) and it has proved a

[2]Probably the best introduction to Prior's views is Prior (1967).

valuable tool for temporal reasoning in computer science (indeed, computer scientists usually treat *Until* as the fundamental temporal modality). However the *Until* operator is *not* definable in ML. As both the non-definability of *Until* and irreflexivity follows from the fact that ML formulas are preserved under *bisimulations*, and as we shall later have to make use of a special form of bisimulation called *quasi-injective bisimulations*, it will be useful to prove these non-definability results here.

A *bisimulation* between two models $\mathcal{M}_1 = (S_1, R_1, V_1)$ and $\mathcal{M}_2 = (S_2, R_2, V_2)$ is a non-empty binary relation Z between S_1 and S_2 such that:

1. For all states s_1 in S_1 and s_2 in S_2, if $s_1 Z s_2$ then s_1 and s_2 satisfy the same propositional symbols.

2. For all states s_1, s_1' in S_1 and s_2 in S_2, if $s_1 R_1 s_1'$ and $s_1 Z s_2$ then there is a state s_2' in S_2 such that $s_2 R_2 s_2'$ and $s_1' Z s_2'$.

3. For all state s_2, s_2' in S_2, and s_1 in S_1, if $s_2 R_2 s_2'$ and $s_1 Z s_2$ then there is a state s_1' in S_1 such that $s_1 R_1 s_1'$ and $s_1' Z s_2'$.

The fundamental result concerning bisimulations (which follows straightforwardly by induction on the structure of ML formulas) is that if Z is a bisimulation between models $\mathcal{M}_1$ and $\mathcal{M}_2$ and $s_1 Z s_2$ then s_1 and s_2 satisfy exactly the same ML formulas.

It follows that neither *Until* nor irreflexivity is definable — indeed the following counterexample (which we believe is due to Johan van Benthem) establishes both points simultaneously. Let $\mathcal{M}_1$ be an irreflexive model containing just two states s_1 and s_2, let $s_1 R s_2$ and $s_2 R s_1$, and suppose all propositional symbols are true at both states. Let $\mathcal{M}_2$ be an reflexive model containing just one state s, and suppose all propositional symbols are true at s. Clearly the relation Z which links both s_1 and s_2 to s and vice-versa is a bisimulation, hence all states in both models satisfy exactly the same ML formulas. So, as $\mathcal{M}_1$ is irreflexive and $\mathcal{M}_2$ reflexive, it follows that no ML formula succeeds in distinguishing irreflexive and reflexive states. Moreover, observe that *Until*$(\top, \bot)$ is false in $\mathcal{M}_1$ (at both s_1 and s_2) but true in $\mathcal{M}_2$. It follows that the *Until* operator cannot be expressed in ML.

Thus, ML has expressive weaknesses that are relevant to temporal reasoning, and one of the key goals of this paper will be to repair them by hybridisation. But what should a hybrid temporal language look like? It is time to draw up a wish-list.

First, we would like our hybrid language to be *local*. Second, we would like our hybrid language to be expressive enough to detect irreflexivity and define *Until*. Third, we are interested in finding hybrid languages in which the central ideas of modal and classical proof systems can be clearly combined. (Indeed, ultimately we hope to find hybrid languages that exhibit a certain 'synergy' between modal and classical ideas; where the whole, so to speak, offers more than the sum of its parts. It remains to be seen how realistic this hope is.)

These goals are not easy to reconcile. As has already been mentioned, the hybrid language $\mathrm{ML}+\forall$ does not meet the locality demand, while $\mathrm{ML}+\Downarrow^0$ cannot define *Until* and seems complex to axiomatise. Much of the work that follows will be devoted to showing that introducing a new local binder $\Downarrow^1$ is a reasonably neat way of resolving these tensions.

3 Three hybrid binders

We now hybridise ML. Syntactically, this involves making two changes. First, we *sort* the atomic symbols of ML. Instead of having just one kind of atom (namely the symbols in PROP) we add a second kind which we call *state symbols*. For reasons we shall soon explain, it is convenient to divide state symbols into two subcategories: *state variables* and *nominals*. Second, we add *binders*. The binders will be used to bind state variables (but not nominals or propositional symbols). For most of this paper we shall work with the binders $\forall$, $\Downarrow^0$, $\Downarrow^1$ and their duals.

Let PROP be as described before. Assume we have denumerably infinite set SVAR of state variables (whose elements we typically write as u, v, w, x, y and z), and a denumerably infinite set NOM of nominals (whose elements we typically write as c, d, e and k). We assume that PROP, SVAR and NOM are pairwise disjoint. We call SVAR$\cup$NOM the set of *state symbols*, and PROP$\cup$SVAR$\cup$NOM the set of *atoms*. Further, choose $\mathcal{B}$ to be one of $\{\forall\}$, $\{\Downarrow^0\}$, $\{\Downarrow^1\}$, or $\{\Downarrow^0, \Downarrow^1\}$. We build the well-formed formulas of the hybrid language (over PROP, SVAR, NOM, and $\mathcal{B}$) as follows:

$$\mathrm{WFF}\ \varphi \ := \ a \mid \neg\varphi \mid \varphi \wedge \psi \mid \Box\varphi \mid Bx\varphi.$$

Here $a \in$ ATOM, $x \in$ SVAR, and $B \in \mathcal{B}$. The different choices of $\mathcal{B}$ give rise to the languages $\mathrm{ML}+\forall$, $\mathrm{ML}+\Downarrow^0$, $\mathrm{ML}+\Downarrow^1$, and $\mathrm{ML}+\Downarrow^0+\Downarrow^1$ respectively. (Strictly speaking, different choices of PROP, SVAR and NOM give rise to different hybrid languages, but we ignore this whenever possible.) Incidentally, while we could choose $\mathcal{B}$ to be $\{\forall, \Downarrow^0, \Downarrow^1\}$, $\{\forall, \Downarrow^0\}$, or $\{\forall, \Downarrow^1\}$, we don't bother as the resulting languages are expressively equivalent to $\mathrm{ML}+\forall$.

A full discussion of the syntax of these languages would need to define such concepts as 'free', 'bound', 'substitutible for', and so on. However, as experience with classical logic is a reliable guide, and as the relevant definitions may be found in Blackburn and Tzakova (1997), we'll simply remark that a *sentence* is a formula containing no free variables, and that we use the notation $\varphi[\mathbf{s}/\mathbf{v}]$ to denote the formula obtained by substituting the state symbol **s** for all free occurrences of the state variable **v** in φ.

As promised in the introduction, our hybrid languages use formulas as terms: in the semantics presented below, both state variables and nominals will be satisfied at exactly one state in any model. Now, the role of the state variables should be clear; but what is the point of having nominals? Simply this: it is convenient to have a supply of 'labels' that cannot be bound by the binders; this simplifies some of the technicalities, for is saves us having to worry about accidental binding. In short, nominals are reminiscent of the 'parameters' used in classical proof theory.

Now for the semantics. The key idea is straightforward: we are going to insist that state symbols are interpreted by singleton subsets of models. We'll also need a smooth way to handle the fact that state variables may become bound, whereas this is not possible for nominals or propositional symbols. But there is an obvious way to do this: we'll let the state variables be handled by a separate assignment function in the manner familiar from classical logic.

Definition 1 (Standard models and assignments) *Let $\mathcal{L}$ be a hybrid language over* PROP, SVAR *and* NOM. *A model $\mathcal{M}$ for $\mathcal{L}$ is a triple (S, R, V) such that S is a non-empty set, R a binary relation on S, and $V : \mathrm{PROP} \cup \mathrm{NOM} \longrightarrow Pow(S)$. A model is called* standard *iff for all nominals $c \in \mathrm{NOM}$, $V(c)$ is a singleton subset of S.*

An assignment for $\mathcal{L}$ on $\mathcal{M}$ is a mapping $g : \mathrm{SVAR} \longrightarrow Pow(S)$. An assignment is called standard *iff for all state variables $x \in \mathrm{SVAR}$, $g(x)$ is a singleton subset of S. The notation $g' \overset{x}{\sim} g$ (g' is an x-variant of g) means that g' and g are standard assignments (on some model $\mathcal{M}$) such that g' agrees with g on all arguments save possibly x.*

Let $\mathcal{M} = (S, R, V)$ be a standard model, and g a standard assignment. For any atom a, let $[V, g](a) = g(a)$ if a is a state variable, and $V(a)$ otherwise. Then interpretation of our hybrid languages is carried out using the following definition:

$$
\begin{array}{lll}
\mathcal{M}, g, s \models a & \text{iff} & s \in [V, g](a), \text{ where } a \in \mathrm{ATOM} \\
\mathcal{M}, g, s \models \neg\varphi & \text{iff} & M, g, s \not\models \varphi \\
\mathcal{M}, g, s \models \varphi \wedge \psi & \text{iff} & \mathcal{M}, g, s \models \varphi \;\&\; \mathcal{M}, g, s \models \psi \\
\mathcal{M}, g, s \models \Box\varphi & \text{iff} & \forall s'(sRs' \;\Rightarrow\; \mathcal{M}, g, s' \models \varphi). \\
\mathcal{M}, g, s \models \forall x\varphi & \text{iff} & \forall g'(g' \overset{x}{\sim} g \;\Rightarrow\; \mathcal{M}, g', s \models \varphi) \\
\mathcal{M}, g, s \models \Downarrow^0_x\varphi & \text{iff} & \mathcal{M}, g', s \models \varphi, \text{ for all } g' \overset{x}{\sim} g \text{ such that } g'(x) = \{s\} \\
\mathcal{M}, g, s \models \Downarrow^1_x\varphi & \text{iff} & \mathcal{M}, g', s \models \varphi, \text{ for all } g' \overset{x}{\sim} g \text{ such that } sRg'(x)
\end{array}
$$

Let $\mathcal{M}$ be a standard model. We say that φ is *valid* on $\mathcal{M}$ iff for all standard assignments g on $\mathcal{M}$, and all states s in $\mathcal{M}$, $\mathcal{M}, g, s \models \varphi$, and if this is the case we write $\mathcal{M} \models \varphi$. In what follows we will be interested in those formulas valid on all standard strictly partially ordered models; for any of our hybrid languages, we call this set of validities the *minimal temporal logic* in that language.

Lemma 2 (Substitution lemma) *Let $\mathcal{M}$ be a standard model, let g be an assignment on $\mathcal{M}$, and let φ be a formula of any of the hybrid languages defined above. Then, for every state s in $\mathcal{M}$, if y is a variable that is substitutable for x in φ and c is a nominal then:*

1. $\mathcal{M}, g, s \models \varphi[y/x]$ *iff* $\mathcal{M}, g', s \models \varphi$, *where* $g' \stackrel{x}{\sim} g$ *and* $g'(x) = g(y)$.
2. $\mathcal{M}, g, s \models \varphi[c/x]$ *iff* $\mathcal{M}, g', s \models \varphi$, *where* $g' \stackrel{x}{\sim} g$ *and* $g'(x) = V(c)$.

Proof. By induction on the complexity of φ. $\dashv$

This concludes the preliminaries. It is time to take a closer look at the various binders.

The binders $\forall$ and $\exists$

The $\forall$ binder is the strongest, most classical, of our binders. Indeed it's just the familiar universal quantifier in a modal setting. Note that it we define $\exists x\varphi$ to be the dual binder $\neg\forall x\neg\varphi$, then:

$$\mathcal{M}, g, s \models \exists x\varphi \text{ iff } \exists g'(g' \stackrel{x}{\sim} g \ \&\ \mathcal{M}, g', s \models \varphi).$$

ML + $\forall$ is a powerful language. As was pointed out in the introduction, it can distinguish irreflexive from reflexive states. Moreover it can define the *Until* operator:

$$\textit{Until}(\varphi, \psi) \ := \ \exists x \Diamond((x \wedge \varphi) \wedge \Box(\Diamond x \rightarrow \psi)).$$

This definition says: it is possible to bind the variable x to a successor state in such a way that (1) φ holds at the state labeled x, and (2) ψ holds at all successors of the current state that precede this x-labeled state.

In addition, the minimal temporal logic of ML + $\forall$ has a simple axiomatisation that can be proved complete reasonably straightforwardly. We won't present the full axiomatisation here, but for comparison with our later work on ML + $\Downarrow^0$ + $\Downarrow^1$ it will be useful to note two of its components.

First, ML + $\forall$ validates two of the schemas standardly used to axiomatise first-order logic ($\mathbf{v}$ and $\mathbf{s}$ are used as metavariables over state variables and state symbols respectively):

$Q1 \quad \forall \mathbf{v}(\varphi \rightarrow \psi) \rightarrow (\varphi \rightarrow \forall \mathbf{v}\psi)$

$Q2 \quad \forall \mathbf{v}\varphi \rightarrow \varphi[\mathbf{s}/\mathbf{v}]$

(In *Q1*, φ cannot contain free occurrences of $\mathbf{v}$; and in *Q2*, $\mathbf{s}$ must be substitutable for $\mathbf{v}$ in φ.) As $\mathrm{ML} + \forall$ also validates the classical rule of generalisation (if φ is provable then so is $\forall x\varphi$) it is clear that it has a complete classical core; this makes much of the completeness proof straightforward.

Second, $\mathrm{ML} + \forall$ also validates the Barcan schema:

Barcan $\qquad \forall \mathbf{v}\Box\varphi \rightarrow \Box\forall \mathbf{v}\varphi$

Now, the *Barcan* schema is familiar from first-order modal logic (that is, the study of languages containing quantifiers that range over the individuals in some underlying collection of first-order models). In first-order modal logic, *Barcan* is a convenient (if rather dubious) principle. Its status is $\mathrm{ML} + \forall$, however, is beyond dispute: all instances are valid in all standard models as the reader can easily check.

Barcan is important because it builds a robust bridge between the modalities and the binders. As it allows us to permute the relative scopes of $\Box$ and $\forall$, it is straightforward to combine the key techniques of modal and classical completeness proofs. From an axiomatic perspective, the presence of *Barcan* is what most sharply distinguishes $\mathrm{ML} + \forall$ from the weaker local languages.

All in all, $\mathrm{ML} + \forall$ is a lovely language. There's just one problem: it isn't local. To see that satisfaction need not be preserved under the formation of generated submodels, consider the following counterexample (taken from Blackburn and Seligman (1995)). Let $\mathcal{M}$ be the following two-element model where $S = \{s, t\}$, and $R = \{(s, s)\}$:

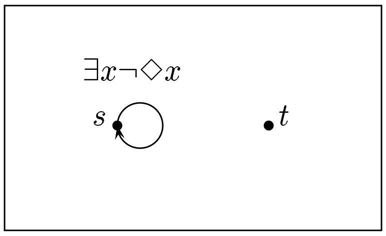

Then $\exists x \neg \Diamond x$ is true at s in $\mathcal{M}$, for we can assign the state t to x and $(s, t) \notin R$. However it is *not* true at s in the submodel $\mathcal{M}^s$ generated by s, for as $\mathcal{M}^s$ contains only the state s, all assignments assign s to x. As s is reflexive, $\neg\Diamond x$ will always be false.

If you want a strong hybrid language and are not interested in maintaining locality, then $\mathrm{ML}{+}\forall$ is probably an excellent choice. Indeed, you may wish consider working with a hybrid language even *less* local, namely ML + $\forall$ enriched with the *universal modality* A.[3] The universal modality has the following satisfaction definition: $\mathcal{M}, s \models A\varphi$ iff $\mathcal{M}, s' \models \varphi$ for all states $s' \in \mathcal{M}$. It is not hard to see that adding the universal modality yields a hybrid language with first-order expressive power (see Blackburn and Seligman (1995) for details). Moreover, the A and $\forall$ work extremely well together and make elegant axiomatisations possible. But while such rich systems are interesting, they are far removed from the local languages we wish to explore.

The binders $\Downarrow^0$ and $\downarrow^0$

If one is interested in local hybrid languages, the $\Downarrow^0$ binder is the most natural starting point. Quite simply, $\Downarrow^0$ binds a variable to the current state; it creates a label for the here-and-now. Let's look at it more closely.

First, observe that if we define $\downarrow^0_x\varphi$ to be the dual binder $\neg\Downarrow^0_x\neg\varphi$, we obtain a binder with the following satisfaction definition:

$$\mathcal{M}, g, s \models \downarrow^0_x\varphi \quad \text{iff} \quad \mathcal{M}, g', s \models \varphi, \text{ where } g' \overset{x}{\sim} g \text{ and } g'(x) = \{s\}$$

However, as there is exactly one current state, it is clear that the $\downarrow^0$ binder is nothing but $\Downarrow^0$ binder in disguise — that is, $\Downarrow^0$ is self-dual. In what follows we regard $\Downarrow^0$ as the official notation for this "label-the-here-and-now" binder, but, we sometimes make use of the $\downarrow^0$ notation as well (this will help us emphasise certain similarities with $\downarrow^1$, the dual of the $\Downarrow^1$ binder).

Next, note that $\Downarrow^0_x\varphi$ is definable in ML+$\forall$. We can define it either as $\forall x(x \to \varphi)$ or $\exists x(x \wedge \varphi)$. Thus ML + $\Downarrow^0$ is a fragment of ML + $\forall$, and in many ways it is a well-behaved fragment. For a start, sentences of ML + $\Downarrow^0$ *are* preserved under the formation of generated submodels. (We leave the very simple proof of this fact to the reader. Essentially it just boils down to the observation that the only states that $\Downarrow^0$ can bind to variables in the course of evaluation must be states in the generated submodel.) Moreover, adding the $\Downarrow^0$ binder boosts the expressive power of ML in temporally interesting ways. In particular, note that the sentence

$$\Downarrow^0_x\Box\neg x$$

is true in a model at a state s iff s is irreflexive. Finally, the minimal temporal logic of ML + $\Downarrow^0$ is decidable (this can be proved by a Rabin-Gabbay style argument; see Blackburn and Seligman (1997) for details).

Unfortunately, ML + $\Downarrow^0$ has two drawbacks. First, for many purposes it simply isn't expressive enough. Second, the only known way of obtaining completeness

[3] Virtually the entire literature on hybrid languages is devoted to such systems. For example, both Bull (1970) and Passy and Tinchev (1991) make use of both $\forall$ and A.

results for this language relies on the use of rather complex rules of proof. Let us examine both issues more closely.

Although adding $\Downarrow^0$ increases the expressive power, *Until* still isn't definable. To see why, we will have to make use of the concept of *quasi-injective bisimulations* introduced in Blackburn and Seligman (1997). Let us say that states s and s' in a model $\mathcal{M} = (S, R, V)$ are *mutually inaccessible* iff s is not in the submodel generated by s' and s' is not in the submodel generated by s. We then define:

Definition 3 (Quasi-injective bisimulations) *Let Z be a bisimulation between $\mathcal{M}_1$ and $\mathcal{M}_2$. Then Z is a* quasi-injective bisimulation *iff:*

1. *For all states s_1, $s_1{}'$ in $\mathcal{M}_1$, and s_2 in $\mathcal{M}_2$, if $s_1 Z s_2$ and $s_1{}' Z s_2$, and $s_1 \neq s_1{}'$ then s_1 and $s_1{}'$ are mutually inaccessible, and*

2. *For all states s_2, $s_2{}'$ in $\mathcal{M}_2$, and s_1 in $\mathcal{M}_1$, if $s_1 Z s_2$ and $s_1 Z s_2{}'$, and $s_2 \neq s_2{}'$ then s_2 and $s_2{}'$ are mutually inaccessible.*

Now, $\mathrm{ML} + \Downarrow^0$ sentences need not be preserved under arbitrary bisimulations (the fact that the sentence $\Downarrow^0_x \Box \neg x$ picks out the irreflexive states shows this), but Blackburn and Seligman show that they *are* preserved under quasi-injective bisimulations. That is:

Proposition 1 *Let Z is a quasi-injective bisimulation between models $\mathcal{M}_1$ and $\mathcal{M}_2$, and let s_1 and s_2 be states in $\mathcal{M}_1$ and $\mathcal{M}_2$ respectively such that $s_1 Z s_2$. Then for all sentences of* $\mathrm{ML} + \Downarrow^0$, *$\mathcal{M}_1, s_1 \models \varphi$ iff $\mathcal{M}_2, s_2 \models \varphi$.*

We can use this result to show that no sentence of $\mathrm{ML} + \Downarrow^0$ defines the *Until* operator. To be more specific, let p and q be propositional symbols. Then, even over strictly partially ordered models, there is no sentence $\varphi^{U(p,q)}$ of $\mathrm{ML} + \Downarrow^0$ that is satisfied in a model $\mathcal{M}$ at a state s iff *Until*(p, q) is satisfied in $\mathcal{M}$ at s. To see this consider the following two models:

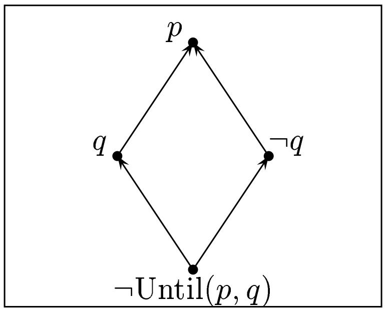

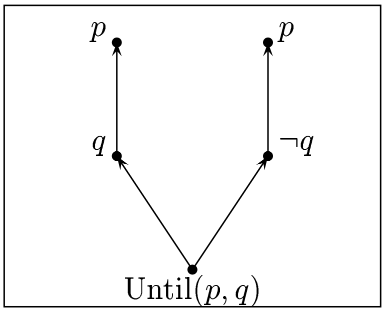

(In both models, the relation we are interested in is the transitive closure of the relation indicated by the arrows, thus both models are strict partial orders.) Note that *Until*(p,q) is false in the left-hand model at the root node, and true in the right-hand model at the root. Hence if some sentence $\varphi^{U(p,q)}$ of $\mathrm{ML} + \Downarrow^0$ expressed *Until*(p,q), it would be false at the root of the left-hand model, and true at the root of the right-hand side one. But this is impossible, for the obvious 'unravelling' relation between the two models is a quasi-injective bisimulation. Hence *Until*(p,q) is not expressible.

There is a further difficulty with $\mathrm{ML} + \Downarrow^0$: even over strictly partially ordered models, there is no obvious way to provide a complete axiomatisation without resorting to fairly complex rules of proof. For a start, $\Downarrow^0_x \Box\varphi \to \Box\Downarrow^0_x\varphi$, the Barcan analog in $\mathrm{ML} + \Downarrow^0$ is unsound over strict partial orders. (For example, consider $\Downarrow^0_x \Box\neg x \to \Box\Downarrow^0_x \neg x$; in any strictly partially ordered model, this is false at any state that has a successor.) Moreover, at present no alternative axioms are known which build a suitable bridge between the modalities and the $\Downarrow^0$ binder. Now, as is shown in Blackburn and Tzakova (1997), there is a way round the problem: by making use of COV (see Gargov and Goranko (1993)) a collection of additional rules of proof, it is possible to prove a completeness result.[4] Unfortunately, the COV rules are rather complex. Given that $\mathrm{ML} + \Downarrow^0$ is too expressively weak to define the *Until* operator, it seems natural to search for a nicer hybrid language for temporal logic: a language which is local, capable of defining *Until*, and for which use of the COV rule can be avoided. Indeed, it was the quest for such a language that led us to the $\Downarrow^1$ and $\downarrow^1$ binders.

The binders $\Downarrow^1$ and $\downarrow^1$

Question: given that the current state is the most modally significant state, what are the next most important states? Answer: the states accessible from the current state, of course! Clearly the $\Downarrow^1$ binder is simply a universal quantifier across such states. Note that it we define $\downarrow^1_x\varphi$ to be the dual binder $\neg\Downarrow^1_x\neg\varphi$, then:

$$\mathcal{M}, g, s \models \downarrow^1_x\varphi \quad \text{iff} \quad \mathcal{M}, g', s \models \varphi, \text{ for some } g' \overset{x}{\sim} g \text{ such that } sRg'(x)$$

Thus $\downarrow^1_x\varphi$ is an existential quantifier across accessible states. The reader should think of $\Downarrow^1$ and $\downarrow^1$ as a pair of binders that "match" the actions of $\Box$ and $\Diamond$ respectively. Observe that both $\Downarrow^1$ and $\downarrow^1$ are definable in $\mathrm{ML} + \forall$, for $\Downarrow^1_x\varphi$ is simply $\forall x(\Diamond x \to \varphi)$ and $\downarrow^1_x\varphi$ is $\exists x(\Diamond x \wedge \varphi)$.

[4]The earliest work on axiomatic systems for $\Downarrow^0$ seems to be that of Goranko (1994) and Goranko (1996a). However Goranko's investigations have little bearing on the concerns of the present paper, for Goranko investigates a language containing both the universal modality and $\Downarrow^0$, and the universal modality plays a central role in his axiomatisation. Note that the $\forall$ binder is definable in this language by $\forall x\varphi := \Downarrow^0_y A \Downarrow^0_x A(y \to \varphi)$, thus this language has full first-order expressive power.

Like $\Downarrow^0$, $\Downarrow^1$ is an intrinsically local binder; it is obvious that its sentences are preserved under the formation of generated submodels. Moreover, over strict partial orders, it is strong enough to define *Until*:

$$\mathit{Until}(\varphi,\psi) \;:=\; \downarrow^1_x(\Diamond(x\wedge\varphi)\wedge\Box(\Diamond x\rightarrow\psi)).$$

So it only remains to check that the $\Downarrow^1$ binder allows us to give a smooth axiomatisation of the minimal temporal logic.

4 The axiomatisation

We now introduce a logic called $\mathcal{H}(I4)$ axiomatically. Our goal is to show that $\mathcal{H}(I4)$ consists of precisely the $\mathrm{ML}+\Downarrow^0+\Downarrow^1$ formulas valid on strict partial orders, and the next few pages will be devoted to proving this.

$\mathcal{H}(I4)$ is an extension of *K4*, the modal logic of strict partial orders. Recall that *K4* is the smallest set of (ordinary modal) formulas containing all propositional tautologies, all instances of $\Box(\varphi\rightarrow\psi)\rightarrow(\Box\varphi\rightarrow\Box\psi)$, and all instances of $\Box\varphi\rightarrow\Box\Box\varphi$, that is closed under *modus ponens* (if φ and $\varphi\rightarrow\psi$ are both provable, then so is ψ) and *necessitation* (if φ is provable then so is $\Box\varphi$).

Suppose we have fixed a (countable) language $\mathcal{L}$ of $\mathrm{ML}+\Downarrow^0+\Downarrow^1$. By $\mathcal{H}(I4)$ we mean the smallest set of $\mathcal{L}$ formulas that contains all instances of the *K4* axiom schemas, and all instances of the axiom schemas listed below, that is closed under the following rules of proof: modus ponens, necessitation, and *state variable generalisation for both* $\Downarrow^0$ *and* $\Downarrow^1$ (that is, if φ is provable then so are $\Downarrow^0_x\varphi$ and $\Downarrow^1_x\varphi$, for all state variables x). We assume the usual notion of formal proof, and if φ is provable we write $\vdash\varphi$.

Our axiom schemas fall naturally into four groups. The first group reflects the basic quantificational powers of $\Downarrow^0$ and $\Downarrow^1$.

Q1	$\Downarrow^0_{\mathbf{v}}(\varphi\rightarrow\psi)\rightarrow(\varphi\rightarrow\Downarrow^0_{\mathbf{v}}\psi)$	$\Downarrow^1_{\mathbf{v}}(\varphi\rightarrow\psi)\rightarrow(\varphi\rightarrow\Downarrow^1_{\mathbf{v}}\psi)$
Q2	$\Downarrow^0_{\mathbf{v}}\varphi\rightarrow(\mathbf{s}\rightarrow\varphi[\mathbf{s}/\mathbf{v}])$	$\Downarrow^1_{\mathbf{v}}\varphi\rightarrow(\Diamond\mathbf{s}\rightarrow\varphi[\mathbf{s}/\mathbf{v}])$
Q3	$\Downarrow^0_{\mathbf{v}}(\mathbf{v}\rightarrow\varphi)\rightarrow\Downarrow^0_{\mathbf{v}}\varphi$	$\Downarrow^1_{\mathbf{v}}(\Diamond\mathbf{v}\rightarrow\varphi)\rightarrow\Downarrow^1_{\mathbf{v}}\varphi$
Self-Dual	$\Downarrow^0_{\mathbf{v}}\varphi\leftrightarrow\downarrow^0_{\mathbf{v}}\varphi$	

(Here $\mathbf{v}$ is a metavariable over state variables, $\mathbf{s}$ a metavariables over state symbols, and φ and ψ metavariables over arbitrary wffs. In *Q1*, φ cannot contain free occurrences of $\mathbf{v}$; in *Q2*, $\mathbf{s}$ must be substitutable for $\mathbf{v}$ in φ.)

Q1 and *Q2* are obvious analogs of familiar first-order axiom schemas, and indeed, the schemas given earlier for ML + $\forall$. The major difference is that the present version of *Q2* only lets us substitute state symbols for binders when the appropriate conditions are fulfilled: either $\mathbf{s}$ must be true in the current state (for the $\Downarrow^0$ binder) or it must be true at a successor state (for the $\Downarrow^1$ binder). These restrictions motivate the introduction of *Q3*. *Q3* allows us to eliminate conditional occurrences of state variables (usually those introduced by *Q2*) that have subsequently become bound; we shall shortly see some examples. Finally, while it is clear that much of the logic of our two binders is similar, $\Downarrow^0$ has one property that $\Downarrow^1$ lacks: it is self-dual. The *Self-Dual* axiom, which is included only for $\Downarrow^0$, reflects this. Summing up: both $\Downarrow^0$ and $\Downarrow^1$ legitimate a *restricted* version of classical quantificational reasoning. When proving completeness we shall have to find a way to manouevre past these restrictions.

The second group consists of analogs of *Barcan*:

Barcan$_{10}$ $\quad \Downarrow^1_{\mathbf{v}}\Box\varphi \rightarrow \Box\Downarrow^0_{\mathbf{v}}\varphi$

Barcan$_{11}$ $\quad \Downarrow^1_{\mathbf{v}}\Box\varphi \rightarrow \Box\Downarrow^1_{\mathbf{v}}\varphi$

These two permutation principles play an absolutely crucial role in the completeness proof. One point is worth noting: the soundness of *Barcan*$_{11}$ depends on the fact that we are working with *transitive* models.

The third group consists of a single schema, and reflects the fact that state variables are essentially 'labels' which uniquely identify nodes. In a sense, it is a modal counterpart of the first-order theory of equality.

Nom $\quad \Downarrow^1_{\mathbf{v}}[\Diamond(\mathbf{v} \wedge \varphi) \rightarrow \Box(\mathbf{v} \rightarrow \varphi)]$

The fourth group consists of a single schema reflecting the fact that we are working with irreflexive models.

I $\quad \Downarrow^0_{\mathbf{v}}\Box\neg\mathbf{v}$

In many ways, $\mathcal{H}(I4)$ behaves in ways familiar from classical logic. For example, the *normality schemas* for $\Downarrow^0$ and $\Downarrow^1$, *replacement of equivalents*, and *α-conversion* for the state variables, can all be straightforwardly established.

Lemma 4 (Normality) *For all formulas φ and ψ we have:*

$$\vdash \Downarrow^0_x(\varphi \rightarrow \psi) \rightarrow (\Downarrow^0_x\varphi \rightarrow \Downarrow^0_x\psi) \qquad \vdash \Downarrow^1_x(\varphi \rightarrow \psi) \rightarrow (\Downarrow^1_x\varphi \rightarrow \Downarrow^1_x\psi).$$

Proof. Note that $\Downarrow^1_x(\varphi \rightarrow \psi) \rightarrow (\Diamond x \rightarrow (\varphi \rightarrow \psi))$ is an instance of *Q2*, as is $\Downarrow^1_x\varphi \rightarrow (\Diamond x \rightarrow \varphi)$. Hence $\vdash (\Downarrow^1_x(\varphi \rightarrow \psi) \wedge \Downarrow^1_x\varphi)) \rightarrow (\Diamond x \rightarrow \psi)$. Use generalisation to prefix this formula with $\Downarrow^1_x$, and then *Q1* to distribute $\Downarrow^1_x$ over the main implication to get $\vdash (\Downarrow^1_x(\varphi \to \psi) \wedge \Downarrow^1_x\varphi) \to \Downarrow^1_x(\Diamond x \to \psi)$. Note that $\Downarrow^1_x(\Diamond x \to \psi) \to \Downarrow^1_x\psi$ is an instance of *Q3*, so we can simplify the consequent and so obtain the result. (Using *Q3* to simplify the conditionals produced by applications of *Q2* is fairly common in $\mathcal{H}(I4)$ proofs.) The proof for $\Downarrow^0$ is analogous. $\dashv$

Lemma 5 (Replacement of equivalents) *If $\vdash \xi \leftrightarrow \chi$, and the formula $\varphi(\xi)$ differs from the formula $\varphi(\chi)$ only in having ξ at zero or more places where $\varphi(\chi)$ has χ, then $\vdash \varphi(\xi) \leftrightarrow \varphi(\chi)$.*

Proof. If we have that $\vdash \varphi \leftrightarrow \psi$ implies $\vdash \Downarrow^0_x\varphi \leftrightarrow \Downarrow^0_x\psi$, $\vdash \Downarrow^1_x\varphi \leftrightarrow \Downarrow^1_x\psi$, and $\vdash \Box\varphi \leftrightarrow \Box\psi$, the lemma follows by a straightforward induction. But these properties are immediate by the previous lemma and the axiom $\Box(\varphi \to \psi) \to (\Box\varphi \to \Box\psi)$. $\dashv$

Lemma 6 (α-conversion) *If y is substitutable for x in φ and φ has no free occurrences of y, then $\vdash \Downarrow^0_x\varphi \leftrightarrow \Downarrow^0_y\varphi[y/x]$ and $\vdash \Downarrow^1_x\varphi \leftrightarrow \Downarrow^1_y\varphi[y/x]$.*

Proof. Straightforward. Use *Q1*, *Q2*, *Q3* and the generalisation rules. $\dashv$

So far, so good — but we really do need to be careful! Hybrid languages (and in particular, the $\Downarrow^1$ operator) have a distinctly modal flavour, and assuming 'obvious' classical principles can be dangerous. A simple example concerns vacuous occurrences of binders. In classical logic, we can simply add or discard vacuous quantifiers. What about our hybrid language? For the most part, the expected vacuity principles hold:

Lemma 7 (Vacuity principles) *Let φ be any formula containing no free occurrences of x. Then: $\vdash \varphi \to \Downarrow^0_x\varphi$, $\vdash \varphi \to \Downarrow^1_x\varphi$ and $\vdash \Downarrow^0_x\varphi \to \varphi$.*

Proof. The fact that $\vdash \varphi \to \varphi$ immediately yields (via generalisation and *Q1*) both $\vdash \varphi \to \Downarrow^0_x\varphi$ and $\vdash \varphi \to \Downarrow^1_x\varphi$. Using *Self-Dual*, we then obtain $\vdash \Downarrow^0_x\varphi \to \varphi$. $\dashv$

But what about $\Downarrow^1_x\varphi \to \varphi$? In fact, this is *not* provable, and we don't want it be, for it is not valid! To see this, consider $\Downarrow^1_x \bot \to \bot$. In any model, the antecedent will be true in any state that has no R-successors — but the consequent, of course, is false. It may be helpful to consider the contraposed and dualised form of this formula: $\top \to \downarrow^1_x\top$. The consequent demands the existence of an R-successor (in effect, the $\downarrow^1$ operator contains a 'hidden $\Diamond$') but the trivially true antecedent obviously doesn't guarantee that a successor exists.

A more significant example, which will affect our later work, is the following: when φ contains no free occurrences of x, then we have that $\vdash (\varphi \to \downarrow^0_x\psi) \to \downarrow^0_x(\varphi \to \psi)$, but we certainly do *not* have $\vdash (\varphi \to \downarrow^1_x\psi) \to \downarrow^1_x(\varphi \to \psi)$ (consider what happens when φ and ψ are $\bot$ and $\top$ respectively; again, the consequent demands the existence of an R-successor).

With the help of the Substitution Lemma it is routine to show that $\mathcal{H}(I4)$ is sound with respect to the class of all strictly partially ordered models, so let's turn to the question of its completeness.

5 Completeness

First some preliminaries. A set of formulas Σ is *consistent* iff for all formulas σ, if σ is a conjunction of (finitely many) formulas from Σ, then $\nvdash \sigma \to \bot$; otherwise Σ is *inconsistent*. A set of $\mathcal{L}$-formulas Σ is a *maximal consistent set* in $\mathcal{L}$ (an $\mathcal{L}$-MCS) iff it is consistent, and any set of $\mathcal{L}$-formulas that properly extends it is inconsistent. As $\mathcal{H}(I4)$ is an extension of classical propositional logic, *Lindenbaum's Lemma* holds: any consistent set of $\mathcal{L}$-formulas can be extended to an $\mathcal{L}$-MCS.

We want to prove the following completeness theorem: *any consistent set of formulas can be satisfied in a strictly partially ordered model*. So, given such a set, how do we build a model? The natural approach, of course, is to blend the modal technique of *canonical models*, with the first-order technique of Henkin proofs via *witnessed sets*, and this is what we shall do.

Definition 8 (Canonical Models) *For a countable language $\mathcal{L}$, the* canonical model *$\mathcal{M}^c$ is (S^c, R^c, V^c), where S^c is the set of all $\mathcal{L}$-MCSs; R^c is the binary relation on S^c defined by $\Gamma R^c \Delta$ iff $\Box\varphi \in \Gamma$ implies $\varphi \in \Delta$, for all $\mathcal{L}$-formulas φ; and V^c is the valuation defined by $V^c(a) = \{\Gamma \in S^c \mid a \in \Gamma\}$, where a is a proposition symbol or nominal.*

Definition 9 (Witnessed sets) *An MCS Γ is called* $\downarrow^0$-witnessed *iff for all formulas $\downarrow^0_x\varphi \in \Gamma$, there is some nominal c such that $(\varphi[c/x] \wedge c) \in \Gamma$. It is called* $\downarrow^1$-witnessed *iff for all formulas $\downarrow^1_x\varphi \in \Gamma$ there is a nominal c such that $(\varphi[c/x] \wedge \Diamond c) \in \Gamma$. It is called* witnessed *iff it is both $\downarrow^0$- and $\downarrow^1$-witnessed.*

Lemma 10 (Extended Lindenbaum's Lemma) *Let $\mathcal{L}^0$ and $\mathcal{L}^n$ be countable languages such that $\mathcal{L}^n$ is $\mathcal{L}^0$ enriched with a countably infinite set of new nominals. Then every consistent set of $\mathcal{L}^0$-formulas Δ can be extended to a witnessed $\mathcal{L}^n$-MCS Θ.*

Proof. Enumerate the nominals that are in $\mathcal{L}^n$ but not in $\mathcal{L}^0$. Let Δ be a consistent set of $\mathcal{L}^0$ formulas. We now inductively extend Δ to a witnessed $\mathcal{L}^n$-MCS. Let $\Theta^0 = \Delta$. Note that this set contains no nominals from the enumeration. Define Θ^{n+1} as follows. If $\Theta^n \cup \{\varphi_n\}$ is *inconsistent*, then $\Theta^{n+1} = \Theta^n$. Otherwise:

1. $\Theta^{n+1} = \Theta^n \cup \{\varphi_n\}$, if φ_n is not of the form $\downarrow^0_x\psi$ or $\downarrow^1_x\psi$.
2. $\Theta^{n+1} = \Theta^n \cup \{\varphi_n\} \cup \{\psi[c/x] \wedge c\}$, if φ_n is of the form $\downarrow^0_x\psi$.
3. $\Theta^{n+1} = \Theta^n \cup \{\varphi_n\} \cup \{\psi[c/x] \wedge \Diamond c\}$, if φ_n is of the form $\downarrow^1_x\psi$.

(In clauses 2 and 3, c is the first nominal in the enumeration that has not been used in the definitions of Θ^i, for all $i \leq n$, and that is not in φ_n.)

Let $\Theta = \bigcup_{n \geq 0} \Theta^n$. Clearly this set is maximal and witnessed. Moreover, it is consistent; the argument required is essentially the classical argument used to show the consistency of Henkin expansions. $\dashv$

Definition 11 (Witnessed models) *Let Σ be a witnessed MCS in some countable language $\mathcal{L}$, let $\mathcal{M}^c = (S^c, R^c, V^c)$ be the canonical model in $\mathcal{L}$, let S^Σ be those states in S^c that belong to the submodel generated by Σ, and let $Wit(S^c)$ be the set of all witnessed MCSs in S^c. Then the* witnessed model *$\mathcal{M}^w$ yielded by Σ is the triple (S^w, R^w, V^w), where $S^w = S^\Sigma \cap Wit(S^c)$, and R^w and V^w are the restrictions of R^c and V^c respectively to S^w.*

Lemma 12 *Let $\mathcal{M}^w$ be the witnessed model yielded by a witnessed MCS Σ. Then (1) R^w is transitive, (2) every MCS in S^w contains at least one nominal, (3) R^w is irreflexive, and (4) no state symbol occurs in more than one MCS in S^w.*

Proof. For (1), note it is a standard result that R^c is transitive (because all MCSs contain all instances of $\Box\varphi \to \Box\Box\varphi$, the transitivity axiom), hence as R^w is a subrelation of R^c it is transitive too.

For (2), note that $\vdash \downarrow^0_x x$ (this follows easily from the fact that $\vdash x \to x$) so every MCS in S^w contains $\downarrow^0_x x$. But as all MCSs in S^w are witnessed MCSs, they are $\downarrow^0$-witnessed, hence every MCS contains a witness for this formula. Such a witness is simply a nominal.

For (3), we have just observed that any MCS Γ contains some nominal, say c. But Γ also contains the axiom $\Downarrow^0_x \Box\neg x$. Hence, by *Q2*, Γ contains $\Box\neg c$ as well. Thus Γ is not R^c-related to itself (that is, R^c is irreflexive), and as R^w is a subrelation of R^c, R^w is irreflexive too.

For (4), suppose first that there are two distinct successors of Σ in $\mathcal{M}$, say Γ and Δ, that both contain the same state symbol $\mathbf{s}$. Since Γ and Δ are distinct, there is a formula δ that distinguishes them; that is, a formula δ such that $(\mathbf{s} \wedge \delta) \in \Gamma$ and $(\mathbf{s} \wedge \neg\delta) \in \Delta$. It follows (using the fact that R^w is a restriction of R^c) that $\Diamond(\mathbf{s} \wedge \delta) \in \Sigma$ and $\Diamond(\mathbf{s} \wedge \neg\delta) \in \Sigma$. Let z be any variable that does not occur in δ and is distinct from $\mathbf{s}$; then $\Downarrow^1_z(\Diamond(z \wedge \delta) \to \Box(z \to \delta))$ is an instance of the *Nom* axiom. As $\Diamond\mathbf{s} \in \Sigma$, it follows by *Q2* that $\vdash \Diamond(\mathbf{s} \wedge \delta) \to \Box(\mathbf{s} \to \delta)$. Thus, as $\Diamond(\mathbf{s} \wedge \delta) \in \Sigma$, $\Box(\mathbf{s} \to \delta) \in \Sigma$. But then, as $\Diamond(\mathbf{s} \wedge \neg\delta) \in \Sigma$, $\Diamond(\mathbf{s} \wedge \neg\delta \wedge \delta) \in \Sigma$, which is impossible. We conclude that no two distinct R^w-successors of Σ contain the same state symbol.

So next assume that some state symbol $\mathbf{s}$ is in both Σ and a successor Γ of Σ. As $\mathbf{s} \in \Sigma$, using the irreflexivity axiom and *Q2* we have that $\Box\neg\mathbf{s} \in \Sigma$. On the other hand, as $\mathbf{s}$ belongs to a successor of Σ, we also have that $\Diamond\mathbf{s} \in \Sigma$; as

Σ is consistent, this is impossible. Thus no state symbol occurs in both Σ and a successor of Σ, and we have shown that no state symbol occurs in more than one MCS in S^w. $\dashv$

Note, however, that there is no guarantee that every state symbol occurs in at least one witnessed MCS in a witnessed model. For example, the set $\{\neg c, \Box\neg c\}$ is consistent and thus can be extended to a witnessed MCS. But any witnessed MCS extending $\{\neg c, \Box\neg c\}$ yields a witnessed model in which c does not occur in any MCS, and in such models the natural valuation V^w is not a standard valuation as $V^w(c) = \{\}$. As we need to build a standard model (and a standard assignment) we need to ensure that all state symbols denote some state, so we shall glue a new root node $*$ onto our witnessed models. This new node will serve as a denotation for any symbol not contained in MCSs. More precisely:

Definition 13 (Completed models and assignments) *Given a witnessed MCS Σ, let $\mathcal{M}^w = (S^w, R^w, V^w)$ be the witnessed model yielded by Σ. The* completion *of $\mathcal{M}^w$ is a triple $\mathcal{M}^+ = (S^+, R^+, V^+)$, where $S^+ = S^w \cup \{*\}$ (where $*$ is an entity that is not an MCS); $R^+ = R^w \cup \{(*, \Gamma) \mid \Gamma \in S^w\}$; for all propositional variables p, $V^+(p) = V^w(p)$; and for all nominals c, $V^+(c) = \{\Gamma \in \mathcal{M}^w \mid c \in \Gamma\}$ if this set is non-empty, and $V^+(c) = \{*\}$ otherwise.*

If $\mathcal{M}^+ = (S^+, R^+, V^+)$ is a completion of a witnessed model $\mathcal{M}^w$, then the completed assignment g^+ on $\mathcal{M}^+$ is defined as follows: for all variables x, $g^+(x) = \{\Gamma \in \mathcal{M}^w \mid x \in \Gamma\}$ if this set is non-empty, and $g(x) = \{*\}$ otherwise.

Clearly (with the help of Lemma 12) completions of witnessed models are standard models, completed assignments are standard assignments, and R^+ is a strict partial ordering. Thus all theorems of the logic $\mathcal{H}(I4)$ are true in completed models with respect to the relevant complete assignment. Completions of witnessed models are well-behaved structures, and we shall use them to prove our Truth Lemma.

But we are not yet ready to do this. First we have to establish a crucial fact: that (completions of) witnessed models contain all the information required to cope with the *modalities*. That is, we need an Existence Lemma which tells us that if $\Diamond\varphi$ belongs to a *witnessed MCS* Δ, then Δ has a R^w-successor Γ containing φ. This is not obvious. The ordinary Existence Lemma for modal logic tells us that Δ has a R^c-successor Γ that contains φ — but it does not guaranteed that Γ is a *witnessed* MCS, and hence we have no idea whether or not Γ is an R^w-successor of Δ. To put it another way, we formed the witnessed model by throwing away the non-witnessed MCSs present in the canonical model. How do we know that we didn't throw away all the φ-containing R^c-successors of Δ?

The obvious approach to this problem is to tackle it head on by inductively constructing the required witnessed successors; this is the approach taken in Blackburn and Tzakova (1997) for $\mathrm{ML} + \forall$. However there is no obvious way to adapt this argument to $\mathrm{ML} + \Downarrow^0 + \Downarrow^1$ because of the following technical problem: as we have already observed, $(\varphi \to \downarrow^1_x \psi) \to \downarrow^1_x(\varphi \to \psi)$ is *not* a theorem of $\mathcal{H}(I4)$ for it is not sound. However $(\varphi \to \exists x \psi) \to \exists x(\varphi \to \psi)$, the corresponding principle in $\mathrm{ML} + \forall$, *is* sound and in fact plays an (easily overlooked, but important) role in the completeness proof for $\mathrm{ML} + \forall$. We might sum up the situation as follows: although $\mathrm{ML}+\Downarrow^0+\Downarrow^1$ is rich enough to support *Barcan* analogs, the restrictions on its classical component seem to make the obvious route to completeness difficult.

But there is an elegant solution — and a rather modally flavoured one at that. Let us call an MCS *named* iff it contains at least one nominal, and call a nominal in an MCS a *name* for that MCS. We are going to ask one simple, but important, question: *could we have thrown away any named* MCS? As we shall show, the answer is "no": all named successors of witnessed MCS are themselves witnessed. Why is this useful? Because it means that we are free to think in terms of named MCSs, rather than just in terms of witnessing — and, as will become clear, because we have the *Nom* schema at our disposal, named MCSs are very easy to work with. In particular, this approach leads to a relatively straightforward proof of the required Existence Lemma.

We shall need the following syntactical preliminary:

Lemma 14 *For any formulas φ and ψ such that φ contains no free occurrences of x we have that $\vdash (\varphi \wedge \downarrow^1_x \psi) \to \downarrow^1_x(\varphi \wedge \psi)$.*

Proof. Because of our restriction on φ, $\Downarrow^1_x(\varphi \to \neg\psi) \to (\varphi \to \Downarrow^1_x \neg\psi)$ is an instance of *Q1*, hence $\vdash \Downarrow^1_x(\varphi \to \neg\psi) \to (\neg\varphi \vee \Downarrow^1_x \neg\psi)$. Now, by replacement of equivalents we have $\vdash \Downarrow^1_x(\neg\varphi \vee \neg\psi) \leftrightarrow \Downarrow^1_x(\varphi \to \neg\psi)$, hence it follows that $\vdash \Downarrow^1_x(\neg\varphi \vee \neg\psi) \to (\neg\varphi \vee \Downarrow^1_x \neg\psi)$. Taking the contrapositive yields $\vdash \neg(\neg\varphi \vee \Downarrow^1_x \neg\psi) \to \neg\Downarrow^1_x(\neg\varphi \vee \neg\psi)$ and the required result follows. $\dashv$

Now for the key observation:

Lemma 15 *Let Σ be a witnessed MCS and $\mathcal{M}^\Sigma$ the submodel of the canonical model it generates. Suppose that Γ is a named MCS in $\mathcal{M}^\Sigma$. Then Γ is witnessed, and hence $\Sigma R^w \Gamma$.*

Proof. If Γ is Σ there is nothing to prove. So suppose that $\Gamma \neq \Sigma$ and that Γ is named by c. Note that $\Sigma R^c \Gamma$, as $\mathcal{M}^\Sigma$ is transitive.

We need to show that Γ is both $\downarrow^0$- and $\downarrow^1$-witnessed. Now, it is easy to see that Γ is $\downarrow^0$-witnessed. For suppose that $\downarrow^0_x \varphi \in \Gamma$. Then, by *Self-Dual*, $\Downarrow^0_x \varphi \in \Gamma$.

But by *Q2* we also have that $\Downarrow^0_x\varphi \rightarrow (c \rightarrow \varphi[c/x]) \in \Gamma$, hence $\varphi[c/x] \in \Gamma$ and Γ is $\downarrow^0$-witnessed.

So suppose $\downarrow^1_x\varphi \in \Gamma$. We have to show that there is a nominal k such that $\varphi[k/x] \wedge \Diamond k \in \Gamma$. Now, from the definition of the canonical relation, we have $\Diamond(c \wedge \downarrow^1_x\varphi) \in \Sigma$. So, suppose that the following formula was provable:

$$\textit{Named-Witness} \qquad \Diamond(c \wedge \downarrow^1_x\varphi) \rightarrow \downarrow^1_u\Diamond(c \wedge \varphi[u/x] \wedge \Diamond u),$$

for some variable u not in $\downarrow^1_x\varphi$. Then the required result would follow easily, for we would have that $\downarrow^1_u\Diamond(c \wedge \varphi[u/x] \wedge \Diamond u) \in \Sigma$, and hence, as Σ is $\downarrow^1$-witnessed, that $\Diamond(c \wedge \varphi[k/u] \wedge \Diamond k) \in \Sigma$, for some nominal k. Using *Nom* we could conclude that $\Box(c \rightarrow (\varphi[k/u] \wedge \Diamond k)) \in \Sigma$, and it would follow that $\varphi[k/u] \wedge \Diamond k \in \Gamma$ since c names Γ. Thus Γ would be $\downarrow^1$-witnessed as required.

And, indeed, *Named-Witness* is provable. Using *Q3* and α-conversion we have $\vdash \downarrow^1_x\varphi \rightarrow \downarrow^1_u(\varphi[u/x] \wedge \Diamond u)$, where u is a variable not occurring in $\downarrow^1_x\varphi$. Hence by propositional calculus and easy modal reasoning $\vdash \Diamond(c \wedge \downarrow^1_x\varphi) \rightarrow \Diamond(c \wedge \downarrow^1_u(\varphi[u/x] \wedge \Diamond u))$. Hence, by the previous lemma we have $\vdash \Diamond(c \wedge \downarrow^1_x\varphi) \rightarrow \Diamond\downarrow^1_u(c \wedge (\varphi[u/x] \wedge \Diamond u))$, and using $\textit{Barcan}_{11}$ to give $\downarrow^1_u$ scope over the consequent, we get *Named-Witness*. $\dashv$

Lemma 16 (Existence Lemma for Witnessed Models) *Let Σ be a witnessed MCS, and let $\mathcal{M}^w = (S^w, R^w, V^w)$ be the witnessed model yielded by Σ. Then for any MCS Δ in S^w, if $\Diamond\varphi \in \Delta$, then there is $\Theta \in S^w$ such that $\Delta R^w \Theta$ and $\varphi \in \Theta$.*

Proof. Easy. We are simply going to use the following formula to reduce the problem to the standard modal Existence Lemma:

$$\textit{Paste} \qquad \Diamond\varphi \rightarrow \downarrow^1_x\Diamond(x \wedge \varphi),$$

where x is a variable not occurring in φ. To see that *Paste* is provable, note that $\vdash \varphi \rightarrow (x \rightarrow (x \wedge \varphi)$. By generalisation and normality we get $\vdash \downarrow^0_x\varphi \rightarrow \downarrow^0_x(x \rightarrow (x \wedge \varphi)$, hence using *Q3* we get $\vdash \downarrow^0_x\varphi \rightarrow \downarrow^0_x(x \wedge \varphi)$. Lemma 7 tells us that $\varphi \rightarrow \downarrow^0_x\varphi$, as x does not occur in φ, hence $\vdash \varphi \rightarrow \downarrow^0_x(x \wedge \varphi)$. Standard modal reasoning yields $\vdash \Diamond\varphi \rightarrow \Diamond\downarrow^0_x(x \wedge \varphi)$, and $\textit{Barcan}_{10}$, then yields *Paste*.

But now the lemma is virtually immediate. If $\Diamond\varphi \in \Delta$, then $\downarrow^1_u\Diamond(u \wedge \varphi) \in \Delta$; hence as Δ is witnessed, for some nominal c, $\Diamond(c \wedge \varphi) \in \Delta$. Now, by the standard modal Existence Lemma, Δ has an R^c-successor containing c and φ; call this c-containing successor Θ. Moreover, as Θ is a named MCS in the submodel of the canonical model generated by Σ, by the previous lemma it is witnessed, and hence $\Delta R^w \Theta$. $\dashv$

Lemma 17 (Truth Lemma) *Let $\mathcal{M}$ be the completion of a witnessed model in some countable language $\mathcal{L}$, g the completed $\mathcal{M}$-assignment, and Δ an $\mathcal{L}$-MCS in $\mathcal{M}$. For every formula φ we have $\varphi \in \Delta$ iff $\mathcal{M}, g, \Delta \models \varphi$.*

Proof. By induction on the complexity of φ. If φ is a state symbol or a propositional variable the equivalence follows from the definition of the model $\mathcal{M}$ and the assignment g. The Boolean cases follow from obvious properties of MCSs. For the modal case, the Existence Lemma gives us the information required to drive through the left to right direction. The right to left direction is more or less immediate, though there is a small point the reader should observe: if $\mathcal{M}, g, \Delta \models \Diamond\psi$, then there is a successor of Δ that satisfies ψ. Since (by definition) no MCS precedes $*$, we conclude that $*$ cannot be this successor of Δ. Thus the successor to Δ that satisfies ψ is itself an MCS, and so we can apply the inductive hypothesis.

So suppose $\downarrow_x^1\psi \in \Delta$. Since Δ is witnessed, there is a nominal c such that $\psi[c/x] \wedge \Diamond c \in \Delta$. By the inductive hypothesis $\mathcal{M}, g, \Delta \models \psi[c/x]$. Moreover, by the Existence Lemma, Δ is related to some witnessed MCS containing c, hence $\mathcal{M}, g, \Delta \models \Diamond c$. Thus, by the contrapositive of the *Q2* axiom, $\mathcal{M}, g, \Delta \models \downarrow_x^1\psi$.

For the other direction assume $\mathcal{M}, g, \Delta \models \downarrow_x^1\psi$. That is, $\mathcal{M}, g', \Delta \models \psi$, where $g' \stackrel{x}{\sim} g$ such that $g'(x) = \{\Gamma\}$, where Γ is a witnessed MCS such that $\Delta R^+\Gamma$. By clause 2 of Lemma 12, Γ contains a nominal, say c. Now, by clause 2 of the Substitution Lemma, $\mathcal{M}, g, \Delta \models \psi[c/x]$, hence by the inductive hypothesis $\psi[c/x] \in \Delta$. Moreover, as $\Delta R^+\Gamma$, $\Diamond c \in \Delta$. So, by the contrapositive of the *Q2* axiom, $\downarrow_x^1\psi$ is in Δ as required.

The argument $\downarrow_x^0\psi \in \Delta$ is similar (in fact, slightly simpler as no appeals to the Existence Lemma are required) and is left to the reader. $\dashv$

Theorem 18 (Completeness) *Every $\mathcal{H}(I4)$-consistent set of formulas in a countable language $\mathcal{L}^0$ is satisfiable in a countable, rooted, strictly partially ordered, standard model with respect to a standard assignment function.*

Proof. Suppose Σ is a consistent set of $\mathcal{L}^0$-formulas. By the Extended Lindenbaum's Lemma we expand it to a witnessed MCS Σ^+ in the countable language $\mathcal{L}^n$. Let $\mathcal{M}^+$ be the completed model yielded by Σ^+, and let g^+ be the completed assignment on this model. $\mathcal{M}^+$ is a rooted strictly partially ordered standard model, and g^+ is a standard assignment. It is countable since every MCS contains a nominal, and $\mathcal{L}^n$ is countable. By the Truth Lemma, $\mathcal{M}^+, g^+, \Sigma^+ \models \Sigma^+$, and so $\mathcal{M}^+, g^+, \Sigma^+ \models \Sigma$. $\dashv$

6 Undecidability

In Blackburn and Seligman (1997) a Rabin-Gabbay style argument is used to show the decidability of the ML + $\Downarrow^0$-satisfiability problem over strict partial order. However, as we shall now show, adding the $\Downarrow^1$ operator leads to undecidability. In fact, even ML + $\Downarrow^1$ has an undecidable satisfiability problem over strict partial orders; the $\Downarrow^0$ operator is not needed. We shall show this by a reduction from the *unbounded tiling problem*.

A tile t is a 1×1 square, of fixed orientation, with coloured edges. The colours are drawn from some denumerable set, and for any tile t we'll use the notation $\mathit{right}(t)$, $\mathit{left}(t)$, $\mathit{up}(t)$, and $\mathit{down}(t)$ to indicate the colours of its four edges. The unbounded tiling problem is this: given a finite set $\mathcal{T}$ of tile types, can one cover $\omega \times \omega$ with tiles of these types in such a way that adjacent tiles have the same colour on their adjacent edges? More precisely, does there exist a function *tile* from $\omega \times \omega$ to $\mathcal{T}$ such that $\mathit{right}(\mathit{tile}(n,m)) = \mathit{left}(\mathit{tile}(n+1,m))$, and $\mathit{up}(\mathit{tile}(n,m)) = \mathit{down}(\mathit{tile}(n,m+1))$? This problem is known to be undecidable; see Harel (1983) for further discussion.

Theorem 19 *The ML + $\Downarrow^1$-satisfiability problem over strict partial order is undecidable.*

Proof. We shall construct, for any finite set of tiles $\mathcal{T} = \{t_1, \ldots, t_i, \ldots t_k\}$, a formula $\varphi^{\mathcal{T}}$ such that $\varphi^{\mathcal{T}}$ is satisfiable on an SPO iff $\mathcal{T}$ tiles $\omega \times \omega$. As the unbounded tiling problem is undecidable, it follows the ML + $\Downarrow^1$ satisfiability problem over strict partial orders is too.

Given a finite set of tiles $\mathcal{T} = \{t_1, \ldots, t_i, \ldots t_k\}$ we take $\mathcal{T} \cup \{0,1,2\}$ as our set of propositional symbols. We make heavy use of the following abbreviation: $[+]\psi \ := \ \psi \wedge \Box\psi$. Note that if $[+]\psi$ is satisfied at a state s in a strictly partially ordered model, then ψ holds at s and at every state in the submodel generated by s.

The formula $\varphi^{\mathcal{T}}$ we construct will have the following form:

$$\varphi^{\textit{2-imm}} \wedge \varphi^{\textit{share}} \wedge \varphi^{\textit{mod-3}} \wedge \varphi^{\textit{tiles}}.$$

The first two conjuncts together force the existence of a generated submodel isomorphic to $\omega \times \omega$, the third draws the distinction between vertical and horizontal, while the last covers this submodel with colour-matching tiles.

Defining $\varphi^{\textit{2-imm}}$. Let φ^{2} be the following sentence:

$$\begin{aligned}\downarrow^1_y\downarrow^1_z[&\Box(y \to \neg z) \wedge \Box(z \to \neg y) \wedge \neg\Diamond\Diamond y \wedge \neg\Diamond\Diamond z \\ &\wedge \Downarrow^1_x(\neg\Diamond\Diamond x \to \Box(x \to (y \vee z)) \\ &\wedge \Downarrow^1_w(\Diamond(w \wedge \neg y \wedge \neg z) \to (\Diamond(y \wedge \Diamond w) \vee \Diamond(z \wedge \Diamond w)))]\end{aligned}$$

If φ^{2} is satisfied at some state in an SPO then this state has exactly two immediate successors (namely the states labeled by y and z), and every other successor w is preceded by at least one of these states. As we want every state in the generated submodel to have such immediate successors, we define $\varphi^{\textit{2-imm}}$ to be $[+]\varphi^{\mathit{2}}$.

Defining φ^{share}. It will be helpful to define a modality that accesses the immediate successors, so we define: $\langle ims\rangle\psi \;:=\; \downarrow^1_w(\Diamond(w \wedge \psi) \wedge \neg\Diamond\Diamond w)$. With the help of this modality, we know stipulate that immediate successors share a common immediate successor. As a first step, we define φ^s to be

$$\downarrow^1_y\downarrow^1_z\downarrow^1_v(\langle ims\rangle(y \wedge \neg z \wedge \langle ims\rangle v) \wedge \langle ims\rangle(z \wedge \neg y \wedge \langle ims\rangle v)).$$

If this formula is satisfied at some state, then its two distinct immediate successors (labeled y and z) share a common minimal successor (the state labeled v). Again, we want this property to hold throughout the generated submodel, so we define $\varphi^{share} \;:=\; [+]\varphi^s$.

Defining $\varphi^{mod\text{-}3}$. Although the first two conjuncts force a gridlike structure, it is too symmetric: there is no obvious distinction between vertical and horizontal. However, by "labeling nodes modulo 3" (a technique we borrow from Spaan (1993)) we can impose the required distinction. So, let $\varphi^{0\text{-}or\text{-}1\text{-}or\text{-}2}$ be a sentence that is satisfied at a state iff that state satisfies exactly one of the propositional variables 0, 1 or 2. Further, suppose that a state s in our model is labeled by the number a, where $0 \leq a \leq 2$. Let $+_3$ denote addition mod 3. Essentially, we're going to insist that one of s's immediate successors is labeled with $a{+_3}1$, the other with $a{+_3}2$. We'll then think of the move from s to the immediate successor labeled $a{+_3}1$ as a *vertical* move, and a move from s to the immediate successor of s labeled $a{+_3}2$ as a *horizontal* move.

To carry this out, we first define two useful modalities:

$$\langle up\rangle\psi \quad := \quad \textstyle\bigwedge_{0\leq a\leq 2}(a \rightarrow \langle ims\rangle((a{+_3}1) \wedge \psi))$$

$$\langle right\rangle\psi \quad := \quad \textstyle\bigwedge_{0\leq a\leq 2}(a \rightarrow \langle ims\rangle((a{+_3}2) \wedge \psi))$$

We then define $\varphi^{mod\text{-}3}$ to be

$$[+](\varphi^{0\text{-}or\text{-}1\text{-}or\text{-}2} \wedge \langle up\rangle\top \wedge \langle right\rangle\top).$$

This guarantees that each state is labeled by a number less than three, and that each state has both a vertical and horizontal immediate successor. (Note that if a state s is labeled by the number a, then the shared common successor of s's two immediate successors is also labeled a.)

Defining φ^{tiles}. We can now force the desired tiling. First, define φ^t as follows:

$$\textstyle\bigvee_{1\leq i\leq k} t_i \wedge \bigwedge_{1\leq i\leq k}(t_i \rightarrow \neg(\bigvee_{j\neq i} t_j))$$

$$\wedge \textstyle\bigwedge_{1\leq i\leq k}(t_i \rightarrow \langle up\rangle(\bigvee_{up(t_i)=down(t_j)} t_j))$$

$$\wedge \textstyle\bigwedge_{1\leq i\leq k}(t_i \rightarrow \langle right\rangle(\bigvee_{right(t_i)=left(t_j)} t_j))$$

This formula straightforwardly forces the propositional variables t_i to act like matched tiles: every state satisfies exactly one tile, and the adjacent tiles have matching colours both vertically and horizontally. We want this to hold throughout the generated submodel, so we define: $\varphi^{tiles} := [+]\varphi^4$.

We have all we need. Let $\varphi^{\mathcal{T}}$ be

$$\varphi^{2\text{-}imm} \wedge \varphi^{share} \wedge \varphi^{mod\text{-}3} \wedge \varphi^{tiles}.$$

Clearly, if $\mathcal{T}$ tiles $\omega \times \omega$ then $\varphi^{\mathcal{T}}$ is satisfiable, for the tiling induces a suitable model in the obvious way. Moreover, as should be clear from the above discussion, if $\varphi^{\mathcal{T}}$ is satisfiable on a strictly partially ordered model, then the model embodies all the information needed to construct a tiling. The theorem follows. $\dashv$

7 Hybrid path languages

Our discussion of ML$+\Downarrow^0+\Downarrow^1$ should have given the reader a reasonable impression of what hybrid languages are, and we hope it has awakened interest in them as tools for temporal reasoning. But it hardly needs stressing that much remains to be done. To close our discussion we would like to pose the following question: what has hybridisation to offer temporal reasoning based on *paths* (or *courses of history*)? We won't be able to fully answer this question. As we shall see, while it is quite easy to devise hybrid path languages, and even to provide sensible looking axiomatisations that build on our previous work, at present it is unclear how to prove completeness (at least, with respect to standard models). We illustrate the issues involved by adding a local path binder $\Downarrow^\pi$ to ML $+ \Downarrow^0 + \Downarrow^1$.

By a path in a strict partial order (S, R) we mean a subset π of S that is convex, maximal under both R-successorship and R-predecessorship, and linear. We denote the set of paths in (S, R) by $\Pi(S, R)$. If $\pi \in \Pi(S, R)$ and $s \in \pi$ then we say that π passes through s. If (S, R) is a partial order, then $\Pi(S, R)$ is not empty; indeed at least one path passes through every state.

The basic strategy for dealing with paths in hybrid languages should be fairly clear. The first step is to add a third sort, the sort of *path symbols*. As with state symbols, path symbols should be divided into two subcategories, namely path variables (which will be open to binding) and path nominals (which will not). So we choose PVAR to be a countably infinite set of path variables (whose elements we typically write as ρ and ρ') and PNOM to be a countably infinite set of path nominals (whose elements we typically write as τ and τ'), and of course we choose these sets to be disjoint from each other and from PROP, SVAR, and NOM. We define the set of atoms of our enriched language to be PROP $\cup$ SVAR $\cup$ NOM $\cup$ PVAR $\cup$ PNOM.

The second step is to add a binder. We shall add a local binder called $\Downarrow^\pi$. As the notation should suggest, this will be a universal quantifier over paths through the current state ('local paths'). We call the resulting language $\mathrm{ML}+\Downarrow^0+\Downarrow^1+\Downarrow^\pi$. The wffs of this language are defined in the expected way, as are such concepts as free and bound path variables, so let us proceed straight to the semantics.

Definition 20 (Standard models and assignments) *Let $ML+\Downarrow^0+\Downarrow^1+\Downarrow^\pi$ be a hybrid language built over* PROP, SVAR, NOM, PVAR *and* PNOM*. A model $\mathcal{M}$ for this language is a triple (S,R,V) such that (S,R) is a strict partial order, and $V:\mathrm{PROP}\cup\mathrm{NOM}\cup\mathrm{PNOM}\longrightarrow Pow(S)$. A model is called* standard *iff for all nominals $c\in\mathrm{NOM}$, $V(c)$ is a singleton subset of S, and for all path nominals $\tau\in\mathrm{PNOM}$, $V(\tau)\in\Pi(S,R)$.*

An assignment on $\mathcal{M}$ is a mapping $g:\mathrm{SVAR}\cup\mathrm{PVAR}\longrightarrow Pow(S)$. An assignment is called standard *iff for all state variables $x\in\mathrm{SVAR}$, $g(x)$ is a singleton subset of S, and for all path variables $\rho\in\mathrm{PNOM}$, $V(\rho)\in\Pi(S,R)$.*

Now to interpret the language. The atomic clause is automatically taken care of by our $[V,g]$ notation, and the clauses for the Booleans and modalities are unchanged. So it only remains to interpret $\Downarrow^\pi$:

$$\mathcal{M},g,s\models\Downarrow^\pi_\rho\varphi \quad \text{iff} \quad \mathcal{M},g',s\models\varphi, \text{ for all } g'\overset{\rho}{\sim}g \text{ such that } s\in g'(\rho)$$

That is, $\Downarrow^\pi$ is indeed a universal quantifier over local paths. Note that the binder $\downarrow^\pi_\rho\varphi$ defined by $\neg\Downarrow^\pi_\rho\neg\varphi$ has the following satisfaction definition:

$$\mathcal{M},g,s\models\downarrow^\pi_\rho\varphi \quad \text{iff} \quad \mathcal{M},g',s\models\varphi, \text{ for some } g'\overset{\rho}{\sim}g \text{ such that } s\in g'(\rho)$$

Clearly this is an existential quantifier over local paths.

It is easy to see that sentences of this language are preserved under generated submodels. Moreover, some interesting concepts are expressible. For example, it is straightforward to define a fairness operator:

$$Fair(\varphi) \;:=\; \Downarrow^\pi_\rho\Downarrow^1_x\Box((x\wedge\rho)\to\Diamond(\varphi\wedge\rho)).$$

This operator guarantees that φ is true infinitely often along every branch.

And indeed, extending $\mathcal{H}(I4)$ to cover the new binder seems reasonably straightforward at first sight. We add a rule of *path variable generalisation* (if φ is provable then so is $\Downarrow^\pi_\rho\varphi$, for any path variable ρ) and all instances of the following three groups of axiom schemas.

Q1 $\quad \Downarrow^{\pi}_{\boldsymbol{\rho}}(\varphi \to \psi) \to (\varphi \to \Downarrow^{\pi}_{\boldsymbol{\rho}}\psi)$

Q2 $\quad \Downarrow^{\pi}_{\boldsymbol{\rho}}\varphi \to (\mathbf{p} \to \varphi[\mathbf{p}/\boldsymbol{\rho}])$

Q3 $\quad \Downarrow^{\pi}_{\boldsymbol{\rho}}(\boldsymbol{\rho} \to \psi) \to \Downarrow^{\pi}_{\boldsymbol{\rho}}\psi$

Local-Path $\quad \downarrow^{\pi}_{\boldsymbol{\rho}}\boldsymbol{\rho}$

(Here $\boldsymbol{\rho}$ and $\mathbf{p}$ are used as metavariables across path variables and path symbols respectively. In *Q1*, $\boldsymbol{\rho}$ must not be free in φ; and in *Q2*, $\mathbf{p}$ must be substitutable for $\boldsymbol{\rho}$ in φ.) In short, the basic quantificational powers of $\Downarrow^{\pi}$ described by *Q1* – *Q3* are analogous to those of $\Downarrow^{0}$ and $\Downarrow^{1}$. Moreover we know a little bit extra: as *Local-Path* tells us, at least one path passes through every state.

Next we add:

$$Barcan_{\pi} \qquad \Downarrow^{\pi}_{\boldsymbol{\rho}}\Box\varphi \to \Box\Downarrow^{\pi}_{\boldsymbol{\rho}}\varphi$$

The contraposed and dualised form $\Diamond\downarrow^{\pi}_{\boldsymbol{\rho}}\varphi \to \downarrow^{\pi}_{\boldsymbol{\rho}}\Diamond\varphi$ is perhaps easier to grasp. Essentially this says: "if we can select a suitable path at a successor state, then we can select a suitable path at the current state".

Finally we add three schemas reflecting the geometry of paths (we use $\boldsymbol{\tau}$ as a metavariable over path nominals and $\mathbf{c}$ and $\mathbf{d}$ as metavariable over state nominals):

P1 $\quad \Diamond\boldsymbol{\tau} \to \boldsymbol{\tau}$

P2 $\quad \boldsymbol{\tau} \wedge \Diamond\top \to \Diamond\boldsymbol{\tau}$

P3 $\quad \Diamond(\mathbf{c} \wedge \boldsymbol{\tau}) \wedge \Diamond(\mathbf{d} \wedge \boldsymbol{\tau}) \to \Diamond(\mathbf{c} \wedge \Diamond\mathbf{d}) \vee \Diamond(\mathbf{c} \wedge \mathbf{d}) \vee \Diamond(\mathbf{d} \wedge \Diamond\mathbf{c})$

Clearly *P1* reflects convexity, *P2* reflects R-maximality under successorship, and *P3* reflects linearity.

The soundness of this axiomatisation should be clear — but what about its completeness? Let's consider what happens if we try to extend the completeness proof for ML + $\Downarrow^{0}$ + $\Downarrow^{1}$ to this richer language.

First, we extend our concept of witnessing by adding the following demand: if $\downarrow^{\pi}_{\rho}\varphi \in \Gamma$ then there is a path nominal τ such that $\downarrow^{\pi}_{\rho}\varphi \to (\varphi[\tau/\rho] \wedge \tau) \in \Gamma$. Extending our earlier Extended Lindenbaum's Lemma to cover this richer notion of witnessing is routine.

Second, in our completeness proof ML + $\Downarrow^{0}$ + $\Downarrow^{1}$ we worked with witnessed models, so let us do so here. The turning point of our earlier work was Lemma 15 in which we proved that named MCSs in witnessed models were themselves witnessed.

This lemma extends straightforwardly to $\mathrm{ML} + \Downarrow^0 + \Downarrow^1 + \Downarrow^\pi$. Because we have $Barcan_\pi$, it is easy to prove the following analog of *Named-Witness*:

$$\textit{Named-Path-Witness} \qquad \Diamond(\tau \wedge \downarrow^\pi_\rho \varphi) \rightarrow \downarrow^\pi_{\rho'} \Diamond(\tau \wedge \varphi[\rho'/\rho] \wedge \rho')$$

(Here ρ' is not free in $\Diamond(\tau \wedge \downarrow^\pi_\rho \varphi)$.) With this at our disposal, the required extension of Lemma 15 is an obvious analog of our earlier work.

This has three more-or-less immediate consequences. First, the Existence Lemma follows exactly as in our earlier work. Second, some path nominal is true at every state. Third, for any path symbol, we can prove that the subset of a witnessed model consisting of all MCSs that contain that symbol is convex, maximal under R-successorship and R-predecessorship, and linear. There are one or two refinements of the earlier model construction that need to be made (for example, to complete the model we not only need to glue on an extra root node, we need to glue on a 'dummy successor' to the dummy root to ensure the R-maximality of the path symbol interpretations). However, the required ideas are pretty straightforward; our earlier work provides a scaffolding on which it possible to build a natural looking model.

So near — and yet so far. In spite of all these pleasant properties, the completeness result does not go through, for we cannot prove the Truth Lemma. The problem is this. We don't have any guarantee that *all* convex R-maximal linear subsets of witnessed models are 'named' by some path symbol. Nothing in our model construction guarantees this, and without it, the proof of the Truth Lemma does not go through.

This does not seem to be an easy problem to fix. Roughly speaking, what is needed is an 'infinitary extension' of the *Local-Path* schema. From *Local-Path* we can deduce that there is a path through the current state; what we also need is a principle that ensures that given a sequence of states (one of which is the current state) that satisfies the convexity, R-maximality, and linearity principles, then there is a path nominal that is true at all the states in this sequence.[5]

Once remark is worth making. This difficulty does not seem to reflect our insistence on working with local languages. As we mentioned in the introduction, Robert Bull in his pioneering article on hybrid languages introduced the idea of path symbols and path binders. Now, Bull worked with a non-local language: he used the universal modality A, the binder $\forall$ over states (thus he had access to the full power of first-order quantification over states) and, instead of the local

[5]A recent paper by Goranko on hybrid languages strong enough to embed CTL* (see Goranko (1996b)) may be relevant to this problem. Goranko's language doesn't contain path binders. Roughly speaking, he uses the universal modality and ordinary nominals to simulate the action of path nominals. Goranko proves a completeness result that makes use of some interesting rules of proof. At present, however, the relation between the two lines of work is unclear.

$\Downarrow^\pi$ binder he used the non-local binder $\forall^\pi$ that quantifies over arbitrary paths. In spite of this power, Bull runs into the same problem. He does not provide a solution but comments (see footnote 5 on page 292) that although not every path is the interpretation of some path symbol, his model:

> *... does provide enough paths V(u) to give a reasonable interpretation.*

Now, for many purposes Bull is probably right. Perhaps our path binders don't need access to all paths, but merely to some (application dependent) set of "well-behaved" ones. But clearly more precision is required here, and it would be nice to have more insight into what can and cannot be done with the standard interpretation. On this inconclusive note we close the paper.

References

Blackburn, P. and Seligman, J. (1995), Hybrid languages, *Journal of Logic, Language and Information* **4**, 251–272.

Blackburn, P. and Seligman, J. (1997), What are hybrid languages?, *in* M. Kracht, M. de Rijke, H. Wansing and M. Zakharyaschev (eds), *Advances in Modal Logic '96*, CSLI Publications, Stanford University. To appear.

Blackburn, P. and Tzakova, M. (1997), Hybrid completeness, *Journal of the IGPL* . To appear.

Bull, R. (1970), An approach to tense logic, *Theoria* **36**, 282–300.

Burgess, J. (1984), Basic tense logic, *in* D. Gabbay and F. Guenthner (eds), *Handbook of Philosophical Logic*, Vol. 2, Reidel, Dordrecht, pp. 135–165.

Gargov, D. and Goranko, V. (1993), Modal logic with names, *Journal of Philosophical Logic* **22**(6), 607–636.

Goranko, V. (1994), Temporal logic with reference pointers, *in* D. Gabbay and H. J. Ohlbach (eds), *Proceedings of the 1st International Conference on Temporal Logic*, Vol. 827 of *LNAI*, Springer, Berlin, pp. 133–148.

Goranko, V. (1996a), Hierarchies of modal and temporal logics with reference pointers, *Journal of Logic, Language and Information* **5**(1), 1–24.

Goranko, V. (1996b), An interpretation of computational tree logics into temporal logics with reference pointers, *Technical Report 2/96*, Verslagreeks van die Department Wiskunde, RAU, Department of Mathematics, Rand Afrikaans University, Johannesburg, South Africa.

Harel, D. (1983), Recurring dominoes: Making the highly undecidable highly understandable, *in* M. Karpinski (ed.), *Foundations of Computing Theory. Proceedings of the 1983 International FCT-Conference*, Springer Lecture Notes in Computer Science 158, Springer, Berlin, pp. 177–194.

Passy, S. and Tinchev, T. (1991), An essay in combinatory dynamic logic, *Information and Computation* **93**, 263–332.

Prior, A. (1967), *Past, Present and Future*, Oxford University Press, Oxford.

Prior, A. N. and Fine, K. (1977), *Worlds, Time and Selves*, University of Massachusetts Press.

Seligman, J. (1994), The logic of correct descriptions, *in* M. de Rijke (ed.), *Advances in Intensional Logic*, Applied Logic Series, Kluwer, Dordrecht. To apear.

Spaan, E. (1993), *Complexity of Modal Logic*, PhD thesis, Department of Mathematics and Computer Science, University of Amsterdam.