7

Extended Modal Logic

As promised in the preface, this chapter is the party at the end of the book. We've chosen six of our favorite topics in extended modal logic, and we're going to tell you a little about them. There's no point in offering detailed advice here: simply read these introductory remarks and the following Chapter Guide and turn to whatever catches your fancy.

Roughly speaking, the chapter works it's way from fairly concrete to more abstract. A recurrent theme is the interplay between modal and first-order ideas. We start by introducing a number of important *logical modalities* (and learn that we've been actually been using logical modalities all through the book). We then examine languages containing the *since* and *until* operators, and show that first-order expressive completeness can be used to show modal deductive completeness. We then explore two contrasting strategies, namely the strategy underlying *hybrid logic* (import first-order ideas into modal logic, notably the ability to refer to worlds) and the strategy that leads to the *guarded fragment* of first-order logic (export the modal locality intuition to classical logic). Following this we discuss *multi-dimensional modal logic* (in which evaluation is performed at a sequence of states), and see that first-order logic itself can be viewed as modal logic. We conclude by proving a *Lindström Theorem* for modal logic.

Chapter guide

Section 7.1: Logical Modalities (Basic track). Logical modalities have a fixed interpretation in every model. We introduce two of the most important (the *global modality*, and the *difference operator*) and briefly discuss *Boolean Modal Logic* (a system which contains an entire *algebra of diamonds*).

Section 7.2: Since and Until (Basic track). We introduce the since and until operators (and their stronger cousins, the *Stavi connectives*), discuss the expressive completeness results they give rise to, and use expressive completeness to prove deductive completeness.

Section 7.3: Hybrid Logic (Basic track). Hybrid languages are modal languages which can refer to worlds. They do so using atomic formulas called *nominals* which are true at exactly one world in any model. We introduce the basic hybrid language and discuss its completeness theory.

Section 7.4: The Guarded Fragment (Advanced track). As is clear from the standard translation, modal operators perform a 'guarded' form of quantification across states. What happens when this idea is exported to first-order logic and generalized? This section provides some answers.

Section 7.5: Multi-Dimensional Modal Logic (Advanced track). By viewing assignments as possible worlds and quantifiers as diamonds, one can treat first-order logic itself as a modal formalism. In fact, orthodox Tarskian semantics for first-order logic provides a prime example of multi-dimensional modal logic: formulas are evaluated at a sequence of points.

Section 7.6: A Lindström Theorem for Modal Logic (Advanced track). As a famous theorem due to Lindström tells us, any logic satisfying completeness, compactness, and Löwenheim Skolem is essentially first-order logic. Is there an analogous abstract characterization of modal logic?

7.1 Logical Modalities

Pure first-order logic has a significant expressive weakness: it's not strong enough to express the concept of equality in arbitrary structures. But because equality is such an important relation, logicians introduce a special binary relation symbol (namely =) and *stipulate* that it denotes the equality relation. As the interpretation of = is fixed, and as the relation it denotes is so fundamental, the equality symbol is called a *logical predicate*.

Logical modalities trade on the same idea. Are there important relations which ordinary modal languages cannot express? Very well then: let's add new modalities and stipulate that they be interpreted by the relation in question. In this section we'll discuss two of the most important logical modalities: the *global modality* (which is interpreted by the relation $W \times W$) and the *difference operator* (which is interpreted by $\neq$, the inequality relation). We'll also make a few remarks about *Boolean Modal Logic* (BML), a system containing an entire family of logical modalities.

But before going any further, let's get one thing absolutely clear: *we've been using logical modalities all through the book*. Here's the simplest example. Suppose we are working with the basic modal language. Now, for many purposes we may be happy simply using $\Diamond$ to talk about the relation R — but sometimes we may want to talk about $R^{\smile}$, the converse of R, as well. Now, we know (see Exercise 2.1.2) that this can't be done in the basic modal language, so we have to add a new backward-looking modality as a primitive; doing so, of course, gives us the basic temporal language. But note: we *don't* have to bring in the concept of time to justify this extension. If a binary relation R is important, its converse is likely to be too — so it's simply common sense to consider adding a diamond for $R^{\breve{}}$. In short, the 'temporal operator' P is really a logical modality.

The other important example is PDL. To motivate PDL we told a story about programs and transition systems — but a more abstract motivation is not only possible, it's more satisfying. The point is this. As soon as we fix a collection of relations R_α, regular algebra is staring us in the face: we can combine these relations using union and composition, and form transitive closures. Any model containing the initial R_α relations implicitly contains many other interesting relations as well — so it's natural to add extra modalities to deal with them explicitly, and doing so yields PDL. As this example shows, we can go way beyond the idea of adding a single new logical modality: we can add an entire *algebra of diamonds*. We'll see another example of this when we discuss BML.

The global modality

Throughout the book we've emphasized the locality of modal logic, and for many purposes local languages are ideal. For example, suppose we're working with a modal language for talking about computer networks, and in this language ϕ means `Server 1 is active` and ψ means `Server 2 is active`. Then we can check whether the network makes it possible for `Server 1` to be active by checking whether ϕ is satisfiable, and we can check whether it is possible for `Server 2` to be inactive by testing for the satisfiability of $\neg\psi$.

But suppose we want to know if *whenever* `Server 1` is active, then so is `Server 2`. There's no obvious way to test this. Testing for the satisfiability of $\phi \to \psi$ does *not* answer this question: if $\phi \to \psi$ is satisfiable, this only means that there is a state where either ϕ is false or ψ is true. We want to know whether *every* state that makes ϕ true is also a state that makes ψ true. This is clearly a global query. What are we to do?

Here's an elegant answer: enrich the language language with the *global modality*. To keeps things simple, suppose we're working in the basic modal language over some fixed choice of proposition letters; let's call this language *ML*$(\Diamond)$. We'll now add a second diamond, written E, and call the resulting language *ML*$(\Diamond, \mathrm{E})$. The interpretation of E is *fixed*: in any model $\mathfrak{M} = (W, R, V)$, E must be interpreted using the relation $W \times W$. That is:

$$\mathfrak{M}, w \Vdash \mathrm{E}\phi \text{ iff there is a } u \in W \text{ such that } \mathfrak{M}, u \Vdash \phi.$$

Thus E scans the entire model for a state that satisfies ϕ. Its dual $\mathrm{A}\phi := \neg\mathrm{E}\neg\phi$ has the following interpretation:

$$\mathfrak{M}, w \Vdash \mathrm{A}\phi \text{ iff } \mathfrak{M}, u \Vdash \phi, \text{for all } u \in W.$$

That is, $\mathrm{A}\phi$ asserts that ϕ holds at *all* points in the model. In effect, A brings the metatheoretic notion of global truth in a model down into the object language: for any model $\mathfrak{M}$, and any formula ϕ, we have that $\mathfrak{M} \Vdash \phi$ iff $\mathrm{A}\phi$ is satisfiable in $\mathfrak{M}$. We'll call E the *global diamond*, and A the *global box*. When it's irrelevant whether we mean E or its dual, we'll simply say *global modality*.

It should now be clear how to handle the computer network problem: to test whether `Server 2` is active whenever `Server 1` is, we test the satisfiability not of $\phi \to \psi$, but of $\mathrm{A}(\phi \to \psi)$. This query has exactly the global force required.

Well — this looks appealing. But what are the properties of this (obviously richer) new language? Maybe introducing the global modality destroys the properties that make model logic attractive in the first place! We've made an important change, and we need to take a closer look at the consequences.

Now, we could begin by discussing the sublanguage *ML*(E) — but this is not very interesting (it's easy to see that E is just an **S5** modality). Anyway (as our server example shows) the main reason for adding logical modalities is to have them available as *additional* tools. So the real question is: what does *ML*$(\Diamond, \mathrm{E})$ offer that *ML*$(\Diamond)$ doesn't? The most obvious answer is *expressivity*. Let's first consider expressivity at the level of frames:

$(R = W^2)$	$\mathrm{E}p \to \Diamond p$
$(R \neq \varnothing)$	$\mathrm{E}\Diamond\top$
$(\exists x \forall y \neg Rxy)$	$\mathrm{E}\Box\bot$
$(\forall x \exists y Ryx)$	$p \to \mathrm{E}\Diamond p$
$(\lvert W \rvert = 1)$	$\mathrm{E}p \to p$
$(\lvert W \rvert \leq n)$	$\bigwedge_{i=1}^{n+1} \mathrm{E}p_i \to \bigvee_{i \neq j} \mathrm{E}(p_i \wedge p_j)$
(R *is trichotomous*)	$(p \wedge \Box q) \to \mathrm{A}(q \vee p \vee \Diamond p)$
($\breve{R}$ *is well-founded*)	$\mathrm{A}(\Box p \to p) \to p$

None of the frame classes listed is definable in *ML*$(\Diamond)$, but (as we ask the reader to check in Exercise 7.1.1) the *ML*$(\Diamond, \mathrm{E})$ formulas to their right do define the corresponding property.

Where does this extra frame expressivity come from? From trivializing the notion of generated submodel (generating on $W \times W$ always yields $W \times W$) and rendering inapplicable the notion of disjoint union (for any disjoint frames (W, R) and (W', R'), $(W \times W) \uplus (W' \times W') \neq (W \uplus W') \times (W \uplus W')$). By insisting that E be interpreted using $W \times W$, we've trashed two of the classic modal preservation results and thereby bought ourselves more expressivity. How much more? For first-order definable frame classes, the answer is elegant:

Theorem 7.1 *A first-order definable class of frames is definable in ML*$(\Diamond, \mathrm{E})$ *iff it is closed under taking bounded morphic images, and reflects ultrafilter extensions.*

This is *exactly* the Goldblatt-Thomason Theorem — minus closure under disjoint unions and generated subframes.

There is also a gain of expressivity at the level of models (the server example makes this clear, and we already know from Section 2.1 that the global modality is not definable in the basic modal language). Moreover, we can measure the gain using our old friends: bisimulations. It's an easy exercise to adapt the definition of bisimulation for the basic modal language to $ML(\Diamond, \mathrm{E})$, and a rather more demanding one to prove a van Benthem style characterization result for the language. The reader is asked to attend to these matters in Exercises 7.1.3 and 7.1.4.

What about completeness? The set of valid $ML(\Diamond, \mathrm{E})$ formulas can be axiomatized as follows. Take the minimal normal logic in $\Diamond$ and E (that is, apply Definition 4.13 to this two-diamond similarity type), and add the following axioms:

(*reflexivity*) $p \to \mathrm{E}p$
(*symmetry*) $p \to \mathrm{AE}p$
(*transitivity*) $\mathrm{EE}p \to \mathrm{E}p$
(*inclusion*) $\Diamond p \to \mathrm{E}p$

Note that first three axioms are the familiar T, B, and 4 axioms (written in E and A rather than $\Diamond$ and $\Box$). We discussed *Inclusion* in Example 1.29(4). We'll call this logic $\mathbf{K}_g$.

Theorem 7.2 $\mathbf{K}_g$ *is strongly complete with respect to the class of all frames.*

This theorem says that to lift the minimal logic $\mathbf{K}$ (for the basic modal language) to $ML(\Diamond, \mathrm{E})$, we need merely treat the global modality as a normal operator that satisfies four further axioms. In fact, we can lift *any* canonical $ML(\Diamond)$ logic in this way. If $\mathbf{K\Gamma}$ is a normal modal logic in $ML(\Diamond)$, let $\mathbf{K}_g\mathbf{\Gamma}$ be the normal modal logic in $ML(\Diamond, \mathrm{E})$ obtained by treating E as a normal operator and adding the four axioms listed above. Then:

Theorem 7.3 *Let Γ be a set of $ML(\Diamond)$ formulas, and let F be the class of frames that Γ defines. If $\mathbf{K\Gamma}$ is canonical, then $\mathbf{K}_g\mathbf{\Gamma}$ is strongly complete with respect to* F.

Proof. Let $\mathfrak{M} = (W, R_\Diamond, R_{\mathrm{E}}, V)$ be the canonical model for $\mathbf{K}_g\mathbf{\Gamma}$. Note that as $\mathbf{K\Gamma} \subseteq \mathbf{K}_g\mathbf{\Gamma}$, we have that $(W, R_\Diamond)$ belongs to F, for $\mathbf{K\Gamma}$ is canonical. Indeed, *any* generated subframe of $(W, R_\Diamond)$ belongs F, for validity in the basic modal language is closed under generated subframes.

Given a $\mathbf{K}_g\mathbf{\Gamma}$-consistent set of sentences Σ, use Lindenbaum's Lemma to expand it to an $\mathbf{K}_g$-MCS Σ^+. By the Canonical Model Theorem, $\mathfrak{M}, \Sigma^+ \Vdash \Sigma$. Now, (*reflexivity*), (*symmetry*), and (*transitivity*) are canonical formulas, thus R_{E} is an equivalence relation. And although there is no guarantee that R_{E} is $W \times W$, this is easy to correct: let $\mathfrak{M}' = (W', R'_{\Diamond}, R'_{\mathrm{E}}, V)$ be the submodel of $\mathfrak{M}$ generated by Σ^+ using the R_{E}-relation. Then $R'_{\mathrm{E}} = W' \times W'$, so we have the global relation we need. Furthermore, because of *Inclusion*, $R_{\Diamond} \subseteq R_{\mathrm{E}}$, thus $\mathfrak{M}'$ is also a generated submodel of $\mathfrak{M}$ with respect to $R_{\Diamond}$, hence $\mathfrak{M}', \Sigma^+ \Vdash \Sigma$. It only remains to observe that (by our initial remarks) $(W', R'_{\Diamond})$ is in F, hence the result follows. (Theorem 7.2 is the special case in which $\Gamma = \varnothing$.) $\dashv$

Example 7.4 Suppose we're working with *ML*$(\Diamond)$ over transitive frames (so the relevant logic is **K4**, which is canonical). Now, we may want to state global constraints on models, or insist that certain information holds somewhere or other, and of course we can do this if we add the global modality. But how do we obtain a complete logic for transitive frames in the enriched language?

Simply enrich **K4** by treating the global modality as a normal operator and adding the (*reflexivity*), (*transitivity*), (*symmetry*), and (*inclusion*) axioms. Doing so yields $\mathbf{K}_g\mathbf{4}$, and by the theorem just proved this logic is strongly complete with respect to the class of transitive frames. $\dashv$

What about decidability and complexity? We briefly met the global modality in Section 6.5, and we saw that its global reach makes it possible to force the existence of gridlike models. This led to undecidability results for languages containing several diamonds, and it's not difficult to adapt these arguments to find frame classes with decidable *ML*$(\Diamond)$ logics and undecidable *ML*$(\Diamond, \mathrm{E})$ logics (we give such an example in Exercise 7.1.5). Moreover, although undecidability does not strike over the class of all frames, $\mathbf{K}_g$ is probably more complex than **K**, for $\mathbf{K}_g$ has an EXPTIME-complete satisfiability problem (the reader was asked to prove this in Exercises 6.8.1 and 6.8.2) while **K** is PSPACE-complete (see Section 6.7). On the other hand, there is a rather nice transfer result concerning the filtration method: if we can prove the decidability of a *ML*$(\Diamond)$ logic by using filtrations to establish establish the strong finite frame property, then we can also do so after adding the global modality. For example, it follows that the logic $\mathbf{K}_g\mathbf{4}$ (see Example 7.4) is decidable. We'll state and prove a stronger version of this result when we discuss the difference operator.

All in all, the global modality is a strikingly natural extension of modal logic — and at first glance this seems surprising. How can something so obviously global blend so well with the locality of modal logic? Basically, because the enriched language still takes an *internal* perspective on relational structure. Although we now have a global operator at our disposal, we still place formulas *inside* models and evaluate them at a particular state. To put it another way, the intuition that a modal formula is an automaton scanning accessible states is remarkably robust: even if we add a special automaton programmed to regard *all* states as accessible, we retain much of the characteristic flavor of ordinary modal logic.

A lot more could be said about the global modality. For a start, it's natural when viewed from an algebraic perspective (it gives rise to *discriminator varieties*). Moreover, the global modality can be added to many richer modal systems, including **PDL** and the hybrid and multi-dimensional logics discussed later in the chapter, often without raising the computational complexity (for example **PDL** is EXPTIME-complete, and adding E doesn't change this). But for more information the reader will have to consult the Notes and Exercises, for it's time to discuss an even more powerful logical modality.

The difference operator

At the bottom of every toolbox lies a heavy cast-iron hammer. It's not the sort of tool we use every day — for delicate jobs it's inappropriate, and we may feel slightly embarrassed about using it at all. Still, there'll always come a time when something simply won't budge, and then we find ourselves reaching for it. Think of the difference operator as that hammer.

Once again, we'll start with *ML*$(\Diamond)$. We'll add a second diamond D, the *difference operator*, and call the resulting language *ML*$(\Diamond, \mathrm{D})$. The interpretation of D is *fixed*: in any model $\mathfrak{M} = (W, R, V)$, D must be interpreted using the inequality relation $\neq$. That is:

$$\mathfrak{M}, w \Vdash \mathrm{D}\phi \text{ iff there is a } u \neq w \text{ such that } \mathfrak{M}, u \Vdash \phi.$$

Thus the difference operator scans the entire model looking for a *different* state that satisfies ϕ. Its dual $\overline{\mathrm{D}} := \neg\mathrm{D}\neg\phi$ has the following interpretation

$$\mathfrak{M}, w \Vdash \overline{\mathrm{D}}\phi \text{ iff } \mathfrak{M}, u \Vdash \phi \text{ for all } u \neq w.$$

In what follows we discuss *ML*$(\Diamond, \mathrm{D})$, but the sublanguage *ML*(D) is quite interesting in its own right, and we ask the reader is asked to explore it in Exercise 7.1.6.

Using the difference operator, we can define the global modality: $\mathrm{E}\phi := \phi \vee \mathrm{D}\phi$. Thus all our earlier examples of frame classes definable in *ML*$(\Diamond, \mathrm{E})$ are definable in *ML*$(\Diamond, \mathrm{D})$ too. But *ML*$(\Diamond, \mathrm{D})$ can define even more:

(*irreflexivity*)	$\Diamond p \to \mathrm{D}p$
(*antisymmetry*)	$(p \wedge \neg\mathrm{D}p) \to \Box(\Diamond p \to p)$
($\exists xy(x \neq y)$)	$\mathrm{D}\top$
($\lvert W\rvert > n$)	$\mathrm{A}(\bigvee_{1 \leq i \leq n} p_i) \to \mathrm{E}\bigvee_{1 \leq i \leq n}(p_i \wedge \mathrm{D}p_i)$

None of these frame classes is closed under bounded morphic images hence (by Theorem 7.1) none of them is definable in *ML*$(\Diamond, \mathrm{E})$; but it is easy to see that the listed *ML*$(\Diamond, \mathrm{D})$ formulas successfully capture them. Incidentally, we have already seen that *ML*$(\Diamond, \mathrm{E})$ can define $\lvert W\rvert \leq n$, thus as *ML*$(\Diamond, \mathrm{D})$ can define $\lvert W\rvert > n$, the difference operator can count states, at least as far as *frames* are concerned; in

Exercise 7.1.7 we ask the reader to investigate whether it can count over *models* as well. Furthermore, note the $p \wedge \neg \mathrm{D}p$ antecedent in the definition of antisymmetry. This is only true when p is true at exactly one state in the model: in effect we are using the power of D to force p to act as 'name' for a state; we'll put this power to good use shortly.

What about completeness? The set of valid $ML(\Diamond, \mathrm{D})$ formulas can be axiomatized as follows. Take the minimal normal logic in $\Diamond$ and D, and add the following axioms:

(*symmetry*)	$p \to \overline{\mathrm{D}}\mathrm{D}p$
(*pseudo-transitivity*)	$\mathrm{DD}p \to (p \vee \mathrm{D}p)$
(*D-inclusion*)	$\Diamond p \to p \vee \mathrm{D}p$

We'll call this logic $\mathbf{K}_d$. Now, it's not particularly difficult to prove the completeness of $\mathbf{K}_d$ (we ask the reader to do so in Exercise 7.1.8) — but it's harder than with $\mathbf{K}_g$ (we have to do more than simply take a generated submodel) and the result doesn't extend to stronger logics so easily (there's no obvious analog of Theorem 7.3). Moreover, it's easy to find frame incompleteness results, indeed we can even find them in the sublanguage $ML(\mathrm{D})$! Things aren't looking too good ...

Enter the hammer. When we discussed rules for the undefinable (Section 4.7) we learned that proof rules which rely on 'names' can lead to general frame completeness results. And as we noted above, the difference operator is powerful enough to simulate state names, thus we can formulate the following rule of proof (the D-rule):

$$\frac{\vdash (p \wedge \neg Dp) \to \theta}{\vdash \theta}$$

(Here p is a proposition letter that doesn't occur in θ. The intuitions underlying this rule are analogous to those underlying the IRR rule discussed in Section 4.7, and we'll leave it to the reader to verify that it preserves validity.) And now for a remarkable result. The D-rule neatly meshes with our earlier work on Sahlqvist formulas to yield one of the most general completeness results known in modal logic, the D-*Sahlqvist theorem*.

Here we only formulate a version in the basic temporal language. Consider the language with operators F, P and D; let, for a set Σ of axioms in this logic, $\mathbf{K}_{td}\Sigma$ be the normal modal logic generated by the axioms of basic temporal logic, the D-axioms and D-rule given above, and the formulas in Σ.

Theorem 7.5 *Let Σ be a collection of Sahlqvist formulas in the basic temporal language. Then $\mathbf{K}_{td}\boldsymbol{\Sigma}$ is strongly sound and complete with respect to the class of bidirectional frames defined by (the first-order frame correspondents of) the axioms in Σ.*

Proof. We will prove weak completeness only. The first step of the proof is to prove the existence of a collection W of maximal consistent sets such that

(i) each Γ in W contains a name, that is, a formula of the form $\phi \wedge \neg \mathrm{D}\phi$,
(ii) for each Γ in W and each formula $F\psi \in \Gamma$, there is a Δ in W such that Γ and Δ are in the canonical accessibility relation R^c_F for F; and likewise, for the operators P and D.
(iii) for each pair of distinct points Γ and Δ in W we have $R^c_{\mathrm{D}}\Gamma\Delta$.

All of this can be proved in the style of Proposition 4.71.

It easily follows from (i) and (iii) above that R^c_{D} is the inequality relation on W. But then the model on W given by $V(p) = \{\Gamma \in W \mid p \in \Gamma\}$ is *named*; that is, for every point in the model there is a formula which is true *only* at this point, see Definition 4.76. However, condition (ii) allows us to prove a *Truth Lemma* which implies that all axioms of the logic are true throughout the model. But then it follows from Theorem 4.77 that the Sahlqvist axioms are valid on the underlying frame as well. $\dashv$

The pinch of Theorem 7.5 lies in the fact that the first-order frame correspondents it mentions use *inequality* for the ‘relation symbol’ referring to the accessibility relation of D. This means that we can automatically axiomatize frame properties like irreflexivity or antisymmetry. The reader doubt the usefulness of this: isn’t the logic of the class of irreflexive frames is identical to the logic of the class of all frames? True, but this may change when we consider irreflexivity in *addition* with other properties. Conditions like irreflexivity, undefinable in themselves, may nevertheless have ‘side effects’ so to speak. What we mean is that there are frame classes K such that the logic of K *differs* from the logic of the irreflexive frames in K. In such cases the above theorem can be of tremendous help.

In a surprisingly large number of cases we find ourselves in the situation that over a certain class of frames, the difference operator is *definable* in the underlying modal language. For example, over the class of strict linear orders, the temporal formula $Fp \vee Pp$ holds at a point if and only if p holds at a *different* point. In general, we say that a formula $\delta(p)$ *acts as* D on a frame $\mathfrak{F}$ if $\mathfrak{F} \Vdash \delta(p) \leftrightarrow \mathrm{D}p$; if $\delta(p)$ acts as the difference operator on every frame in a class K then we say that δ *defines* D *over* K.

Definability of the difference operator is of great use for axiomatizability, as the following result shows. For a formula $\delta(p)$, let $\mathbf{K}_{t\delta}\Sigma$ be the ‘δ’-version of $\mathbf{K}_{td}$, that is, the logic in the language without the D-operator obtained by replacing, in all axioms and derivation rules of $\mathbf{K}_{td}$ every formula $\mathrm{D}\phi$ with $\delta(\phi)$.

Theorem 7.6 *Let Σ be a collection of Sahlqvist formulas. Then $\mathbf{K}_{t\delta}\boldsymbol{\Sigma}$ is strongly sound and complete with respect to the class of those bidirectional frames on which Σ is valid and on which δ acts as the difference operator.*

In the section on multi-dimensional modal logic we will see an application of this theorem; for a proof, we refer the reader to Exercise 7.1.9. We will examine another name-driven proof rule (called PASTE) in detail when we discuss hybrid logic. First we turn to decidability issues concerning the difference operator.

ML$(\Diamond, \mathrm{D})$ is a strong language. As it can define the global modality, $\mathbf{K}_d$ must have an EXPTIME-hard satisfiability problem (in fact, the problem is EXPTIME-complete; see Exercise 7.1.10) and it is even easier to find undecidable logics than in *ML*$(\Diamond, \mathrm{E})$. Nonetheless, decidability is often retained. In particular, if the *ML*$(\Diamond)$ logic of a class of frames can be proved decidable by using a filtration argument to establish the strong finite frame property, then the *ML*$(\Diamond, \mathrm{D})$ logic of that same frame class can be proved decidable in the same way. Let's prove this.

Definition 7.7 Let Λ be a logic, and let F be a class of frames for Λ. We say that Λ *admits filtrations on* F if for any model $\mathfrak{M}$ which is based on a frame in F, and for any finite subformula closed set Σ of *ML*$(\Diamond)$ formulas, there is a filtration $\mathfrak{M}^f$ of $\mathfrak{M}$ through Σ which is based on a frame in F. $\dashv$

Theorem 7.8 *Suppose that* F *is a class of frames, and that* Λ_{F} *(the set of all* *ML*$(\Diamond)$*-formulas valid on* F*) admits filtrations on* F*. Then the logic* Λ^d_{F} *(the set of all* *ML*$(\Diamond, \mathrm{D})$*-formulas valid on* F*) has the strong finite frame property with respect to* F*.*

Proof. Let ξ be a *ML*$(\Diamond, \mathrm{D})$-formula satisfiable in a model $\mathfrak{M} = (W, R, V)$ of which the underlying frame (W, R) is in F. We want to show that ξ is satisfiable in an F-frame of bounded size.

Let Σ be the set of subformulas of ξ. First consider the relation $\equiv_\Sigma$ which holds between two points if they satisfy the same formulas in Σ. As the points of our finite model we would like to take the equivalence classes of this relation but this would not work out well (it is instructive to see how the proof of the filtration lemma fails in the inductive step of the difference operator). The key idea of the proof of the theorem is to solve this problem by splitting each equivalence class in two parts — unless the original class is a singleton. To achieve this we add a new proposition letter d to the language and we make d true at exactly one point of each equivalence class. We would then like to filtrate the new model according the equivalence relation $\equiv_{\Sigma \cup \{d\}}$.

There is still a problem however: we can only guarantee that the underlying frame of the filtrated model is in F if we filtrate through a set of *ML*$(\Diamond)$ formulas. But Σ may contain formulas with occurrences of D. In order to get rid of these, we employ a little technical trick. For every formula of the form $\mathrm{D}\psi$ in Σ, choose a distinct propositional variable q_ψ that does not occur in any formula in Σ. Let V' be the valuation that differs from V, if at all, only in that $V'(q_\psi) = \{w \mid \mathfrak{M}, w \Vdash \mathrm{D}\psi\}$ and that $V'(d)$ is as indicated above. Let $\mathfrak{M}'$ be the model (W', R', V').

Now define the set Σ' as follows. It is not difficult to see that for every $\phi \in \Sigma$ there is a *unique* *ML*$(\Diamond)$ formula ϕ' such that ϕ can be obtained from ϕ' by replacing in ϕ' every proposition letter q_ψ by $\mathrm{D}\psi$. Put

$$\Sigma' = \{\phi' \mid \phi \in \Sigma\} \cup \{d, q_\psi \mid \mathrm{D}\psi \in \Sigma\}.$$

Observe that the formulas in Σ' are D-free and that Σ' is subformula closed. The model $\mathfrak{M}'$ is (or can be seen as) an *ML*$(\Diamond)$-model satisfying

$$\mathfrak{M}, s \Vdash \phi \text{ iff } \mathfrak{M}', s \Vdash \phi' \tag{7.1}$$

for all formulas ϕ in Σ. Let $\equiv_{\Sigma'}$ hold between two points iff they satisfy the same formulas in Σ'; it is easy to see that every $\equiv_\Sigma$-equivalence class $|s|$ splits into either one or two $\equiv_{\Sigma'}$-equivalence classes, depending on whether $|s|$ has one or more elements.

In any case, it follows from the assumption in the theorem that there is a filtration $\mathfrak{M}^f$ through Σ' which is based on a frame in F. Note that by definition, the points of $\mathfrak{M}^f$ are the $\equiv_{\Sigma'}$-equivalence classes. We claim that this model $\mathfrak{M}^f$ satisfies the following property for all *ML*$(\Diamond, \mathrm{D})$-formulas ϕ in Σ and all states s in $\mathfrak{M}$:

$$\mathfrak{M}, s \Vdash \phi \text{ iff } \mathfrak{M}^f, |s| \Vdash \phi. \tag{7.2}$$

From this, the theorem is almost immediate.

The proof of (7.2) proceeds by a formula induction of which we omit the standard inductive steps concerning the boolean operators; the clauses for $\Diamond$ are fairly easy as well — but note that for one direction, one needs (7.1). For the case that ϕ is of the form $\mathrm{D}\psi$ we also omit the easy right-to-left direction of (7.2). For the other direction, suppose that $\mathfrak{M}, s \Vdash \mathrm{D}\psi$. Then there is a point $s' \neq s$ such that $\mathfrak{M}, s' \Vdash \psi$. If $|s|$ and $|s'|$ are distinct then we are finished, so suppose otherwise. But from $s \equiv_{\Sigma'} s'$ it follows on the one hand that $\mathfrak{M}, s \Vdash d$ iff $\mathfrak{M}, s' \Vdash d$, and on the other hand, that s and s' belong to the same $\equiv_\Sigma$-equivalence class. Since we chose *exactly* one point in each $\equiv_\Sigma$-class to satisfy d, this means that neither s nor s' can be this special point. Hence, there must be *another* point s'' in this $\equiv_\Sigma$-equivalence class which does make d true. From $s' \equiv_\Sigma s''$ it follows that $\mathfrak{M}, s'' \Vdash \psi$, so by the inductive hypothesis we have that $\mathfrak{M}^f, |s''| \Vdash \psi$. But $|s''|$ is distinct from $|s|$ since d holds at s'' and not at s. This gives that $\mathfrak{M}^f, |s| \Vdash \mathrm{D}\psi$, as required. $\dashv$

How does decidability follow? Any logic Λ that admits filtrations on F has the strong finite frame property with respect to F — so if F is recursive we can apply Theorem 6.7 and conclude that Λ_{F} is decidable. But then by the result just proved, we know that Λ^d_{F} also has the strong finite frame property with respect to F, so we can apply the model enumeration idea underlying the proof of Theorem 6.7 to formulas of the richer languages. As D is always interpreted by the inequality relation, and as this relation is obviously computable on finite structures, the decidability of Λ^d_{F} follows.

A great deal more could be said about the difference operator (in particular, bisimulations are easily adapted to cope with D, and a van Benthem style characterization result is forthcoming; see Exercises 6.8.1 and 6.8.2) but it's time to take a brief look at a system containing a whole family of logical modalities.

Boolean modal logic

As we have remarked, as soon as we fix a collection of relations R_α, we can form the regular algebra over this base; building an algebra of diamonds corresponding to these leads to PDL. But an even more obvious algebra demands attention: we can also form the *boolean algebra* over base relations R_α. Why not define an algebra of diamond corresponding to 1, $-$, $\cap$, and $\cup$? Doing so leads to Boolean Modal Logic (BML).

We define the language of BML as follows. As with PDL, we fix a set of primitive relation symbols a, b, c, $\ldots$, and in addition a distinguished relation symbol 1. From these we build complex relations using the relation constructors $-$, $\cap$ and $\cup$: that is, if α and β are relation symbols, then so are $\neg\alpha$, $\alpha \cap \beta$, and $\alpha \cup \beta$. BML is the modal language containing a diamond $\langle\alpha\rangle$ for each relation symbol α. In principle we can interpret BML on any model of appropriate similarity type — that is triples $\mathfrak{M} = (W, \{R_\alpha \mid \alpha \text{ is a relation symbol }\}, V)$ — but most such models are inappropriate. We are only interested in *boolean models*, the models in which $R_1 = W \times W$, and such that, for all relation symbols α and β, $R_{-\alpha} = \overline{R_\alpha}$ (that is, $(W \times W) \setminus R_\alpha$), $R_{\alpha\cap\beta} = R_\alpha \cap R_\beta$, and $R_{\alpha\cup\beta} = R_\alpha \cup R_\beta$.

BML is an expressive language — for a start, it contains the global modality — and it may seem that we've bitten off more than we can chew. While the $\cup$ constructor is well behaved (in particular $\mathfrak{F} \Vdash \langle\alpha\cup\beta\rangle p \leftrightarrow \langle\alpha\rangle\phi \vee \langle\beta\rangle p$ iff $R_{\alpha\cup\beta} = R_\alpha \cup R_\beta$), the $\cap$ constructor is difficult to work with. However, as we will now see, with the help of the $-$ constructor we can get an exact grip on the relations of interest.

First we define the following operator (often called *window*): for any relation symbol α:

$$[\![\alpha]\!]\phi \;:=\; [-\alpha]\neg\phi.$$

That is:

$$\mathfrak{M}, w \Vdash [\![\alpha]\!]\phi \text{ iff } \forall u(\mathfrak{M}, u \Vdash \phi \;\Rightarrow\; R_\alpha wu).$$

Window is an extremely natural operator — once you've seen it, you wonder how you ever managed without it. For example, if we read $[\alpha]\phi$ as saying that *all* executions of program α lead to a ϕ state, then $[\![\alpha]\!]\phi$ says that *only* executions of program α can lead to a ϕ state, and it has other useful readings too (see the Notes) But what concerns us here is the following result: window allows very smooth definitions of the relations we are interested in.

Proposition 7.9 *Let $\mathfrak{F}$ be a frame $(W, \{R_\alpha \mid \alpha$ is a relation symbol $\})$. Then:*

(i) $\mathfrak{F} \Vdash [-\alpha]p \leftrightarrow \llbracket\alpha\rrbracket\neg p$ *iff* $R_{\overline{\alpha}} \subseteq \overline{R_\alpha}$

(ii) $\mathfrak{F} \Vdash [\alpha]\neg p \leftrightarrow \llbracket -\alpha\rrbracket p$ *iff* $\overline{R_\alpha} \subseteq R_{\overline{\alpha}}$

(iii) $\mathfrak{F} \Vdash \llbracket\alpha \cap \beta\rrbracket p \leftrightarrow \llbracket\alpha\rrbracket p \wedge \llbracket\beta\rrbracket p$ *iff* $R_{\alpha\cap\beta} = R_\alpha \cap R_\beta$.

Proof. We prove the third claim. The right to left direction is trivial. For the left to right direction, assume that $\mathfrak{F} \Vdash \llbracket\alpha \cap \beta\rrbracket p \leftrightarrow \llbracket\alpha\rrbracket p \wedge \llbracket\beta\rrbracket p$. We need to show that $R_{\alpha\cap\beta} = R_\alpha \cap R_\beta$. To see that $R_{\alpha\cap\beta} \subseteq R_\alpha \cap R_\beta$, suppose that $R_{\alpha\cap\beta}wu$, and let V be any valuation on $\mathfrak{F}$ such that $V(p) = \{u\}$. Then $(\mathfrak{F}, V), w \Vdash \llbracket\alpha \cap \beta\rrbracket p$. As $\mathfrak{F} \Vdash \llbracket\alpha \cap \beta\rrbracket p \leftrightarrow \llbracket\alpha\rrbracket p \wedge \llbracket\beta\rrbracket p$ we have $(\mathfrak{F}, V), w \Vdash \llbracket\alpha\rrbracket p \wedge \llbracket\beta\rrbracket p$. But u is the only point satisfying p, hence $R_\alpha wu$ and $R_\beta wu$. A similar argument shows that $R_\alpha \cap R_\beta \subseteq R_{\alpha\cap\beta}$. $\dashv$

In a sense, the relations are divided into two kingdoms: the ordinary $[\alpha]$ modalities govern relations built with $\cup$, the widow modalities $\llbracket\alpha\rrbracket$ govern the relations built with $\cap$, and the $-$ constructor acts as a bridge between the two realms. Moreover the bridging function of $-$ also finds expression in a new rule of proof, BR. Unlike the other additional rules discussed in this book, BR is not name-driven:

$$\frac{\vdash [\alpha]p \to ([\beta]p \to [\gamma]p)}{\vdash [\alpha]p \to (\llbracket\gamma\rrbracket\neg p \to \llbracket\beta\rrbracket\neg p)} \qquad (\text{BR})$$

While it is possible to prove a completeness result for BML without using BR, its use leads to an elegant axiomatization, for it enables us to thread negations through the structured modalities.

A final surprise is in store. In Theorem 6.31 we showed that the fragment containing the $\cap$ constructor and the global modality was undecidable over deterministic frames. Nonetheless, the minimal logic in BML actually turns out to be *decidable*. All in all, BML is a fascinating system. For more information, see the Notes.

Exercises for Section 7.1

7.1.1 We listed numerous frame conditions definable in *ML*$(\Diamond, \mathrm{E})$ and *ML*$(\Diamond, \mathrm{D})$ which were not definable in *ML*$(\Diamond)$. Show that these definability claims are correct.

7.1.2 Show that *ML*$(\Diamond, \mathrm{E})$ validity is preserved under bounded morphisms and reflects ultrafilter extensions. (That is, show the easy direction of the Goldblatt-Thomason style result for *ML*$(\Diamond, \mathrm{E})$ stated in Theorem 7.1.) Can you prove the (far more demanding) converse?

7.1.3 Extend the standard translation to the global modality and the difference operator. Extend the notion of bisimulation for the basic modal language to *ML*$(\Diamond, \mathrm{E})$ and *ML*$(\Diamond, \mathrm{D})$, and show prove that your definition leads to an invariance result.

7.1.4 Building on the previous exercise, characterize the expressivity of *ML*$(\Diamond, \mathrm{E})$ and *ML*$(\Diamond, \mathrm{D})$ over models.

7.1.5 Let 2-3 be the class of frames (W, R) such that every state has 2 R-successors, and 3 R-successors of R-successors. First show that the satisfiability problem in *ML*$(\Diamond)$ over 2-3 is *decidable* (note: this *cannot* be proved using a filtration argument). Then show that the satisfiability problem in *ML*$(\Diamond, \mathrm{E})$ over 2-3 is undecidable. (It may be helpful to note that this exercise is related to Exercise 6.5.2.)

7.1.6 Show that a class of frames is definable in *ML*(D) if and only if it is definable in the first-order language over $=$ (that is, the first-order language of equality). What is the complexity of the satisfiability problem for *ML*(D)?

7.1.7 Clearly we can define in *ML*$(\Diamond, \mathrm{D})$ an operator Q with the following satisfaction definition: for any model $\mathfrak{M}$, any state w in $\mathfrak{M}$, and any formula ϕ, $\mathfrak{M}, w \models Q\phi$ iff there is exactly one state u in $\mathfrak{M}$ such that $\mathfrak{M}, u \models Q\phi$. But it is also possible in to define modalities $Q_2\phi$, $Q_3\phi$, $Q_3\phi$, and so on, that are satisfied when ϕ holds at precisely Q_n states $(n \geq 2)$ in the model?

7.1.8 Show that $\mathbf{K}_d$ is complete with respect to the class of all frames. (No need to try anything fancy here — just fiddle with the canonical model.)

7.1.9 Prove Theorem 7.6. That is, let Σ be a collection of Sahlqvist formulas in the basic modal language. Show that $\mathbf{K}_{t\delta}\Sigma$ is strongly sound and complete with respect to the class of those frames on which Σ is valid and on which δ acts as the difference operator.
(Hint: use an auxiliary logic $\mathbf{K}_{t\delta}\Sigma^+$ in the temporal language expanded with the difference operator. Simply define this logic as having *both* the D *and* the δ versions of the D-axioms and rules. Now first use Theorem 7.5 to prove that this logic is sound and strongly complete with respect to the class of Σ-frames on which δ acts as the difference operator. Then, prove that $\mathbf{K}_{t\delta}\Sigma^+$ is conservative over $\mathbf{K}_{t\delta}\Sigma$; that is, show that for every purely temporal formula ϕ, we have that ϕ belongs to $\mathbf{K}_{t\delta}\Sigma$ iff it belongs to $\mathbf{K}_{t\delta}\Sigma^+$.)

7.1.10 Use an elimination of Hintikka sets argument to show that the $\mathbf{K}_d$ satisfiability problem is solvable in EXPTIME.

7.2 Since and Until

The modal operators considered in previous chapters all have satisfaction definitions involving only existential or only universal quantifiers. In this section we look at a popular temporal logic whose operators are based on modalities with more complex satisfaction definitions: S (since) and U (until). The main reason for considering these modalities is, again, to achieve an increase in expressive power. We'll first give some examples demonstrating why the increased expressivity is useful. We'll then learn that (over Dedekind complete frames) we have actually achieved expressive *completeness*: any expression in the first-order correspondence language (in one free variable) has an equivalent in the modal language in S and U. Finally, we'll show that this (first-order) *expressive* completeness leads to (modal) *deductive* completeness.

Basic definitions

The basic operators needed for temporal reasoning seem to be F and P. These allow us to say things like 'Something good will happen' and 'Something bad has happened.'

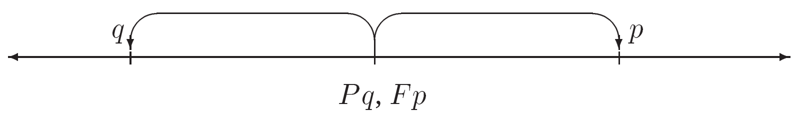

But in several application areas this is not enough. For example, in the semantics of concurrent programs one often needs to be able to express properties of executions of programs that have the general format 'Something good is going to happen, *and until that time nothing bad will happen*.' Or, more concretely: p will be the case, and until that time q will hold:

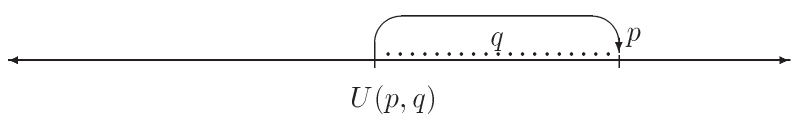

Such properties are sometimes called *guarantee properties* in the computational literature. To state them, the binary *until* operator U can be used; its satisfaction definition reads:

$t \Vdash U(\phi, \psi)$ iff

there is a $v > t$ such that $v \Vdash \phi$ and for all s with $t < s < v$: $s \Vdash \psi$.

The mirror image of U is the *since* operator S:

$t \Vdash S(\phi, \psi)$ iff

there is a $v < t$ such that $v \Vdash \phi$ and for all s with $v < s < t$: $s \Vdash \psi$.

That's the basic idea — but before going further, let's make our discussion a little more precise. The set of S, U*-formulas* is built up from a collection Φ of proposition letters, the usual boolean connectives, and the *binary* operators S and U. The *mirror image* of a formula ϕ is obtained by simultaneously substituting S for U and U for S in ϕ.

S, U-formulas are interpreted on frames of the form $\mathfrak{F} = (T, <)$, where T is a set of time points and $<$ is a binary relation on T. U looks forward along $<$, and S looks backwards. We use the notation $(T, <)$ for frames (rather than our usual (T, R)) because here we are primarily interested in the temporal interpretation of S and U. In fact, will be working with frames $(T, <)$ such that $<$ is a Dedekind complete order — more on this below. To emphasize our interest in the temporal interpretation, we will often refer to frames as *flows of time*. As usual, a valuation is a function assigning subsets of T to the proposition letters in the language.

How does the language in S and U relate to the basic temporal language? First, observe that F and P are definable in the language with S and U: we can define $F\phi := U(\phi, \top)$, $P\phi := S(\phi, \top)$, $G\phi := \neg F \neg \phi$ and $H\phi := \neg P \neg \phi$. Thus the language with S and U is at least as strong as the basic temporal language. In fact, it is strictly stronger. For a start, we saw in Exercise 2.2.4 that the basic temporal language couldn't define U. Moreover, as the following proposition shows, even if we restrict attention to models based on the real numbers, the basic temporal language still isn't strong enough to define U.

Proposition 7.10 *U is not definable over $(\mathbb{R}, <)$ using F and P.*

Proof. We will give two models that agree on all formulas in the language with F and P only, but that can be distinguished using the until operator. Consider the following model $\mathfrak{M}_1$ based on the reals:

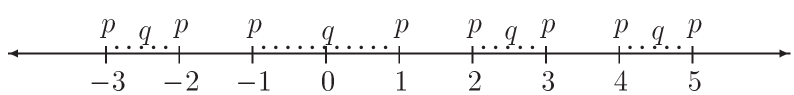

$0 \Vdash U(p, q)$

So, $V_1(p) = \{r \mid r \in \mathbb{Z}\}$, and $V_1(q) = \{0\} \cup \{r \mid \exists n \in \mathbb{N}(-2n - 1 < r < -2n)\} \cup \{r \mid \exists n \in \mathbb{N}(2n < r < 2n + 1)\}$.

Next, consider the model $\mathfrak{M}_2$ given by the following picture:

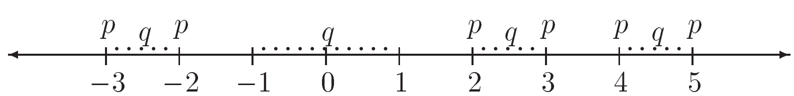

$0 \not\Vdash U(p, q)$

We leave it to the reader to show that the models $\mathfrak{M}_1$ and $\mathfrak{M}_2$ agree on all formulas in F and P, but that $\mathfrak{M}_1, 0 \Vdash U(p, q)$, whereas $\mathfrak{M}_2, 0 \not\Vdash U(p, q)$ (see Exercise 7.2.1). $\dashv$

So the temporal language in S and U is expressive — but just how expressive is it? To answer such questions we need a correspondence language and a standard translation of S and U into the correspondence language. Let Φ be a collection of proposition letters, and let $\mathcal{L}^1_<(\Phi)$, or simply $\mathcal{L}^1_<$, be the first-order language with unary predicate symbols corresponding to the proposition letters in Φ, and with $=$ and $<$ as binary relation symbols. We use $\mathcal{L}^1_<(x)$ to denote the set of $\mathcal{L}^1_<$ formulas having one free variable x. Note: this is the familiar correspondence language for the basic temporal language, except that we are using $<$ rather than R as the binary relation symbol.

The *standard translation* ST_x for the until operator U is

$$ST_x(U(\phi,\psi)) \quad = \quad \exists z\,(x < z \wedge ST_z(\phi) \wedge \forall y\,(x < y < z \rightarrow ST_y(\psi))).$$

The standard translation of S is the mirror image of that of U. Observe that we need 3 variables to specify the translation of since and until! We only needed 2 variables to specify the translation of the basic modal operators (see Proposition 2.49).

Let K be a class of models, ML a modal or temporal language, and $\mathcal{L}$ a classical language. Then ML is *expressively complete over* K, if every $\mathcal{L}^1_<(x)$-formula has an equivalent (over K) in the modal language ML. The study of expressive completeness is an important theme in temporal logics with since and until because of the following remarkable result: the language with S and U is expressively complete over the class of all Dedekind complete flows of time (we will define this class shortly). Moreover, below we will define an even richer temporal language that is expressively complete for the class of *all* linear flows of time. In the remainder of this section we will briefly explain these expressive completeness results, and use them to obtain a deductive completeness result for since and until over well-ordered flows of time.

Further preliminaries

A flow of time is called *Dedekind complete* if every subset with an upper bound has a least upper bound. The standard examples are the reals $(\mathbb{R},<)$ and the natural numbers $(\mathbb{N},<)$. A flow of time is *well-ordered* if every non-empty subset has a smallest element; the canonical example here is $(\mathbb{N},<)$.

To arrive at our goal of axiomatizing the well-ordered flows of time, we make a detour through a still richer temporal language built using the *Stavi connectives*.

Definition 7.11 (The Stavi Connectives) To introduce the Stavi connectives we need the notion of a gap. A *gap* of a frame $\mathfrak{F} = (T,<)$ is a proper subset $g \subset T$ which is downward closed (that is, $t \in g$ and $s < t$ implies $s \in g$), and which does not have a supremum. One can think of a gap as a hole in a Dedekind-*in*complete flow of time; see Figure 7.1 Now, $U'(\phi,\psi)$ holds at a point t if the situation depicted in the above figure holds; that is, if

(i) there are a point s and a gap g such that $t \in g$ and $s \notin g$;
(ii) ψ holds between t and g;

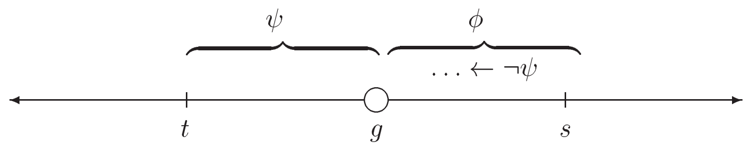

Fig. 7.1. The Stavi connectives

(iii) ϕ holds between s and g; and
(iv) $\neg\psi$ is true arbitrarily soon after g.

$S'(\phi, \psi)$ is the mirror image of $U'(\phi, \psi)$.

The above informal second-order definition (we quantify over gaps, and hence over sets) can be replaced by a first-order definition; see Exercise 7.2.2. $\dashv$

Theorem 7.12 (Expressive Completeness)

(i) U, S *is complete over Dedekind complete flows of time.*
(ii) U, S, U', S' *are complete over all linear flows of time.*

Next, we need an complete axiom system for the class of linear flows of time:

Definition 7.13 Consider the following collection of axioms:

(A1a) $G(p \to q) \to (U(p, r) \to U(q, r))$
(A2a) $G(p \to q) \to (U(r, p) \to U(r, q))$
(A3a) $p \wedge U(q, r) \to U(q \wedge S(p, r), r)$
(A4a) $U(p, q) \wedge \neg U(p, r) \to U(q \wedge \neg r, q)$
(A5a) $U(p, q) \to U(p, q \wedge U(p, q))$
(A6a) $U(q \wedge U(p, q), q) \to U(p, q)$
(A7a) $U(p, q) \wedge U(r, s) \to$
$\quad U(p \wedge r, q \wedge s) \vee U(p \wedge s, q \wedge s) \vee U(q \wedge r, q \wedge s)$
(Aib) the mirror images of (A1a)–(A7a)
(D) $(F\top \to U(\top, \bot)) \wedge (P\top \to S(\top, \bot))$
(L) $H\bot \vee PH\bot$
(W) $Fp \to U(p, \neg p)$
(N) D $\wedge$ L $\wedge$ $F\top$ $\dashv$

Axioms (D), (L), (W), and (N) are discussed in Lemma 7.14 and Exercise 7.2.3 below. As to the other axioms, (A1a) and (A2a) can be viewed as counterparts of the familiar distribution or K axiom $\Box(p \to q) \to (\Box p \to \Box q)$. (A3a) captures the fact that U and S explore relations that are each other's converse. (A4a) and (A5a) connect the current and the future point (at which something good is going to happen) on the one hand with the points in between on the other hand. (A6a) expresses transitivity of the flow of time, and, finally, (A7a) forces the flow of time to be linearly ordered.

Lemma 7.14 *Let $\mathfrak{F}$ be a linear flow of time. Then*

(i) $\mathfrak{F} \models \mathrm{D}$ *iff* $\mathfrak{F}$ *is a discrete ordering.*
(ii) $\mathfrak{F} \models \mathrm{W} \wedge \mathrm{L}$ *iff* $\mathfrak{F}$ *is a well-ordering.*
(iii) $\mathfrak{F} \models \mathrm{W} \wedge \mathrm{N}$ *iff* $\mathfrak{F} \cong (\mathbb{N}, <)$.

The proof of Lemma 7.14 is left as Exercise 7.2.3.

Next, we define three axiom systems: **B**, **BW**, and **BN**. The set of axioms of **B** consists of all classical tautologies, (A1a)–(A7a), and (A1b)–(A7b). **BW** extends **B** with W, and **BN** extends **BW** with N. All three derivation systems have modus ponens, temporal generalization, and uniform substitution as derivation rules:

(MP) If $\vdash \phi$ and $\vdash \phi \to \psi$, then $\vdash \psi$.
(TG) If $\vdash \phi$, then $\vdash G\phi$ and $\vdash H\phi$.
(SUB) If $\vdash \phi$, then $\vdash [\psi/p]\phi$.

A model $\mathfrak{M}$ is called an **X**-*model* if it has $\mathfrak{M} \models \phi$ for all **X**-theorems ϕ, where $\mathbf{X} \in \{\mathbf{B}, \mathbf{BW}, \mathbf{BN}\}$.

For future use we state the following axiomatic completeness result:

Theorem 7.15 *For all sets of S, U-formulas Σ and formulas ϕ: $\Sigma \vdash_{\mathbf{B}} \phi$ iff $\Sigma \models_{\mathbf{B}} \phi$.*

We need one more preliminary result, on definable properties. By Exercise 7.2.4, well-foundedness is a condition on linear frames which cannot be expressed in first-order logic: it involves an essential second-order quantification over all subsets of the universe. However, to arrive at our expressive completeness result we can get by with less, namely the condition that every *first-order definable* non-empty subset must have a smallest element; one can show that definably well-ordered models are sufficiently similar to genuine well-ordered models.

The following definition and lemma capture what we need.

Definition 7.16 Let α be a first-order formula in $\mathcal{L}^1_<(x)$, $\mathfrak{M} = (T, <, V)$ a model for $\mathcal{L}^1_<$. Define X_α to be the set defined by α, that is, $X_\alpha := \{t \in T \mid \mathfrak{M} \models \alpha[t]\}$. Then, $\mathfrak{M}$ is called *definably well-ordered* if for all $\alpha(x) \in \mathcal{L}^1_<$, the set X_α has a smallest element.

Two $\mathcal{L}^1_<$-models $\mathfrak{M}_1$ and $\mathfrak{M}_2$ are called n-*equivalent*, notation $\mathfrak{M}_1 \equiv^n_{FOL} \mathfrak{M}_2$, if for all first-order sentences $\alpha \in \mathcal{L}^1_<$ of quantifier depth at most n, $\mathfrak{M}_1 \models \alpha$ iff $\mathfrak{M}_2 \models \alpha$. $\dashv$

Proviso. For the remainder of this section we will assume that our collection of proposition letters Φ is finite. This is not an essential restriction, but it simplifies some of the arguments below (see Exercise 7.2.5 for a way of circumventing the assumption).

Lemma 7.17 *Let $n \in \mathbb{N}$. Them every definably well-ordered linear model is n-equivalent to a fully well-ordered model.*

Proof. Let $\mathfrak{M} = (T, <, V)$ be a definably well-ordered linear model. For $a, b \in T$ such that $b < a$, define $[b, a) = \{t \in T \mid b \leq t < a\}$, and $(\infty, a) = \{t \in T \mid t < a\}$. Obviously, we can view such sets — with the ordering and valuation induced by $\mathfrak{M}$ — as linear $\mathcal{L}^1_<$-models in their own right. Define

$$Z := \{a \in T \mid \forall b < a\,([b, a) \text{ has a well-ordered } n\text{-equivalent})\}.$$

By Exercise 7.2.6 there are only finitely many first-order formulas $\alpha(x, y)$ of quantifier depth at most n, say $\alpha_1(x, y), \ldots, \alpha_m(x, y)$. Let $\beta_1(x, y), \ldots, \beta_k(x, y) \in \{\alpha_1(x, y), \ldots, \alpha_m(x, y)\}$ be such that if $\mathfrak{M} \models \beta_i(x, y)[ab]$ then $[b, a)$ has a well-ordered n-equivalent. Then Z is defined by the formula

$$\alpha(x) := \forall y \left(y \leq x \to \bigvee_{i \leq k} \beta(x, y) \right).$$

As a consequence, $T \setminus Z$ (the complement of Z in $\mathfrak{M}$) is definable as well. We will now show that $T \setminus Z$ is empty. For, suppose otherwise. Then Z must have a smallest element a (as $\mathfrak{M}$ is definably well-ordered). Distinguish the following cases:

(i) a is the first element of T,
(ii) a has an immediate successor, and
(iii) there exists an ascending sequence $(b_\xi)_{\xi<\lambda}$, which is cofinal in $[b, a)$ and such that $b_0 = b$. (That is, $b_0 = b$, $b_i < b_j$ whenever $i < j$, and for all $c \in [b, a)$ there exists a $b_i > c$.)

It is easy to see that the first two cases lead to contradictions. As to the third case, since a is the minimal element of $T \setminus Z$, all b_ξ are in Z. So, by definition, every interval $[b_\xi, b_{\xi+1})$ has a well-ordered n-equivalent $\mathfrak{M}_\xi$. By Exercise 7.2.7 the lexicographic sum $\sum_{\xi<\lambda} \mathfrak{M}_\xi$ is well-ordered and an n-equivalent to $[b, a)$. But then $a \in Z$ — a contradiction.

Therefore $T \setminus Z = \varnothing$, and hence $Z = T$, so every interval $[b, a)$ of T has an n-equivalent well-ordered model. By using Exercise 7.2.7 again, we see that $\mathfrak{M}$ must have a well-ordered n-equivalent, as required. $\dashv$

Completeness via completeness

With the above preliminaries out of the way, we are now in a position to use the expressive completeness result recorded in Theorem 7.12 to arrive at an axiomatic completeness result for **BW** over well-ordered flows of time.

We need the following lemma.

Lemma 7.18 *Every linear* **BW**-*model is definably well-ordered.*

Proof. Let $\mathfrak{M}$ be a linear model satisfying all instances of the **BW**-theorems. We will prove that every non-empty $\mathcal{L}^1_<$-definable subset of T has a smallest element via detour using the Stavi connectives S' and U'.

Let X be a non-empty $\mathcal{L}^1_<$-definable subset of T. By Theorem 7.12.2 it follows that X has a defining formula ϕ in the language with S, U, S', U'. If we can show that ϕ does in fact belong to the sublanguage with S and U, then we are done, because then we can use the validity of the axioms W and L to show that there must be a *minimal* element in X.

It suffices to show that *every* formula in the language with S, U, S', U' is equivalent to an S, U-formula over $\mathfrak{M}$. To this end we argue by induction of formulas in the richer language. The only non-trivial case is for formulas of the form $U'(\phi, \psi)$ (and their mirror images), where ϕ and ψ are already assumed to equivalent to S, U formulas by the induction hypothesis. So assume $\mathfrak{M}, t \Vdash U'(\phi, \psi)$. Then there is a gap g after t such that (i) ψ holds everywhere between t and g, and (ii) ψ is false arbitrarily soon after g. Now (i) implies that $\mathfrak{M}, t \Vdash F\psi$, so by the validity of the W axiom in $\mathfrak{M}$ it follows that $\mathfrak{M}, t \Vdash U(\neg\psi, \psi)$. But this contradicts (ii). ⊣

Theorem 7.19 **BW** *is (weakly) complete for the class of all well-ordered flows of time.*

Proof. Let ϕ be a **BW**-consistent formula. Construct a maximal **BW**-consistent set Δ with $\phi \in \Delta$. As **BW** extends **B**, Δ must also be **B**-consistent. By Theorem 7.15 there exists a linear model $\mathfrak{M} = (T, <, V)$ in which Δ is satisfiable. Clearly, for every S, U-formula ψ, the formula $H\mathrm{W}(\psi) \wedge \mathrm{W}(\psi) \wedge G\mathrm{W}(\psi)$ is in Δ, where $\mathrm{W}(\psi)$ is the W axiom instantiated for ψ. Thus $\mathfrak{M}$ is a **BW**-model, and hence, by Lemma 7.18 it is definably well-ordered.

Now, for the final step, let n be the quantifier rank of $ST(\phi)$. By Lemma 7.17 there is well-ordered model $\mathfrak{M}'$ that is $n+1$-equivalent to $\mathfrak{M}$. Therefore, $\mathfrak{M}' \models \exists x\, ST(\phi)(x)$, and we are done. ⊣

Using Theorem 7.19 it is easy to obtain a further completeness result, for the temporal logic of the natural numbers.

Theorem 7.20 **BN** *is weakly complete for* $(\omega, <)$*, the natural numbers with their standard ordering.*

The proof of Theorem 7.20 is left as Exercise 7.2.8.

Exercises for Section 7.2

7.2.1 Supply the missing details for the proof of Proposition 7.10.

7.2.2 Give a first-order definition for the Stavi connectives introduced in Definition 7.11 — you may assume that we are working on linear flows of time.

7.2.3 Prove Lemma 7.14. That is, show that D defines discrete orderings, that W $\wedge$ L, defines well-orderings, and that W$\wedge$N picks out the natural numbers in their usual ordering up to isomorphism.

7.2.4 Show that well-foundedness is a condition on linear frames which cannot be expressed in first-order logic.

7.2.5 Throughout this section we assumed that the collection of proposition symbols that we are working with is finite. Show that this assumption can be lifted.

7.2.6 Show that, over a finite vocabulary, there are at only finitely many non-equivalent first-order formulas $\alpha(x, y)$ of quantifier depth at most n

7.2.7 Show that the lexicographic sum of a collection of structures that are well-ordered and n-equivalent to a given structure $\mathfrak{M}$, is again well-ordered and n-equivalent to $\mathfrak{M}$.

7.2.8 Prove Theorem 7.20: show that **BN** is weakly complete for $(\omega, <)$, the natural numbers with their standard ordering.

7.3 Hybrid Logic

An oddity lurks at the heart of modal logic: although states are the cornerstone of modal semantics, they are not directly reflected in modal syntax. We evaluate formulas inside models, at some state, and use the modalities to scan accessible states. But modal syntax offers no grip on the states themselves: it does not let us name them, and it does not let us reason about state equality. Modal syntax and semantics dance to different tunes.

For many applications, this is a drawback. As we mentioned in Example 1.17, both feature and description logics can be viewed as modal logics — or at least, they can up to a point. Real feature logics contain mechanisms for asserting that two sequences of transitions lead to the same state, and description logics allow us to name and reason about individuals. Such capabilities (which are crucial) take us beyond the kinds of modal language we have considered so far. Similarly, it is often important to reason about what is going on at particular times, and the temporal formalisms used in artificial intelligence usually provide expressions such as *Holds*(i, ϕ), asserting that the information ϕ holds at the time named by i, to make this possible. The modal logics considered so far contain no analogs of these important tools.

In their simplest form, *hybrid languages* are modal languages which put this right. Hybrid languages treat states as first class citizens, and they do so in a particularly simple way. The key idea is simply to *sort* the atomic formulas, and to use one sort of atom — the *nominals* — to refer to states. Because this mechanism is so simple, may of the attractive properties of modal logic (such as robust decidability) are unaffected. Indeed, in certain respects hybrid logics are arguably better behaved than their ordinary modal counterparts: their completeness theory is particularly straightforward, and they are proof theoretically natural.

In this section we examine one of the simplest hybrid languages: a two-sorted system with names for states. To build such a language, take a basic modal language (built over propositional variables p, q, r, and so on) and add a second sort of atomic formula. These new atoms are called *nominals*, and are typically written i, j and k. Both types of atom can be freely combined to form more complex formulas in the usual way. For example,

$$\Diamond(i \wedge p) \wedge \Diamond(i \wedge q) \to \Diamond(p \wedge q)$$

is a well formed formula. And now for the key idea: insist that *each nominal be true at exactly one state in any model*. Thus a nominal names a state by being true there and nowhere else. This simple idea gives rise to richer logics. Note, for example, that the previous formula is valid: if the antecedent is satisfied at a state w, then the unique state named by i must be accessible from w, and both p and q must be true there. And note that the use of the nominal i is crucial: if we substituted the ordinary propositional variable r for i, the resulting formula could be falsified.

Actually, what we call the *basic hybrid language* offers more than this: it also enables us to build formulas of the form $@_i\phi$, where i is a nominal. The composite symbol $@_i$ is called a *satisfaction operator*, and it has the following interpretation: $@_i\phi$ is true at any state in a model if and only if ϕ is satisfied at the (unique) state named by i (so $@_i\phi$ is analogous to *Holds*(i, ϕ)). Satisfaction operators play an important role in hybrid proof theory.

Our discussion of basic hybrid logic is largely confined to a single topic: the link between frame definability and completeness. We will show that when *pure* formulas are used as axioms they always yield systems which are complete with respect to the class of frames they define. Now, a pure formula is simply a formula whose only atoms are nominals, so in effect this result tells us that frame completeness is automatic for axioms constructed solely out of names. Our discussion will center on a proof rule called PASTE which is related to the IRR rule discussed in Section 4.7 and the D-rule of Section 7.1.

The basic hybrid language

Given a basic modal language built over propositional variables $\Phi = \{p, q, r, \ldots\}$, let $\Omega = \{i, j, k, \ldots\}$ be a nonempty set disjoint from Φ. The elements of Ω are called *nominals*; they are a second sort of atomic formula which will be used to name states. We call $\Phi \cup \Omega$ the set of *atoms* and define *basic hybrid language* (over $\Phi \cup \Omega$) as follows:

$$\phi ::= i \mid p \mid \bot \mid \neg\phi \mid \phi \wedge \psi \mid \Diamond\phi \mid @_i\phi.$$

For any nominal i, the symbol $@_i$ is called a satisfaction operator. Note that, syntactically speaking, the basic hybrid language is simply a multimodal language (the modalities being $\Diamond$ and all the $@_i$), whose atomic symbols are subdivided into two sorts. If a formula contains no propositional variables (that is, if its only atoms are nominals) we call it a *pure* formula. In what follows we assume that we are working with a fixed basic hybrid language $\mathcal{L}$ in which both Φ and Ω are countably infinite.

The basic hybrid language is interpreted on models. As usual, a model $\mathfrak{M}$ is a triple (W, R, V), where (W, R) is a frame, and V is a valuation. But although the definition of a frame is unchanged, we want nominals to act as names, so we will insist that a valuation V on a frame (W, R) is a function with domain $\Phi \cup \Omega$ and range $\mathcal{P}(W)$ such that for all $i \in \Omega$, $V(i)$ is a *singleton* subset of W. That is, as usual we place no restrictions on the interpretation of ordinary propositional variables, but we insist that a valuation makes each nominal true at a *unique* state. We call the unique state w that belongs to $V(i)$ the *denotation* of i under V. We interpret the basic hybrid language by adding the following two clauses to the satisfaction definition for the basic modal language:

$$\begin{array}{rll} \mathfrak{M}, w \Vdash i & \text{iff} & w \in V(i) \\ \mathfrak{M}, w \Vdash @_i\phi & \text{iff} & \mathfrak{M}, d \Vdash \phi \text{ where } d \text{ is the denotation of } i \text{ under } V. \end{array}$$

As usual, $\mathfrak{M} \Vdash \phi$ means that ϕ is true at all states in $\mathfrak{M}$, $\mathfrak{F} \Vdash \phi$ means that ϕ is valid on the frame $\mathfrak{F}$, and $\Vdash \phi$ means that ϕ is valid on all frames.

Note that a formula of the form $@_i j$ expresses the *identity* of the states named by i and j. Further, note that a formula of the form $@_i\Diamond j$ says that the state named by i has as an R-successor the state named by j.

Although it allows us to refer to states, and talk about state equality, the basic hybrid language is very much a modal language. Nominals name, but they are simply a second sort of atomic formula. Moreover, satisfaction operators are *normal* modal operators: note that for every nominal i, $\Vdash @_i(\phi \to \psi) \to (@_i\phi \to @_i\psi)$, is valid; and if $\Vdash \phi$, then $\Vdash @_i\phi$.

Moreover, the basic hybrid language is quite a simple modal language. For example, its satisfiability problem is known to be no more complex than the satisfiability problem for the basic modal language:

Theorem 7.21 *The satisfiability problem for the basic hybrid logic is PSPACE-complete.*

But in spite of its simplicity the basic hybrid language is surprisingly strong when it comes to frame definability. For a start, many properties definable in the basic modal language can be defined using pure formulas:

(*reflexivity*)	$i \to \Diamond i$
(*symmetry*)	$i \to \Box\Diamond i$
(*transitivity*)	$\Diamond\Diamond i \to \Diamond i$
(*density*)	$\Diamond i \to \Diamond\Diamond i$
(*determinism*)	$\Diamond i \to \Box i$

Moreover, pure formulas also enable us to define many properties *not* definable in the basic modal language, as the reader can easily verify:

(*irreflexivity*)	$i \to \neg\Diamond i$
(*asymmetry*)	$i \to \neg\Diamond\Diamond i$
(*antisymmetry*)	$i \to \Box(\Diamond i \to i)$
(*intransitivity*)	$\Diamond\Diamond i \to \neg\Diamond i$
(*universality*)	$\Diamond i$
(*trichotomy*)	$@_j\Diamond i \vee @_j i \vee @_i\Diamond j$
(*at most 2 states*)	$@_i(\neg j \wedge \neg k) \to @_j k$

All the frame properties defined above are first-order. This is no coincidence: all pure formulas define first-order frame conditions. This is easy to prove: there is a natural way of extending the Standard Translation to cover nominals and satisfaction operators which explains why (see Exercise 7.3.1).

But not only do pure formulas define first-order properties, when used as axioms they are automatically *complete* with respect to the class of frames they define. More precisely, there is a proof system called $\mathbf{K}_h + \text{RULES}$ such that for any set of pure formulas Π:

> If **P** is the normal hybrid logic (which we will shortly define) obtained by adding the formulas in Π as axioms to $\mathbf{K}_h + \text{RULES}$, then **P** is complete with respect to the class of frames defined by P.

The rest of the section is devoted to proving this, but before diving into the technicalities it is worth noting that the result hinges on a rather simple observation. Let us say that a model (W, R, V) is *named* if every state in the model is the denotation of some nominal (that is, for all $w \in W$ there is some nominal i such that $V(i) = \{w\}$). Furthermore, if ϕ is a pure formula, we say that ψ is a *pure instance* of ϕ if ψ is obtained from ϕ by uniformly substituting nominals for nominals. Then we have:

Lemma 7.22 *Let* $\mathfrak{M} = (\mathfrak{F}, V)$ *be a named model and* ϕ *a pure formula. Suppose that for all pure instances* ψ *of* ϕ, $\mathfrak{M} \Vdash \psi$. *Then* $\mathfrak{F} \Vdash \phi$.

Proof. Exercise 7.3.3. $\dashv$

That is, for named models and pure formulas the gap between truth in a model and validity in a frame is non-existent. So if we had a way of building named models, we wouldn't need to appeal to relatively complex syntactic criteria (such as being a Sahlqvist formula) to obtain general completeness results: *any* pure formula would give rise to strongly complete logic for the class of frames it defined. In essence, the work that follows can be summed as follows: we are going to isolate the logic $\mathbf{K}_h + \text{RULES}$ and show that we can build named models from its MCSs and prove an Existence Lemma. Once this is done, a wide range of frame completeness results will be immediate by appeal to Lemma 7.22.

Pure extensions of $\mathbf{K}_h$ + RULES

Let's first say what a normal hybrid logic is:

Definition 7.23 A set of formulas Λ in the basic hybrid language is a *normal hybrid logic* if it contains all tautologies, $\Box(p \to q) \to (\Box p \to \Box q)$, $\Diamond p \leftrightarrow \neg\Box\neg p$, the axioms listed below, and it is closed under the following rules of proof: modus ponens, generalization, $@_i$-*generalization* (if ϕ is provable then so is $@_i\phi$, for any nominal i) and *sorted substitution* (if $\phi \in \Lambda$, and θ results from ϕ by uniformly replacing *propositional letters by arbitrary formulas*, and *nominals by nominals*, then $\theta \in \Lambda$). We call the smallest normal hybrid logic $\mathbf{K}_h$. $\dashv$

The motivation for the sorted substitution rule should be clear: while propositional variables are placeholder for arbitrary information, nominals are names, and substitution must respect the distinction.

The axioms needed to complete our definition of $\mathbf{K}_h$ fall into three groups. The first identifies the basic logic of satisfaction operators:

($\mathrm{K}_@$)	$@_i(p \to q) \to (@_i p \to @_i q)$
(*self-dual*)	$@_i p \leftrightarrow \neg @_i \neg p$
(*introduction*)	$i \wedge p \to @_i p$

As satisfaction operators are normal modal operators, the inclusion of $K_@$ should come as no surprise. As for *self-dual*, note that self-dual modalities are those whose transition relation is a *function*: given the jump-to-the-named-state interpretation of satisfaction operators, this is exactly the axiom we would expect. *Introduction* tells us how to place information under the scope of satisfaction operators. Actually, it also tells us how to get hold of such information, for if we replace p by $\neg p$, contrapose, and make use of *self-dual*, we obtain $(i \wedge @_i p) \to p$; we call this the *elimination* formula.

The second group is a modal theory of naming (or to put it another way, a modal theory of state equality):

(*ref*) $@_i i$
(*sym*) $@_i j \leftrightarrow @_j i$
(*nom*) $@_i j \wedge @_j p \to @_i p$
(*agree*) $@_j @_i p \leftrightarrow @_i p$

Note that the transitivity of naming follows from the *nom* axiom; for example, substituting the nominal k for the propositional variable p yields $@_i j \wedge @_j k \to @_i k$.

The final axiom pins down the interaction between @ and $\Diamond$:

(*back*) $\Diamond @_i p \to @_i p$

Note that $\Diamond i \wedge @_i p \to \Diamond p$ is another valid @-$\Diamond$ interaction principle; it is called *bridge* and we will use it when we prove the Existence Lemma. However *bridge* is provable in $\mathbf{K}_h$ as the reader is asked to show in Exercise 7.3.4.

The soundness of these axioms is clear — but what about completeness? Let us say that a $\mathbf{K}_h$-MCS is *named* if and only if it contains a nominal, and call any nominal belonging to a $\mathbf{K}_h$-MCS a *name* for that MCS. Now, $\mathbf{K}_h$ is strong enough to prove a lemma which is fundamental to our later work: hidden inside any $\mathbf{K}_h$-MCS are a collection of named MCSs with a number of desirable properties:

Lemma 7.24 *Let Γ be a $\mathbf{K}_h$-MCS. For every nominal i, let Δ_i be $\{\phi \mid @_i \phi \in \Gamma\}$. Then:*

(i) *For every nominal i, Δ_i is a $\mathbf{K}_h$-MCS that contains i.*
(ii) *For all nominals i and j, if $i \in \Delta_j$ then $\Delta_j = \Delta_i$.*
(iii) *For all nominals i and j, $@_i \phi \in \Delta_j$ iff $@_i \phi \in \Gamma$.*
(iv) *If k is a name for Γ, then $\Gamma = \Delta_k$.*

Proof. (i) First, for every nominal i we have the *ref* axiom $@_i i$, hence $i \in \Delta_i$. Next, Δ_i is consistent. For assume for the sake of a contradiction that it is not. Then there are $\delta_1, \ldots, \delta_n \in \Delta_i$ such that $\vdash \neg(\delta_1 \wedge \cdots \wedge \delta_n)$. By $@_i$-necessitation, $\vdash @_i \neg(\delta_1 \wedge \cdots \wedge \delta_n)$, hence $@_i \neg(\delta_1 \wedge \cdots \wedge \delta_n)$ is in Γ, and thus by *self-dual* $\neg @_i(\delta_1 \wedge \cdots \wedge \delta_n)$ is in Γ too. On the other hand, as $\delta_1, \ldots, \delta_n \in \Delta_i$, we have $@_i \delta_1, \ldots, @_i \delta_n \in \Gamma$. As $@_i$ is a normal modality, $@_i(\delta_1 \wedge \cdots \wedge \delta_n) \in \Gamma$ as well, contradicting the consistency of Γ. So Δ_i is consistent.

Is Δ_i maximal? Assume it is not. Then there is a formula χ such that neither χ nor $\neg\chi$ is in Δ_i. But then both $\neg @_i \chi$ and $\neg @_i \neg \chi$ belong to Γ, and this is impossible: if $\neg @_i \chi \in \Gamma$, then by self-duality $@_i \neg \chi \in \Gamma$ as well. We conclude that Δ_i is a $\mathbf{K}_h$-MCS named by i.

(ii) Suppose $i \in \Delta_j$; we will show that $\Delta_j = \Delta_i$. As $i \in \Delta_j$, $@_j i \in \Gamma$. Hence, by *sym*, $@_i j \in \Gamma$ too. But now the result is more-or-less immediate. First, $\Delta_j \subseteq \Delta_i$. For if $\phi \in \Delta_j$, then $@_j \phi \in \Gamma$. Hence, as $@_i j \in \Gamma$, it follows by *nom* that $@_i \phi \in \Gamma$, and hence that $\phi \in \Delta_i$ as required. A similar *nom*-based argument shows that $\Delta_i \subseteq \Delta_j$.

(iii) By definition $@_i \phi \in \Delta_j$ iff $@_j @_i \phi \in \Gamma$. By *agree*, $@_j @_i \phi \in \Gamma$ iff $@_i \phi \in \Gamma$. (We call this the *@-agreement property*; it plays an important role in the completeness proof.)

(iv) Suppose Γ is named by k. Let $\phi \in \Gamma$. Then as $k \in \Gamma$, by *Introduction* $@_k \phi \in \Gamma$, and hence $\phi \in \Delta_k$. Conversely, if $\phi \in \Delta_k$, then $@_k \phi \in \Gamma$. Hence, as $k \in \Gamma$, by *elimination* we have $\phi \in \Gamma$. $\dashv$

In what follows, if Γ is a $\mathbf{K}_h$-MCS and i is a nominal, then we will call $\{\phi \mid @_i \phi \in \Gamma\}$ a *named set yielded by* Γ.

We have reached an important crossroad. It is now reasonably straightforward to prove that $\mathbf{K}_h$ is the minimal hybrid logic. We would do so as follows. Given a $\mathbf{K}_h$-consistent set of sentences Σ, use the ordinary Lindenbaum's Lemma to expand it to a $\mathbf{K}_h$-MCS Σ^+, and build a model by taking the submodel of the ordinary canonical model generated by $\Sigma^+ \cup \{\Delta_i \mid \Delta_i \text{ is a named set yielded by } \Sigma^+\}$. The reader is asked to do this in Exercise 7.3.5.

But we have a more ambitious goal in mind: we don't want to build just any model, we want a named model. This will enable us to apply Lemma 7.22 and prove the completeness of pure axiomatic extensions. However we face two problems. The first is this. Given a $\mathbf{K}_h$-consistent set of formula, we can certainly expand it to an MCS using Lindenbaum's Lemma — but nothing guarantees that this MCS will be named. The second problem is much deeper. Suppose we overcame the first problem and learned how to expand any consistent set of sentences Σ to a named MCS Σ^+. Now, as we want to build a named model, this pretty much dictates that only the named MCSs yielded by Σ^+ should be used in the model construction. And now for the tough part: nothing we have seen so far guarantees that there are enough MCSs here to support an Existence Lemma. Incidentally, note that the completeness-via-generation method sketched in the previous paragraph doesn't face this problem: generation automatically gives us all successor MCSs, so we can make use of the ordinary modal Existence Lemma. Unfortunately, not all these successor MCSs need be named, so the generation method won't help with the stronger result we have in mind.

But these difficulties are similar to those we faced when discussing rules for the undefinable, and this suggests a solution. In Section 7.22 we simulated names using tense operators, and used the forward-and-backwards interplay of F and P to create a coherent network of named MCSs which supported a suitable Existence Lemma. Moreover, simulated names were used to define the D-rule mentioned in

Section 7. But nominals are genuine names, and satisfaction operators are an excellent way of enforcing coherence — surely it must be possible to define analogous proof rules for the basic hybrid language? Indeed it is:

$$(\text{NAME})\ \frac{\vdash j \to \theta}{\vdash \theta} \qquad\qquad (\text{PASTE})\ \frac{\vdash @_i\Diamond j \land @_j\phi \to \theta}{\vdash @_i\Diamond\phi \to \theta}$$

In both rules, j is a nominal distinct from i that does not occur in ϕ or θ. The NAME rule is going to solve our first problem, the PASTE rule our second. These rules are clearly close cousins of the IRR rule and the D-rule, but let's defer further discussion till the end of the section, and put them to work right away.

Let $\mathbf{K}_h$ + RULES be the logic obtained by adding the NAME and PASTE rules to $\mathbf{K}_h$. We say that an $\mathbf{K}_h$ + RULES-MCS Γ is *pasted* iff $@_i\Diamond\phi \in \Gamma$ implies that for some nominal j, $@_i\Diamond j \land @_j\phi \in \Gamma$. And now for the key observation: our new rules guarantee we can extend any $\mathbf{K}_h$ + RULES-consistent set of sentences to a *named and pasted* $\mathbf{K}_h$ + RULES-MCS, provided we enrich the language with new nominals:

Lemma 7.25 (Extended Lindenbaum Lemma) *Let Ω' be a (countably) infinite collection of nominals disjoint from Ω, and let $\mathcal{L}'$ be the language obtained by adding these new nominals to $\mathcal{L}$. Then every $\mathbf{K}_h$ + RULES-consistent set of formulas in language $\mathcal{L}$ can be extended to a named and pasted $\mathbf{K}_h$ + RULES-MCS in language $\mathcal{L}'$.*

Proof. Enumerate Ω'. Given a consistent set of $\mathcal{L}$-formulas Σ, define Σ_k to be $\Sigma \cup \{k\}$, where k is the first new nominal in our enumeration. Σ_k is consistent. For suppose not. Then for some conjunction of formulas θ from Σ, $\vdash k \to \neg\theta$. But as k is a new nominal, it does not occur in θ; hence, by the NAME rule, $\vdash \neg\theta$. But this contradicts the consistency of Σ, so Σ_k must be consistent after all.

We now paste. Enumerate all the formulas of $\mathcal{L}'$, define Σ^0 to be Σ_k, and suppose we have defined Σ^m, where $m \geq 0$. Let ϕ_{m+1} be the $m+1$-th formula in our enumeration of $\mathcal{L}'$. We define Σ^{m+1} as follows. If $\Sigma^{m+1} \cup \{\phi_{m+1}\}$ is inconsistent, then $\Sigma^{m+1} = \Sigma^m$. Otherwise:

(i) $\Sigma^{m+1} = \Sigma^m \cup \{\phi_{m+1}\}$ if ϕ_{m+1} is not of the form $@_i\Diamond\phi$. (Here i can be any nominal.)

(ii) $\Sigma^{m+1} = \Sigma^m \cup \{\phi_{m+1}\} \cup \{@_i\Diamond j \land @_j\phi\}$, if ϕ_{m+1} is of the form $@_i\Diamond\phi$. (Here j is the first nominal in the new nominal enumeration that does not occur in Σ^m or $@_i\Diamond\phi$.)

Let $\Sigma^+ = \bigcup_{n\geq 0} \Sigma^n$. Clearly this set is named (by k), maximal, and pasted. Furthermore, it is consistent, for the only non-trivial aspects of the expansion is that defined by the second item, and the consistency of this step is precisely what the PASTE rule guarantees. Note the similarity of this argument to the standard completeness proof for first-order logic: in essence, PASTE gives us the deductive power required to use nominals as Henkin constants. $\dashv$

And now we can define the models we need. In fact, we're basically going to use the named sets examined in Lemma 7.24, but with one small but crucial change: instead of starting with an arbitrary $\mathbf{K}_h$-MCS, we'll insist on using the named sets yielded by a *named and pasted* $\mathbf{K}_h + $ RULES-MCS.

Definition 7.26 Let Γ be a named and pasted $\mathbf{K}_h + $ RULES-MCS. The *named model yielded by* Γ, is $\mathfrak{M}^\Gamma = (W^\Gamma, R^\Gamma, V^\Gamma)$. Here W^Γ is the set of all named sets yielded by Γ, R is the restriction to W^Γ of the usual canonical relation between MCSs (so $R^\Gamma uv$ iff for all formulas ϕ, $\phi \in v$ implies $\Diamond\phi \in u$) and V^Γ is the usual canonical valuation (so for any atom a, $V^\Gamma(a) = \{w \in W^\Gamma \mid a \in w\}$). $\dashv$

Note that $\mathfrak{M}^\Gamma$ really is a *model*: by items (i) and (ii) of Lemma 7.24, V^Γ assigns every nominal a *singleton* subset of W^Γ. And, because we insisted that Γ be named and pasted, we can prove the Existence Lemma we require:

Lemma 7.27 (Existence Lemma) *Let Γ be a named and pasted $\mathbf{K}_h + $ RULES-MCS, and let $\mathfrak{M} = (W, R, V)$ be the named model yielded by Γ. Suppose $u \in W$ and $\Diamond\phi \in u$. Then there is a $v \in W$ such that Ruv and $\phi \in v$.*

Proof. As $u \in W$, for some nominal i we have that $u = \Delta_i$. Hence as $\Diamond\phi \in u$, $@_i\Diamond\phi \in \Gamma$. But Γ is pasted so for some nominal j, $@_i\Diamond j \wedge @_j\phi \in \Gamma$, and so $\Diamond j \in \Delta_i$ and $\phi \in \Delta_j$. If we could show that $R\Delta_i\Delta_j$, then Δ_j would be a suitable choice of v. So suppose $\psi \in \Delta_j$. This means that $@_j\psi \in \Gamma$. By @-agreement (item (iii) of Lemma 7.24) $@_j\psi \in \Delta_i$. But $\Diamond j \in \Delta_i$. Hence, by *Bridge*, $\Diamond\psi \in \Delta_i$ as required. $\dashv$

In short, we have successfully blended the first-order idea of Henkin constants with the modal idea of canonical models, and it's plain sailing all the way to the desired completeness result.

Lemma 7.28 (Truth Lemma) *Let $\mathfrak{M} = (W, R, V)$ be the named model yielded by a named and pasted $\mathbf{K}_h + $RULES-MCS Γ, and let $u \in W$. Then, for all formulas ϕ, $\phi \in u$ iff $\mathfrak{M}, u \Vdash \phi$.*

Proof. Induction on the structure of ϕ. The atomic, boolean, and modal cases are obvious (we use the Existence Lemma just proved for the modalities). What about the satisfaction operators? Suppose $\mathfrak{M}, u \Vdash @_i\psi$. This happens iff $\mathfrak{M}, \Delta_i \Vdash \psi$ (for by items (i) and (ii) of Lemma 7.24, Δ_i is the only MCS containing i, and hence, by the the atomic case of the present lemma, the only state in $\mathfrak{M}$ where i is true) iff $\psi \in \Delta_i$ (inductive hypothesis) iff $@_i\psi \in \Delta_i$ (using the fact that $i \in \Delta_i$ together with *Introduction* for the left-to-right direction and *elimination* for the right-to-left direction) iff $@_i\psi \in u$ (@-agreement). $\dashv$

Theorem 7.29 (Completeness) *Every* $\mathbf{K}_h + \text{RULES}$*-consistent set of formulas in language* $\mathcal{L}$ *is satisfiable in a countable named model. Moreover, if* Π *is a set of pure formulas (in* $\mathcal{L}$*), and* **P** *is the normal hybrid logic obtained by adding all the formulas in* Π *as extra axioms to* $\mathbf{K}_h + \text{RULES}$*, then every* **P***-consistent set of sentences is satisfiable in a countable named model based on a frame which validates every formula in* Π*.*

Proof. For the first claim, given a $\mathbf{K}_h + \text{RULES}$-consistent set of formulas Σ, use the Extended Lindenbaum Lemma to expand it to a named and pasted set Σ^+ in a countable language $\mathcal{L}'$. Let $\mathfrak{M} = (W, R, V)$ be the named model yielded by Σ^+. By item (iv) of Lemma 7.24, because Σ^+ is named, $\Sigma^+ \in W$. By the Truth Lemma, $\mathfrak{M}, \Sigma^+ \Vdash \Sigma$. The model is countable because each state is named by some $\mathcal{L}'$ nominal, and there are only countably many of these.

For the 'moreover' claim, given a **P**-consistent set of formulas Ξ, use the Extended Lindenbaum Lemma to expand it to a named pasted **P**-MCS Ξ^+. The named model $\mathfrak{M}^\Xi$ that Ξ^+ gives rise to will satisfy Ξ at Ξ^+; but in addition, as every formula in Π belongs to every **P**-MCS, we have that $\mathfrak{M}^\Xi \Vdash \Pi$. Hence, by Lemma 7.22, the frame underlying $\mathfrak{M}^\Xi$ validates Π. $\dashv$

Example 7.30 We know that $i \to \neg\Diamond i$ defines irreflexivity and $\Diamond\Diamond i \to \Diamond i$ defines transitivity, hence adding these formulas as axioms to $\mathbf{K}_h + \text{RULES}$ yields a logic (let's call it **I4**) which is complete with respect to the class of strict preorders. Hence $\Diamond\Diamond p \to \Diamond p$, the ordinary modal transitivity axiom, must be **I4**-provable. Furthermore, as $i \to \neg\Diamond\Diamond i$ is valid on any asymmetric frame, and $i \to \Box(\Diamond i \to i)$ is valid on any antisymmetric frame, these must be **I4**-provable too. The reader is asked to supply **I4**-proofs in Exercise 7.3.6. $\dashv$

The PASTE rule has played an pivotal role in our work; is there anything we can say about it apart from 'Hey, it works!'? There is. As we will now see, PASTE is actually a lightly-disguised sequent rule.

A *sequent* is an expression of the form $\Gamma \longrightarrow \Theta$, where Γ and Θ are multisets of formulas (that is, Γ and Θ may contain multiple occurrences of the same formula). Note that the sequent arrow $\longrightarrow$ is longer than the material implication arrow $\to$. Sequents can be read as follows: whenever all the formulas in Γ are true at some state in a model, at least one formula in Θ is true at that state too. A *sequent rule* takes a sequent as input, and returns another sequent as output.

Now, here's PASTE as we stated it above:

$$\frac{\vdash @_i\Diamond j \wedge @_j\phi \to \theta}{\vdash @_i\Diamond\phi \to \theta}$$

Let's get rid of the $\vdash$ symbols and replace the implications by sequent arrows:

$$\frac{@_i\Diamond j \wedge @_j\phi \longrightarrow \theta}{@_i\Diamond\phi \longrightarrow \theta}$$

Splitting the formula in the top line into two simpler formulas yields:

$$\frac{@_i\Diamond j, @_j\phi \longrightarrow \theta}{@_i\Diamond\phi \longrightarrow \theta}$$

This rule works in arbitrary deductive contexts, so let's add a left-hand multiset Γ, and turn θ into a right-hand multiset Θ, thus obtaining:

$$\frac{@_i\Diamond j, @_j\phi, \Gamma \longrightarrow \Theta}{@_i\Diamond\phi, \Gamma \longrightarrow \Theta}$$

But this is just a sequent rule, and a useful one at that. Let's read it from bottom to top: to prove Θ given the information $@_i\Diamond\phi$ and Γ (that's the bottom line) introduce a brand new nominal j and try to prove Θ from $@_i\Diamond j$, $@_j\phi$ and Γ (that's the top line). That is, we should search for a proof by decomposing the formula $@_i\Diamond\phi$ into a near-atomic formula $@_i\Diamond j$ and simpler formula $@_j\phi$. In fact, this decomposition is the key idea needed to define sequent calculi, tableaux, and natural deduction systems for hybrid logics, and several systems which work this way have been developed (see the Notes for details). In effect, such systems discard $\mathbf{K}_h$ from $\mathbf{K}_h$ + RULES (after all, why bother keeping the clumsy Hilbert-style part?) and strengthen the RULES component so it can assume full deductive responsibility.

To conclude, a general remark. As should now be clear (especially if you have already done Exercises 7.3.1, 7.3.2, and 7.3.3), the basic hybrid language is a genuine hybrid between first-order and modal logic: it makes available a number of key first-order capabilities (such as names for states and state-equality assertions) in a decidable (indeed, PSPACE-complete) propositional modal logic. But now that we are used to viewing names as formulas, it is easy to go even further. For example, instead of thinking of nominals as names, we could think of them as variables over states and bind them with quantifiers. For example, we could allow ourselves to form expressions such as

$$\exists x\,(x \wedge \Diamond\exists y(y \wedge \phi \wedge @_x\Box(\Diamond y \to \psi))).$$

This expression captures the effect of the until operator: it says $U(\phi, \psi)$. Note that in this example the $\exists$ quantifier is only used to bind nominals to the *current* state. This is such an important operation that a special notation, $\downarrow$, has been introduced for it. Using this notation the definition of $U(\phi, \psi)$ can be written as

$$\downarrow x(x \wedge \Diamond\downarrow y(y \wedge \phi \wedge @_x\Box(\Diamond y \to \psi))).$$

It turns out that when the basic hybrid language is enriched only with $\downarrow$ (that is, not with the full power of $\exists$) then the resulting language picks out *exactly* the fragment of the first-order correspondence language that is invariant under generated submodels. See the Notes for more details.

Exercises for Section 7.3

7.3.1 Extend the standard translation to the basic hybrid language by adding clauses for nominals and satisfaction operators. Use your translation to show that all classes of frames defined by pure formulas are first-order definable. (Hint: translate nominals to free first-order variables.)

7.3.2 For any $n \geq 1$, let $R^n xy$ be the first-order formula $\exists z_1 \cdots \exists z_n(Rxz_1 \wedge Rz_1z_2 \wedge \cdots \wedge Rz_ny)$. Let ψ be a first-order formula that is a boolean combination of formulas of the form $R^n xy$, Rxy, and $x = y$. Show that the class of frames defined by the universal closure of ψ is definable in the basic hybrid language. (Hint: look at the way we defined trichotomy.)

7.3.3 Prove Lemma 7.22. That is, if $\mathfrak{M} = (\mathfrak{F}, V)$ is a named model and ϕ is a pure formula and for all pure instances ψ of ϕ we have that $\mathfrak{M} \Vdash \psi$, then $\mathfrak{F} \Vdash \phi$.

7.3.4 Show that $\Diamond i \wedge @_i p \to \Diamond p$, the *Bridge* formula, is provable in $\mathbf{K}_h$. (Hint: prove the contraposed form $\Diamond i \wedge \Box p \to @_i p$ with the help of $\Diamond q \wedge \Box p \to \Diamond(q \wedge p)$, *Introduction*, and *Back*.)

7.3.5 Prove that $\mathbf{K}_h$ is the minimal hybrid logic by fleshing out the completeness-via-generation argument sketched in the text.

7.3.6 Find **I4**-proofs of $\Diamond\Diamond p \to \Diamond p$, $i \to \neg\Diamond\Diamond i$, and $i \to \Box(\Diamond i \to i)$. (The logic **I4** was introduced in Example 7.30.)

7.3.7 The PASTE rule makes crucial use of @-operators. Prove an analog of Theorem 7.29 for the @-free sublanguage of the basic hybrid language. (Hint: you need to simulate the satisfaction operators using the modalities. So for all $n, m \geq 0$, add the axiom $\Diamond^n(i \wedge p) \to \Box^m(i \to p)$. Furthermore, let $\Diamond_i \phi$ be shorthand for $\Diamond(i \wedge \phi)$, and add all rules of the form

$$\frac{\vdash \Diamond_k \cdots \Diamond_i \Diamond_j \phi \to \theta}{\vdash \Diamond_k \cdots \Diamond_i \Diamond \phi \to \theta}$$

Here j is a nominal distinct from $k, \cdots, i$ that does not occur in ϕ or θ.)

7.3.8 Let **I4D** be the normal hybrid logic obtained by adding the axiom $\Diamond(i \vee \neg i)$ to **I4**. Clearly **I4D** lacks the finite frame property. Show that it possesses the finite *model* property (and hence that Theorem 3.28 fails for hybrid languages). Exploit this by proving the decidability of **I4D** using a filtration argument.

7.3.9 Add the global diamond E to the basic hybrid language. Use a filtration argument to show that the satisfiability problem the resulting language is decidable. What is its complexity? (Note that $@_i\phi$ can be defined to be $\mathrm{E}(i \wedge \phi)$, so you don't have to deal explicitly with the satisfaction operators.) Show that a class of frames is definable in this language if and only if it is definable in the basic modal language enriched with the D-operator. (Here 'definable' means definable by an arbitrary formula, not just a pure formula.)

7.4 The Guarded Fragment

In Chapter 2 we saw that modal languages can be viewed as fragments of first-order logic, and in Chapter 6 we discovered that these fragments have some nice computational properties. It thus seems natural to try and see how far we can generalize these properties to larger fragments of first-order logic. This will be the main aim of this section: we will define and discuss two extensions of the modal fragment with reasonably nice computational behavior.

In order to isolate such fragments, what properties of the modal fragment of first-order logic should we concentrate on? In particular, what makes modal logic decidable? If we confine ourselves to the basic modal language, is it perhaps the fact that the standard translation can be carried out entirely within the two variable fragment of first-order logic (which has a decidable satisfiability problem)? This argument immediately breaks down if we consider languages with modal operators of higher arity: while giving rise to decidable logics as well, these languages have standard translations that really need *more* than two variables. But as soon as we are considering n-variable fragments of first-order logic with $n > 2$, we face an *undecidable* satisfiability problem.

Rather, it seems to be the fact that the modal fragment of first-order logic allows quantification only in a very restricted form, as is obvious from the modal clause in the definition of the standard translation function:

$$ST_x(\Diamond\phi) = \exists y\,(Rxy \land ST_y(\phi)). \tag{7.3}$$

It is this restricted form of quantification which ensures that modal logic is the bisimulation invariant fragment of first-order logic, and bisimulation invariance of modal truth was critical in the first method of proving the finite model property for the basic modal language (see Section 2.3). Recall that the starting point of this method was the observation that modal logic has the *tree model property* (meaning that every satisfiable modal formula is satisfiable on a tree model), and that bisimulation invariance is pivotal in proving this result. In short, there seems to be a direct line from the restricted quantifier pattern in (7.3), via bisimulation invariance and the tree model property, to the finite model property and decidability.

This provides our first search direction: look for first-order fragments characterized by restricted quantification. It turns out that one can easily relax many constraints applying to the (basic) modal fragment. For example, we do not have to confine ourselves to formulas using two variables only, to formulas having precisely one free variable, or to formulas with predicates of arity at most two. Relaxing these constraints naturally leads to the so-called *guarded fragment* of first-order logic; the idea here is that quantifiers may appear only in the following form:

$$\exists\overline{y}\,(G(\overline{x},\overline{y}) \land \psi(\overline{x},\overline{y})) \tag{7.4}$$

in which $G(\overline{x},\overline{y})$ is an atomic formula that we will call the *guard* of the quantification (or, of the formula). The crucial ingredient that we *keep* from (7.3) is that all free variables of ψ are also free in the guard $G(\overline{x}, \overline{y})$. And indeed, it can be shown that the guarded fragment has various nice properties, such as a decidable satisfiability problem and the finite model property.

However, there are some very natural modal-like languages, or alternative but intuitive interpretations for standard modal languages, that correspond to a decidable fragment of first-order logic as well, but are not covered by this definition. For example, consider the language with the since and until operators: it is straightforward to turn the truth definitions for these operators into a standard translation to first-order logic. The interesting clauses are

$$ST_x(U(\phi, \psi)) = \exists y\,(Rxy \wedge ST_y(\phi) \wedge \forall z\,((Rxz \wedge Rzy) \rightarrow ST_z(\psi))), \quad (*)$$

and a similar one for the since operator. We can prove that this kind of clause takes us outside the guarded fragment of first-order logic: the problem concerns the 'betweenness conjunct' $\forall z\,((Rxz \wedge Rzy) \rightarrow ST_z(\psi))$ which has a 'composite' guard, $(Rxz \wedge Rzy)$. Nevertheless, the language with since and until has a decidable satisfiability problem; apparently, *some* composite guards are admissible as well.

Examples such as $(*)$ lead to extensions of the guarded fragment to fragments in which one is more liberal in the precise conditions imposed on the guard. One *can* be a bit more liberal here because in the 'direct line' mentioned earlier there are some steps that could be skipped on the way. In particular, if we are interested in *decidability* rather than the finite model property, we could just as well settle for fragments of first-order logic to which we may apply the *mosaic method* of Section 6.4. Recall that the mosaic method is a way of proving decidability by 'deconstructing' a model into a finite number of finite pieces, and then using such finite toolboxes for constructing models again, models that usually hang together quite loosely (in a sense to be made precise later). This provides the second direction in our quest: try to find fragments of first-order logic to which the mosaic method applies, leading to a *loose model property*. Implementing this idea one naturally finds quantifier restrictions of the form

$$\exists \overline{y}\,(\pi(\overline{x}, \overline{y}) \wedge \psi(\overline{x}, \overline{y})) \qquad (7.5)$$

in which there are constraints on the presence of variables in certain subformulas of the guard π. For such fragments one may find a direct line from the restricted quantifier pattern in (7.5), via an appropriate notion of bisimulation invariance and the *loose* model property, to some finite *mosaic* property and decidability.

The particular extension that we discuss in this section is that of the *packed fragment*; it fits very nicely in the mosaic approach. On a first reading of the section the reader may choose to skip the parts referring to this packed fragment, and concentrate on the guarded fragment.

The guarded and the packed fragment

We need some preliminaries. The first-order language that we will be working in is purely relational, with equality; the language contains neither constants nor function symbols. For a sequence of variables $\overline{x} = x_1, \ldots, x_n$, we frequently write $\exists\overline{x}\phi$, which, as usual, has the same meaning as $\exists x_1 \ldots \exists x_n \phi$. However, in this section we view $\exists\overline{x}$ not as an abbreviation, but as a primitive operator. In particular this means that the subformulas of $\exists\overline{x}\,\phi$ are just $\exists\overline{x}\,\phi$ itself, together with the subformulas of ϕ. As usual, by writing $\phi(\overline{x})$ we indicate that the free variables of ϕ are among $x_1, \ldots, x_n$.

Definition 7.31 We say that a formula ϕ *packs* a set of variables $\{x_1, \ldots, x_k\}$ if (i) $\mathit{Free}(\phi) = \{x_1, \ldots, x_k\}$ and (ii) ϕ is a conjunction of formulas of the form $x_i = x_j$ or $R(x_{i_1}, \ldots, x_{i_n})$ or $\exists\overline{y}\, R(x_{i_1}, \ldots, x_{i_n})$ such that (iii) for every $x_i \neq x_j$, there is a conjunct in ϕ in which x_i and x_j both occur free.

The *packed fragment* PF is defined as the smallest set of first order formulas which contains all atomic formulas and is closed under the boolean connectives and under *packed quantification*. That is, whenever ψ is a packed formula, π packs $\mathit{Free}(\pi)$, and $\mathit{Free}(\psi) \subseteq \mathit{Free}(\pi)$, then $\exists\overline{x}\,(\pi \wedge \psi)$ is packed as well; π is called the *guard* of this formula. The *guarded fragment* GF is the subfragment of PF in which we only allow *guarded quantification* as displayed in (7.4); that is, packed quantification in which the guard π is an *atomic* formula.

PF_n and GF_n denote the restrictions to n variables and at most n-ary predicate symbols of PF and GF, respectively. $\dashv$

Examples of guarded formulas are

(i) the standard translation of any modal formula (in any language),
(ii) the standard translation of any formula in the basic temporal language,
(iii) formulas like $\forall xy\,(Rxy \to Ryx)$, $\exists xy\,(Rxy \wedge Ryx \wedge (Rxx \vee Ryy))$, $\ldots$

For an example of a packed formula which is not guarded, consider $\exists xyz\,((Rxy \wedge Rxz \wedge Ryz) \wedge \neg Cxyz)$. For another example, first consider the standard translation $\exists y\,(Rxy \wedge Py \wedge \forall z\,((Rxz \wedge Rzy) \to Qz))$ of the formula $U(p,q)$. This formula is not packed *itself*, because the guard of the subformula $\forall z\,((Rxz \wedge Rzy) \to Qz))$ has no conjunct in which the variables x and y occur together. But of course, the formula is *equivalent* to

$$\exists x\,(Rxy \wedge Py \wedge \forall z\,((Rxz \wedge Rzy \wedge Rxy) \to Qz))$$

which is packed. It is not hard to convert this example into a proof showing that *every* formula in the since and until language is equivalent to a packed formula.

Second, note that the notion of packedness only places meaningful restrictions on pairs of *distinct* variables: since the formula $x = x$ packs the set of variables

$\{x\}$, the formula $\exists x(x = x \wedge \psi(x))$, (that is, with a *single* quantification over the variable x) is a packed formula, at least, provided that $\psi(x)$ is packed. Since the given formula is equivalent to $\exists x\, \psi(x)$ this shows that packedness allows a fairly mild form of ordinary quantification, namely over formulas with one free variable only. A nice corollary of this is that we may perform the standard translation of the global diamond E within the two variable guarded fragment:

$$ST_x(\mathsf{E}\phi) = ST_y(\mathsf{E}\phi) = \exists x\, (ST_x(\phi)) \equiv \exists x\, (x = x \wedge ST_x(\phi)).$$

Finally, not *all* formulas are packed, or equivalent to a packed formula. For example, the *transitivity* formula $\forall yz\, ((Rxy \wedge Ryz) \rightarrow Rxz)$ is not packed, and neither is the standard translation of the difference operator: $\exists y\, (x \neq y \wedge Py)$.

Nice properties

Having defined the packed and the guarded fragment of first-order logic, let us see now what we can *prove* about these fragments. To start with, for each of the two fragments we can find a suitable notion of bisimulation which characterizes the fragment in the same way as the ordinary bisimulation characterizes the modal fragment of first-order logic. Unfortunately we do not have the space to go into detail here. Nevertheless, we will show that both fragments have what we call the *loose model property*: in Theorem 7.33 we will show that every satisfiable packed formula can be satisfied on a loose model. What, then, is a loose model?

Definition 7.32 Let $\mathfrak{A} = (A, I)$ be a first-order structure. A tuple $(a_1, \ldots, a_n)$ of objects in A is called *live in* $\mathfrak{A}$ if either $a_1 = \cdots = a_n$ or $(a_1, \ldots, a_n) \in I(P)$ for some predicate symbol P. A subset X of A is called *guarded* if there is some live tuple $(a_1, \ldots, a_n)$ such that $X \subseteq \{a_1, \ldots, a_n\}$. In particular, singleton sets are always guarded; note also that guarded sets are always finite. X is *packed* or *pairwise guarded* if it is finite and each of its two-element subsets is guarded.

We say that $\mathfrak{A}$ is a *loose model of degree* $k \in \mathbb{N}$ if there is some acyclic connected graph $\mathfrak{G} = (G, E)$ and a function f mapping nodes of $\mathfrak{G}$ to subsets of A of size not exceeding k such that for every live tuple $\bar{s}$ from $\mathfrak{A}$, the set $L(\overline{s}) = \{g \in G \mid s_i \in f(g) \text{ for all } s_i\}$, is a non-empty and connected subset of $\mathfrak{G}$. $\dashv$

In words, we call a model $\mathfrak{A} = (A, I)$ loose if we can associate a connected graph $\mathfrak{G} = (G, E)$ with it in the following way. Each node t of the graph corresponds to a *small* subset $f(t)$ of the model; a good way of thinking about this is that t 'describes' $f(t)$. One then requires that the graph 'covers' the entire model in the sense that any $a \in A$ belongs to one of these sets (this follows from the fact that for any $a \in A$, the 'tuple' a is live). The fact that each set $L(\overline{a})$ is connected whenever $\overline{a}$ is live, implies that various nodes of the graph will not give contradictory descriptions of the model. Finally, the *looseness* of the model intuitively stems from the acyclicity of $\mathfrak{G}$ and the connectedness of the sets $L(\overline{a})$; for, this ensures that in walking through the graph we may describe different parts of the model, *but we never have to worry about returning to the same part once we have left it*. Summarizing, we may see the graph as a loose, coherent collection of descriptions of local submodels of the model. Loose models are the ones for which we can find such a graph.

The following result states that the packed fragment of first-order logic has the *loose model property*.

Theorem 7.33 *Every satisfiable packed formula can be satisfied on a loose model (of degree at most the number of $\exists\overline{x}$ subformulas of ξ).*

But the big question is of course whether following this looseness principle we have indeed arrived at a decidable fragment of first-order logic. The next theorem states that we have.

Theorem 7.34 *The satisfiability problems for the guarded and the packed fragment are decidable; both problems are DEXPTIME-complete (complete for doubly exponential time). However, for a fixed natural number n, the satisfiability problem for formulas in the packed fragment PF_n is decidable in EXPTIME.*

And finally, what about the finite model property? Will every satisfiable packed formula have a finite model? Here as well, the packed fragment displays very nice behavior. Unfortunately, we do not have the space for a proof of the finite model property for the packed fragment — suffice it to say that it involves some quite advanced techniques from finite model theory. For some further information the reader is referred to the Notes at the end of the section.

Mosaics

The remainder of the section is devoted to proving the Theorems 7.34 and 7.33. The main idea behind the proof is to use the *mosaic method* that we met in Chapter 6. Roughly speaking, this method is based on the idea of deconstructing models into a finite collection of finite submodels, and conversely, of building up new, ‘loose’, models from such parts. We will see that the packed fragment is in a sense tailored towards making this idea work.

The proof is structured as follows. We start by formally defining mosaics and some related concepts. After that we state the main result concerning the mosaic method, namely the *Mosaic Theorem* stating that a packed formula is satisfiable if and only if there is a so-called *linked* set of mosaics for it, of bounded size. This equivalence enables us to define our decision algorithm and establish the complexity upper bounds mentioned in Theorem 7.34. We then continue to prove the Mosaic Theorem. In doing so we obtain the loose model property for the packed fragment as a spin-off.

For a formal definition of the concept of a mosaic we first need some syntactic preliminaries. Given a first-order formula ξ, we let $Var(\xi)$ and $Free(\xi)$ denote the sets of variables and free variables occurring in ξ, respectively. Let V be a set of variables. A V*-substitution* is any partial map $\sigma : V \to V$. The result of performing the substitution σ on the formula ψ is denoted by ψ^σ. (We can and may assume that such substitutions can be carried out without increasing the total number of variables involved; more precisely, we assume that if $Var(\psi) \subseteq V$ then $Var(\psi^\sigma) \subseteq V$.)

As usual, we will employ a notion of closure to delineate a finite set of *relevant* formulas, that is formulas that for some reason critically influence the truth of a given formula ξ. Let the *single negation* $\sim\phi$ of a formula ϕ denote the formula ψ if ϕ is of the form $\neg\psi$; otherwise, $\sim\phi$ is the formula $\neg\phi$; we say that a set Σ of formulas is closed under single negations if $\sim\phi \in \Sigma$ whenever $\phi \in \Sigma$.

Definition 7.35 Let Σ be a set of packed formulas in the set V of variables. We call Σ V*-closed* if it is closed under subformulas, single negations and V-substitutions (that is, if ψ belongs to Σ, then so does ψ^σ for every V-substitution σ). With $Cl_g(\xi)$ we denote the smallest $Var(\xi)$-closed set of formulas containing ξ. $\dashv$

For the remainder of this section, we fix a packed formula ξ — all definitions to come should be understood as being relativized to ξ. The number of variables occurring in ξ (free or bound) is denoted by k; that is, k is the size of $Var(\xi)$. It can easily be verified that the sets of guarded and packed formulas are both closed under taking subformulas; hence, the set $Cl_g(\xi)$ consists of guarded (packed, respectively) formulas. An easy calculation shows that the cardinality of $Cl_g(\xi)$ is bounded by $k^k \cdot (2|\xi|)$.

The following notion is the counterpart of the atoms that we have met in earlier decidability proofs (see Lemma 6.29, for instance). All three defining conditions are fairly obvious.

Definition 7.36 Let $X \subseteq Var(\xi)$ be a set of variables. An X*-type* is a set $\Gamma \subseteq Cl_g(\xi)$ with free variables in X satisfying, for all formulas $\phi \wedge \psi$, $\sim\phi$, ϕ in $Cl_g(\xi)$ with free variables in X, the conditions (i) $\phi \wedge \psi \in \Gamma$ iff $\phi \in \Gamma$ and $\psi \in \Gamma$, (ii) $\phi \notin \Gamma$ iff $\sim\phi \in \Gamma$ and (iii) if $\phi, x_i = x_j \in \Gamma$ then $\phi^\sigma \in \Gamma$ for any substitution σ mapping x_i to x_j and/or x_j to x_i, while leaving all other variables fixed. $\dashv$

The next definition introduces our key tool in proving the decidability of the packed fragment: mosaics and linked sets of them. Basically, a mosaic consists of a subset X of $Var(\xi)$ together with a set Γ encoding the relevant information on some small part of a model. Here 'small' means that its size is bounded by the number of objects that can be named using variables in X, and 'relevant' refers to all formulas in $Cl_g(\xi)$ whose free variables are in X. It turns out that a finite set of such mosaics contains sufficient information to construct a model for ξ provided that the set links the mosaics together in a nice way. Here is a more formal definition.

Definition 7.37 A *mosaic* is a pair (X, Γ) such that $X \subseteq Var(\xi)$ and $\Gamma \subseteq Cl_g(\xi)$. A mosaic is *coherent* if it satisfies the following conditions:

(C1) Γ is an X-type,
(C2) if $\psi(\overline{x}, \overline{z})$ and $\pi(\overline{x}, \overline{z})$ are in Γ, then so is $\exists\overline{y}\,(\pi(\overline{x}, \overline{y}) \wedge \psi(\overline{x}, \overline{y}))$, (provided that the latter formula belongs to $Cl_g(\xi)$).

A *link* between two mosaics (X, Γ) and (X', Γ') is a renaming (that is, an injective substitution) σ with $\mathrm{dom}\,\sigma \subseteq X$ and $\mathrm{range}\,\sigma \subseteq X'$ which satisfies, for all formulas $\phi \in Cl_g(\xi)$: $\phi \in \Gamma$ iff $\phi^\sigma \in \Gamma'$.

A *requirement* of a mosaic is a formula of the form $\exists\overline{y}\,(\pi(\overline{x}, \overline{y}) \wedge \psi(\overline{x}, \overline{y}))$ belonging to Γ. A mosaic (X', Γ') *fulfills* the *requirement* $\exists\overline{y}\,(\pi(\overline{x}, \overline{y}) \wedge \psi(\overline{x}, \overline{y}))$ of a mosaic (X, Γ) *via* the link σ if for some variables $\overline{u}$, $\overline{v}$ in X' we have that $\sigma(\overline{x}) = \overline{u}$ and $\pi(\overline{u}, \overline{v})$ and $\psi(\overline{u}, \overline{v})$ belong to Γ'. A set S of mosaics is *linked* if every requirement of a mosaic in s is fulfilled via a link to some mosaic in S. S is a linked set of mosaics for ξ if it is linked and $\xi \in \Gamma$ for some (X, Γ) in S. $\dashv$

Note that a mosaic (X, Γ) may fulfill its own requirements, either via the identity map or via some other map from X to X.

The key result concerning mosaics is the following Mosaic Theorem.

Theorem 7.38 (Mosaic Theorem) *Let ξ be a packed formula. Then ξ is satisfiable if and only if there is a linked set of mosaics for ξ.*

Proof. The hard, right to left, direction of the theorem is treated in Lemma 7.39 below; here we only prove the other direction.

Suppose that ξ is satisfied in the model $\mathfrak{A} = (A, I)$. In a straightforward way we can 'cut out' from $\mathfrak{A}$ a linked set of mosaics for ξ. Consider the set of partial assignments of elements in A to variables in $Var(\xi)$. For each such α, let $(X_\alpha, \Gamma_\alpha)$ be the mosaic given by $X_\alpha = \mathrm{dom}\,\alpha$ and

$$\Gamma_\alpha = \{\phi \in Cl_g(\xi) \mid \mathfrak{A} \models \phi[\alpha]\}.$$

We leave it to the reader to verify that this collection forms a linked set of mosaics for ξ. $\dashv$

When establishing the hard direction of this proposition we will in fact prove something stronger: starting from a linked set of mosaics for a formula ξ we will show, via a step by step argument, that there is a *loose* or *tree-like* model for ξ. First however, we want to show that the Mosaic Theorem is the key towards proving the decidability of the packed fragment, and also for finding an upper bound for its complexity.

The decision algorithm and its complexity

The mosaic theorem tells us that any packed formula ξ is satisfiable if and only if there is a linked set of mosaics for ξ. Thus an algorithm answering the question whether a linked set of mosaics exists for ξ, also decides whether ξ is satisfiable. By providing such an algorithm we establish the upper complexity bound for the satisfiability problem of the packed fragment.

Recall that k denotes the number of variables occurring in ξ. The following observations are fairly straightforward consequences of our definitions:

(i) up to isomorphism there are at most $2^k \cdot 2^{2|\xi| \cdot k^k}$ mosaics. Using the big O notation, this is at most $2^{\mathcal{O}(|\xi|) \cdot 2^{k \log k}}$.
(ii) given sets X, Γ with $|X| \leq k$ and $\Gamma \subseteq Cl_g(\xi)$ it is decidable in time polynomial in k^k and $|\xi|$ whether (X, Γ) is a coherent mosaic.
(iii) given a set X of coherent mosaics and a requirement $\phi(\overline{x})$ it is decidable in time polynomial in $|X|$ and $|\phi(\overline{x})|$ whether X fulfills the requirement $\phi(\overline{x})$.

Using methods similar to the elimination of Hintikka sets that we saw in the decidability proof for propositional dynamic logic (see Section 6.8), we now give an algorithm which decides the existence of a linked set of mosaics for ξ. Let S_0 be the set of *all* coherent mosaics. By the observations above, S_0 contains at most $2^{\mathcal{O}(|\xi|) \cdot 2^{k \log k}}$ elements and can be constructed in time polynomial in $|S_0|$. We now inductively construct a sequence of sets of mosaics $S_0 \supseteq S_1 \supseteq S_2 \supseteq S_3 \cdots$. If every requirement of a mosaic μ in a set S_i is fulfilled we call μ *happy*. If every mosaic in S_i is happy then return 'there is a linked set of mosaics for ξ' if S_i contains a mosaic (X, Γ) with $\xi \in \Gamma$, and return 'there is no linked set of mosaics for ξ' otherwise. If, on the other hand, S_i contains *unhappy* mosaics, let S_{i+1} consist of all happy mosaics in S_i and continue the construction. Since our sets decrease in size at every step, the construction must halt after at most $|S_0|$ many stages. By the observations above, computing which states in S_i are happy can be done in time polynomial in ξ and $|S_i|$. Thus the entire computation can be performed in time polynomial in $|S_0|$. Clearly the algorithm is correct.

Hence, if we consider a formula ξ in a packed fragment with a *fixed number of variables*, $|S_0|$ is exponential in $|\xi|$. In general however, the number of variables occurring in a formula depends on the formula's length and hence in general, $|S_0|$

is doubly exponential in $|\xi|$. Thus, pending the correctness of Lemma 7.39 below, this establishes the complexity upper bounds in the Theorem 7.34.

Loose models

Finally, we show the hard direction of the Mosaic Theorem; as a spin-off we establish the 'loose model property' mentioned in Theorem 7.33.

Lemma 7.39 *Let ξ be a packed formula. If there is a linked set of mosaics for ξ, then ξ is satisfiable in a loose model of degree $|Var(\xi)|$.*

Proof. Assume that S is a linked set of mosaics for ξ. Using a step-by-step construction, we will build a loose model for ξ, together with an acyclic graph associated with the model. At each stage of the construction we will be dealing with some kind of approximation of the final model and tree; these approximations will be called networks and are slightly involved structures.

A *network* is a quintuple $(\mathfrak{A}, \mathfrak{G}, \mu, \alpha, \sigma)$ such that $\mathfrak{A} = (A, I)$ is a model for the first-order language; $\mathfrak{G} = (G, E)$ is a connected, directed and acyclic graph; $\mu : G \to S$ is a map associating a mosaic $\mu_t = (X_t, \Gamma_t)$ in S with each node t of the graph; α is a map associating an assignment $\alpha_t : X_t \to A$ with each node t of the graph; and finally, σ is a map associating with each edge (t, t') of the graph a link $\sigma_{tt'}$ from μ_t to $\mu_{t'}$ (we will usually simplify our notation by writing σ instead of $\sigma_{tt'}$).

The idea is that each mosaic μ_t is meant to give a complete description of the relevant requirements that we impose on a small part of the model-to-be. Which part? This is given by the assignment α_t. And the word 'relevant' refers to the fact that we are only interested in the formulas influencing the truth of ξ; that is, the formulas in $Cl_g(\xi)$. The links between neighboring mosaics are there to ensure that distinct mosaics agree on the part of the model that they both have access to.

Now obviously, if we want all of this to work properly we have to impose some conditions on networks. In order to formulate these, we need some auxiliary notation. For a subset $Q \subseteq A$, let $L(Q)$ denote the set of nodes in $\mathfrak{G}$ that have 'access' to Q; formally, we define $L(Q) = \{t \in G \mid A \subseteq \mathrm{range}(\alpha_t)\}$. For a tuple $\overline{a} = (a_1, \ldots, a_n)$ of elements in A we set $L(\overline{a}) = L(\{a_1, \ldots, a_n\})$. Now a network is called *coherent* if it satisfies the following conditions (all to be read universally quantified):

(C1) $P\overline{x} \in \Gamma_t$ iff $\mathfrak{A} \models P\overline{x}[\alpha_t]$,
(C2) $x_i = x_j \in \Gamma_t$ iff $\alpha_t(x_i) = \alpha_t(x_j)$,
(C3) $L(Q)$ is non-empty for every guarded set $Q \subseteq A$,
(C4) $L(Q)$ is connected for every guarded set $Q \subseteq A$,
(C5) if Ett' then $\sigma_{tt'}(x) = x'$ iff $\alpha_t(x) = \alpha_{t'}(x')$.

A few words of explanation about these conditions: (C1) and (C2) ensure that every mosaic is a complete description of the atomic formulas holding in the part of the model it refers to. Condition (C3) states that no live tuple of the model remains unseen from the graph, while the conditions (C4) and (C5) are the crucial ones making that remote parts of the graph cannot contain contradictory information about the model — how this works precisely will become clear further on. Note that condition (C5) has two directions: the left-to-right direction states that neighboring mosaics have common access to part of the model, while the other direction ensures that they agree on their requirements concerning this common part.

The motivation for using these networks is that in the end we want any formula $\phi(\overline{x}) \in Cl_g(\xi)$ to hold in $\mathfrak{A}$ under the assignment α_t if and only $\phi(\overline{x})$ belongs to Γ_t. Coherence on its own is not sufficient to make this happen. A *defect* of a network consists of a formula $\exists\overline{y}\,(\pi(\overline{x},\overline{y}) \wedge \psi(\overline{x},\overline{y}))$ which is a requirement of the mosaic μ_t for some node t while there is no neighboring node t' such that $\mu_{t'}$ fulfills $\exists\overline{y}\,(\pi(\overline{x},\overline{y}) \wedge \psi(\overline{x},\overline{y}))$ via the link $\sigma_{tt'}$. A coherent network $\mathfrak{N}$ is *perfect* if it has no defects. We say that $\mathfrak{N}$ is a network *for* ξ if for some $t \in G$, $\mu_t = (X_t, \Gamma_t)$ is such that $\xi \in \Gamma_t$.

Claim 1 *If $\mathfrak{N} = (\mathfrak{A}, \mathfrak{G}, \mu, \alpha, \sigma)$ is a perfect network, then*

(i) *$\mathfrak{A}$ is a loose model of degree $|\,Var(\xi)|$, and*

(ii) *for all formulas $\phi(\overline{x}) \in Cl_g(\xi)$ and all nodes t of $\mathfrak{G}$: $\phi \in \Gamma_t$ iff $\mathfrak{A} \models \phi[\alpha_t]$.*

Proof of Claim. For part (i) of the claim, let $\mathfrak{N} = (\mathfrak{A}, \mathfrak{G}, \mu, \alpha, \sigma)$ be the perfect network for ξ. Let $\mathfrak{A} = (A, I)$. As the function f mapping nodes of $\mathfrak{G}$ to *subsets* of A, simply take the map that assigns the *range* of α_t to the node t. Since the domain of each map α_t is always a subset of $Var(\xi)$, it follows immediately that $f(t)$ will always be a set of size at most $|\,Var(\xi)|$. Now take an arbitrary live tuple $\overline{s}$ in $\mathfrak{A}$; it follows from (C3) and (C4) that $L(\overline{s})$ is a non-empty and connected part of the graph $\mathfrak{G}$. Thus $\mathfrak{A}$ is a loose model of degree $|\,Var(\xi)|$.

We prove part (ii) of the claim by induction on the complexity of ϕ. For atomic formulas the claim follows by conditions (C1) and (C2), and the boolean case of the induction step is straightforward (since Γ_t is an X–type) and left to the reader. We concentrate on the case that $\phi(\overline{x})$ is of the form $\exists\overline{y}\,(\pi(\overline{x},\overline{y}) \wedge \psi(\overline{x},\overline{y}))$.

First assume that $\phi(\overline{x}) \in \Gamma_t$. Since $\mathfrak{N}$ is perfect there is a node t' in G and variables $\overline{u}$, $\overline{v}$ in $X_{t'}$ such that Ett', $\pi(\overline{u},\overline{v})$ and $\psi(\overline{u},\overline{v})$ belong to $\Gamma_{t'}$, while the link σ from μ_t to $\mu_{t'}$ maps $\overline{x}$ to $\overline{u}$. By the induction hypothesis we find that

$$\mathfrak{A} \models \pi(\overline{u},\overline{v}) \wedge \psi(\overline{u},\overline{v})[\alpha_{t'}]. \tag{7.6}$$

But from condition (C5) it follows that $\alpha_{t'}(\overline{x}) = \alpha_t(\overline{u})$, whence (7.6) implies that

$$\mathfrak{A} \models \exists\overline{y}\,(\pi(\overline{x},\overline{y}) \wedge \psi(\overline{x},\overline{y}))[\alpha_t],$$

which is what we were after.

Now suppose, in order to prove the converse direction, that $\mathfrak{A} \models \phi(\overline{x})[\alpha_t]$. Let $\overline{a}$ denote $\alpha_t(\overline{x})$, then there are $\overline{b}$ in A such that $\mathfrak{A} \models \pi(\overline{x}, \overline{y})[\overline{a}\overline{b}]$ and $\mathfrak{A} \models \psi(\overline{x}, \overline{y})[\overline{a}\overline{b}]$. Our first aims are to prove that

$$L(\overline{a}\overline{b}) \neq \varnothing, \tag{7.7}$$

and

$$L(Q) \text{ is connected for every } Q \subseteq \{\overline{a}, \overline{b}\}. \tag{7.8}$$

Note that if we are working in the guarded fragment, then $\pi(\overline{x}, \overline{y})$ is an atomic formula, whence it follows from $\mathfrak{A} \models \pi(\overline{x}, \overline{y})[\overline{a}\overline{b}]$ that $\overline{a}\overline{b}$ is live. Thus $\{\overline{a}, \overline{b}\}$ is guarded, and hence (7.7) is immediate by condition (C3). In fact, *every* $Q \subseteq \{\overline{a}, \overline{b}\}$ is guarded in this case, so (7.8) is immediate by condition (C4).

In the more general case of the packed fragment we have to work a bit harder. First, observe that it *does* follow from $\mathfrak{A} \models \pi(\overline{x}, \overline{y})[\overline{a}\overline{b}]$ and the conditions on $\pi(\overline{x}, \overline{y})$ in the definition of packed quantification, that $\{c, d\}$ is guarded, and thus, $L(c, d) \neq \varnothing$, for every *pair* (c, d) of points taken from $\overline{a}\overline{b}$. It follows from (C4) that $\{L(c, d) \mid c, d \text{ taken from } \overline{a}\overline{b}\}$ is a collection of non-empty, connected, pairwise overlapping subgraphs of the acyclic graph $\mathfrak{G}$. It is fairly straightforward to prove, for instance, by induction on the size of the graph $\mathfrak{G}$, that any such collection must have a non-empty intersection. From this, (7.7) and (7.8) are almost immediate.

Thus, we may assume the existence of a node t' in $\mathfrak{G}$ such that $\{\overline{a}, \overline{b}\} \subseteq \text{range}\,\alpha_{t'}$. Let $\overline{u}$ and $\overline{v}$ in $X_{t'}$ be the variables such that $\alpha_{t'}(\overline{u}) = \overline{a}$ and $\alpha_{t'}(\overline{v}) = \overline{b}$. The induction hypothesis implies that $\pi(\overline{u}, \overline{v})$ and $\psi(\overline{u}, \overline{v})$ belong to $\Gamma_{t'}$, whence $\phi(\overline{u}) \in \Gamma_{t'}$ by coherence of $\mu_{t'}$. Since both t and t' belong to $L(\overline{a})$, it follows from (7.8) that there is a path from t to t' *within* $L(\overline{a})$, say $t' = s_0 E s_1 E \ldots E s_n = t$. Let σ_i be the link between the mosaics of s_i and s_{i+1}, and define ρ to be the composition of these maps. It follows by an easy inductive argument on the length of the path that ρ is a link between $\mu_{t'}$ and μ_t such that $\rho(\overline{u}) = \overline{x}$. Hence, by definition of a link we have that $\phi(\overline{x}) \in \Gamma_{t'}$. $\dashv$

By Claim 1, in order to prove the Lemma it suffices to construct a perfect network for ξ. This construction uses a step-by-step argument; to start the construction we need *some* coherent network for ξ.

Claim 2 *There is a coherent network for* ξ.

Proof of Claim. By our assumption on ξ there is a coherent mosaic $\mu = (X, \Gamma)$ such that $\xi \in \Gamma$. Without loss of generality we may assume that X is the set $\{x_1, \ldots, x_n\}$ (otherwise, take an isomorphic copy of μ in which X does have this form). Let $a_1, \ldots, a_n$ be a list of objects such that for all i and j we have that $a_i = a_j$ if and only if the formula $x_i = x_j$ belongs to Γ. Define $A = \{a_1, \ldots, a_n\}$

and put the tuple $(a_{i_1}, \ldots, a_{i_k})$ in the interpretation $I(P)$ of the k-ary predicate symbol P precisely if $Px_{i_1} \ldots x_{i_n} \in \Gamma$. Let $\mathfrak{A}$ be the resulting model (A, I) and define $\mathfrak{G}$ as the trivial graph with one node 0 and no edges. Let $\mu(0)$ be the mosaic μ; $\alpha_0 : X \to A$ is given by $\alpha(x_i) = a_i$; and finally, σ_{00} is the identity map from X to X.

We leave it to the reader to verify that the quintuple $(\mathfrak{A}, \mathfrak{G}, \mu, \alpha, \sigma)$ is a coherent network for ξ. $\dashv$

The crucial step of this construction will be to show that any defect of a coherent network can be repaired.

Claim 3 *For any coherent network $\mathfrak{N} = (\mathfrak{A}, \mathfrak{G}, \mu, \alpha, \sigma)$ and any defect of $\mathfrak{N}$ there is a coherent network $\mathfrak{N}^+$ extending $\mathfrak{N}$ and lacking this defect.*

Proof of Claim. Suppose that $\phi(\overline{x})$ is a defect of $\mathfrak{N}$ because it is a requirement of the mosaic μ_t and not fulfilled by any neighboring mosaic $\mu_{t'}$. We will define an extension $\mathfrak{N}^+$ of $\mathfrak{N}$ in which this defect is repaired.

Since S is a linked set of mosaics and μ_t belongs to S, μ_t is linked to a mosaic $(X', \Gamma') \in S$ in which the requirement is fulfilled via some link ρ. Let Y be the set of variables in X' that do not belong to the range of ρ; suppose that $Y = \{y_1, \ldots, y_k\}$ (with all y_i being distinct). For the sake of a smooth presentation, assume that Γ' contains the formulas $\neg x' = y$ for all variables $x' \in X'$ and $y \in Y$ (this is not without loss of generality — we leave the general case as an exercise to the reader). Take a set $\{c_1, \ldots, c_k\}$ of fresh objects (that is, no c_i is an element of the domain A of $\mathfrak{A}$), and let γ be the assignment with domain X' defined as follows:

$$\gamma(x') \;=\; \begin{cases} \alpha_t(x) & \text{if } x' = \rho(x), \\ c_i & \text{if } x' = y_i, \end{cases}$$

and let t' be an object not belonging to G. Now define the network $\mathfrak{N}^+ = (\mathfrak{A}^+, \mathfrak{G}^+, \mu^+, \alpha^+, \sigma^+)$ as follows:

$$\begin{aligned} A^+ &= A \cup \{c_1, \ldots, c_k\}, \\ I^+(P) &= I(P) \cup \{\overline{d} \mid \text{ for some } \overline{x}, \overline{d} = \gamma(\overline{x}) \text{ and } P\overline{x} \in \Gamma'\}, \\ G^+ &= G \cup \{t'\}, \\ E^+ &= E \cup \{(t, t')\}, \end{aligned}$$

while μ^+, α^+ and σ^+ are given as the obvious extensions of μ, α and σ, namely by putting $\mu^+_{t'} = (X', \Gamma')$, $\alpha^+_{t'} = \gamma$ and $\sigma_{tt'} = \rho$.

Since the interpretation I^+ agrees with I on 'old' tuples it is a straightforward exercise to verify that the new network $\mathfrak{N}^+$ satisfies the conditions (C1)–(C3) and (C5).

In order to check that condition (C4) holds, take some guarded subset Q from

A^+; we will show that $L^+(Q)$ is a connected subgraph of $\mathfrak{G}^+$. It is rather easy to see that $L^+(Q)$ is identical to either $L(Q)$ or $L(Q) \cup \{t'\}$; hence by the connectedness of $L(Q)$ it suffices to prove, on the assumptions that $t' \in L^+(Q)$ and $L(Q) \neq \varnothing$, that $t \in L(Q)$. Hence, suppose that $t' \in L^+(Q)$; that is, each $a \in Q$ is in the range of γ. But if $L(Q) \neq \varnothing$, each such point a must be old; hence, by definition of γ, each $a \in Q$ must belong to $\mathrm{range}\,\alpha_t$. This gives that $t \in L(Q)$, as required. $\dashv$

As in our earlier step-by-step proofs, the previous two claims show that using some standard combinatorics we can construct a chain of networks such that their *limit* is a perfect network. This completes the proof of the lemma. $\dashv$

Exercises for Section 7.4

7.4.1 In the *loosely guarded fragment* the following quantification patterns are allowed: $\exists\overline{x}(\pi(\overline{x},\overline{y}) \wedge \psi(\overline{x},\overline{y}))$ is a loosely guarded formula if $\psi(\overline{x},\overline{y})$ is loosely guarded, $\pi(\overline{x},\overline{y})$ is a conjunction as in the packed fragment, and any pair z, z' of distinct variables from $\overline{xy}$ occurs free in some conjunct of the guard π, unless z and z' are both from $\overline{y}$. For example, $\exists x\,((Ryx \wedge Rxy') \wedge \neg Cxyy')$ is loosely guarded, but not packed since there is no conjunct having both y and y' free.

Show that for every loosely guarded *sentence* ξ there exists an equivalent packed sentence ξ' in the same language.

7.4.2 Define the *universal packed fragment* as the fragment of first-order logic that is generated from atoms, negated atoms, conjunction, disjunction, ordinary existential quantification, and packed universal quantification. (With the latter we mean that $\forall\overline{x}(\pi \rightarrow \psi)$ is in the fragment if ψ is universally packed, π packs its own free variables, and $Free(\psi) \subseteq Free(\pi)$.)

Show that *satisfiability* is decidable for the universal packed fragment.

7.4.3 Fix a natural number n, and suppose that we are working in an n-*bounded* first-order signature; that is, all predicate symbols have arity at most n. Prove that in such a signature, every guarded sentence is equivalent to a guarded sentence using at most n variables. Does this hold for packed sentences as well? What are the consequences for the complexity of the respective satisfiability problems?

7.4.4 Let ξ be a packed formula, and suppose that ξ is satisfiable. Prove that ξ is satisfiable in a loose model with an associated graph $\mathfrak{G}$ of which the *out-degree* is bounded by some recursive function on ξ. In particular, this out-degree should be *finite*. (The out-degree of a node k of a graph (G, E) is defined as the number of its neighbors, or, formally, as the size of the set $\{k' \in G \mid kEk'\}$; the out-degree of a graph is defined as the supremum of the out-degrees of the individual nodes.)

7.5 Multi-Dimensional Modal Logic

In Chapter 2 we backed up our claim that logical formalisms do not live in isolation by developing the correspondence theory of modal logic: we studied modal languages as fragments of first-order languages. In this section we will turn the looking glass around and examine first-order logic as if it were a modal formalism. The basic observations enabling this perspective are that we may view *assignments* (the functions that give first-order variables their value in a first-order structure) as *states* of a modal model, and that this makes standard first-order *quantifiers* behave just like modal diamonds and boxes. First-order logic thus forms an example of a *multi-dimensional* modal system. Multi-dimensional modal logic is a branch of modal logic dealing with special relational structures in which the states, rather than being abstract entities, have some inner structure. More specifically, these states are tuples or sequences over some base set, in our case, the domain of the first-order structure. Furthermore, the accessibility relations between these states are (partly) determined by this inner structure of the states.

Reverse correspondence theory

To simplify our presentation, in this section we will not treat modal versions of first-order logic in general, but restrict our attention to certain finite variable fragments. A precise definition of these fragments will be given later on (see Definition 7.40). For the time being, we fix a natural number $n \geq 2$ and invite the reader to think of a first-order language with equality, but without constants or function symbols, in which all predicates are n-adic. Consider the basic declarative statement in first-order logic concerning the truth of a formula in a model under an assignment s:

$$\mathfrak{M} \models \phi\ [s]. \tag{7.9}$$

The basic observation underlying our approach, is that we can read (7.9) from a modal perspective as: 'the formula ϕ is true in $\mathfrak{M}$ *at state* s'. But since we have only n variables at our disposal, say $v_0, \ldots, v_{n-1}$, we can identify assignments with maps: $n\ (= \{0, \ldots, n-1\}) \to U$, or equivalently, with n-tuples over the domain U of the structure $\mathfrak{M}$ — we will denote the set of such n-tuples by U^n. But then we find ourselves in the setting of multi-dimensional modal logic: the universe of our modal models will be of the form U^n for some base set U. Now recall that the truth definition of the quantifiers reads as follows:

$$\mathfrak{M} \models \exists v_i\, \phi[s] \text{ iff there is an } u \in U \text{ such that } \mathfrak{M} \models \phi\ [s^i_u],$$

where s^i_u is the assignment defined by $s^i_u(k) = u$ if $k = i$ and $s^i_u(k) = s(k)$ otherwise. We can replace the above truth definition with the more 'modal' equivalent,

$$\mathfrak{M} \models \exists v_i \phi[s] \text{ iff there is an assignment } s' \text{ with } s \equiv_i s' \text{ and } \mathfrak{M} \models \phi\ [s'],$$

where $\equiv_i$ is given by

$$s \equiv_i s' \text{ iff for all } j \neq i,\ s_j = s'_j. \tag{7.10}$$

In other words: existential quantification behaves like a modal *diamond*, having $\equiv_i$ as its *accessibility relation*.

Since the semantics of the boolean connectives in the predicate calculus is the same as in modal logic, this shows that the inductive clauses in the truth definition of first-order logic neatly fit a modal mould. So let us now concentrate on the atomic formulas. To start with, we observe that *equality* formulas do not cause any problem: the formula $v_i = v_j$, with truth definition

$$\mathfrak{M} \models v_i = v_j[s] \text{ iff } s \in Id_{ij},$$

can be seen as a modal *constant*. Here Id_{ij} is defined by

$$s \in Id_{ij} \text{ iff } s_i = s_j. \tag{7.11}$$

The case of the other atomic formulas is more involved, however. Since we confined ourselves to the calculus of n-adic relations and do not have constants or function symbols, our atomic predicate formulas are of the form $Pv_{\sigma(0)} \dots v_{\sigma(n-1)}$. Here σ is an n-*transformation*, that is, a map: $n \to n$. In the model theory of first-order logic the predicate symbol P will be interpreted as a subset of U^n; but this is precisely how modal valuations treat propositional variables in models where the universe is of the form U^n! Therefore, we can identify the set of propositional variables of the modal formalism with the set of predicate symbols of our first-order language. In this way, we obtain a modal reading of (7.9) for the case where ϕ is the atomic formula $Pv_0 \dots v_{n-1}$: $\mathfrak{M} \models Pv_0 \dots v_{n-1}[s]$ iff s belongs to the interpretation of P. However, as a consequence of this approach our set-up will not enjoy a one-to-one correspondence between atomic first-order formulas and atomic modal ones: the atomic formula $Pv_{\sigma(0)} \dots v_{\sigma(n-1)}$ will correspond to the modal atom p only if σ is the identity function on n. For the cases where σ is not the identity map we still have to find some kind of solution. There are many options here.

Since we are working in a first-order language with equality, atomic formulas with a *multiple* occurrence of a variable can be rewritten as formulas with only 'unproblematic' atomic subformulas, for instance

$$\begin{aligned} Pv_1v_0v_0 &\sim \exists v_2\,(v_2 = v_0 \,\wedge\, Pv_1v_2v_2) \\ &\sim \exists v_2\,(v_2 = v_0 \,\wedge\, \exists v_0\,(v_0 = v_1 \,\wedge\, Pv_0v_2v_2)) \\ &\sim \exists v_2\,(v_2 = v_0 \,\wedge\, \exists v_0\,(v_0 = v_1 \,\wedge\, \exists v_1\,(v_1 = v_2 \,\wedge\, Pv_0v_1v_2))). \end{aligned}$$

This leaves the case what to do with atoms of the form $Pv_{\sigma(0)} \dots v_{\sigma(n-1)}$, where σ is a permutation of n, or in other words, atomic formulas where variables have been substituted *simultaneously*. The previous trick does not work here: for example, to write an equivalent of the formula $Pv_1v_0v_2$ one needs *extra* variables as buffers, for instance, when replacing $Pv_1v_0v_2$ by

$$\exists v_3 \exists v_4 (v_3 = v_0 \wedge v_4 = v_1 \wedge \exists v_0 \exists v_1 (v_0 = v_4 \wedge v_1 = v_3 \wedge Pv_0v_1v_2)).$$

One might consider a solution where a predicate P is translated into *various* modal propositional variables p_σ, one for every permutation σ of n, but this is not very elegant. One might also forget about simultaneous substitutions and confine oneself to a *fragment* of n-variable logic where all atomic predicate formulas are of the form $Pv_0 \dots v_{n-1}$ — this fragment of *restricted* first-order logic is defined below. A third solution is to take substitution seriously, so to speak, by adding special 'substitution operators' to the language. The crucial observation is that for any transformation $\sigma \in n^n$, we have that

$$\mathfrak{M} \models Pv_{\sigma(0)} \dots v_{\sigma(n-1)}[s] \text{ iff } \mathfrak{M} \models Pv_0 \dots v_{n-1}\, [s \circ \sigma], \tag{7.12}$$

where $s \circ \sigma$ is the composition of σ and s (recall that s is a map: $n \to U$). So, if we define the relation $\bowtie_\sigma \subseteq U^n \times U^n$ by

$$s \bowtie_\sigma t \text{ iff } t = s \circ \sigma, \tag{7.13}$$

we have rephrased (7.12) in terms of an accessibility relation (in fact, a function):

$$\begin{aligned} &\mathfrak{M} \models Pv_{\sigma(0)} \dots v_{\sigma(n-1)}[s] \text{ iff} \\ &\quad \mathfrak{M} \models Pv_0 \dots v_{n-1}\, [t] \text{ for some } t \text{ with } s \bowtie_\sigma t. \end{aligned}$$

So if we add an operator $\bigcirc_\sigma$ to the modal language for every n-transformation σ in n^n, with $\bowtie_\sigma$ as its intended accessibility relation, we have found the desired modal equivalent for any atomic formula $Pv_{\sigma(0)} \dots v_{\sigma(n-1)}$, namely in the form $\bigcirc_\sigma p$. (As a special case, for the formula $Pv_0 \dots v_{n-1}$ one can take the identity map on n.)

Definition 7.40 Let n be an arbitrary but fixed natural number. The alphabet of L_n and of L_n^r consists of a set of variables $\{v_i \mid i < n\}$, a countable set of n-adic relation symbols $(P_0, P_1, \dots)$, equality (=), the boolean connectives $\neg, \vee$ and the quantifiers $\exists v_i$. The collection of formulas is defined as usual in first-order logic, with the restriction that the atomic formulas of L_n^r are of the form $v_i = v_j$ or $P_l(v_0 \dots v_{n-1})$; for L_n we allow all atomic formulas (but note that all predicates are of arity n).

A *first-order structure for L_n (L_n^r)* is a pair $\mathfrak{M} = (U, V)$ such that U is a set called the domain of the structure and V is an interpretation function mapping every P to a subset of U^n. The notion of a formula ϕ being *true* in a first-order structure $\mathfrak{M}$ under an assignment s is defined as usual. For instance, given our notation we have, for any atomic formula:

$$\begin{aligned} \mathfrak{M} \models P(v_0 \dots v_{n-1})\, [s] &\quad \text{if} \quad s \in V(P), \\ \mathfrak{M} \models P(v_{\sigma(0)} \dots v_{\sigma(n-1)})\, [s] &\quad \text{if} \quad s \circ \sigma\ (= (s_{\sigma(0)} \dots s_{\sigma(n-1)})) \in V(P). \end{aligned}$$

An L_n-formula ϕ is *true in* $\mathfrak{M}$ (notation: $\mathfrak{M} \models \phi$), if $\mathfrak{M} \models \phi\ [s]$ for all $s \in U^n$;

it is *valid* (notation: $\models_{fo} \phi$), if it is true in every first-order structure of L_n. The same definition applies to L_n^r. $\dashv$

From now on, we will concentrate on the *modal* versions of L_n^r and L_n, which are given in the following definition:

Definition 7.41 Let n be an arbitrary but fixed natural number. MLR_n (short for: *modal language of relations*) is the modal similarity type having constants $\iota\delta_{ij}$ and diamonds $\Diamond_i$, $\bigcirc_\sigma$ (for all $i, j < n, \sigma \in n^n$). CML_n, the similarity type of *cylindric modal logic*, is the fragment of MLR_n-formulas in which no substitution operator $\bigcirc_\sigma$ occurs.

A first-order structure $\mathfrak{M} = (U, V)$ can be seen as a modal model based on the universe ${}^nU$, and formulas of these modal similarity types are interpreted in such a structure in the obvious way; for instance, we have

$$\begin{array}{rll} \mathfrak{M}, s \Vdash \iota\delta_{ij} & \text{iff} & s_i = s_j \\ \mathfrak{M}, s \Vdash \bigcirc_\sigma \phi & \text{iff} & \mathfrak{M}, s \circ \sigma \Vdash \phi \\ & (\text{iff} & \text{there is a } t \text{ with } s \bowtie_\sigma t \text{ and } \mathfrak{M}, t \Vdash \phi) \\ \mathfrak{M}, s \Vdash \Diamond_i \phi & \text{iff} & \text{there is a } t \text{ with } s \equiv_i t \text{ and } \mathfrak{M}, t \Vdash \phi. \end{array}$$

If an MLR_n-formula ϕ holds throughout any first-order structure, we say that it is *first-order valid*, notation: $\mathsf{C}_n \Vdash \phi$ (this notation will be clarified further on). $\dashv$

The modal disguise of L_n in MLR_n and of L_n^r in CML is so thin, that we give the translations mapping first-order formulas to modal ones without further comments.

Definition 7.42 Let $(\cdot)^t$ be the following translation from L_n to MLR_n:

$$\begin{array}{rcl} (Pv_{\sigma(0)} \dots v_{\sigma(n-1)})^t & = & \bigcirc_\sigma p \\ (v_i = v_j)^t & = & \iota\delta_{ij} \\ (\neg\phi)^t & = & \neg\phi^t \\ (\phi \vee \psi)^t & = & \phi^t \vee \psi^t \\ (\exists v_i \phi)^t & = & \Diamond_i \phi^t. \quad \dashv \end{array}$$

This translation allows us to see L_n^r and CML_n as syntactic variants: $(\cdot)^t$ is easily seen to be an *isomorphism* between the formula algebras of L_n^r and CML_n. Note that in the case of L_n versus MLR_n, we face a different situation: where in MLR the simultaneous substitution of two variables for each other is a *primitive* operator, in first-order logic it can only be defined by induction. Nevertheless, we could easily define a translation mapping MLR_n-formulas to equivalent L_n^r-formulas. In any case, the following proposition shows that we really have developed a reverse correspondence theory; we leave the proof as an exercise to the reader.

Proposition 7.43 *Let ϕ be a formula in L_n, then*

(i) *for any first-order structure* $\mathfrak{M}$, *and any* n*-tuple/assignment* s, *we have that* $\mathfrak{M} \models \phi[s]$ *if and only if* $\mathfrak{M}, s \Vdash \phi^t$;

(ii) *as a corollary, we have that* $\models_{fo} \phi \iff \mathsf{C}_n \Vdash \phi^t$.

Let us now put the modal machinery to work and see whether we can find out something new about first-order logic.

Degrees of validity

Perhaps the most interesting aspect of this modal perspective on first-order logic is that it allows us to generalize the semantics of first-order logic, and thus offers a wider perspective on the standard Tarskian semantics. The basic idea is fairly obvious: now that we are talking about *modal* languages, it is clear that the first-order structures of Definition 7.41 are *very specific* modal models for these languages. We may abstract from the first-order background of these models, and consider modal models in which the universe is an *arbitrary* set and the accessibility relations are *arbitrary* relations (of the appropriate arity).

Definition 7.44 A MLR_n*-frame* is a tuple $(W, T_i, E_{ij}, F_\sigma)_{i,j<n,\sigma\in n^n}$ such that every E_{ij} is a subset of the universe W, and such that every T_i and every F_σ is a binary relation on W. A MLR_n*-model* is a pair $\mathfrak{M} = (\mathfrak{F}, V)$ with $\mathfrak{F}$ a MLR_n-frame and V a *valuation*, that is, a map assigning subsets of W to propositional variables. CML_n-models and frames are defined likewise. $\dashv$

For such models, *truth* of a formula at a state is defined via the usual modal induction, for instance:

$$\mathfrak{M}, w \Vdash \bigcirc_\sigma \phi \text{ iff there is a } v \text{ with } F_\sigma wv \text{ and } \mathfrak{M}, v \Vdash \phi.$$

In this very general semantics, states (that is, elements of the universe) are no longer real assignments, but rather, abstractions thereof. First-order logic now really has become a poly-modal logic, with quantification and substitution diamonds. It is interesting and instructive to see how familiar laws of the predicate calculus behave in this new set-up. For example, the axiom schema $\phi \to \exists v_i \phi$ will be valid only in n-frames where T_i is a reflexive relation (this follows from the fact that the modal formula $p \to \Diamond_i p$ corresponds to the frame condition $\forall x T_i xx$). Likewise, the axiom schemes $\exists v_i \exists v_i \phi \to \exists v_i \phi$ and $\phi \to \forall v_i \exists v_i \phi$ will be valid only in frames where the relation T_i is transitive and symmetric, respectively.

Later on we will see more of such correspondences; the point to be made here is that the abstract perspective on the semantics of first-order logic imposes a certain 'degree of validity' on well-known theorems of the predicate calculus. Some theorems are valid in *all* abstract assignment frames, like distribution:

$$\forall v_i(\phi \to \psi) \to (\forall v_i \phi \to \forall v_i \psi),$$

which is nothing but the modal K-axiom. Other theorems of the predicate calculus, like the ones mentioned above, are only valid in *some* classes of frames. Narrowing down the class of frames means increasing the set of valid formulas, and vice versa. In particular, we now have the option to look at classes of frames that are only slightly more general than the standard first-order structures, but have much nicer computational properties. This new perspective on first-order logic, which was inspired by the literature on algebraic logic, provides us with enormous freedom to play with the semantics for first-order logic. In particular, consider the fact that first-order structures can be seen as frames of the form $(U^n, \equiv_i, Id_{ij}, \bowtie_\sigma)_{i,j<n,\sigma\in n^n}$ where *all* assignments $s \in U^n$ are available. But why not study a semantics where states are still real assignments on the base set U, but *not all such assignments are available?*

There are at least two good reasons to make such a move. First, it turns out that the logic of such generalized assignment frames has much nicer meta-properties than the logic of the cubes such as decidability, see for instance Theorem 7.46 below. These logics will provide less laws than the usual predicate calculus, but their supply of theorems may be sufficient for particular applications. Note for instance, that the schemes $\phi \to \exists v_i \phi$, $\exists v_i \exists v_i \phi \to \exists v_i \phi$ and $\phi \to \forall v_i \exists v_i \phi$ are still valid in every generalized assignment frame, since $\equiv_i \restriction_W$ is always an equivalence relation.

In some situations it may even be *useful* not to have all familiar validities. Consider for instance the schema

$$\exists v_i \exists v_j\, \phi \to \exists v_j \exists v_i\, \phi. \tag{7.14}$$

It follows from correspondence theory that (7.14) is valid in a frame $\mathfrak{F}$ iff (7.15) below holds in $\mathfrak{F}$.

$$\forall xz\, (\exists y\, (T_i xy \wedge T_j yz) \to \exists u\, (T_j xu \wedge T_i uz)). \tag{7.15}$$

The point is that the schema (7.14) disables us to make the dependency of variables explicit in the language (that is, whether v_j is dependent of v_i or the other way around), while these dependencies play an important role in some proof-theoretical approaches. So, the second motivation for generalizing the semantics of first-order logic is that it gives us a finer sieve on the notion of equivalence between first-order formulas. Note for instance that (7.14) is not valid in frames with assignment 'holes': take $n = 2$. In a square (that is, 2-cubic) frame we have $(a, b) \equiv_0 (a', b) \equiv_1 (a', b')$, but if (a, b') is not an available tuple, then there is no s such that $(a, b) \equiv_1 s \equiv_0 (a', b')$ — hence this frame will not satisfy (7.15). So, the schema (7.14) will not be valid in this frame.

In this new paradigm, a whole landscape of frame classes and corresponding logics arises. In the most general approach, *any* subset of U^n may serve as the universe of a multi-dimensional frame, but it seems natural to impose restrictions on the set of available assignments. Unfortunately, for reasons of space limitations we cannot go into further detail here, confining ourselves to the following definition.

Definition 7.45 Let U be some set, and W a set of n-tuples over U, that is, $W \subseteq U^n$. The *cube over* U or *full assignment frame over* U is defined as the frame

$$\mathfrak{C}_n(U) = (U^n, \equiv_i, Id_{ij}, \bowtie_\sigma)_{i,j<n,\sigma\in n^n}.$$

The W*-relativized cube over* U or W-assignment frame on U is defined as the frame

$$\mathfrak{C}_n^W(U) = (W, \equiv_i\restriction_W, Id_{ij} \cap W, \bowtie_\sigma\restriction_W)_{i,j<n,\sigma\in n^n}.$$

C_n and R_n are the classes of cubes and relativized cubes, respectively. $\dashv$

Observe that this definition clarifies our earlier notation '$\mathsf{C}_n \Vdash \phi$' for the fact that the modal formula ϕ is 'first-order valid'.

Decidability

As we already mentioned, one of the reasons for developing the abstract and generalized assignment semantics is to 'tame' first-order logic by looking for core versions with nicer computational behavior. This idea is substantiated by the following theorem.

Theorem 7.46 *It is decidable in exponential time whether a given MLR_n-formula is satisfiable in a given relativized cube. As a corollary, the problem whether a given first-order formula in L_n can be satisfied in a general assignment frame is also decidable in exponential time.*

Proof. This theorem can be proved directly by using the *mosaic method* that we encountered in Section 6.4 — in fact, the mosaic method was developed for this particular proof! However, space limitations prevent us from giving the mosaic argument here. Therefore, we prove the theorem by a reduction of the R_n satisfiability problem to the satisfiability problem of the n-variable *guarded fragment* of Section 7.4.

This reduction is quite interesting in itself: the key idea is that we find a syntactic counterpart to the semantic notion of restricting the set of *available* assignments. There is in fact a very simple way of doing so, namely by introducing a special n-adic predicate G that will be interpreted as the collection of available assignments. One can then translate modal formulas (or L_n-formulas) into first-order ones, with the proviso that this translation is *syntactically relativized* to G. The formula $Gv_0 \ldots v_{n-1}$ so to speak acts as a *guard* of the translated formula, and indeed, it will be easily seen that the range of this translation formally falls inside the guarded fragment.

Now for the technical details. Given a collection Φ of propositional variables, assume that with each $p \in \Phi$ we have an associated n-adic predicate symbol P. Also, fix a *new* n-adic predicate symbol G; let Φ^+ denote the expanded signature $\{P \mid p \in \Phi\} \cup \{G\}$. Consider the following translation $(\cdot)^\bullet$ mapping MLR_n-formulas to first-order formulas:

$$\begin{array}{rcl} p^\bullet & = & Pv_0 \dots v_{n-1} \\ \iota\delta_{ij}^\bullet & = & v_i = v_j \\ (\neg\phi)^\bullet & = & Gv_0 \dots v_{n-1} \wedge \neg\phi^\bullet \\ (\phi \wedge \psi)^\bullet & = & \phi^\bullet \wedge \psi^\bullet \\ (\bigcirc_\sigma \phi)^\bullet & = & (Gv_0 \dots v_{n-1} \wedge \phi^\bullet)^\sigma \\ (\Diamond_i \phi)^\bullet & = & \exists v_i (Gv_0 \dots v_{n-1} \wedge \phi^\bullet) \end{array}$$

Here, for a given transformation σ, $(\cdot)^\sigma$ denotes the corresponding syntactic substitution operation on first-order formulas.

We want to show the following claim.

Claim 1 *For any MLR_n-formula ϕ we have that $\mathsf{R}_n \Vdash \phi$ if and only if the formula $Gv_0 \dots v_{n-1} \to \phi^\bullet$ is a first order validity.*

Proof of Claim. In order to prove this claim, we need a correspondence between modal models and first-order models for the new language. Given a relativized assignment model $\mathfrak{M} = (\mathfrak{C}_n^W(U), V)$, define the corresponding first-order model $\mathfrak{M}^\bullet$ as the structure (U, I) where $I(P) = V(p)$ for every propositional variable p, and $I(G) = W$. Conversely, given a first-order structure $\mathfrak{A} = (A, I)$ for the expanded first-order signature Φ, let $\mathfrak{A}_\bullet$ be the relativized cube model $(\mathfrak{C}_n^{I(G)}(A), V)$, where the valuation V is given by $V(p) = I(P)$.

For any relativized assignment model $\mathfrak{M}$, and any available assignment s, we have

$$\mathfrak{M}, s \Vdash \phi \text{ iff } \mathfrak{M}^\bullet \models \phi^\bullet[s]. \tag{7.16}$$

This suffices to prove Claim 1, because of the following. First suppose that the modal formula ϕ is satisfiable in some relativized cube model $\mathfrak{M}$, say at state s. Since s is an available tuple, it follows from (7.16) that $\phi^\bullet$ is satisfiable in the first-order structure $\mathfrak{M}^\bullet$ under the assignment s; but also, since s is available we have $\mathfrak{M}^\bullet \models Gv_0 \dots v_{n-1}[s]$. This shows that $\phi^\bullet \wedge Gv_0 \dots v_{n-1}$ is satisfiable.

Conversely, if the latter formula is satisfiable, there is some first-order structure $\mathfrak{A}$ for the language Φ^+, and some assignment s such that $\mathfrak{A} \models \phi^\bullet \wedge Gv_0 \dots v_{n-1}[s]$. It is not difficult to see that $(\mathfrak{A}_\bullet)^\bullet = \mathfrak{A}$. Since $\mathfrak{A} \models Gv_0 \dots v_{n-1}[s]$, it follows by definition that s is an available assignment of $\mathfrak{A}_\bullet$. But then we may apply (7.16) which yields that $\mathfrak{A}_\bullet, s \Vdash \phi$; in particular, ϕ is satisfiable in R_n. The proof of (7.16) proceeds by a standard induction, which we leave to the reader. $\dashv$

Finally, we leave it to the reader to verify that the range of $(\cdot)^\bullet$ indeed falls entirely inside the n-variable guarded fragment $\mathfrak{F}_n$. From Claim 1 and this observation the theorem is immediate. $\dashv$

Axiomatization

To finish off the section we will sketch how to prove completeness for the class of cube models. For simplicity we confine ourselves to the similarity type of cylindric modal logic — but observe that this completeness result will immediately transfer to the restricted n-variable fragment L^r_n.

Multi-dimensional modal logic is an area with a very interesting completeness theory. For instance, if one only admits the standard modal derivation rules (modus ponens, necessitation and uniform substitution), then *finite* axiomatizations are few and far between. For instance, concerning the CML_n-theory of the class C_n, Andréka proved that if Σ is a set of CML_n-formulas axiomatizing C_n, then for each natural number m, Σ contains infinitely many formulas that contain all diamonds $\Diamond_i$, at least one diagonal constant $\iota\delta_{ij}$ and at least m propositional variables... However, if we allow special derivation rules, in the style of Section 4.7, then a nice finite axiomatization can be obtained, as we will see now. A key role in our axiomatization and in our proof will be played by a defined operator $\mathrm{D}_n p$ which acts as the *difference operator* on the class of cube frames, see Section 7.1. For its definition we need some auxiliary operators:

$$\begin{array}{rcll} \bigcirc_{ij}\phi & = & \Diamond_i(\iota\delta_{ij} \wedge \phi) & \qquad (i \neq j) \\ \mathrm{E}^n_i\phi & = & \Diamond_0 \ldots \Diamond_{i-1}\Diamond_{i+1} \ldots \Diamond_{n-1}\phi & \\ \mathrm{D}_n\phi & = & \bigvee_{j\neq i} \bigcirc_{ji}\Diamond_i(\neg\iota\delta_{ij} \wedge \mathrm{E}^n_i\phi). & \end{array}$$

The definition of D_n may look fairly complex, but it is directly based on the observation that two n-tuples s and t are *distinct* if and only for some coordinate i, s_i is distinct from t_i.

Proposition 7.47 D_n *acts as the difference operator on the class of cubes.*

Proof. Let $\mathfrak{M} = (\mathfrak{C}_n(U), V)$ be a cube model. We will show that

$$\mathfrak{M}, s \Vdash \mathrm{D}_n p \text{ iff } \mathfrak{M}, t \models p \text{ for some } t \text{ such that } s \neq t. \tag{7.17}$$

For the sake of a clear exposition we assume that $n = 3$, so that we may write $s = (s_0, s_1, s_2)$.

For the left to right direction of (7.17), suppose that $\mathfrak{M}, s \Vdash \mathrm{D}_n p$. Without loss of generality we may assume that $s \Vdash \bigcirc_{10}\Diamond_0(\neg\iota\delta_{01} \wedge \mathrm{E}^n_0 p)$. By definition of $\bigcirc_{10}$ it follows that $(s_0, s_0, s_2) \Vdash \Diamond_0(\neg\iota\delta_{01} \wedge \mathrm{E}^n_0 p)$. This in its turn implies that there is some s'_0 such that $(s'_0, s_0, s_2) \Vdash \neg\iota\delta_{01}$ and $(s'_0, s_0, s_2) \Vdash \mathrm{E}^n_0 p$. It is easily seen that the meaning of E^n_0 is given by

$$\mathfrak{M}, u \Vdash \mathrm{E}^n_i\psi \text{ iff } \mathfrak{M}, v \models \psi \text{ for some } v \text{ such that } u_i = v_i,$$

so $(s'_0, s_0, s_2) \Vdash \mathrm{E}^n_0 p$ means that there is some n-tuple t such that $t \Vdash p$ and $s'_0 = t_0$. But it follows from $(s'_0, s_0, s_2) \Vdash \neg\iota\delta_{01}$ that $s_0 \neq s'_0$, so that we find that $t_0 \neq s_0$. But then, indeed, t is distinct from s. We leave it to the reader to prove the right to left direction of (7.17). $\dashv$

However, the connection between D_n and the class of cubes is far tighter than this Proposition suggests. In fact, the cubes are the *only* frames on which D_n acts as the difference operator, at least, against the right background of the class $\mathsf{HCF_n}$ of *hypercylindric frames*.

Definition 7.48 A CML_n-frame is called *hypercylindric* if the following formulas are valid on it:

$(CM1_i)$ $\quad p \to \Diamond_i p$
$(CM2_i)$ $\quad p \to \Box_i \Diamond_i p$
$(CM3_i)$ $\quad \Diamond_i \Diamond_i p \to \Diamond_i p$
$(CM4_{ij})$ $\quad \Diamond_i \Diamond_j p \to \Diamond_j \Diamond_i p$
$(CM5_i)$ $\quad \iota\delta_{ii}$
$(CM6_{ij})$ $\quad \Diamond_i(\iota\delta_{ij} \wedge p) \to \Box_i(\iota\delta_{ij} \to p)) \quad (i \neq j)$
$(CM7_{ijk})$ $\quad \iota\delta_{ij} \leftrightarrow \Diamond_k(\iota\delta_{ik} \wedge \iota\delta_{kj}) \quad (k \notin \{i, j\})$
$(CM8_{ij})$ $\quad (\iota\delta_{ij} \wedge \Diamond_i(\neg p \wedge \Diamond_j p)) \to \Diamond_j(\neg\iota\delta_{ij} \wedge \Diamond_i p) \quad (i \neq j)$ $\dashv$

All these axioms are Sahlqvist formulas and thus express first-order properties of frames. Clearly, the axioms *CM1–3* together say that each T_i is an equivalence relation. $CM6_{ij}$ then means that in every T_i-equivalence class there is *at most one* element on the diagonal E_{ij} ($i \neq j$). One can combine this fact with the (first-order translations of) $CM5_j$ and $CM7_{jji}$ to show that every T_i-equivalence class contains *exactly* one representative on the E_{ij}-diagonal. Apart from this effect, the contribution of *CM7* is rather technical. Finally, the meaning of *CM4* and *CM8* is best made clear by Figure 7.2 below, where the straight lines represent the antecedent of the first-order correspondents, and the dotted lines, the relations holding of the 'old' states and the 'new' ones given by the succedent.
The key theorem in our completeness proof is the following.

Theorem 7.49 *For any frame $\mathfrak{F}$ in* $\mathsf{HCF_n}$, D_n *acts as the difference operator on $\mathfrak{F}$ if* and only if *$\mathfrak{F}$ is a cube.*

Proof. We have already proved the left to right direction of this equivalence in Proposition 7.47. The proof of the other direction is technically rather involved and falls outside the scope of this book. $\dashv$

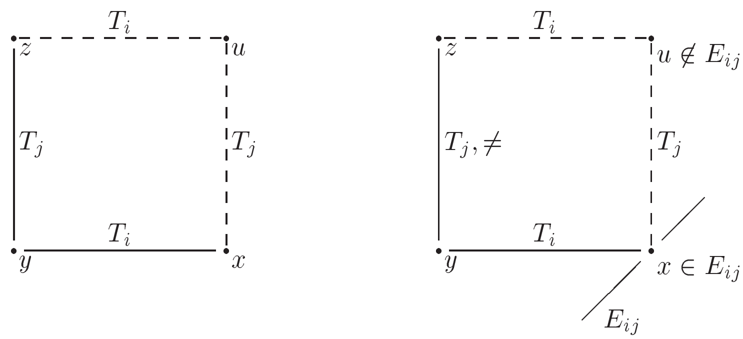

Fig. 7.2. The meaning of $CM4_{ij}$ (left) and $CM8_{ij}$ (right)

In fact, with Theorem 7.49 we have all the material in our hands to prove the desired completeness result.

Definition 7.50 Consider the following modal derivation system Ω_n. Its axioms are (besides the ones of the minimal modal logic for the similarity type CML_n), the formulas $CM1$–8; as its derivation rules we take, besides the standard ones, also the D_n-rule:

$$\frac{\vdash (p \wedge \neg \mathrm{D}_n p) \to \theta}{\vdash \theta}$$

As usual, Ω_n will also denote the *logic* generated by this derivation system. ⊣

Theorem 7.51 *Ω_n is sound and strongly complete with respect to the class C_n.*

Proof. It follows immediately from Theorem 7.6 and Theorem 7.49 that we obtain a complete axiomatization for C_n if we extend Ω_n with the D_n-versions of the axioms *Symmetry*, *Pseudo-transitivity* and D-*Inclusion*. However, as its turns out, these axioms are valid on the class of hypercylindric frames, so they are already derivable in Ω_n (even without the use of the D_n-rule). From this, the theorem is immediate. ⊣

Exercises for Section 7.5

7.5.1 Let n and m be natural numbers such that $n < m$, and consider a CML_n-formula ϕ. First, observe that ϕ is also a CML_m-formula. Prove that $\mathsf{C}_n \Vdash \phi$ iff $\mathsf{C}_m \Vdash \phi$. Conclude that our definition of an MLR_n-formula being *first order valid*, is unambiguous.

7.5.2 Prove that the formula $\Diamond_0 \cdots \Diamond_{n-1} p$ acts as the global modality on the class of hypercylindric frames. That is, show that for any model $\mathfrak{M}$ based on such a frame we have that

$$\mathfrak{M}, s \Vdash \Diamond_0 \cdots \Diamond_{n-1} p \text{ iff } \mathfrak{M}, t \Vdash p \text{ for some } t \text{ in } \mathfrak{M}.$$

Which of the axioms CM1–8 are actually needed for this?

7.5.3 Let L_n^- denote the equality-free fragment of L_n^r; that is, all atomic formulas are of the form $Pv_0 \dots v_{n-1}$. In an obvious way we can define relativized assignment frames for this language. Prove that the satisfiability problem for L_n^- in this class of frames can be solved in PSPACE.

7.5.4 Prove that every hypercylindric CML_2-frame is the bounded morphic image of a square frame (that is, a 2-cube). Use this fact to find a complete axiomatization for the class C_2 that only uses the standard modal derivation rules.

7.5.5 Let CF_n be the class of cylindric frames, that is, those CML_n-frames that satisfy the axioms CM1–7. The class of *n-dimensional cylindric algebras* is defined as $\mathsf{CA}_n = \mathbb{SP}\mathbf{Cm}\mathsf{CF}_n$. The classes HCF_n and HCA_n are defined similarly, now using all axioms CM1–8.

(a) Prove that CA_n and HCA_n are canonical, that is, closed under taking canonical embedding algebras.
(b) Prove that CA_n and HCA_n are varieties.

7.5.6 A *full n-dimensional cylindric set algebra* is an algebra of the form

$$(\mathcal{P}(U^n), \cup, -, \varnothing, C_i, Id_{ij})_{i,j<n}.$$

Here the *i-th cylindrification* is defined as the map $C_i : \mathcal{P}(U^n) \to \mathcal{P}(U^n)$ given by

$$C_i(X) = \{s \in U^n \mid t \in X \text{ for some } t \text{ in } X \text{ with } s \equiv_i t\}.$$

If we close the class of these algebras under products and subalgebras, we arrive at the variety RCA_n of *representable n-dimensional cylindric algebras*.

(a) Prove that every representable n-dimensional cylindric algebra is a boolean algebra with operators.
(b) Prove that RCA_n is contained in the classes CA_n and HCA_n of the previous exercise.
(c) Prove that RCA_n is canonical. (Hint: use Theorem 7.49 to show that the class C_n of n-dimensional cubes is first-order definable in the frame language of CML_n.)

7.6 A Lindström Theorem for Modal Logic

Throughout this book we have seen many examples of modal languages, especially in the present chapter. To get a clear picture of the emerging spectrum, these languages may be classified according to their expressive power or their semantic properties. But what — if any — is the special status of the familiar modal languages defined in Chapter 1. If we focus on characteristic semantic properties, then clearly their invariance under bisimulations must be a key feature. But what else is needed to single the out (standard) modal languages?

The answer to this question is a modal analogue of a classic result in first-order model theory: Lindström's Theorem. It states that, given a suitable explication of what 'classical logic' is, first-order logic is the strongest logic to possess the

Compactness and Löwenheim-Skolem properties. To prove an analogous characterization result for modal logic we need to agree on a number of things:

- What will be the distinguishing property of the logic that we want to characterize (on top of its invariance for bisimulations)? To answer this question we will exploit the notion of degree introduced in Definition 2.28.
- What is a suitable notion of an abstract modal logic? To answer this question we will introduce some bookkeeping properties from the formulation of the original Lindström Theorem for first-order logic, and add a further property having to do with invariance under bisimulations.

Our plan for this section is to discuss each of the above items, one after the other, and to conclude with a Lindström Theorem for modal logic.

Background material

Throughout this section models for modal languages are *pointed models* of the form $(\mathfrak{M}, w)$, where $\mathfrak{M}$ is a relational structure and w is an element of $\mathfrak{M}$ (its *distinguished point*) at which evaluation takes place.

Our main reasons for adopting this convention are the following. First, the basic semantic unit in modal logic simply *is* a structure together with a distinguished node at which evaluation takes place. Second, some of the results below admit smoother formulations when we adopt the *local* perspective of pointed models.

Bisimulations between pointed models $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ are required to link the distinguished points w and v.

Definition 7.52 (In-degree) Let τ be a modal similarity type, and let $\mathfrak{M}$ be a τ-model. The *in-degree* of a state u in $\mathfrak{M}$ is the number of times u occurs as an non-first argument in a relation: $Rw \ldots u \ldots$. More formally, it is defined as

$$|\{\vec{w} \in \mathfrak{M}^{<\omega} \mid \text{ for some } R \text{ and } i > 1,\, u = w_i \text{ and } R^{\mathfrak{M}} w_1 \ldots w_i \ldots w_n)\ \}|. \quad \dashv$$

In addition to the in-degree of an element of a model, we will also need to use the notion of *height* as defined in Definition 2.32.

Below we will want to get models that have nice properties, such as a low in-degree or finite height for each of its elements. To obtain such models, the notion of forcing comes in handy. Fix a similarity type τ. A property P of models is $\underline{\leftrightarrow}_\tau$*-enforceable*, or *enforceable*, iff for every pointed τ-model $(\mathfrak{M}, w)$, there is a pointed τ-model $(\mathfrak{N}, v)$ with $(\mathfrak{M}, w) \mathrel{\underline{\leftrightarrow}_\tau} (\mathfrak{N}, v)$ and $(\mathfrak{N}, v)$ has P.

For example, the property 'every element has finite height' is enforceable. To see this, let $(\mathfrak{M}, w)$ be a pointed τ-model; we may assume that $\mathfrak{M}$ is generated by w. Let $(\mathfrak{N}, w)$ be the submodel of $\mathfrak{M}$ whose domain consists of all elements of finite height. Then $(\mathfrak{M}, w) \mathrel{\underline{\leftrightarrow}_\tau} (\mathfrak{N}, w)$.

Proposition 7.53 below generalizes the *unraveling construction* from the standard modal language to arbitrary vocabularies.

Proposition 7.53 *The following properties of models are enforceable:*

(i) *tree-likeness, and*
(ii) *the conjunction of 'having a root with in-degree 0' and 'every element (except the root) has in-degree at most 1'.*

Proof. Item (ii) follows from item (i). A proof of item (i) for similarity types only involving diamonds is given in Proposition 2.15; for the general case, consult Exercise 2.1.7. ⊣

We will characterize modal logic (in the sense of Definitions 1.12 and 1.23) by showing that it is the only modal logic satisfying a modal counterpart of the original Lindström conditions: having a notion of *finite degree* which gives a fixed upper bound on the height of the elements that need to be considered to verify a formula; recall Definition 2.28 for the definition.

To wrap up our discussion of background material needed for our Lindström Theorem, let us briefly recall some basic facts related to degrees and height. Here's the first of these facts; recall that $((\mathfrak{M}, w) \upharpoonright n, w)$ denotes the submodel of $\mathfrak{M}$ that is generated from w and that only has states of height at most n.

Proposition 7.54 *Let ϕ be a modal formula with $\deg(\phi) \leq n$. Then $(\mathfrak{M}, w) \Vdash \phi$ iff $((\mathfrak{M}, w) \upharpoonright n, w) \Vdash \phi$.*

Next, recall from Proposition 2.29 that, up to logical equivalence, there are only finitely many non-equivalent modal formulas with a fixed finite degree over a finite similarity type.

We say that $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ are *n-equivalent* if w and v satisfy the same modal formulas of degree at most n.

Proposition 7.55 *Let τ be a finite similarity type. Let $(\mathfrak{M}, w)$, $(\mathfrak{N}, v)$ be two rooted models such that the roots have in-degree 0, every element different from the root has in-degree at most 1, all nodes have and height at most n.*

If $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ are $n+1$-equivalent, then $(\mathfrak{M}, w) \leftrightarroweq (\mathfrak{N}, v)$.

Proof. Define $Z \subseteq A \times B$ by xZy iff:

$$\mathrm{height}(x) = \mathrm{height}(y) = m \text{ and } (\mathfrak{M}, x) \text{ and } (\mathfrak{N}, y) \text{ are } (n-m)\text{-equivalent.}$$

We claim that $Z : (\mathfrak{M}, w) \leftrightarroweq (\mathfrak{N}, v)$. To prove this, we only show the forth condition. Assume xZy and $R^{\mathfrak{M}} x x_1 \ldots x_k$, where $\mathrm{height}(x) = \mathrm{height}(y) = m$. Then $n - m \geq 1$. Let $\triangle$ be the modal operator whose semantics is based on R.

As τ is finite, there are only finitely many non-equivalent formulas of degree at most $n-m-1$. Let ψ_i be the conjunction of all non-equivalent modal formulas of at most this degree that are satisfied at x_i ($1 \leq i \leq k$). Then $(\mathfrak{M}, x) \Vdash \triangle(\psi_1, \ldots, \psi_k)$, and $\triangle(\psi_1, \ldots, \psi_k)$ has degree $n-m$. Hence, as xZy, $(\mathfrak{N}, y) \Vdash \triangle(\psi_1, \ldots, \psi_k)$. So there are $y_1, \ldots, y_k$ in $\mathfrak{N}$ such that $R^{\mathfrak{N}}yy_1 \ldots y_k$ and $(\mathfrak{N}, y_i) \Vdash \psi_i$ ($1 \leq i \leq k$).

Now, as all states have in-degree at most 1, $\mathrm{height}(x_i) = \mathrm{height}(y_i) = m+1$, and $(\mathfrak{M}, x_i)$ and $(\mathfrak{N}, y_i)$ ($1 \leq i \leq k$) are $(n-(m+1))$-equivalent. Hence, $(\mathfrak{M}, x_i) \leftrightarroweq_\tau (\mathfrak{N}, y_i)$. This proves the forth condition. $\dashv$

Abstract modal logic

The original Lindström Theorem for first-order logic starts from a definition of an abstract classical logic as a pair $(\mathcal{L}, \models_\mathcal{L})$ consisting of a set of formulas $\mathcal{L}$ and a satisfaction relation $\models_\mathcal{L}$ between $\mathcal{L}$-structures and $\mathcal{L}$-formulas that satisfies three bookkeeping conditions, an Isomorphism property, and a Relativization property which allows one to consider definable submodels. Then, an abstract logic extending first-order logic coincides with first-order logic if, and only if, it satisfies the Compactness and Löwenheim-Skolem properties. We will now set up our modal analogue of Lindström's Theorem along similar lines.

The definition runs along the same lines as the definition of an abstract classical logic. An abstract modal logic is characterized by three properties: two book keeping properties, and a Bisimilarity property to replace the Isomorphism property.

Definition 7.56 (Abstract Modal Logic) By an *abstract modal logic* we mean a pair $(\mathcal{L}, \Vdash_\mathcal{L})$ with the following properties (here $\mathcal{L}$ is the set of formulas, and $\Vdash_\mathcal{L}$ is its satisfaction relation, that is, a relation between (pointed) models and $\mathcal{L}$-formulas):

(i) ***Occurrence property.*** For each ϕ in $\mathcal{L}$ there is an associated finite language $\mathcal{L}(\tau_\phi)$. The relation $(\mathfrak{M}, w) \Vdash_\mathcal{L} \phi$ is a relation between $\mathcal{L}$-formulas ϕ and structures $(\mathfrak{M}, w)$ for languages $\mathcal{L}$ containing $\mathcal{L}(\tau_\phi)$. That is, if ϕ is in $\mathcal{L}$, and $\mathfrak{M}$ is an $\mathcal{L}$-model, then the statement $(\mathfrak{M}, w) \Vdash_\mathcal{L} \phi$ is either true or false if $\mathcal{L}$ contains $\mathcal{L}(\tau_\phi)$, and undefined otherwise.

(ii) ***Expansion property.*** The relation $(\mathfrak{M}, w) \Vdash_\mathcal{L} \phi$ depends only on the reduct of $\mathfrak{M}$ to $\mathcal{L}(\tau_\phi)$. That is, if $(\mathfrak{M}, w) \Vdash_\mathcal{L} \phi$ and $(\mathfrak{N}, w)$ is an expansion of $(\mathfrak{M}, w)$ to a larger language, then $(\mathfrak{N}, v) \Vdash_\mathcal{L} \phi$.

(iii) ***Bisimilarity property.*** The relation $(\mathfrak{M}, w) \Vdash_\mathcal{L} \phi$ is preserved under bisimulations: if $(\mathfrak{M}, w) \leftrightarroweq_\tau (\mathfrak{N}, v)$ and $(\mathfrak{M}, w) \Vdash_\mathcal{L} \phi$, then $(\mathfrak{N}, v) \Vdash_\mathcal{L} \phi$. $\dashv$

If we compare the above definition to the list of properties defining an abstract classical logic, we see that it's the Bisimilarity property that determines the *modal* character of an abstract modal logic.

Obviously, ordinary modal formulas provide an example of an abstract modal logic, but so does propositional dynamic logic. In contrast, the language of basic temporal logic provides an example of a logic that is *not* an abstract modal logic, as formulas from basic temporal logic are not preserved under bisimulations.

Next, we need to say what we mean by '$(\mathcal{L}, \Vdash_{\mathcal{L}})$ extends basic modal logic' and by closure under negation.

Definition 7.57 We say that $(\mathcal{L}, \Vdash_{\mathcal{L}})$ *extends modal logic* if for every basic modal formula there exists an equivalent $\mathcal{L}$-formula, that is, if for each basic modal formula ϕ there exists an $\mathcal{L}$-formula ψ such that for any model $(\mathfrak{M}, w)$ we have $(\mathfrak{M}, w) \Vdash \phi$ iff $(\mathfrak{M}, w) \Vdash_{\mathcal{L}} \psi$.

Also, $(\mathcal{L}, \Vdash_{\mathcal{L}})$ is *closed under negation* if for all $\mathcal{L}$-formulas ϕ there exists an $\mathcal{L}$-formula $\neg\phi$ such that for all models $(\mathfrak{M}, w)$, $(\mathfrak{M}, w) \Vdash \phi$ iff $(\mathfrak{M}, w) \nVdash \neg\phi$. $\dashv$

Of course, propositional dynamic logic is an example of an abstract modal logic that extends (basic) modal logic.

Logics in the sense of Definition 7.56 deal with the same class of pointed models as (basic) modal logic, and only the formulas and satisfaction relation may be different. This implies, for example, that intuitionistic logic or the hybrid logics considered in Section 7.3 are not abstract modal logics: their models need to satisfy special constraints. The original Lindström characterization of first-order logic suffers from similar limitations (by not allowing ω-logic as a logic, for example).

As a final step in our preparations, we need to say what the notion of degree means in the setting of an abstract modal logic.

Definition 7.58 (Notion of Finite Degree) An abstract modal logic has a *notion of finite degree* if there is a function $\deg_{\mathcal{L}} : \mathcal{L} \to \omega$ such that for all $(\mathfrak{M}, w)$, all ϕ in $\mathcal{L}$,

$$(\mathfrak{M}, w) \Vdash_{\mathcal{L}} \phi \quad \text{iff} \quad ((\mathfrak{M}, w) \restriction \deg_{\mathcal{L}}(\phi)), w \Vdash_{\mathcal{L}} \phi.$$

If $\mathcal{L}$ extends (basic) modal logic, we assume that $\deg_{\mathcal{L}}$ behaves regularly with respect to standard modal operators and proposition letters. That is, if $\triangle$ is a modal operator (see Definition 1.12), then $\deg_{\mathcal{L}}(p) = 0$ and $\deg_{\mathcal{L}}(\triangle(\phi_1, \ldots, \phi_n)) = 1 + \max\{\deg_{\mathcal{L}}(\phi_i) \mid 1 \leq i \leq n\}$.

Finally, two models $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ for the same language are *$\mathcal{L}$-equivalent* if for every ϕ in $\mathcal{L}$, $(\mathfrak{M}, w) \Vdash \phi$ iff $(\mathfrak{N}, v) \Vdash \phi$. $\dashv$

Having a finite degree is a very restrictive property, which is not implied by the finite model property (FMP). To see this recall that propositional dynamic logic has the FMP: it has the property that every satisfiable formula ϕ is satisfiable on a model of size at most $|\phi|^3$, where ϕ is the length of ϕ. However, it does not have a notion of finite degree. To see this, consider the model (ω, R_a, V), where R_a is the successor relation and V is an arbitrary valuation, and let $\phi = [a^*]\langle a\rangle\top$; clearly $(\omega, R_a, V), 0 \Vdash \phi$. But for no $n \in \omega$ does the restriction $(\omega, R_a, V) \upharpoonright n$ satisfy ϕ at 0. It follows that **PDL** does not have a notion of finite degree.

Characterizing modal logic

We are almost ready now to prove our characterization result. The following lemma is instrumental.

Lemma 7.59 *Let $(\mathcal{L}, \Vdash_{\mathcal{L}})$ be an abstract modal logic which is closed under negation. Assume $\mathcal{L}$ has a notion of finite degree $\deg_{\mathcal{L}}$. Let ϕ be an $\mathcal{L}$-formula with $\deg_{\mathcal{L}}(\phi) = n$. Then, for any two models $(\mathfrak{M}, w)$, $(\mathfrak{N}, v)$ such that $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ are n-equivalent, we have that $(\mathfrak{M}, w) \Vdash_{\mathcal{L}} \phi$ implies $(\mathfrak{N}, v) \Vdash_{\mathcal{L}} \phi$.*

Proof. Assume that the conclusion of the lemma does not hold. Let $(\mathfrak{M}, w)$, $(\mathfrak{N}, v)$ be such that $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ are n-equivalent, but $(\mathfrak{M}, w) \Vdash_{\mathcal{L}} \phi$ and $(\mathfrak{N}, v) \Vdash_{\mathcal{L}} \neg\phi$.

By the Occurrence and Expansion properties we may assume that $\mathcal{L} = \mathcal{L}(\tau_\phi)$, where $\mathcal{L}(\tau_\phi)$ is the finite language in which ϕ lives.

By Proposition 7.53 we can assume that $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ are rooted such that the roots have in-degree 0, while all other nodes have in-degree at most 1. Then $((\mathfrak{M}, w) \upharpoonright n, w)$ and $((\mathfrak{N}, v) \upharpoonright n, v)$ are n-equivalent, and $((\mathfrak{M}, w) \upharpoonright n, w) \Vdash_{\mathcal{L}} \phi$ but $((\mathfrak{N}, v) \upharpoonright n, v) \Vdash_{\mathcal{L}} \neg\phi$. In addition $((\mathfrak{M}, w) \upharpoonright n, w)$ and $((\mathfrak{N}, v) \upharpoonright n, v)$ both have in-degree 1 and roots of in-degree 0. By Proposition 7.55 it follows that $((\mathfrak{M}, w) \upharpoonright n, w)$ and $((\mathfrak{N}, v) \upharpoonright n, v)$ are bisimilar — but now we have a contradiction with the Bisimilarity property as $((\mathfrak{M}, w) \upharpoonright n, w)$ and $((\mathfrak{N}, v) \upharpoonright n, v)$ are bisimilar but don't agree on ϕ. $\dashv$

Theorem 7.60 *Let $(\mathcal{L}, \Vdash_{\mathcal{L}})$ extend modal logic. If $(\mathcal{L}, \Vdash_{\mathcal{L}})$ has a notion of finite degree, then it is equivalent to the modal language as defined in Definition 1.12.*

Proof. We must show that every $\mathcal{L}$-formula ϕ is $\mathcal{L}$-equivalent to a basic modal formula ψ, that is, for all $(\mathfrak{M}, w)$, $(\mathfrak{M}, w) \Vdash_{\mathcal{L}} \phi$ iff $(\mathfrak{M}, w) \Vdash_{\mathcal{L}} \psi$. As before, by the Occurrence and Expansion properties we may restrict ourselves to a finite language. Moreover, ϕ has a basic modal equivalent iff it has such an equivalent with the same degree; so we have to locate the equivalent we are after among the basic modal formulas whose degree equals the $\mathcal{L}$-degree of ϕ.

Assume $n = \deg_{\mathcal{L}}(\phi)$. By Proposition 2.29 there are only finitely many (non-equivalent) basic modal formulas whose degree equals n; assume that they are all contained in Γ_n. It suffices to show the following

if $(\mathfrak{M}, w)$ and $(\mathfrak{N}, v)$ agree on all formulas in Γ_n, then they agree on ϕ. (7.18)

For then, ϕ will be equivalent to a Boolean combination of formulas in Γ_n. To see this, reason as follows. The relation 'satisfies the same formulas in Γ_n' is an equivalence relation on the class of all models; as Γ_n is finite, there can only be finitely many equivalence classes. Choose representatives $(\mathfrak{M}_1, w_1)$, ..., $(\mathfrak{M}_m, w_m)$, and for each i, with $1 \leq i \leq m$, let ψ_i be the conjunction of all formulas in Γ_n that are satisfied by $(\mathfrak{M}_i, w_i)$. Then ϕ is equivalent to $\bigvee\{\psi_i \mid (\mathfrak{M}_i, w_i) \Vdash_{\mathcal{L}} \phi\}$.

Now to conclude the proof of the theorem we need only observe that condition (7.18) is exactly the content of Lemma 7.59. $\dashv$

To conclude this section a few remarks are in order. First, the property of having a notion of finite degree can be characterized algebraically in terms of preservation under ultraproducts over the natural numbers; Theorem 7.60 can then be reformulated accordingly.

Second, in the proof of the Lindström Theorem the basic modal formula ψ that is found as the equivalent of the abstract modal formula ϕ is in the same vocabulary as ϕ. This means, for example, that the only abstract modal logic over a binary relation that has a notion of finite degree is the standard modal logic with a single modal operator $\Diamond$.

Here, we have only covered the modal logics as defined in Definition 1.12; in some cases extensions beyond this pattern can easily be obtained. As a first example, consider the basic temporal language with operators F and P, where $x \Vdash Fp$ ($x \Vdash Pp$) iff for some y, Rxy and $y \Vdash \phi$ (Ryx and $y \Vdash \phi$). Consider *temporal bisimulations* in which one not only looks forward along the binary relation, but also backward, and adopt the notion of height accordingly. Given the obvious definition of an *abstract temporal logic*, standard temporal logic is the only temporal logic over a single binary relation that has a notion of finite degree.

7.7 Summary of Chapter 7

- *Logical Modalities*: Logical modalities receive a fixed interpretation in every model. Simple examples are the past tense operator P, the global diamond E, and the difference operator D. As well a enhancing expressivity, some of them (notably P and D) make it possible to prove general completeness theorems using additional rules of proof.
- *Algebra of Diamonds*: Some modal languages offer not just a single logical modality but an entire algebra of diamonds. Good examples are PDL and BML.
- *Since and Until*: The since and until operators are interesting in applied logic because they enable us to specify guarantee properties. They are mathematically

interesting because they are expressively complete over Dedekind complete total orders.

- *Completeness-via-Completeness*: While deductive completeness of since and until logic can be proved using standard modal techniques, for Dedekind complete total order there is an interesting alternative: taking a detour via expressive completeness.
- *Hybrid Logic*: The basic hybrid language lets us refer to states using nominals, atomic symbols true at exactly one state in every model. Some stronger hybrid languages allow us to bind nominals.
- *Hybrid Proof Theory*: We can define a rule of proof called PASTE in the basic hybrid language. This rule is essentially a sequent rule lightly disguised. With its help, a frame completeness result covering all pure formulas can be proved fairly straightforwardly.
- *Guarded fragment*: As the standard translation shows, modalities are essentially macros which permit restricted forms of quantification. Abstracting from this insight leads to the guarded fragment, a decidable fragment of first-order logic with the final model property.
- *Packed Fragment*: By taking this observation even further, and noting that the mosaic method suffices to prove decidability, it is possible to isolate an even larger decidable fragment of first-order logic: the packed fragment. This fragment also has the finite model property.
- *Multi-Dimensional Modal Logic*: Multi-dimensional modal logic is essentially modal logic in which evaluation is performed at a sequence of states, rather than at a single state. By viewing variable assignments as sequence of states, it is possible to view first-order logic itself as a multi-dimensional modal logic.
- *Lindström's Theorem*: Given a suitable (bisimulation centered) explication of what an abstract modal logic is, our Lindström Theorem for modal logic says that the general modal languages defined in Definition 1.12 are the strongest ones to have a notion of finite degree.
- *Extended Modal Logic*: In many ways, this chapter is badly named. Among other things, we've just seen that not only it is possible to introduce globality, more complex quantifier alternations in satisfaction definitions, names for states, and evaluation at sequences of states, but we can do so without losing the properties that made modal logic attractive in the first place. So forget the 'extended'. As we said in the Preface: *it's all just modal logic*!

Notes

A really serious guide to extended modal logic would have to cover the (vast) literature on temporal logics, fixed point logics, and variants of PDL discussed in the theoretical computer science literature, plus formalisms such as feature and description logic, and much else besides. We don't have space to do all that, and the following Notes stick to the six topics discussed in the text. Nonetheless, with the help of the following remarks (coupled with a little judicious reference chasing) the reader should be able to form a coherent map of territory.

Logical Modalities. It's hard to precise about when the idea of adding fixed interpretation operators to modal languages came to be seen as standard. Certainly the writings of Johan van Benthem (for example, his book on temporal logic, his 'manual' on intensional logic, and his influential survey of correspondence theory) played an important role. So did the new applications of modal logic, particularly in computer science (once you've seen PDL it's hard to believe that the basic modal language is the be-all and end-all of modal logic). At any rate, by the end of the 1980s the idea that modal languages are abstract tools for talking about relational structures — tools that it was not only legitimate, but actually *interesting* to extend — was well established in both Amsterdam and Bulgaria. Nowadays this view is taken for granted by many (perhaps most) modal logicians, and given this perspective the use of logical modalities is as natural as breathing.

Of course, many of the operators we now call 'logical' have been around a lot longer than that. In a way, the global modality has always been there (after all its just a plain old **S5** operator). But when did it first emerge as an *additional* operator? We're not sure. Prior used it on a number of occasions (see, for example, [369, Appendix B4]), though sometimes Prior's global modality is actually the master modality $\boxast$ discussed in Section 6.5 (that is, sometimes Prior views globality as the reflexive transitive closure of the underlying relation).

But it seems fair to say that it was the Bulgarian-school who first exploited it systematically: it's the Swiss Army knife underlying their investigation of BML, and their work on hybrid logic. Goranko and Passy [198] is a systematic study of the global modality as an additional operator, and is the source of Theorem 7.1, the Goldblatt-Thomason theorem for $ML(\Diamond, \mathrm{E})$. The operator has also been studied from an algebraic angle, being closely connected to the notion of a discriminator variety; these classes display nice algebraic behavior and have been intensively investigated in universal algebra. For, in the context of boolean algebra with operators, having the global modality is equivalent to having a so-called discriminator term; this is why in algebraic circles this modality is sometimes dubbed a 'unary discriminator term'; see Jipsen [253] for some information. The basic complexity results for the global modality were proved in Hemaspaandra's thesis [412]. Incidentally, the global modality is usually referred to as the 'universal' modality in the literature. However the word 'universal' suggests that we are working with a box, so we prefer the term 'global', which is appropriate for both boxes and diamonds.

The history of the difference operator is harder to untangle. It is probably due to von Wright [457] (who viewed it as a 'logic of elsewhere') and Segerberg gave an axiomatization in a *festschrift* for von Wright (see [399]). Segerberg's axiomatization, together with a more detailed completeness proof, was later published in [401]. But Segerberg treats D as an isolated modality. The use of D as an additional modality seems to have been proposed independently by Koymans [276, 277] and Sain [389]. The difference operator is also discussed in Goranko [195]. For a systematic investigation of D as an additional, logical modality, see de Rijke [104]. The D-Sahlqvist theorem in the text is due to Venema [439]. Theorem 7.8 is an unpublished result due to Szabolcs Mikulás.

BML is a Bulgarian school invention. The system is first described in Gargov, Passy and Tinchev [173] (as part of a wide ranging discussion of extended modal logic) and Gargov and Passy [172] concentrates on BML and gives proofs of the key completeness and decidability results. See also the results on modal definability in Goranko [195]. All these papers view modal languages as general tools for talking about structures, very much in the spirit of the present book. The window operator has an interesting independent history: van Benthem [37] used it as part of a logic of permissions and obligations, Goldblatt [182] used it to define negation in quantum logic, Humberstone [242] used it in a discussion of inaccessible worlds, while Gargov, Passy and Tinchev [173] view it as a 'logic of sufficiency' that balances the usual 'logic of necessity' provided by $\Box$. Complexity-theoretic aspects of BML have been studied and surveyed by Lutz and Sattler [310], while resolution-based decision procedures for extensions of BML and related languages are explored by Hustadt and Schmidt [244].

As we pointed out in the text, both BML and PDL are examples of modal languages equipped with highly structured collections of modal operators. The dynamic modal logic of De Rijke [112] is a further example, and many description logics allow for the construction of complex roles (that is, accessibility relations) by means of some or all of the booleans, converse, and sometimes even transitive closure and least fixed point constructors; see Donini *et al.* [123].

The algebraic counterparts of modal languages with structured collections of modal operators can best be phrased in terms of multi-sorted algebras, where the (algebraic counterparts of the) modal operators provide the links between the sorts. Kleene algebras [278] and Peirce algebras [108, 111] are two important examples. The former provide an algebraic semantics for PDL and consist of a boolean algebra and a regular algebra together with systematic links between them that are used to interpret the diamonds. The latter provide an algebraic semantics of dynamic modal logic and consist of a boolean algebra and a relation algebra together various links between that are, again, used to interpret the modalities in the language.

Since and Until. The invention of since and until logic was a major breakthrough in the study of modal logic. Hans Kamp tells the story this way. In a semester-long course Arthur Prior gave on tense logic at UCLA in the fall of 1965, when Kamp had just started his PhD, Prior stressed that the P and F operators operators were strictly *topological*, and asked whether it was possible to develop some notion of *metric* time within the framework of tense logic. Now, a first requirement on such an enterprise is that it can express what it is for some proposition q to have been true since the last time some periodically true proposition p was true. Trying to find a genuinely topological tense logic in which these kinds of relations could be expressed lead Kamp to the definitions of since and until. As the technical interest of the new operators became clear, the original topological motivation seems to have been shelved (Kamp, personal communication, remarks that 'The question of how to embed a logic of metric temporal notions within a topological tense logic unfortunately never got properly off the ground.'). Kamp first showed that P and F were not capable of expressing since and until, and eventually succeeded in proving Theorem 7.12(i), the expressive completeness of since and until logic over Dedekind complete total orders (see his thesis [263]). At that time, deductive completeness was the dominant interest in modal logic. Kamp's result showed that the neglected topic of modal expressivity deserved further attention, and can be regarded as a precursor to the study of correspondence theory that emerged in the 1970s.

The next step was taken by Dov Gabbay. Kamp's result was clearly important, but his direct proof was complex, and although Jonathan Stavi [415] succeeded in providing a direct proof of Theorem 7.12(ii), it was not obvious how proceed further. Matter were greatly simplified when Gabbay introduced the notion of *separability* (see [157, 159]). Roughly speaking, a language is separable over a class of models if every formula is equivalent to a boolean combination of atomic formulas, formulas that only talk about the past, and formulas that only talk about the future. This idea drastically simplifies the proofs of Theorem 7.12(i) and Theorem 7.12(ii), and opens the way to more general investigations. Nowadays a variety of techniques are used for proving expressive completeness results for modal (and other) languages; game-based approaches (see Immerman and Kozen [246]) have proved particularly useful. The best introduction to expressive completeness is the encyclopedic Gabbay, Hodkinson, and Reynolds [163]; both separability and game-based proofs are discussed. It also contains many other results on since and until logic and a useful bibliography.

But what really made the until operator so popular is the simple observation made in the text: it offers precisely the what is needed to express guarantee properties (this was first noted in Gabbay, Pnueli, Shelah, and Stavi [167]). Nowadays until may well be the single best known modal operator (at least in computer science) and it occurs in both in its original form, and in a number of variant forms in the study of linear and branching time temporal logics (see Clarke and Emerson [92], Goldblatt [183]).

Good discussions of step-by-step completeness proofs for since and until can be found in Burgess [76] and Xu [458]. The classification of properties of flows of time (in terms of safety, liveness, and guarantees) referred to in Section 7.2 can be found in Manna and Pnueli's textbook [318] on using temporal logic for specifying concurrent and reactive systems. Theorem 7.19 is due to Venema [438]; the strategy of using expressive completeness to obtain axiomatic completeness results goes back at least to Gabbay and Hodkinson [164].

One final remark: in spite of the fact that its satisfaction definition makes use of a more complex patterns of quantification, the since and until operators are genuinely modal. In particular, the notion of bisimulation can be adapted to these operators: the only complication is that, instead of the simple 'complete the square' idea illustrated in Figure 2.3 (65), bisimulations now need to match relational steps *plus intermediate intervals* in suitable ways. Kurtonina and de Rijke [295] contain a solution to this issue as well as a survey of earlier proposals.

Hybrid Logic. Arthur Prior introduced and made systematic use of hybrid logic; see Prior [369] (in particular, Chapter 5 and Appendix B.3), several of the papers in Prior [370], and the posthumously published Prior and Fine [371]. Prior's systems typically allowed explicit quantification over states using $\forall$ and $\exists$, and contained the global modality. Technical aspects of such languages were explored in Bull [71], an important paper, which among other things notes that pure formulas give rise to easy frame completeness results. In the mid 1980s Passy and Tinchev independently reinvented the idea of 'names as formulas'. Their earliest paper [360] added nominals and the global modality to a rich version of PDL; in [361] they considered $\forall$ and $\exists$ (again in the setting of PDL); and [362], their beautiful essay on hybrid languages, remains one of the key papers on hybrid languages.

The subsequent history of hybrid languages revolves around attempts to find well-behaved sublanguages of such strong systems. The most obvious way to do this is one explored in the text: treat nominals as names, rather than variables open to binding, and keep the underlying modal language relatively weak. Early papers which explore this option include Gargov and Goranko [171] (the basic modal language enriched with nominals and the global modality) and Blackburn [52] (the basic tense language enriched with nominals alone). The basic hybrid language discussed in the text can be viewed as an interesting compromise between simply adding nominals to the basic modal language (which makes the axiomatics messier, as Exercise 7.3.7 shows) and adding both nominals and the global modality (which raises the complexity to EXPTIME-complete). A proof of Theorem 7.21 (that the basic hybrid language has a PSPACE-complete satisfiability problem) can be found in Areces, Blackburn and Marx [14]. For a more detailed look at the complexity of hybrid logic, see [13] by the same authors. Theorem 7.29 is a modification of results proved in Blackburn and Tzakova [61]. It simplifies similar a result proved in Gargov and Goranko [171] with the aid of the global modality.

But the idea of binding variables to states turns out to be important. Binding admits a rich expressivity hierarchy. For a start, even if binding with $\forall$ and $\exists$ is allowed, when there are no satisfaction operators in the language, the resulting language does *not* have full first-order expressivity; see Blackburn and Seligman [57]. Moreover, as we mentioned in the text, the $\downarrow$ binder simply binds variables to the *current* state; in effect, it lets us create a name for the here-and-now (see Goranko [196], Blackburn and Seligman [57, 58], Blackburn and Tzakova [61]). If we enrich the basic hybrid language with the $\downarrow$ binder we obtain a hybrid language which corresponds to precisely the fragment of the first-order correspondence language which is invariant under generated submodels. This is proved in Areces, Blackburn and Marx [14] by isolating notions of bisimulation suitable for various hybrid languages and proving a characterization theorem. The paper also links these notions of bisimulation to restricted forms of Ehrenfeucht-Fraïssé games.

Hybrid logic provide a natural setting for modal proof theory. Seligman [404] is the pioneering paper here, and Seligman [405] discusses satisfaction operator based natural deduction and sequent systems. Blackburn [55] defines satisfaction operator driven tableau and sequent systems and uses Hintikka sets to prove an analog of Theorem 7.29. Tzakova [431] combines the use of nominals with the prefix systems of Fitting [145]. Demri [115] defines a sequent calculus for the basic tense language enriched with nominals, and Demri and Gore [116] introduce a display calculus for the basic tense language enriched with nominals and D.

Hybrid logics turn up naturally in a number of applications. The AVMs used in computational linguistics (recall Example 1.17) can be viewed as modal logics: path re-entrancy tags are treated as nominals (see, for example, Blackburn and Spaan [59]). And while it has long been known that description logics are notational variants of modal logics, this relation only holds at the level of concepts. So-called A-Box (or assertional) reasoning — that is, reasoning about how concepts apply to particular individuals — corresponds to a restricted use of satisfaction operators, while the ‘one-of’ operators used in some versions of description logic are essentially disjunctions of nominals; see Blackburn and Tzakova [60], Areces and de Rijke [15], and Areces’s PhD thesis [12]. Nominals also turn up in the Polish tradition of modal logics for information systems and rough-set theory: see Konikowska [274, 275]. They also provide a natural model of tense and other forms of temporal reference in natural language (see Blackburn [54]).

A final remark. The basic hybrid language shows that sorting is interesting in the setting of modal logic — so why not introduce further sorts? In fact, this step was already taken in Bull [71] who introduced a third sort of atomic symbol: *path nominals*, true at precisely the points belonging to some path through the model. For more information on hybrid logic, see the Hybrid Logic home page at `www.hylo.net`. For a recent ‘manifesto’ on hybrid logic that touches on most of the themes just mentioned, see Blackburn [56]

The Guarded Fragment. The guarded fragment was introduced by Andréka, van Benthem and Németi in 1994. The roots of the decidability proof date back to 1986, when Németi [345] showed that the equational theory of the class of so-called relativized cylindric set algebras is decidable. The first-order counterpart of this result is that a certain subfragment of the guarded fragment is decidable.

The importance of this result for first-order logic was realized in 1994 when Andréka, van Benthem and Németi introduced the guarded fragment and showed that many nice properties of the basic modal system **K** generalize to it. In particular, the authors established a characterization in terms of guarded bisimulations, decidability and a kind of tree model property. The journal version of their paper is [9]. Some time later van Benthem, was able to generalize some of the results, introducing the loosely guarded fragment in [433]. The slightly more general packed fragment was introduced in Marx [323] in order to give a semantic characterization in terms of packed bisimulations. (An example of a packed sentence which is not equivalent to a loosely guarded sentence in the same signature is $\exists xyz\,(\exists wCxyw \wedge \exists wCxzw \wedge \exists wCzyw \wedge \neg Cxyz)$.)

The mosaic based decision algorithms of Andréka, van Benthem and Németi were essentially optimal: a result established by Grädel [200]. In this paper, Grädel also defines and establishes the loose model property for the loosely guarded fragment. Our definition of a loose model is based on the definition of a tree model given there. Grädel and Walukiewicz [203] showed that the same bounds obtain when the guarded fragment is expanded with least and greatest fixed point operators. Marx, Mikulás and Schlobach [325] defined a PSPACE-complete guarded fragment with the finite tree model property. This fragment satisfies both locality principles.

The finite model property for the guarded fragment, and several subfragments of the packed fragment, was established in an algebraic setting by Andréka, Hodkinson and Németi [7]. Grädel [200] provides a direct proof for the guarded fragment. The remaining open question for the full packed fragment was solved affirmatively by Hodkinson [236]. All these results are based on variants of a result due to Herwig [228]. The use of Herwig's Theorem to establish the finite model property and to eliminate the need of step-by-step constructions originates with Hirsch *et al.* [232].

Multi-Dimensional Modal Logic. The idea of evaluating modal languages at sequences of points, rather than at the points simpliciter, is extremely natural, so it is no surprise that over the years modal logicians with very diverse interests have devised multi-dimensional systems.

It seems that logicians interested in natural language were first off the mark. Natural language utterances are so context dependent, that evaluating at sequences of points (each coordinate modelling a different aspect of context) proved a useful idea. Evaluation at pairs of points is built into Montague's [342] general framework for natural language semantics. Kamp's [264] classic analysis of the word 'now' uses a second coordinate to keep track of utterance time. Vlach [445] provided an analysis of the word 'then', and in a series of papers, Åqvist and co-workers [11] developed a number of rich multi-dimensional modal logics for analyzing natural language temporal phenomena. Before long, such systems were subjected to rigorous logical investigation: see, for example, Segerberg's elegant decidability and completeness result in [398], and Gabbay's work on expressiveness and other topics (much of which reappeared in the later work by Gabbay *et al.* [163]).

Somewhat later, a rich source of inspiration came from logic itself. Some work here, such as the sorted modal logic PREDBOX of Kuhn [293], fitted in the tradition of Quine-style first-order logic without variables, but most of it was linked, one way or another, with the algebraic logic framework of the Tarskian school (see the Notes of Chapter 5). This certainly applies to the multi-dimensional logics that we presented in Section 7.5. Venema [436], from which our Theorem 7.51 originates, made the connection between modal logic and cylindric algebras. Subsequent research drew on existing ideas on relativized cylindric algebras (see Németi [345]) to use the modal framework to 'tame' first-order logic and its finite variable fragments (see our discussion of the abstract and relativized assignment frames in the text; more information on this program can be found in van Benthem [47] or Mikulás [335]). This line of work is closely related to arrow logic, which is a multi-dimensional modal logic in its own right (see Marx *et al.* [324] for more information) and in fact this strand of work ultimately lead to the isolation of the guarded fragment. All of these (and more) multi-dimensional modal logics are covered in the monograph Marx and Venema [326]; readers interested in complexity results should consult Marx [322].

Computer scientists have different motivations for studying multi-dimensional modal logics. In order to build formal models of an application domain, they need to take account of various features simultaneously. Of the wealth of literature on this topic we'll just mention Fagin *et al.* [133], which concentrates on the combination of temporal and epistemic logics in the context of distributed systems. Such applications have led logicians to study various ways of constructing complex logics from relatively simple ones. A particularly interesting and mathematically non-trivial branch of multi-dimensional modal logics arises if one studies a modal language with various modal operators over a semantics in which the frames are cartesian products of frames for the individual operators. This area of so-called product logics, which has an early predecessor in Shehtman [406], has recently become very active; a monograph Gabbay *et al.* [153] is on its way.

Finally, multi-dimensional modal logic remains one of the most philosophically important branches of modal logic. Important references include Kaplan [269, 270], Stalnaker [414], and Chalmers [88]

The Lindström Theorem for Modal Logic. Theorem 7.60, a Lindström-type characterization of the modal languages defined in Definitions 1.9 and 1.12 is due to De Rijke [107]; the result was obtained as part of a general program to come up with modal counterparts of model-theoretic results in first-order logic [106]. The original first-order version of Lindström's Theorem was first presented in Lindström [309]. The original result states that, given a suitable explication of a 'classical logic', first-order logic is the strongest logic to possess the Compactness and Löwenheim-Skolem properties; it formed an important source of inspiration for the area of model-theoretic logics [25]. Definitions of the abstract notion of a logic can be found in Chang and Keisler [89] and in Barwise [24]. A very accessible presentation of Lindström's Theorem for first-order logic can be found in Doets [119, Chapter 4].