MODAL LOGIC IN TWO GESTALTS

Johan van Benthem, Amsterdam & Stanford

in M. de Rijke, H. Wansing & M. Zakharyashev, eds., *Advances in Modal logic, vol. II*, Uppsala 1998, CSLI Publications, Stanford, 73–100.

Abstract

We develop a translation-based view dual of modal logic as the study of intensional languages that are at the same time interesting expressive and decidable parts of standard logical systems. This tandem approach improves our understanding of modal logic – while at the same time, it extends the range of modal notions and techniques into broader areas of standard logic.

1 Translation as a Way of Life

1.1 Basic modal logic and the modal fragment of FOL

Modal languages as used to-day can be considered a species of their own, inhabiting the realm of Intensional Logic. But they can also be translated into fragments of standard logical languages, mostly first-order, sometimes higher-order or infinitary. These translations reflect the truth conditions for modal operators in possible worlds models. The ur-example is the basic modal language of possibility and necessity, whose standard translation ST inspired Correspondence Theory (van Benthem 1976, 1985):

an existential modality <>p goes to a bounded quantifier $\exists y\,(Rxy \,\&\, Py)$
stating that the current world x has a successor y in which p holds

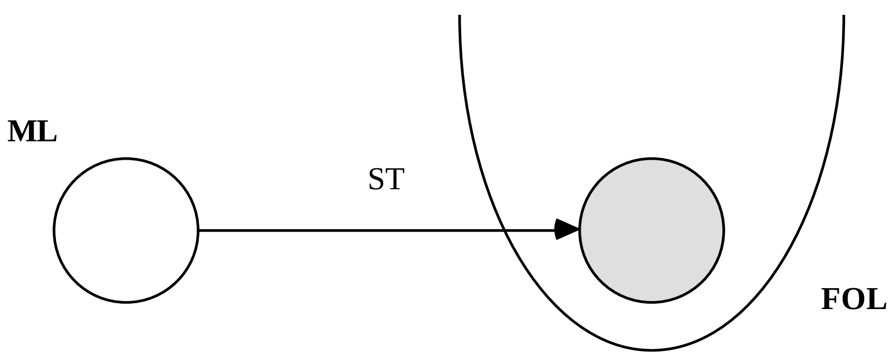

In this manner, the basic modal language transcribes into a fragment of a first-order language over possible worlds models, in the appropriate similarity type with a binary accessibility relation and unary local atomic properties of worlds. The fundamental semantic feature that locates this syntactic modal fragment inside the full first-order language is a semantic invariance property, measuring expressive power with respect to the appropriate structural equivalence 'between models, comparing 'bisimilar nodes':

> <u>Modal Invariance Theorem</u>
> A first-order formula is definable by a modal formula
> if and only if it is invariant for *bisimulation*.

The basic modal fragment has a 'nice package' of properties. It is reasonably expressive for many recurrent purposes. It has a good model theory – based on bisimulation, rather than classical potential isomorphism - enjoying all the classical meta-theorems (compactness, interpolation, Los-Tarski, etcetera). It also has a good proof theory. In addition, however, unlike full first-order logic, the basic modal logic is *decidable*. In this paper, we want to generalize this observation. Our presentation revolves mostly around minimal modal logic of model classes (better called 'universal' or 'central' logic), and on first-order semantics on models (rather than second-order semantics on frames).

1.2 General translation of modal truth conditions

The same translational view can be applied any (more expressive) modal or temporal or dynamic language with a well-defined semantics. Examples may be taken from any presentation at the AiML-II conference. Here are two cases out of many:

intuitionistic modal logic

$A \Rightarrow \langle\rangle B \qquad \forall y\,(x \leq y \rightarrow (Ay \rightarrow \exists z\,(Ryz \,\&\, Bz)))$

interpretability logic

$A \mid\!> B \qquad \forall yz\,(Rxyz \rightarrow (Ay \rightarrow \exists u\,(Szu \,\&\, Bu)))$

Our general approach to modal logic will be to develop these two viewpoints: modal formalisms per se, and their standard counterparts, *in tandem* for purposes of language design and meta-logical analysis. Quite radically, we see this as instances of a desirable Gestalt switch: one should develop the ability to see them as *both*. This tandem approach has several virtues. Seeing your favourite modal language in a broader environment allows for transfer of known results from standard logics, which saves time and effort.

(Of course, not every property of the larger standard language will transfer automatically to its modal fragments. Additional effort is often required...) A broader standard setting may even suggest redesign for recalcitrant modal languages. A case in point is Since/Until temporal logic, whose basic operators are double quantifiers $\exists\forall$ that work better when decomposed into iterated single ones, living in a suitable two-dimensional modal logic. But profitable traffic can also run in the *opposite* direction. The tandem view allows for natural penetration of notions and techniques that were first developed inside modal logic into broader areas of standard logic. Various demonstrations of this occur in this paper. Thus, a systematic translational perspective has both practical and theoretical value when engaging in modal logic.

2 Syntactic Fine-Structure: Quantifier Guards

The results in this section are largely based on Andréka, van Benthem & Németi 1998, to which we refer for missing definitions and proofs, and statements of further results.

2.1 From basic modal logic to the Guarded Fragment

Here is the more general thrust of modal logic in standard first-order logic. One can extend the above pattern for an existential modality to the more general polyadic format

$\exists \mathbf{y}\ (G(\mathbf{x}, \mathbf{y})\ \&\ \phi(\mathbf{x}, \mathbf{y}))$ called *guarded quantification*.

Here, the 'guard atoms' $G(\mathbf{x}, \mathbf{y})$ bounding the existential quantifier may have their variables occurring in any order and multiplicity. Bounded quantification is a known concept, witness the important 'absolute' Δ_0–formulas of set theory. The new idea here is that guards totally restrict the group of objects about which the subsequent matrix statement $\phi(\mathbf{x}, \mathbf{y})$ may make assertions. (Allowing new objects beyond $\mathbf{x}, \mathbf{y}$ in the matrix, as is allowed in standard Δ_0–formulas, leads to undecidability.) The sublanguage of first-order predicate logic with this syntax is the *Guarded Fragment* **GF**.

GF has a semantic characterization much like that for basic modal logic in the Modal Invariance Theorem, but now involving invariance under *guarded bisimulations*. The latter are classical 'potential isomorphisms', i.e. non-empty families **I** of finite partial isomorphisms between two models **M**, **N**, but with only guarded choices of tuples of new elements in the back-and-forth clauses. More precisely, Let $f \in \mathbf{I}$ and let a guard atom $G(\mathbf{d}, \mathbf{e})$ hold in **M** for objects **d** from the domain of F plus new objects **e** . Then there exists a tuple **e'** of objects in the other model **N** verifying $G(f(\mathbf{d}), \mathbf{e}')$ while the partial isomorphism sending the **d** to their old f-values, and objects in **e** to the corresponding objects in the tuple **e'** belongs to **I** . The opposite direction asks the same.

This generalizes basic modal bisimulation in an obvious manner.

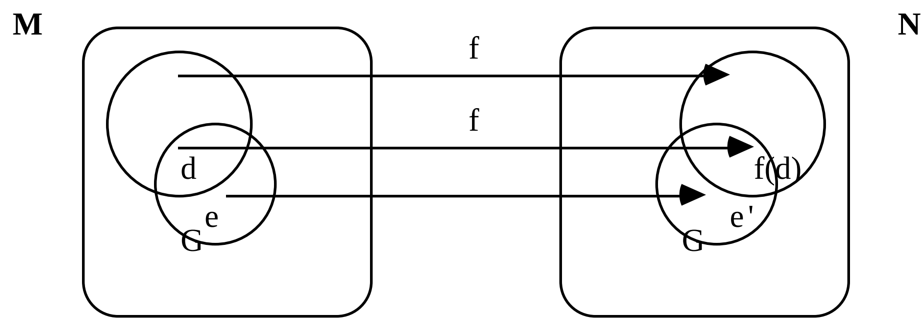

<u>Theorem</u> The first-order formulas invariant for guarded bisimulations are precisely those definable by means of guarded ones.

<u>Proof</u> The modal proof of the MIT lifts to first-order logic.

The main result establishing the interest of **GF** is the following. Its proof, again, generalizes a modal method of reasoning, namely, filtration (using 'quasi-models'):

<u>Theorem</u> **GF** is decidable.

<u>Proof</u> For later reference, we reproduce the basic steps from Andréka, van Benthem & Németi 1997. Their method amounts to a generalization of modal filtration arguments. Any satisfiable **GF**-formula ϕ has a finite so-called 'quasi-model', consisting of 'types' made up of subformulas of ϕ, of some effectively computable size – and conversely, each such quasi-model generates a model for ϕ. Thus, whether a guarded formula is satisfiable is equivalent to its having a finite quasi-model: a decidable property.

From Standard Models to Finite Quasi-Models Suppose that formula ϕ is satisfiable in standard model **M** . Let V be the set of variables occurring in ϕ (free or bound). Henceforth, we restrict attention to the finite set Sub_ϕ consisting of ϕ and all its subformulas, closed under simultaneous substitutions using only variables in V, that do not change syntactic forms. (This is feasible, by the Remark following this proof: simultaneous substitutions do not need bound variables beyond the original supply.) Each variable assignment verifies a 'type' Δ of finitely many formulas from this set. Our quasi-model will have a universe that consists of the finitely many types realized in the model $\mathbf{M}$. Moreover, evidently, in this syntactic structure, for each guarded formula $\exists \mathbf{y} (Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y})) \in \Delta$, there exists a type Δ' with (i) $Q\mathbf{xy}, \psi(\mathbf{x}, \mathbf{y}) \in \Delta'$, (ii) Δ, Δ' agree on all 'unaffected' formulas with only free variables in $\mathbf{x}$.

<u>Definition</u> (i) Let F denote the finite set of all guarded formulas of length $\leq |\phi|$ that use only variables from V . Note that $\phi \in F$ and F is closed under taking subformulas and 'alphabetic variants'. (ii) An F-*type* is a subset Δ of F for which we have

(a)	$\neg \psi \in \Delta$	iff	*not* $\psi \in \Delta$	whenever $\neg \psi \in F$
(b)	$\psi \wedge \xi \in \Delta$	iff	$\psi \in \Delta$ *and* $\xi \in \Delta$	whenever $\psi \wedge \xi \in F$
(c)	$[\mathbf{u}/\mathbf{y}]\psi$	implies	$\exists y\, \psi \in \Delta$	whenever $\exists y\, \psi \in F$

$[\mathbf{u}/\mathbf{y}]\psi$ comes from ψ by replacing each free variable in $\mathbf{y}$ with the corresponding variable in $\mathbf{u}$, simultaneously. (iii) Let $\mathbf{y}$ be a sequence of variables, and Δ, Δ' types. Write $\Delta =_{\mathbf{y}} \Delta'$ if Δ , Δ' have the same formulas with free variables disjoint from $\mathbf{y}$. (iv) A *quasimodel* is a set of F–types S such that, for each $\Delta \in S$ and each guarded formula $\exists \mathbf{y} (Q\mathbf{xy} \wedge \psi) \in \Delta$, there is a type $\Delta' \in S$ with $Q\mathbf{xy}$ and $\psi(\mathbf{x}, \mathbf{y})$ in Δ' and $\Delta =_{\mathbf{y}} \Delta'$. We say that ϕ *holds in a quasi-model* if $\phi \in \Delta$ for some Δ in this model. ■

Clearly, if ϕ is satisfied by some model, then ϕ also holds in some quasi-model.

From Quasi-Models to Standard Models From any quasi-model $\mathbf{M}$, we can define a standard model $\mathbf{N}$. Call π a *path* if $\pi = \langle \Delta_1, \phi_1, ..., \Delta_n, \phi_n, \Delta_{n+1} \rangle$ where Δ_1, Δ_{n+1} are types in $\mathbf{M}$, each formula ϕ_i is of the form $\exists \mathbf{y} (Q\mathbf{xy} \wedge \psi) \in \Delta_i$ and Δ_{i+1} is an alternative type as described above (i.e., $Q\mathbf{xy}$, $\psi(\mathbf{x}, \mathbf{y})$ in Δ_{i+1} and $\Delta_{i+1} =_{\mathbf{y}} \Delta_i$) . We say that the variables in $\mathbf{y}$ *changed their values* from Δ_i to Δ_{i+1} (the others did not). Finally, variable z is called *new in path* π if either $|\pi| = 1$ or z's value was changed at the last round in π . *Objects* in $\mathbf{N}$ are all pairs (π, z) with π a path, z new in π . Next, we *interpret predicates* over these objects. $I(Q)$ holds of the sequence of objects $\langle(\pi_j, x_j)\rangle_{j \in J}$ iff the paths π_j fit into one linear sequence under inclusion, with a maximal path π^* such that (i) the atom $Q\langle x_j \rangle_{j \in J} \in \Delta^*$ (the last type on π^*) and for no

(π_j, x_j) does x_j change its value on the further path to the end of π^* . Finally, we define an *assignment* s_π for each path. We set $s_\pi(x) =_{def} (\pi', x)$ with π' the unique subpath of π^* at whose end x was new, while it remained unchanged afterwards.

The correctness of this model construction shows at *last*(π), the last type on the path π :

Truth Lemma For all paths π in **N** , and all formulas $\psi \in F$,

$\mathbf{N}, s_\pi \models \psi$ iff $\psi \in last(\pi)$.

Proof Induction on ψ . ***Boolean cases*** are immediate by the closure conditions for $\neg$, $\wedge$ on types. ***Atoms***: involve a straightforward calculation, via the linearity condition in the interpretation function I, plus the '$=_\mathbf{y}$ -clause' in quasi-models ensuring transfer of 'unaffected formulas' along paths. For later reference, we repeat the full argument for bounded ***Existential Quantifiers*** $\exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y}))$. (i) First, suppose that $\exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y})) \in last(\pi)$. Then there is an extended path $\pi^+ =_{def} \pi$ *concatenated with* $< \exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y})), \Delta' >$, where Δ' is a successor type for Δ chosen as above with $Q\mathbf{xy}, \psi(\mathbf{x}, \mathbf{y}) \in \Delta'$ (satisfying the transfer condition for unaffected formulas with free variables $\mathbf{x}$). All objects (π^+, y_i) with y_i in $\mathbf{y}$ are new here. By definition, the atomic guard I(Q) holds for the object tuples $s_{\pi+}(\mathbf{y})$, $s_{\pi+}(\mathbf{x})$ ($= s_\pi(\mathbf{x})$) . Also, by the inductive hypothesis, $\mathbf{N}, s_{\pi+} \models \psi(\mathbf{x}, \mathbf{y})$. Therefore, $\mathbf{N}, s_{\pi+} \models \exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y}))$. And by **x**-invariance in the standard model **N** , then, indeed $\mathbf{N}, s_\pi \models \exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y}))$. (ii) Conversely, suppose that $\mathbf{N}, s_\pi \models \exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y}))$. By the truth definition, there are objects $d_i = (\pi_i, u_i)$ with $\mathbf{N}, s_\pi{}^{\mathbf{y}}{}_{\mathbf{d}} \models Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y})$. (Here, $s_\pi{}^{\mathbf{y}}{}_{\mathbf{d}}$ is the assignment which is like s_π except for setting all y_i to d_i .) In particular, I(Q) holds of the objects $s_\pi(\mathbf{x})$, $\mathbf{d}_i$. This leads to a picture of forking paths. The $s_\pi(\mathbf{x})$ were all introduced by stage π^* inside π , and then the $\mathbf{d}_i$ were (either interpolated, or) added to form a maximal sequence π^+ with the atom $Q\mathbf{xy}$ true at the end. The fork is such that **x**-values do not change any more from π^* onward, whether toward π or π^+ .

This is the only case where the atomic guard restriction on our quantifiers comes in essentially.)

We now analyse this 'forking situation' a bit more carefully:

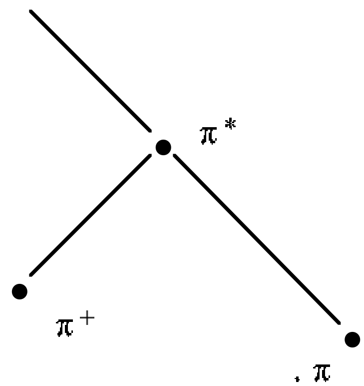

Now, the variables u_i do not have to be the y_i . Say, π^+ has $s_{\pi+}(u_i) = (\pi_i, u_i) = d_i$. Thus, the assignments $s_\pi{}^{\mathbf{y}}{}_{\mathbf{d}}$ and $s_{\pi+}$ agree on $\mathbf{x}$, and for all $y_i \in \mathbf{y}$ we have $s_\pi{}^{\mathbf{y}}{}_{\mathbf{d}}(y_i) = d_i = s_{\pi+}(u_i)$. Then, by $\mathbf{N}, s_\pi{}^{\mathbf{y}}{}_{\mathbf{d}} \models Q\mathbf{xy} \wedge \psi$ and the above observations, we have $\mathbf{N}, s_{\pi+} \models [\mathbf{u}/\mathbf{y}]Q\mathbf{xy}$, $\mathbf{N}, s_{\pi+} \models [\mathbf{u}/\mathbf{y}]\psi$. By the inductive hypothesis, $[\mathbf{u}/\mathbf{y}]\psi \in \text{last}(\pi^+)$. Also, from the initial description of π^+ , we see at once that $[\mathbf{u}/\mathbf{y}]Q\mathbf{xy} \in \text{last}(\pi^+)$ (by the interpretation of atomic predicates). By closure conditions (b), (c) for types, one gets $\exists\mathbf{y}\,(Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y})) \in \text{last}(\pi^+)$. Finally, since no changes in $\mathbf{x}$-values occurred on the fork from π^* , the transfer condition for unaffected formulas along successor types along paths ensures that this same formula is in last(π) . ◆

Thus having a quasi-model implies having a real model, and the Theorem is proved. ◆

Remark *Finite Variable Fragments are closed under Simultaneous Substitutions*

Our proof assumes the finite set of relevant formulas is closed under simultaneous substitutions – without enlarging the set V of relevant variables. To see this, consider any substitution $[\mathbf{x} := f(\mathbf{x})]\,\phi$ in a k-variable fragment with variables $\mathbf{x} = x_1, \ldots, x_k$. Atomic replacements are straightforward. Also, we can push substitutions inside over Booleans. The only interesting case is when we encounter an existential quantifier: $[\mathbf{x} := f(\mathbf{x})]\,\exists x_j\,\psi$. Then, the assignment clause $x_j := f(x_j)$ has no effect, and so it can be omitted. Hence, in the remaining substitution σ , at least one variable x_k is not used at all on the right-hand side in any assignment. But then, the following formula is easily shown to be equivalent to the original one: $\exists x_k\,[x_j := x_k, \sigma]\,\psi$. This gives a simple recursive algorithm computing substitutions inside our fragment. (With function symbols present, however, the result fails: witness the case of $[x := fxy]\,\exists y\,Rxy$.) ◆

This analysis of modal languages is orthogonal to a well-known alternative, via *finite-variable fragments*. Guarded fragments seem better-behaved than finite-variable ones in qua meta-theory and complexity (Andréka, van Benthem & Németi 1998 review pros and cons). Indeed, **GF** offers a new perspective of 'syntactic quantifier bounds', different in its thrust from the usual decidable fragments of first-order logic. The latter go by restrictions on predicate arities, or quantifier prefixes, or numbers of variables (cf. the survey in Börger, Grädel & Gurevich 1996). Thus, the above results about the Guarded Fragment by themselves show how modal logic can inspire classical logic.

2.2 Guarded analysis of modal languages

As an immediate application, modal languages have a decidable minimal logic when their modalities translate into **GF**. Examples are

(1) basic temporal logic: existential past quantifiers $\exists y\ (Ryx\ \&\ Py)$
over predecessors are guarded (just as the future ones)
(2) relevant logic: ternary implications $\forall yz\ (Rxyz \rightarrow (Ay \rightarrow Bz))$
and conjunctions $\exists yz\ (Rxyz\ \&\ Ay\ \&\ Bz)$ are guarded
(3) the algebraic logic CRS of 'cylindric relativised set algebras':
where the relativization supplies uniform guards.

Sometime, a little ingenuity is needed in finding a suitable translation. For instance, as a challenge, consider the 'second-order' *neighbourhood semantics* for the basic modal language. This involves a binary relation R between worlds w and sets of worlds X (serving as 'neighbourhoods'), where one stipulates that

$w \models \langle\rangle\phi$ iff there exists a set X with RwX all of whose members satisfy ϕ

The resulting minimal logic is weaker than minimal modal K, dropping distributivity of possibility over disjunction. Now, the new truth condition may obviously be written as follows in second-order notation: $\exists X\ (RwX\ \&\ \forall y\ (y \in X \rightarrow \phi(y))$. This *is* guarded, however, if we read the formula as a *two-sorted first-order one* – noting at the same time that this move makes no difference to the functioning of neighbourhood semantics. Further working examples of guarded analysis are found in van Benthem 1997B, which analyses various 'Sofia fragments' in extended modal logic.

On top of a minimal logic, one often imposes extra *frame conditions*. Even for the basic modal language, these need not fit into **GF**! (We never promised a miracle cure.) E.g., symmetry is guarded, but transitivity is not – as may be seen by exhibiting two models, one transitive, one non-transitive, with a guarded bisimulation between them. This fact points to a natural division of labour. The Guarded Fragment was invented to explain decidability for general modal formalisms. But there may well be different sources of decidability beyond it, in the special theories of specific well-behaved frame classes. (But read on for some enjoyable subtleties in the range of guarded analysis...) In any case, our main point is not the selling of **GF** as some uniquely preferred modal system. What we advertize is rather the study of a fundamental theme, namely quantifier bounding patterns, and the practical ability of *guarded analysis*, paying attention to the syntactic fine-structure of modalities wherever they occur.

3 Packed Conjunctions and the Edge of Undecidability

The results in this section are drawn largely from the unpublished paper 'Extending the Guarded Fragment to Betweenness and Pair Arrows' occurring in van Benthem 1997A.

3.1 From single guards to conjunctive ones

Returning to general modal languages, there are natural decidable cases beyond **GF**. These often involve *conjunctions of guard atoms*. A typical example is *betweenness* of points in temporal logic, which leads to inherently non-guarded assertions.

<u>Example</u> The modality UNTIL AB says $\exists y\ (x<y \wedge Ay \wedge \forall z\ ((x<z \wedge z<y) \rightarrow Bz))$. Its betweenness clause has a composite guard $x<z \wedge z<y$. This assertion is not in **GF**, even though the minimal temporal logic of UNTIL (and its dual SINCE) is decidable.

Another illustration is the technique of *relativization* in Relational Algebra, which weakens classically undecidable systems to decidable ones. In so-called Arrow Logic, one quantifies over arrows as primary objects, treating relations as unary predicates over these. A typical clause is the truth condition for relational composition – which involves an underlying primitive ternary relation Comp of arrow composition:

$$R \circ S\ (a) \text{ iff } \quad \exists bc\ (R(b)\ \&\ S(c)\ \&\ Comp\ (a, bc))$$

This is guarded, and the decidability of basic Arrow Logic follows immediately from that of **GF**. But algebraic logicians have a slightly more concrete technique for relativizing, letting relations still be set of ordered pairs, but now over arbitrary 'top relations' U (not just full Cartesian squares DxD). That is

<u>Example</u> Pair arrow models have their binary composition defined as follows: $R \circ S =_{def} \lambda xy \bullet \exists z\, ((Uxz \wedge Uzy) \wedge Rxz \wedge Szy)$, with a composite guard $Uxz \wedge Uzy$.

The point here cannot be that *arbitrary conjunctions* of atoms are acceptable guards.

<u>Proposition</u> **GF** extended with arbitrary conjunctions of guards is undecidable.

<u>Proof</u> The 3-variable fragment of first-order logic is known to be undecidable. Here is an effective reduction taking its satisfiability problem into that for **GF** with arbitrary conjunctive guards. Clearly, any 3-variable formula ϕ is satisfiable iff its guarded relativization $(\phi)^U$ to some new ternary predicate U is satisfiable in a *full Cartesian product* U = DxDxD. Now, it suffices to observe that the latter assertion can be expressed as the satisfiability of a formula

$$(\phi)^U \ \& \ CART(U)$$

where $CART(U) =_{def}$

(i) ∃xyz Uxyz & (ii) ∀xyz (Uxyz → **&** U-followed-by "all permutations and identifications among {x, y, z}") & (iii) ∀xyzuvw ((Uxyz & Uuvw) → **&** U-followed-by "all selections of three variables from among {x, y, z, u, v, w}").

Note that the latter formula is indeed in **GF** with added conjunctions of guards. ■

3.2 The Packed Fragment

Here is the proper generalization covering the above two positive examples. We call a quantification *pairwise guarded*, or *packed*, it has the following syntactic format:

$$\exists \mathbf{y}\ (\&Q\mathbf{xy} \wedge \psi(\mathbf{x}, \mathbf{y}))\ ,$$

where $\&Q\mathbf{xy}$ *is a conjunction of atoms with free variables* $\mathbf{y}, \mathbf{x}$ *in which every two variables from* $\mathbf{y} \cup \mathbf{x}$ *co-occur in at least one of the listed atoms.*

An obvious inductive definition gives the *Packed Fragment* **PF**, where the matrix formula $\psi(\mathbf{x}, \mathbf{y})$ itself comes from **PF** . The above temporal and arrow formulas were packed (with UNTIL, the match $x<y$ was given 'on the outside' already). By contrast, the above formula CART(U) is not packed. Another typical non-example is *transitivity*

$$\forall y_1 y_2 y_3 \, ((y_1<y_2 \wedge y_2<y_3) \rightarrow y_1<y_3)$$

without co-occurrence of y_1, y_3 in a guard atom: the point of this relational condition is precisely to *get* a relationship between the latter.

<u>Theorem</u> **PF** is decidable.

<u>Proof</u> We analyse the representation argument for **GF**. The definition of quasi-models carries over without major changes, as does their representation via 'path models'. Here, we now allow path extensions via the new generalized form of bounded quantification. Again, the crucial result is the Truth Lemma, saying that guarded formulas hold under the assignment induced by a path iff they occur in the last set encoded in that path. The step from right to left here is as before. Thus, the key is a combinatoric aspect of the converse direction, whose main step was illustrated in the above picture. The argument for true existential formulas still works with a conjunction of atomic guards like above. We look at the maximal position π^* as before. For each new variable y , again given the truth condition for atomic statements, loose guardedness requires that the path of the new y-value fits linearly with the original path on which the x-values occurred. Therefore, it either lies on the latter, or it extends it starting from π^* . Moreover, the condition also applies to all new values y amongst each other - and hence, these form at worst some linear path π^+ extending π^*, up to some maximal node where the highest new y-value has been introduced. The rest of the argument is as before, since all relevant y-atoms hold at π^+, and no y-values change in going back towards π^*. Cases of mere interpolation of the new y-values on the old path π are merely simpler. (Here, we heavily use the constancy of relevant variable values in an atom along the path up to the highest variable mentioned. This requires some checking of cases.) ∎

There is also a semantic characterization of the Packed Fragment, in terms of invariance under appropriately enriched guarded simulations, which will not be pursued here.

4 Boosting Decidability: Infinitary Languages and Fixed Points

The above does not yet represent the limit of guarded analysis. For there are further ways around apparent failures. An instructive case is the above non-pairwise-guarded *transitivity*. Consider the modal logic K4. Why is it (quite easily) decidable? There are two possible lines of attack here. One tries to extend the syntactic scope of **PF** and its ilk, to find broader decidable logics covering K4. We doubt this is feasible. Transitivity is *dangerous*: it is known to make first-order fragments undecidable (Börger, Grädel & Gurevich 1996). But there is a way-around this difficulty, by an alternative diagnosis of K4's decidability, transcending first-order logic, while still retaining the key role of bisimulation invariance. Recall that propositional dynamic logic PDL (or the modal μ–calculus) is decidable. Now it is easy to see – reinterpreting the usual decidability arguments as they stand – that *K4 is also the logic of any iteration modality* $[a^*]$, on which we impose no special frame restrictions at all. This is a genuinely different strategy. For, the PDL-language cannot define transitivity of models! Like the basic modal language, it is *invariant for bisimulation* (the infinitary conjunctions needed to define iteration do not affect this), while transitivity is not.

<u>Conjecture</u> **GF** extended with fixed points for defining new assertions is decidable.

This result has been claimed already informally, and we have a proof for the special case of fixed points that occur with so-called 'finitely distributive' monotone operators (the latter always stabilize at the ordinal ω). Presumably, the Tree Model Property highlighted in Vardi 1997 for decidable modal logics, which also underlies the decidability of **GF**, will prove the key notion here. (**Added in print, February 1999: Erich Graedel & Igor Walukiewicz have just circulated such a proof, to be presented at LICS 99.**) But there is a further subtlety.

A positive answer to the conjecture may be viewed as a natural generalization of the celebrated decidability of the *modal μ–calculus*. This is poly-modal logic extended with fixed point operators $\mu p \bullet \phi(p)$ defining new propositions (where all occurrences of the proposition letter p are syntactically positive in ϕ). But there is a subtlety here. The μ–calculus has only part of its possible fixed points, viz. those that define *assertions about states*. What it lacks, however, are fixed points that define new *program constructions*, by recursing over transition relations. E.g., a transitive closure $\langle a^*\rangle p$ is mimicked by the fixed-point assertion $\mu q \bullet \langle a\rangle p \vee \langle a\rangle q$. But the natural recursion $a^* = a \cup a;a^*$ over binary relations is not expressed directly.

<u>Question</u> Is the μ–calculus with relational fixed points decidable, too?

This distinction between state assertions and state transitions is a natural one – and it will return in later sections. In particular, it also makes sense for guarded fragments and their ilk. 'State recursion' and 'action recursion' are two different ways of adding fixed points. E.g., finite approximations for state-predicate based fixed point equations remain inside **GF**, but those for action predicates need not. The reason is that substituting an arbitrary guarded formula for a guard atom need not produce a guarded formula (e.g., substitute ¬ Rxy for Axy in ∃y (Axy & Qy)). Only so-called *safe formats* for action expressions have this substitution property, which unpack into iterated guarded quantifications. (A precise definition of safety is not attempted here: cf. van Benthem 1996, Ch. 5, 1998C, or Hollenberg 1998 – or the sketch given in Section 7.) (**Graedel & Janin also show that GF with arbitrary action fixed-points is undecidable.**) The dangers of unbridled action fixed-points show once more in the tiling problems of Section 5, where transitive closure of action predicates North, East gives undecidability.

<u>Remark</u> *Shifting between truth conditons and frame conditions*

Our re-analysis of K4 high-lights a usually implicit division of labour in modal logic: viz. between general truth conditions and special frame conditions. What we see is that the same effects can be obtained by manipulating the 'balance' between these two semantic features. Another example is the 'Brouwer logic' B of symmetric frames, which is also the minimal logic of the existential modality ∃y (Rxy & Ryx & ϕ(y)). This trade-off is far from being understood in its generality.

<u>Remark</u> *Higher-order extensions of decidable fragments*

Decidable fragments of first-order logic are not just weakenings. They may be able to carry 'extra weight' which first-order logic as a whole cannot bear without pain. The earlier fixed-point operators provide one example of this: added to FOL they generate highly exprerssive and complex parts of infinitary languages. But on top of modal languages, they seem more harmless. Another example found at AiML-II was the second-order quantifier "most". On top of FOL, it creates a logic which can define the natural numbers categorically, and hence incurs very high complexity. But the methods of

Ohlbach & Koehler 1998 show that basic modal logic with a new numerical modal operator saying "more a-successors are A than b-successors are B" remains decidable. What about similar extensions of **GF** with non-standard generalized quantifiers?

Remark *Weaker propositional bases*
As in description logics, one can study our main questions over a weaker underlying propositional logic: e.g., with just conjunction plus existential and universal guarded quantifiers. What happens then to the complexity of satisfiability and consequence? And, what are appropriate 'directed bisimulations'? Cf. Kurtonina & de Rijke 1997.

5 A Dynamic Perspective: From Sequential Action to Parallellism

This section is based on the analysis of dynamic logic in the unpublished paper 'Guarded Questions and Variations', which occurs in van Benthem 1997A.

5.1 State predicates versus action predicates

Syntactic decidability analysis without some concrete point of view may become blind. One powerful more focused view takes modal languages *dynamically* as descriptions of actions and their effects across states. That is, one thinks of possible worlds models as *process graphs* or 'labeled transition systems'. Intuitively, **GF** is about states which can be changed by 'guard actions' G(**x**, **y**), going either from **x** to **y**, or from **x** to **x**, **y** . This is still like basic modal logic, or propositional dynamic logic, in that it describes the effects of *sequential actions*. By contrast, conjunctive guards suggest *parallel action*, where different actors (sub-processes) change components of one global state. Here is a typical quantifier pattern which arises with 'collective action':

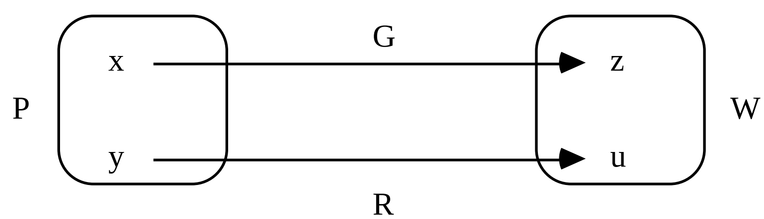

For a concrete interpretation, let P ('poor') describe the sum total of our current possessions, G your action of gambling at the casino, R my action of robbing the bank, while W ('wealthy') describes the sum of our new individual financial states. The typical quantifier pattern describing this outcome is the non-guarded, non-packed

$Pxy \,\&\, \exists zu\, (Gxz \,\&\, Ryu \,\&\, Wzu)$

We can think of this as a 'product action' G**x**H, of a kind studied in process theories. Given its syntactic shape, must we conclude that 'Parallellism implies Undecidability'? Such a clear-cut outcome might be pleasing. Indeed, our negative result in Section 3 on free conjunctive guards said that *unconstrained parallellism leads to undecidability*. But the situation with collective action is more delicate, and much more can be said. For that purpose, we need to backtrack a bit from **GF**.

Basic modal logic has an intuitive distinction between *action predicates* Rxy that jump across accessibility links (from x to y), and *state predicates* Px making static assertions about the current state x. This distinction is obliterated in **GF** syntax, whose atomic predicates can serve indifferently for describing moves between states, or fixed states. But, by maintaining such a distinction, we can be more liberal with quantifier bounds – and in the limit, allow any conjunction at all. Henceforth, we distinguish between state atoms $Q\mathbf{x}$ and action atoms $R\mathbf{x}, \mathbf{y}$. The comma in action atoms separates *input states* on the left from *output states* on the right. The total language has both 'action formulas' and 'state formulas', whose syntax can be manipulated independently – as happens in propositional dynamic logic. Here are some concrete options for languages like this. We start with two sequential action formalisms.

GSAL1 Action formulas $R\mathbf{x},\mathbf{y}$

State formulas $Q\mathbf{x}$, Booleans, $\exists \mathbf{y}\, (R\mathbf{x}, \mathbf{y} \,\&\, \phi(\mathbf{y}))$

This 'guarded state-action language' describes transitions from an old state to a new one, but without any cross-comparison between old states and new ones. The input-output distinction has various effects. E.g., action atoms Rx, y are very different from their converses Ry, x . Moreover, the above restriction to only action-guarded quantifiers has the effect of making every formula depend on some initial tuple of free variables. Thus, all formulas in GSAL1 are 'local': there are no closed sentences. As in ordinary modal logic, 'satisfiability' then refers to local truth at some tuple of states in a model. 'Global satisfiability', truth at all tuples in a model, is a much more powerful notion. Next, if some input states are allowed to to persist as output, we need further atoms like Rx, yx, while quantifiers $\exists \mathbf{y}$ only range over the *new* components of the output state. Naturally, a matrix statement may now refer to these new **y** plus the persistent **x** . These additional syntactic features turn GSAL1 into a more expressive modal action language **GSAL2** . Both are effective parts of **GF**, and thus inherit its decidability. Note that their syntax has no explicit *operations on action predicates*. One may add certain *safe* operations, however (cf. Section 4) – mainly some forms of 'choice' and 'composition' – without increasing the expressive resources of these fragments.

5.2 Modal logics for parallel action

This was all 'sequential' action. Genuinely *parallel* versions enrich the action formulas by (unsafe!) *conjunctions*, while imposing various constraints on quantifier patterns. Quantifiers then collect all output states mentioned in conjunctions of atoms **&**R**x**, **y** . Moreover, to emphasize that the new objects form a coherent state, one may require the occurrnce of an atomic guard, either over the new **y**, or over the new **y** plus the persistent **x**. We list some options. But before proceeding, a warning may be in order. The purpose of all this variation is not to create a boring catalogue of formal languages – but rather, to demonstrate the effect of various expressive resources on decidability.

P-GSAL1	Action formulas	R**x**,**y**, conjunctions
	State formulas	Q**x**, Booleans, $\exists$**y** (**&**R**x**, **y** & ϕ(**y**))
P-GSAL1*	Action formulas	R**x**,**y**, conjunctions
	State formulas	Q**x**, Booleans, $\exists$**y** (**&**R**x**, **y** & Q**y** & ϕ(**y**))

As before, both languages allow only 'local' formulas, describing some tuple of states. The second fragment is obviously a part of the first. **P-GSAL2** and **P-GSAL2*** are defined analogously, but now allowing input states from **x** to reappear as output states. None of these languages lies inside **GF** (even though P-GSAL2* adds strong guards):

$\exists y_1y_2$ (Rx_1, y_1 & Rx_2, y_2 & Qy_1y_2) *is in P-GSAL1*, but not in GF*

$\exists y_1y_2$ (Rx_1, y_1 & Sx_2, x_2y_2 & $Qx_2y_1y_2$) *is in P-GSAL2*, but not in GF*

By a somewhat brute force argument, one can obtain the following result.

Theorem Satisfiability in P-GSAL1* is decidable.

Proof We start again from the decidability proof for **GF**, with a universe of 'types' (sets taken from the finite family of relevant formulas) satisfying suitable closure conditions.

From this, we constructed paths of types recording which formulas are true at any stage. We modify this idea slightly, allowing types that describe desired behaviour on only some subset of the variables. Transitions extending a path are triggered explicitly by existential formulas $\exists\mathbf{y}\,(\&R\mathbf{x}, \mathbf{y}\ \&\ \phi(\mathbf{y}))$ occurring in the last type so far, with the $\mathbf{y}$ 'changing their values' – while the new end-type only has formulas *with free variables among the* $\mathbf{y}$. As a result, the 'life-time' of the input variables $\mathbf{x}$ ends at such a step. In the model construction, we use objects (π, x) as before, where x is among the active variables at the end of the path π. For the interpretation of predicates, we set

(a) a state atom $Q\mathbf{d}$ is only true of a tuple of objects if these lie on the same path, and were introduced *simultaneously* at the final transition, whose result-type contains the atom with the variables of the $\mathbf{d}$ (in the same order)

(b) an action atom $A\mathbf{d}, \mathbf{e}$ is only true if all its objects lie on the same path, and the atom with the corresponding variables plugged in (as in (a)) occurred in the conjunctive action prefix of some transition.

Each path has an associated assignment s_π defined on the variables in the last and one-but-last types of the path, sending x to the object (π^*, x), where π^* is that subpath of π in which x was changed last. Clearly, action atoms will only hold between objects in the one-but-last and last stages. The Truth Lemma then says that

a (relevant) *state formula* ϕ holds under the assignment of a path iff ϕ literally occurs in the last type of that path

As in the original decidability argument for GF, there are two cases of major interest. (**1**) First, consider state *atoms* $Q\mathbf{x}$. If $Q\mathbf{x}$ is in the last type of π, then – by our restriction on result-types of path transitions – its variables were among those affected by the final change. So, we have the above condition for truth of the atom. Conversely, if $Q\mathbf{x}$ is true under s_π, this can only have happened by a simultaneous introduction on π, with $Q\mathbf{x}$ explicitly present. (**2**) Now consider *existential quantifiers* $\exists\mathbf{y}\,(\&R\mathbf{x}, \mathbf{y}\ \&\ Q\mathbf{y}\ \&\ \phi(\mathbf{y}))$. If the latter occurs in the final type, then it is true – by an argument as for GF: one looks at the obvious path extension triggered by the existential formula. The crucial case is when such a formula is true under s_π: while it should occur in the last type of π. Let some tuple $\mathbf{d}$ of objects satisfy the specified action predicates, plus the state guard $Q\mathbf{y}$ and the matrix statement $\phi(\mathbf{y})$. By the definition of true action predicates, the $\mathbf{d}$ must have been introduced following the end of the current path. Moreover, as the state atom $Q\mathbf{y}$

holds, they were introduced together in one transition, resulting in one final type Δ (i.e., they do not lie on separate forks) containing $Q\mathbf{y}$. Call this extended path π^+. Its s-assignment sends the variables $\mathbf{y}$ to the objects $\mathbf{d}$. By the inductive hypothesis then, $\phi(\mathbf{y})$ occurs in Δ, the last type of π^+. But then, by an obvious existential closure condition on quasi-models, $\exists\mathbf{y}\ (\&R\mathbf{x}, \mathbf{y}\ \&\ Q\mathbf{y}\ \&\ \phi(\mathbf{y}))$ occurred in the type before that, which was the final type of π. ♠

We think that P-GSAL1 (without the guard condition on new state tuples) is decidable, too. But the above proof method does not work, since there is no guarantee that the new states introduced by a true existential quantifier $\exists\mathbf{y}$ ($\&R\mathbf{x}$, $\mathbf{y}$ & form a 'simultaneous set' introduced in one parallel action step. (Different y might come from different steps.) On the other hand, various parts of the preceding proof seem to admit of generalisation. As for the two stronger languages P-GSAL2 and P-GSAL2*, we leave their decidability as an open question. Finally, note that the above proof is about local satisfiability only. It does not settle the decidability of *global* satisfiability (truth in all states of a model). This issue will return below.

<u>Remark</u> *Parallel Bisimulation*
Guarded bisimulations for GSAL may be extended to stricter bisimulations for the richer language PGAL. We need additional zig-zags for joint actions, with clauses like

if $\mathbf{a}\ E\ \mathbf{b}$*, and* $R\mathbf{a'c'}$*,* $S\ \mathbf{a''c''}$*, then there must be* $\mathbf{d'}, \mathbf{d''}$
with $R\mathbf{b'd'}$*,* $S\ \mathbf{b''d''}$ *such that* $\mathbf{c'c''}\ E\ \mathbf{d'd''}$

The above parallel languages are a new area for modal analysis. We noted several open questions of decidability. But also, their model theory remains to be explored.

6 The Danger Zone: Grids and Tiling Problems

Let us now approach these issues from a different angle, and see where undecidability strikes for sure. We will use insights on this matter from Spaan 1993, Marx 1997.

6.1 Encoding tiling problems

Consider the embedding of *tiling problems*. The undecidable task is to put coloured tiles on the infinite grid NxN, with some finite set of colours, and tiles having four coloured edges, subject to the constraint that adjacent tiles have the same colour along their boundary. First-order formulas expressing the relevant constraints have a definite P-GSAL flavour, with actions N (*go one step north*), E (*go one step east*) and state predicates Cx for the colours. Here are some examples. Adjacency of colours can be expressed by straightforward universal conditions of the form

$$\forall x: \quad \forall y\ (Nx, y \rightarrow (C_1x \rightarrow \bigvee C_2y))$$
$$\forall x: \quad \forall y\ (Ex, y \rightarrow (C_1x \rightarrow \bigvee C_2y))$$

where the unary predicates C_i describe the various possible kinds of tiles. General behaviour of colours is expressed by conditions of the form

$\forall x$: "at least and at most one C holds of x"

Next, the crucial *grid pattern* seen from x is expressed by the assertions

$$\forall x: \quad \exists y\ Nx, y \qquad\qquad \forall x: \quad \exists y\ Ex, y$$

and more importantly,

$$\forall x: \quad \forall yz\ ((Nx, y\ \&\ Ex, z) \rightarrow \exists u\ (Ey, u\ \&\ Nz, u))$$

These assertions lie in P-GSAL1, modulo one unbounded universal quantifier in front. Let us call their conjunction TILE. Now it is not hard to prove the following

<u>Fact</u> NxN has a tiling iff TILE is satisfiable.

<u>Proof</u> Here is a sketch (for detailed arguments of this kind, cf. Blackburn, de Rijke & Venema 1998). Clearly, if a tiling exists, NxN *itself*, suitably expanded, verifies TILE. Conversely, consider any model for TILE. It is easy to define a map f from NxN, sending the origin to any point in the model, with the following property:

if y is a northern (eastern) neighbour of x, then N f(x), f(y) (E f(x), f(y))

To see this, use the last three formulas above repeatedly to construct a grid of squares x N y E u, x E z N u, which provides all necessary f-values. Then, a colouring for NxN meeting all constraints can be copied from the C-behaviour of the f-values. ●

6.2 Exactly what causes undecidability?

This result tells us that *the expressive power of parallellism comes close to encoding grids*, and hence undecidable problems may arise. But the encoding does not quite lie in P-GSAL1. We need *one unbounded universal quantifier in front* to make TILE work – and the latter's dangers are well-known. Spaan 1993 shows how decidable modal logics can become undecidable with this simple addition. She states this in terms of adding a 'universal modality' to the logic, but also observes that one such modality in front, i.e., our earlier *global satisfiability*, would do the harm already. An alternative would use only those points (in models for TILE) reachable from some fixed origin by a finite number of E, N steps. This uses *transitive closure* of the relation $N \cup E$, which is again outside our fragments – and even more dangerous for decidability, as it can embed the Σ^1_1-hard problem of 'recurrent tiling'. Thus, a mixture of encoding grids plus some weak form of universal prefix quantification will make process logics undecidable. Nevertheless, things remain delicate. Adding one universal quantifier up front to the non-conjunctively-bounded Guarded Fragment does no harm! (Cf. van Benthem 1997B for similar observations on formalisms in extended modal logic.)

<u>Fact</u> Satisfiability in **GF** with one universal prefix quantifier is decidable.

<u>Proof</u> Start with any type containing a few universally quantified guarded formulas x $\forall x\, \phi(x)$. Add all instances $[u/x]\phi$ (for the relevant variables u) to the types in the quasi-model. The original tree-model construction will still work as it stands – and it is easy to show that ϕ will hold for all tuples of 'path objects' of the form (π, u).

Recall that minimal modal logic plus a 'universal modality' remains decidable. Thus, it is the mixture of parallellism and universal quantification that generates undecidability. As to extensions of our observation about **GF**, Marx 1997 presents undecidable modal logics with characteristic *universal Horn* frame conditions. Therefore, allowing universal prefix quantification over larger tuples seems problematic already.

<u>Remark</u> The formulas in TILE did not satisfy the syntactic constraint of the language *PGSAL1** , that new objects in quantification must come simultaneously *guarded* by some state predicate Q. But we can modify the definition of TILE by using a trivial unary predicate P at *all* points, as well as a trivial binary predicate Q at all point pairs:

$$\forall x: \quad Px \qquad\qquad \forall xy: \quad Qxy$$

Without the (double) universal prefix quantifiers allowing this trivial obedience, it is unclear how to modify the necessary grid encoding, and get things right for proper tiling within the syntactic constraint on outputs imposed by $PGSAL1^*$.

Summing up, parallel constructions (with conjunctive guards) flirt with undecidability. On the other hand, they need not do so in general (witness the decidability of PGSAL1), and they seem harmful mainly in league with universal prefix quantifiers. We leave the intermediate possibilities alone here. We hope to have shown at least how guarded analysis can probe the effects of expressive power on decidability in a sensitive manner.

7 Model Theory: Simulations and Splitting Expressive Power

In this section, we explore the outline of a model theory for our extended formalisms. The results stated here are generalisations of ones already known for basic modal logic, and the proofs of the relevant results in van Benthem 1996, chapters 4, 5 largely go through, with some straightforward obvious modifications. Therefore, we omit details.

7.1 The state–action split in model theory

In addition to decidability, the above fragments have other interesting logical features. Consider the central notion of *bisimulation*. First, the split between *state predicates* and *action predicates* may be given a concrete meaning in standard first-order logic by assigning them different roles in guarded bisimulations. Action predicates regulate the picking of suitable object tuples in back-and-forth moves, while state predicates determine the 'quality' of what counts as a 'partial isomorphism'. (This difference of two meaningful roles is of more general interest, as standard first-order logic seems highly uniform in its treatment of non-logical vocabulary.) In the other Gestalt of this paper, one can also design various modal languages incorporating this distinction. Our example is a multi-state version of propositional dynamic logic, to be defined below.

7.2 Case-study: multi-state propositional dynamic logic

Consider sequential actions performed on 'collective states' with many components. This requires a shift from binary transition relations to general *finitary relations* $R\mathbf{xy}$ between finite tuples of individual states. One modal language for this is a many-dimensional one, with two components: state predicates, and action predicates. The new system $\mathbf{PDL}^*$

requires a two-level syntax, as for PDL, plus some book-keeping of arities for both levels (position numbers, or with variables themselves as 'positions').

Assertions

State atoms $P\mathbf{x}$, all Boolean operations, existential modalities $\langle R\rangle_{\mathbf{x},\mathbf{y}}$ taking **y**-state formulas to **x**-state formulas, and 'lifters' $[\phi, T]_{\mathbf{z}}$ (from **x**-state formulas ϕ to **x**+**z**-state ones).

Programs

Action atoms $R_{\mathbf{x},\mathbf{y}}$, relation composition (with arity fit), union (with arity fit), tests (ϕ)**?**, projections $\Pi_{\mathbf{x},\mathbf{y}}$ (from **x** to some subset **y**).

Note that both formulas and programs of this language translate immediately into the Guarded Fragment. So, we can either view it as a modal formalism, or as a piece of first-order logic. Either way, models and the truth definition are obvious. In particular, the lifter holds at an **x**+**z**-tuple if ϕ holds at its **x**-subtuple. It is also straightforward to generalize the original modal bisimulation (which still works for PDL) to an appropriate notion of guarded bisimulation, slightly modifying the back-and-forth clauses stated in Section 2.1 when analysing **GF**. (A small technical feature here: one has to close the relevant families under sub-partial isomorphisms.)

We will briefly discuss the further theory of PDL^*, which may be viewed as either 'modal' or 'classical' model theory demonstrating some interesting themes. To prove an Invariance Theorem for our two-level language, just as with PDL, we also have to identify another basic notion (cf. van Benthem 1996, Chapter 5). As before, we call a first-order formula $\phi(\mathbf{x})$ *invariant for guarded bisimulations* E if,

whenever $\mathbf{a}\ E\ \mathbf{b}$, then $\mathbf{M} \models \phi(\mathbf{a})$ iff $\mathbf{N} \models \phi(\mathbf{b})$

But next, we call a first-order formula $\pi(\mathbf{x}, \mathbf{y})$ *safe for guarded bisimulations* if,

whenever E is a guarded bisimulation (zigzagging for the basic action predicates of the language), the above zigzag clauses hold automatically for the new relation defined by π in the models $\mathbf{M}$, $\mathbf{N}$

Thus, safe formulas define transition relations that 'stay inside' our simulation semantics, i.e. our process realm. The following basic property of PDL^* is proved by a simultaneous induction on formulas and programs.

Proposition

(1) All formulas are invariant for guarded bisimulations.

(2) All programs are safe for guarded bisimulations.

An adaptation of a known argument for modal logic shows a converse result as well.

Invariance Theorem

For all first-order formulas ϕ, the following assertions are equivalent:

(1) ϕ is invariant for guarded bisimulations

(2) ϕ is definable in PDL^*

Another modally inspired proof (cf. van Benthem 1998C) captures the safe operations. This amounts to *expressive completeness* for the key operations in the above language.

Safety Theorem

The safe operations are precisely those definable using
(1) atomic action predicates, (2) tests for arbitrary state formulas,
(3) projections, (4) relation composition, and (5) union.

We can vary a bit on this syntactic description. Instead of having all tests, just atomic ones will do, if one adds an 'impossibility negation' ~ on actions. Essentially, the safe programs describe unions (OR-trees) of finite sequences of multi-states linked by action steps or projections, with test assertions interspersed. The model theory of PDL^* is a blend of 'modal ideas pursued by first-order means'. Guarded bisimulation is like plain bisimulation, though a bit more difficult to visualise, as matches are between finite tuples of states. There is an *unraveling* method creating tree models – involving paths

<atom R**a**, **b**, selected object b_i , atom S**b'**, **c**, etcetera>

This can be used for various purposes, amongst others for interpolation and preservation properties. Here is a sample result, used in proving the Safety Theorem. A formula $\phi(Q)$ is *totally distributive* in the displayed *state predicate* Q if its truth for the union of any family $\{Q_i \mid i \in I\}$ is equivalent to that for some Q_i separately.

<u>Distribution Theorem</u>

A formula is totally distributive in the state predicate $Q\mathbf{x}$ iff it can be defined in the form $<\pi>Q\mathbf{x}$, with π a safe program as above whose test conditions on intermediate states do not involve the predicate Q.

PDL^* is decidable, because **GF** is. It even has an effective Finite Model Property, since it lies inside a simple fragment of **GF** with 'distinguished guards' for which Andréka, van Benthem & Németi 1998 provide an effective decidability argument. Valid principles are much as in PDL itself. Several methods for completeness exist (many-dimensional modal logic, algebraic representation, or proof-theoretic modification of decidability proofs). As with PDL or **GF**, there is also an interest in adding general fixed-point operators, and especially, ones that can be reached in ω steps. In our first-order Gestalt, PDL-style operators suffice for all ω–fixed points $\mu Q \bullet \phi(Q)$ that can be computed with a matrix formula $\phi(Q)$ involving *one* suitable occurrence of the atom $Q\mathbf{x}$. Semantically, general ω-stability follows from *Finite Distribution*, i.e.,

ϕ holds of Q iff it holds of some finite subpredicate Q_0

The latter allows forms of definition with a finite number of suitable occurrences of Q. Full first-order logic has this syntactic normal form for finite-distributive operators:

$\mu Q \bullet \phi(Q)$ where the occurrences of Q-atoms in ϕ lie only in the scope of logical operators $\vee, \wedge, \exists$

For PDL^*, a similar syntactic classification exists, of finite distributivity for state predicates. It involves *finite action trees*, being AND-trees whose steps are safe actions, and whose nodes may carry both Q-free test conditions and atomic tests involving Q.

<u>Finite Distribution Theorem</u>

For state-formulas ϕ, the following two assertions are equivalent:

(1) ϕ is finitely distributive in Q,

(2) ϕ says there exists one out of some set of finite action trees.

We have a simple quasi-model proof on probation to the effect that PDL^* extended with fixed-point operators for state predicates defined by the above operations is decidable. (It generalizes the standard Fisher-Ladner filtration argument for PDL.) But see the earlier positive news about fixed-point extensions of the Guarded Fragment. One open questions is whether they can also be obtained by direct quasi-model-style arguments.

7.3 Further issues in modal model theory

We conclude with some further issues in modal logic that seem to have a more general model-theoretic interest. First, in modal logic, one often encounters two related versions of basic results. For instance, modal interpolation theorems state that

if $\phi \models \psi$*, then there exists an interpolant* α *with* $\phi \models \alpha \models \psi$
which lies in the 'joint language' of ϕ *and* ψ

The latter may either refer to the joint vocabulary of proposition letters, or also to the joint modalities indexed by actions. Also, Los-Tarski theorems may characterize preservation, either when dropping worlds from a model, or when dropping arrows from its accessibility relation. This split between state predicates and action predicates returns in our more general languages. For instance, the above discussion of PDL^* had preservation theorems for semantic distributivity w.r.t. state predicates. But there are similar (open) questions concerning action predicates. This split also has repercussions for other basic semantic notions, like monotonicity. One final example was already mentioned in Section 4. There are two natural kinds of *fixed-point operator*: one for state predicates, and one for action predicates. The two turned out to be different.

Remark *Boosting via bisimulation*

Also, well-known modal representation and completeness theorems suggest new standard notions and results. Consider the 'model surgery' that occurs in many modal completeness arguments. One finds a simple (Henkin) countermodel to some non-theorem ϕ, and then constructs a bisimulation equivalent (where modal ϕ still fails) satisfying some desired extra feature defined by, say, α. Behind this technique lies an *existential preservation property*, different from the usual universal versions:

whenever $\mathbf{M} \models \phi$, there exists a bisimilar model $\mathbf{N} \models \phi \,\&\, \alpha$

'Boosting via bisimulation' is a new notion of general interest (cf. the paper on 'Information Links and Logical Transfer' in van Benthem 1998A).

Another set of open questions arises when we move from *sequential* to *parallel* modal formalisms allowing conjunctive guards. In that case, our simulations must be extended with new clauses, and the above basic model theory of modal invariance and safety (van Benthem 1996, Chapters 4 & 5) needs to be redone. In particular, can one find *expressively complete* sets of natural modal operations for parallel actions?

Digression The notion of 'partial isomorphism' needs change, too, due to the special status of *identity* in our fragments. Identity statements $\exists y\ (Rx_1x_2, y\ \&\ \ldots\ \&\ y=x_1\ \&\ \ldots)$ circumvent the distinction between input and output states, and their effect is hard to predict. But without identity, bisimulation must be adjusted, even for **GF** itself. The basic building blocks will now be binary relations between finite tuples of objects of the same length – or alternatively, binary relations between finite variable assignments.

We conclude with some more general issues behind the above language constructions. There is a general spectrum of *correspondences between simulations and languages*, running from 'modal-logic/bisimulation' to 'first-order-logic/potential-isomorphism'. This needs to be understood more generally. In particular, why are the modal fragments of first-order logic chosen on this spectrum usually so well-behaved? Do the specific choices that people make perhaps obey some implicit *transfer principles for a good meta-theory*? (Caveat. A warning example is the recent discovery reported in Hoogland and Marx 1998 that Craig interpolation fails for **GF**. What is the general picture?)

Even in this discursive format, with more questions than answers, we hope to have shown that modal logic engenders interesting novel themes for standard logic.

8 Proof-Theoretic Alternatives

For the record, we note that generality in modal logic can also come from proof-theoretic, rather than model-theoretic considerations. Here are two illustrations.

8.1 Resolution

Decidability of modal languages may also be analysed in a computational perspective. There are new *resolution strategies* for **GF**, providing a complete but finite search space (De Nivelle 1998), using Skolemisation techniques plus sophisticated proof strategies. A theorem prover 'Bliksem' incorporating these reached second place over-all at the international competition CADE, Konstanz 1998. Here, the emphasis is not so much on the syntax of modal languages as on correctness and termination of specific proof strategies.

This is a really different approach, based on algorithmics rather than syntax or semantics, to what makes modal decidability tick. Our second illustration is in the same vein.

8.2 Contraction

It is easy to show (Andréka, van Benthem & Németi 1998) that basic modal logic can be axiomatized completely with the usual Gentzen introduction rules for the logical operators plus all structural rules *minus Contraction*. This follows from a simple reduction method for valid modalized/atomic sequents. For stronger modal fragments, *effectively finitely bounded* versions of contraction often suffice. This observation again suggests an independent proof-theoretic perspective. As is well-known, in linear logic, one 'shuts off' the contraction rule, and then sees what (decidable) logics remain. What we observed here is that basic modal logic is insensitive to this shift: no validities are lost. Moreover, generalized modal languages can make do with effectively limited contraction without losing validities. So, we ask *just which* fragments of full classical logic can do with effectively limited forms of contraction (keeping the search space finite). Will the outcomes of this query match up with the results of guarded analysis?

9 A Summary of General Themes

What we have advocated are the virtues of *general translation* and adopting a *tandem approach*. We do note that this should be done with care. Our 'standard translations' enshrine one particular view of the semantics for a modal language, and hence, they may encourage undue conservatism. These issues were hotly debated in the seventies: cf. van Benthem 1977 on intrinsic versus translationist views of *temporal logic*. For a contemporary example, in the 'logic of proofs' of Artemov 1998, the box modality [] is not a *universal quantifier* (over all accessible worlds), but an *existential* one (running over available proofs). But when well-defined, such alternative views, too, can always be 'translated'. Also, modal translations need not run into first-order logic. For instance, when translating Beth semantics for intuitionistic logic, one will naturally encounter *second-order* quantification over 'bars' of nodes across a tree. Here too, translation may still be useful, because it forces one to rethink the given semantics. Do we really want this second-order version, or rather a *many-sorted first-order* one treating nodes, bars and branches on a par as first-class semantic citizens? (Van Benthem, van Eijck & Stebletsova 1995 make a similar point concerning process logics with states and paths.) And thus, 'translationism' need not be a conservative force after all.

Next, we have emphasized the duality between *language design* and the search for *characteristic simulations*. There are no qualifications here: this is just a Good Thing. Then, in this language design, we stressed the importance of quantifier fine-structure, especially that involving *guards*. Our claim is not that this gives us a miracle cure explaining every form of decidability in modal logic. Our discussion of minimal logics versus extra frame conditions has shown clear limitations to the guarded approach – but also some surprising extensions (witness the discussion of transitivity and fixed-points in Section 4). Then, we have advocated the use of *concrete metaphors* in extending the range of modal logic, in particular, a *dynamic perspective* with new distinctions between state versus action predicates, and sequential versus parallel actions. Thus we are traveling in a landscape of modal languages, where we want to study general phenomena, rather than enjoy the attractions of any particular spot forever. This landscape also has its exciting features, such as *undecidability thresholds*, occurring in a generally undetectable manner (it is undecidable if a given modal logic is decidable: Chagrov & Zakharyashev 1993), much like the deep cracks in the ice-cap of Antarctica. This perhaps outlandish methodology of 'landscapism' (cf. Moss' 1998 review of van Benthem 1996) puts broad logical phenomena in focus as our real topic of research, rather than – *pace* our Uppsala ancestor Linnaeus – the usual 'botany of modal logics'.

Nevertheless, this paper has offered no *definition* of Modal Logic. The most I will say here is this. Our field is concerned with the balance between expressive power and complexity in designing logical systems. This is not a minor issue. If there are universal conservation principles underlying logic (as I myself believe: cf. van Benthem 1997C), then one must surely be some kind of Golden Rule inversely relating expressive power and complexity. Our investigations in this paper are about just that subtle relationship.

REFERENCES

H. Andréka, J. van Benthem & I. Németi, 1998, 'Modal Logics and Bounded Fragments of Predicate Logic', *Journal of Philosophical Logic* 27:3, 217–274.

S. Artemov, 1998, "Explicit Modal Logic', in *Proceedings AiML-II*, Philosophical Institute, Uppsala, 22 – 31.

J. van Benthem, 1976, *Modal Correspondence Theory*, Ph.D. Thesis, Mathematical Institute, University of Amsterdam.

J. van Benthem, 1977, 'Tense Logic and Standard Logic', *Logique et Analyse* 20, 41-83.

J. van Benthem, 1985, *Modal Logic and Classical Logic*, Bibliopolis, Napoli.

J. van Benthem, 1996, *Exploring Logical Dynamics*, CSLI Publications, Stanford, (distributed by Cambridge University Press).

J. van Benthem, 1997A, 'Dynamic Bits and Pieces', Report LP-97-01, Institute for Logic, Language and Computation, University of Amsterdam.

J. van Benthem, 1997B, 'The Range of Modal Logic', Report ML–97–05, Institute for Logic, Language and Computation, University of Amsterdam. To appear in the *Journal of Applied Non-Classical Logics*, memorial issue for George Gargov.

J. van Benthem, 1997C, 'Wider Still and Wider: resetting the bounds of logic', Report LP–97–08, Institute for Logic, Language and Computation, University of Amsterdam. To appear in A. Varzi, ed., *The European Review of Philosophy*.

J. van Benthem, 1998A, ''Dynamic Odds and Ends'', Report ML-98-08, Institute for Logic, Language and Computation, University of Amsterdam.

J. van Benthem, 1998B, 'Process Operations in Extended Dynamic Logic', *Proceedings Logic in Computer Science (LICS 98)*, Indianapolis, IEEE Publications, Los Alamitos, 1998, 244–250.

J. van Benthem, 1998C, 'Programming Operations that are Safe for Bisimulation', *Logic Colloquium. Clermont-Ferrand 1994*, Studia Logica 60:2, 311-330.

J. van Benthem, J. van Eijck & V. Stebletsova, 1995, 'Modal Logic, Transition Systems and Processes', *Journal of Logic and Computation* 4:5, 811-855.

P. Blackburn, M. de Rijke & Y. Venema, 1998, *Modal Logic*, textbook, ILLC Amsterdam & computer linguistics, Saarbruecken.

E. Börger, E. Grädel & Y. Gurevich, 1996, *The Classical Decision Problem*, Springer, Berlin.

A. Chagrov & M. Zakharyaschev, 1993, 'The Undecidability of the Disjunction Property of Propositional Logics and Other Related Problems', *Journal of Symbolic Logic* 58, 49 – 82.

E. Grädel, 1997, 'On the Complexity of the Guarded Fragment', Department of Informatics & Mathematics, RWTH Aachen.

E. Grädel & I. Walukiewicz, 1999, 'The Guarded Fragment with Fixed Points is Decidable', RWTH Aachen, to be presented at LICS 1999.

M. Hollenberg, 1998, *Logic and Bisimulation*, Ph.D. Thesis, Philosophical Institute, Utrecht.

E.Hoogland & M. Marx, 1998, 'On the Failure of Interpolation for the Guarded Fragment', manuscript, ILLC, University of Amsterdam.

N. Kurtonina & M. de Rijke, 1997, 'Simulating Without Negation', *Journal of Logic and Computation* 7:4, 501-522,

M. Marx, 1997, 'Complexity of Modal Logics of Relations', Report ILLC-ML-97-02, Institute for Logic, Language and Computation, University of Amsterdam.

M. Marx & Y. Venema, 1996, *Multi-Dimensional Modal Logic*, Kluwer, Dordrecht.

L. Moss, 1998, review of "Exploring Logical Dynamics", to appear in *Journal of Logic, Language and Information*.

H. J. Ohlbach & J. Koehler, 1998, 'Modal Logics, Description Logics and Arithmetical Reasoning', in *Proceedings AiML-II*, Philosophical Institute, Uppsala, 231–255.

M. de Rijke, 1993, *Extending Modal Logics*, Ph.D. Thesis, Institute for Logic, Language and Computation, University of Amsterdam.

E. Spaan, 1993, *Complexity of Modal Logics*, dissertation, Institute for Logic, Language and Computation, University of Amsterdam.

M.Y. Vardi, 1997, 'What Makes Modal Logic so Robustly Decidable?', in *Descriptive Complexity and Finite Models*, American Mathematical Society.