Complexity of Modal Logics of Relations

Maarten Marx
Department of Computing, Imperial College, London
email: mm28@doc.ic.ac.uk

October 17, 1997

Abstract

We consider two families of modal logics of relations: arrow logic and cylindric modal logic and several natural expansions of these, interpreted on a range of (relativised) model-classes. We give a systematic study of the complexity of the validity problem of these logics, obtaining price tags for various features as assumptions on the universe of the models, similarity types, and number of variables involved. The general picture is that the process of relativisation turns an undecidable logic into one whose validity problem is exptime-complete. There are interesting deviations to this though, which we also discuss. The numerous results in this paper are all directed to obtain a better understanding why relativisation can turn an undecidable modal logic of relations into a decidable one. We connect the semantic way of "taming logic" by relativisation with the syntactic approach of isolating decidable so-called guarded fragments by showing that validity of loosely guarded formulas is preserved under relativisation.

The author is supported by UK EPSRC grant No. GR/K54946.

Contents

1 Introduction

Modal logics of relations forms a major part of the subject of the recent book (Marx and Venema, 1997). This paper can be seen as an additional chapter to that book in which we are concerned with the computational complexity of these logics. In the book we only looked at decidability vs undecidability. As it turned out, the decision algorithms provided there are in several cases just about right: they provide an upper bound for which we will find a matching lower bound here.

The two main logics of this paper –arrow logic and cylindric modal logic– stem from (Venema, 1989; Venema, 1992). The modal perspective on the algebraic classes of representable relation and cylindric algebras was fully developed there, but the concentration was only on the "square" classical case. The modal perspective however offers a broad range of semantics different from the classical one. Following the modal tradition one immediately tends to study the following lattice:

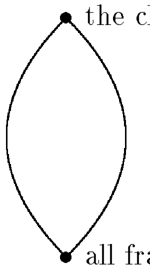

The logics siting at the top of this lattice are almost always undecidable, while the logics at the bottom are –by general modal considerations– always decidable. As it turned out, there are several frame classes below the top –whose logic still rightly can be called a logic of relations– which have a decidable modal theory. The search and study of these classes became known under the headings "relativisation" and "taming logic", a topic to which we turn just below.

This paper can be seen as an attempt to tackle the question *"Why does relativisation work?"*, in the sense that it turns an undecidable logic into a decidable one. We will see that indeed a complexity-analysis of these logics helps in understanding this process.

Modal logic of relations

We briefly describe what we mean by this term; concrete examples are introduced in section 1.2. In "standard" modal logic, formulas are interpreted at worlds/states/points which are elements of some universe endowed with several (accessibility) relations. The meaning of a formula in such a Kripke model is a subset of the universe (the set of worlds where the formula holds). In other words, a modal logical formula describes a *property* of worlds. This show up in the range of the standard translation of modal logic to first order logic: propositional variables go to monadic relation symbols, and all formulas are translated to first order formulas in one free variable.

In modal logic of relations, the meaning of a formula is a relation (i.e., a set of n-tuples, for some n). Let's discuss this for the example of binary relations. The worlds in a Kripke model are then pairs (u, v), so the meaning of a proposition letter p in a model $\mathfrak{M}$ is the set $\{(x, y) \mid \mathfrak{M}, (x, y) \Vdash p\}$, whence it behaves as a genuine binary predicate.

Because the worlds now have internal structure, one can define certain accessibility relations in terms of this structure. We list some examples.

unary identity relation	$I(x, y)$	iff	$x = y$
unary converse function	$f((x, y))$	$=$	(y, x)
unary domain function	$d((x, y))$	$=$	(x, x)
binary chancing coordinate relation	$T_0(x, y), (x', y')$	iff	$y = y'$
ternary composition relation	$C(x, y), (u, v), (z, w)$	iff	$x = u,\ v = z \;\&\; y = w.$

All of these relations have an intuitive geometric interpretation which we will see later on. Let us briefly look at the standard translation of a formula created with the modality $\Diamond_0$ (which has T_0 as its accessibility relation). The obvious choice for translating $\Diamond_0 p$ is $\exists x'y'(T_0(x, y)(x', y') \wedge P(x', y'))$ which is equivalent when we write out the definition of T_0 to $\exists x' P(x', y)$. In other words, $\Diamond_0$ is really the same as the existential quantifier (plus attached variable) $\exists x'$. Moreover, the binary worlds play now the same role as *assignments* in first order logic. In this way we can give a modal account of first order logic, cf. section 1.2. There are much more examples in the literature. For now we delimit the field as follows. A *modal logic of n-ary relations* is a multi-modal logic (where we can also have modalities of rank different from 1) defined by a class of Kripke frames $(W, R_i)_{i \in I}$ which satisfy the following properties

A the elements of W are n-tuples. That is W is a subset of ${}^nU$ for some set U;

B the relations R_i can be retrieved from W, that is they are defined in terms of the tuple-structure of the elements of W. (This implies that two frames are equal iff they have the same universe.)

All accessibility relation defined above satisfy the last criterion. As base cases we have the modal logic of the universal modality and of the difference operator. Sometimes the second criterion is weakened, cf. the products of (Gabbay and Shehtman, 1997). We said that the relations should be defined using the structure of the elements, because we want to exclude examples where we have an implicit structure on the set U (say it is the natural numbers) and we use that structure to define the relations (like $<$). In the cases we deal with here U is always an unstructured set.

Relativisation

Let K be a class of frames of some modal similarity type. Let $ML(\mathsf{K})$ denote the modal theory (modal logic) of K in the language of the same type. Roughly speaking the relativisation of $ML(\mathsf{K})$ is $ML(\mathsf{SK})$, where SK is the class of all *substructures* of elements in K. From now on let K be elementary (i.e., first-order definable). (This is no real restriction, since all frame-classes in this paper are elementary.) Then, also SK is elementary and axiomatized by the set of all universal consequences of (the first-order theory of) K. In modal lingo, $ML(\mathsf{SK})$ is a *subframe logic*.

Is this the reason that relativisation works? That is, are subframe logics decidable? No, (Spaan, 1993) contains an example of an undecidable canonical bimodal subframe logic which is basic elementary (cf., also Theorem 2.5 in here). The logic of Spaan is not a logic of relations, but we do have examples of undecidable relativisations of modal logics of relations (Theorems 2.19 and 2.22).

If we look closer at the decidable relativised classes of modal logics of relations which appear in the literature we see however often the following phenomenon:

$$\begin{aligned} ML(\mathsf{SK}) \;&=\; ML(\mathsf{S}\mathbf{P}\mathsf{K}) \;= \\ &\quad ML(\text{some Sahlqvist axiomatisable universal Horn definable class of frames.}) \end{aligned}$$

(Here $\mathbf{P}$ stands for taking products of structures.) By Corollary B.4.5 in (Marx and Venema, 1997) this implies that $ML(\mathsf{SK})$ has the interpolation property, and that the (modal) variety generated by the complex algebras of frames in SK has the strong amalgamation property. This last notion –which intuitively states that we can always combine two structures in some class which have a common substructure, into a larger structure *in the same class*– is crucial in all the decidability proofs in the literature. We will meet it here in the proof of Lemma 3.1.

Could we perhaps generalise, and say that relativisation leads to decidability for these logics because they are universal Horn Sahlqvist logics? The answer is again no, but less stringent as the previous. Section 2.1 contains several undecidable universal Horn Sahlqvist logics. They are all based on Spaan's example and again not modal logics of relations. We do not have examples of undecidable universal Horn Sahlqvist logics of relations.

So why does it work then? Or at least why does it work for modal logics of relations? It seems that there are just too many parameters involved to give a general answer. (Andréka et al., 1996) gives a criterion on the truth-definition of modalities –being existentially packed– which leads on relativised models to decidability and even the finite model property. Such modalities allow the amalgamation of finite structures. Although this covers most of the studied systems there are still decidable modal logics of relations which do not fall under their heading, but their decidability can be proved using the same "mosaic-amalgamation" technique. Here we will provide a systematic study in terms of complexity of the computational costs of the three main features of modal logics of relations:

- similarity types involved;
- conditions on the universes of our classes (i.e., what is the cost in terms of complexity of allowing some existential conditions on our frame classes?);
- the arity of the relations, and the addities of the modalities involved (in a certain sense, in which finite variable fragment of first order logic is the modal logic "living"?).

We believe this indeed provides insights into the reasons why relativisation leads to decidability and shows the rough borders of the technique. After we have defined our logics we will give an overview of our results along the lines of these three features in section 1.3. In section 3.2 we connect our results with van Benthem's loosely guarded fragment of first order logic. It turns out that for these formulas, validity is unaffected by turning from classical to relativised models.

Overview of paper

We assume a basic knowledge of modal logic and complexity theory. All our modal logics are defined semantically in terms of Kripke frames. We hardly ever deal with axiomatic systems, and when we do we will be sparse in explanations. In general we follow the standard definitions in (multi-)modal logic. For precise definitions of how we use our notions we refer the reader to the appendix of (Marx and Venema, 1997). We give a brief overview of complexity theory in the next subsection.

The rest of the paper is organised as follows. In the remainder of this section we recall the definitions of the two main systems of this paper – arrow logic and cylindric modal logic– from (Marx and Venema, 1997), say a few words about extra connectives, and finish with an overview of our results. The following two sections establish our central complexity results. We divided them into lower and upper bound results. Section 3.2 makes the connection between (relativised) modal logics of relations and first order logic, and shows how the loosely guarded fragment can be seen as a relativised modal logic of relations. Section 4 deals with the rather special case of cylindric modal logic of dimension 2. Even here we have systems in several complexity classes. In the last section we take an axiomatic point of view and show which axioms "cause" undecidability.

1.1 A few words on complexity

Decision problems. Connected with a semantically presented modal logic of a class of frames K are several natural decision problems. We will study the following.

Satisfaction problem: given a formula φ, can φ be satisfied in a state in a model over a frame in K?

Global satisfaction problem: given a formula φ, can φ be satisfied in every state in a model over a frame in K?

Validity problem: given a formula φ, is φ valid, that is does φ hold in every state in every model on every frame in K?

Global consequence problem: given formulas φ and ψ, does φ globally entails ψ, that is, is ψ true in every model on a K-frame where φ is true? This is also called the *valid rule* problem (i.e., is φ/ψ a K-valid inference rule).

The satisfaction problem and the validity problem are each others complement: φ is satisfiable iff $\neg\varphi$ is not valid. The global satisfaction problem is partly connected to the global consequence problem: φ is globally satisfiable iff φ doesn't globally entail $\perp$.

The global satisfaction and consequence problem are not very often studied for modal logics, mainly because the emphasis is on the local consequence relation (φ locally entails ψ if in any model, every state where φ holds, is a state where ψ holds). With the local consequence relation, the implication trivially gives us a deduction theorem, so the validity and the consequence problems collapse into one. With the global consequence relation this need not be the case. In section 2.1 we will see examples of logics with a decidable satisfaction (whence validity) problem, but an undecidable global satisfaction (whence global consequence) problem.

The distinction also naturally arises when considering the standard translation of modal logic to first order logic. The satisfaction problem corresponds to satisfying the first order *formula* with one free variable, while the global satisfaction problem corresponds to satisfying the *sentence* obtained from the translation by universally quantifying over the free variable.

Obviously, for practical applications of a logic it is the complexity of the consequence relation that one is interested in. It might depend on the specific application whether this should be the global or the local one.

Complexity. Complexity analysis for modal logics is a well-developed field and we won't attempt to review it here. Rather we will point to the relevant literature and very briefly recall some basic facts. For general computational complexity theory we refer to (Papadimitriou, 1994). Classical papers on complexity of modal logics are (Ladner, 1977; Fisher and Ladner, 1979; Halpern and Moses, 1992). Especially the last one forms a good introduction for modal logicians. An excellent presentation of complexity in modal logic is given in the book (Blackburn et al., 1997). This is also a good introduction to the more advanced (Spaan, 1993).

We now briefly recall some basic facts, quoting freely from (Halpern and Moses, 1992). In this paper we are interested in the inherent difficulty of deciding for a given formula one of the above mentioned problems. Since all these problems are yes/no problems, we can view them as the difficulty of determining membership in a set. Thus the validity problem is viewed as the problem of determining whether a given formula φ is a member of the set of valid formulas with respect to a certain class K. The difficulty of determining set membership is measured in terms of time and/or (memory)space required to do this, as a function of the size of the input, in our case the length of the formula φ. Here we will only be concerned with complexity classes like P, NP, PSPACE, EXPTIME and NEXPTIME: those sets such that determining whether a given element x is a member of the set can be done in deterministic polynomial time, nondeterministic polynomial time, deterministic polynomial space, deterministic exponential time, and nondeterministic exponential time, respectively (as a function of the size of x). These classes are ordered as P$\subseteq$NP$\subseteq$PSPACE$\subseteq$EXPTIME$\subseteq$NEXPTIME, and it is known that P$\neq$EXPTIME.

Given a complexity class $\mathcal{C}$, the class CO-$\mathcal{C}$ consists of all the sets whose *complement* is a member of $\mathcal{C}$. It is easy to show that for deterministic classes $\mathcal{C}$ =CO-$\mathcal{C}$, while it is believed that for non-deterministic (time)classes they are different. As mentioned above, satisfiability and validity are complementary problems, so determining that satisfiability is in $\mathcal{C}$ immediately implies that validity is in CO-$\mathcal{C}$.

Roughly speaking, a set A is said to be *hard* with respect to a complexity class $\mathcal{C}$ if every set in $\mathcal{C}$ can be *effectively reduced* to A; i.e. for any set $B \in \mathcal{C}$, an algorithm deciding membership in B can be *easily* obtained from an algorithm deciding membership in A. For classes above P, *easily* means that there is a polynomial time computable function which takes any B input x and transforms it into an A-input fx such that $x \in B$ iff $fx \in A$. A set is said to be *complete* with respect to $\mathcal{C}$ if it is both in $\mathcal{C}$ and $\mathcal{C}$-hard.

Usually P is considered the only *tractable* complexity class. For modal logicians this is unfeasible since satisfiability of propositional logic is already NP-complete. The complexity of the satisfaction problem of most well-known modal logics lies between NP and EXPTIME, with many being *polynomial space* (PSPACE) complete. (E.g., **S5** is NP-complete, all of **K,T,K4,S4** are PSPACE-complete, and the expansion of modal logic **K** with the universal modality is EXPTIME-complete.)

As a convention we say that a modal logic is $\mathcal{C}$-complete, for some complexity class $\mathcal{C}$, meaning that the satisfaction problem of that logic is $\mathcal{C}$-complete. If $\mathcal{C}$ is a deterministic class then also the validity problem is $\mathcal{C}$-complete. We use the same convention for the phrases "$\mathcal{C}$-hard" and "in $\mathcal{C}$".

1.2 The logics

1.2.1 Cylindric modal logic

There are several ways to "modalise" first order logic, cf (Marx and Venema, 1997). Here we follow the tradition set out by cylindric algebra theory, which is also the route taken in (Marx and Venema, 1997). Traditional first order logic contains infinitely many variables and relations symbols of every finite rank. This double nature makes a straightforward analogy with modal logic rather cumbersome, and for that reason we will only focus on finite variable fragments of first order logic. Since we are dealing with decision issues in this paper, this is a harmless restriction. We refer to (Marx and Venema, 1997) for a modal treatment of the infinite variable case. Besides dealing with only n-variable fragments, we will also restrict our attention to formulas containing just n-ary relation symbols: again such a single-sorted approach is no real restriction, cf., section 3.2, and it makes the modal counterpart much neater. We first define finite n-variable first order logic. All definitions in this section come from (Marx and Venema, 1997): Chapter 5.

Definition 1.1 *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 $\{v_i \mid i < n\}$, it has a countable set of n-adic relation symbols $(P_0, P_1, \ldots)$, identity $(=)$, the boolean connectives $\neg, \vee$ and the quantifiers $\exists v_i$. Formulas of L_n and L_n^r are 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 $P_l(v_0 \ldots v_n)$; for L_n, we allow all atomic formulas of the form $P_l(v_{\sigma(0)} \ldots v_{\sigma(n-1)})$, where σ is a transformation of n.*

A first-order structure for $L_n^{(r)}$ is a pair $\mathfrak{M} = (U, \mathsf{v})$ 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 ${}^nU$. Truth of a formula in a model is defined as usual: let s be in ${}^nU$, then

$$\begin{array}{ll} \mathfrak{M} \models v_i = v_j\ [s] & \textit{if}\ \ s_i = s_j, \\ \mathfrak{M} \models P(v_0 \ldots v_{n-1})\ [s] & \textit{if}\ \ s \in \mathsf{v}(P), \\ \mathfrak{M} \models P(v_{\sigma 0} v_{\sigma 1} \ldots v_{\sigma n})\ [s] & \textit{if}\ \ s \circ \sigma\ (= (s_{\sigma(0)}, \ldots, s_{\sigma(n-1)})) \in \mathsf{v}(P), \\ \mathfrak{M} \models \exists v_i \varphi\ [s] & \textit{if there is a } t \textit{ with } s \equiv_i t \textit{ and } \mathfrak{M} \models \varphi\ [t], \\ \textit{etc,} & \end{array}$$

where $\equiv_i$ is given by

$$s \equiv_i s' \iff \text{for all } j \neq i,\ s_j = s'_j. \tag{1}$$

An L_n-formula φ is valid in $\mathfrak{M}$ (notation: $\mathfrak{M} \models \varphi$) if $\mathfrak{M} \models \varphi\ [s]$ for all $s \in {}^nU$, first-order valid (notation: $\models_{fo} \varphi$) if it is valid in every first-order structure of L_n. The same definition applies to L_n^r.

The *modal* languages corresponding to L_n^r and L_n, on which we will concentrate from now on, are given in the following definition:

Definition 1.2 *Let n be an arbitrary but fixed natural number. MLR_n is the modal similarity type having constants $\iota\delta_{ij}$ and unary operators $\Diamond_i$, $\odot_\sigma$ (for all $i, j < n, \sigma \in {}^nn$). For a set of propositional variables Q, the language of n-dimensional cylindric/substitution modal formulas in Q, or shortly, MLR_n-formulas (in Q), is built up as usual: the atomic formulas are the (modal or boolean) constants and the propositional variables, and a formula is either atomic or of the form $\neg\varphi$, $\varphi \vee \psi$, $\Diamond_i\varphi$ or $\odot_\sigma\varphi$, where φ, ψ are formulas. We abbreviate $\Box_i\varphi \equiv \neg\Diamond_i\neg\varphi$.*

CML_n is the similarity type of cylindric modal logic, i.e., the fragment of MLR_n-formulas in which no substitution operator $\odot_\sigma$ occurs.

The languages L_n^r and CML_n are syntactic variants by the following translation:

$$\begin{array}{rclrcl} (P_l v_0, \ldots, v_{n-1})^t & = & p_l & (v_i = v_j)^t & = & \iota\delta_{ij} \\ (\neg\varphi)^t & = & \neg\varphi^t & (\exists v_i \varphi)^t & = & \Diamond_i \varphi^t \\ (\varphi \wedge \psi)^t & = & \varphi^t \wedge \psi^t. & & & \end{array}$$

Below we see that when CML_n is interpreted on a certain class of frames, this translation is truth preserving.

The language MLR_n on which we will mostly concentrate is interpreted on structures of the form $(W, \equiv_i, \bowtie_\sigma, Id_{ij})_{i,j<n,\sigma\in{}^n n}$ where W is a set of n-tuples over U for some set U (i.e., $W \subseteq {}^n U$) and the relations are defined for $s, t \in W$ as

$$\begin{array}{lcl} s \in Id_{ij} & \iff & s_i = t_j \\ s \equiv_i t & \iff & \text{for all } j \neq i\text{: } s_j = t_j \\ s \bowtie_\sigma t & \iff & t = s \circ \sigma. \end{array}$$

We call such structures MLR_n-frames. Note that these frames are determined by their universe W. If we endow such a frame with a valuation v sending propositional variables to subsets of W, we obtain an MLR_n-model in which the MLR_n formulas are interpreted as

$$\begin{array}{lcl} \mathfrak{M}, s \Vdash p & \iff & s \in \mathsf{v}(p), \\ \mathfrak{M}, s \Vdash \iota\delta_{ij} & \iff & s \in Id_{ij}, \\ \mathfrak{M}, s \Vdash \neg\varphi & \iff & \mathfrak{M}, s \nVdash \varphi, \\ \mathfrak{M}, s \Vdash \varphi \wedge \psi & \iff & \mathfrak{M}, s \Vdash \varphi \ \& \ \mathfrak{M}, s \Vdash \psi, \\ \mathfrak{M}, s \Vdash \Diamond_i \varphi & \iff & \text{there is a } t \text{ with } s \equiv_i t \ \& \ \mathfrak{M}, t \Vdash \varphi, \\ \mathfrak{M}, s \Vdash \odot_\sigma \varphi & \iff & \text{there is a } t \text{ with } \bowtie_\sigma st \ \& \ \mathfrak{M}, t \Vdash \varphi. \end{array}$$

We employ the usual notion of truth in a model: $\mathfrak{M} \models \varphi$ iff $\mathfrak{M}, s \Vdash \varphi$ for every s in the domain of $\mathfrak{M}$.

Definition 1.3 *We define four classes of frames:* R_n, D_n, LC_n *and* C_n. *Each of these is given as the class of frames of the form* $(W, \equiv_i, \bowtie_\sigma, Id_{ij})_{i,j<n,\sigma\in{}^n n}$ *where* $W \subseteq {}^n U$ *for some set* U *satisfying the conditions listed below.*

$$\begin{array}{ll} \mathsf{R}_n & W \text{ is an arbitrary subset of } {}^n U \\ \mathsf{D}_n & \forall s \in W \forall i, j < n \ \ s \circ [i/j] \in W \\ \mathsf{LC}_n & \forall s \in W \forall \sigma \in {}^n n \ \ s \circ \sigma \in W \\ \mathsf{C}_n & W = {}^n U. \end{array}$$

Frames in R_n *are called generalized assignment frames, frames in* LC_n*, local cubes, and frames in* C_n*, cubes.*

For a class of frames K, we say that φ is valid in K (notation: $\mathsf{K} \models \varphi$ or $\models_{\mathsf{K}} \varphi$) if $\mathfrak{M} \models \varphi$ for every model on every frame in K.

Now we are ready to make the connection between $L_n^{(r)}$ and CML_n and MLR_n. Let $\mathfrak{M} = (U, \mathsf{v})$ be a first order structure for L_n^r. Let $\mathfrak{M}' = ({}^nU, \equiv_i, Id_{ij}, \mathsf{v}')_{i.j<n}$ be a CML_n-model such that $s \in \mathsf{v}(P_l)$ iff $s \in \mathsf{v}'(p_l)$. Then for all L_n^r formulas φ, for all CML_n formulas ψ, and for all $s \in^n U$,

$$\begin{aligned}\mathfrak{M} \models \varphi\,[s] &\iff \mathfrak{M}', s \Vdash \varphi^t\\ \mathfrak{M}', s \Vdash \psi &\iff \mathfrak{M} \models \psi^{t^{-1}}\,[s],\end{aligned}$$

as an easy induction shows. We find a similar one-to-one correspondence between L_n formulas, and MLR_n formulas if we restrict the application of $\odot_\sigma$ to propositional variables.

Thus *when interpreted on cubes*, MLR_n is nothing more but L_n in modal disguise. When we consider larger classes of frames, like LC_n, the logic becomes weaker than L_n, and more interesting from the computational point of view. Interpreted on such frame-classes, MLR_n corresponds to relativised or generalised first order logic (Németi, 1995b; van Benthem, 1996). For more on these connections, cf., section 3.2.

In the proofs to come we will use the following definitions. For $s \in {}^nU$, we define

$$\begin{aligned}\mathcal{R}(s) &= \{s_i \mid i < n\} && (2)\\ \mathcal{C}(s) &= {}^n\mathcal{R}(s). && (3)\end{aligned}$$

$\mathcal{R}(s)$ is simply called the *range* of s, and $\mathcal{C}(s)$ the *cubicle* of s. Note that a generalized assignment frame with universe W is locally cube iff for all $s \in W$, $\mathcal{C}(s) \subseteq W$.

For $s \in {}^nU$ we also define

$$s_u^i(k) = \begin{cases} u & \text{if } k = i\\ s(k) & \text{otherwise.}\end{cases} \tag{4}$$

Finally s_j^i just abbreviates $s_{s(j)}^i$, and we use $\mathsf{s}_j^i\varphi$ as an abbreviation of $\Diamond_i(\varphi \wedge \iota\delta_{ij})$ (alternatively we can define it using the $\odot_\sigma$). Clearly on any n-assignment model $\mathfrak{M} = (V, \mathsf{v})$ and $s \in V$ we have

$$\mathfrak{M}, s \Vdash \mathsf{s}_j^i\varphi \iff \mathfrak{M}, s_j^i \Vdash \varphi.$$

We will also sometimes use the "swapping" operator $\otimes_{ij}$ defined as $\odot_\sigma$ where $\sigma(i) = j, \sigma(j) = i$ and $\sigma(k) = k$, for all $k \notin \{i, j\}$.

1.2.2 Arrow logic

A very interesting sublanguage of first order logic is L^2_3 —first order logic with 3 variables and only binary predicate symbols. Most of the metalogical phenomena of L_n occur already here, and can be more readily understood. In algebraic disguise, L^2_3 existed already before the invention of the predicate calculus, as the calculus of binary relations. Later, Tarski defined the class RRA of representable relation algebras. This class of algebras is precisely as expressive as L^2_3 (Tarski and Givant, 1987). The modal counterpart of relation algebras is arrow logic which originated with Venema and van Benthem. Arrow logic is a modal logic in which the worlds are seen as arrows and there are modalities reflecting the nature of arrows: one for composing arrows, one for taking the inverse, and a constant denoting the identity arrow. For an introduction to arrow logic we refer to the volume (Marx et al., 1996) or to Chapter 3 of (Marx and Venema, 1997). For now we just copy the relevant definitions from that chapter.

Definition 1.4 *(i) An* arrow frame *is a quadruple* $\mathfrak{F} = (V, C, F, I)$ *such that* V *is a set of objects called arrows,* $C \subseteq V \times V \times V$*,* $F \subseteq V \times V$ *and* $I \subseteq V$*.*
(ii) An arrow frame is called a pair frame *if for some base set* U*,* V *is a subset of* $U \times U$*,* C *and* I *are given as in* (5) *and* (7)*, and* F *is a partial function which can be represented functionally as an* f *satisfying* (6) *(whenever it is defined).*

$$\begin{aligned} Cabc &\quad \textit{iff} \quad a_0 = b_0,\ a_1 = c_1\ \&\ b_1 = c_0 && (5)\\ fa &\quad = \quad (a_1, a_0). && (6)\\ Ia &\quad \textit{iff} \quad a_0 = a_1. && (7) \end{aligned}$$

Since a pair frame is completely determined by its universe $V \subseteq U \times U$*, we make this explicit in the relations, and write them as* C_V, I_V *and* f_V*.*

Definition 1.5 *The language of* arrow logic *is a modal language of type* $\{\circ, \otimes, \iota\delta\}$*, with* $\circ$ *a dyadic modality,* $\otimes$ *a monadic one, and* $\iota\delta$ *a modal constant*[1]*. We abbreviate this similarity type with AL. As a convention, we use infix notation for* $\circ$*.*

Definition 1.6 *An* arrow model *is a pair* $\mathfrak{M} = (\mathfrak{F}, \mathsf{v})$ *with* $\mathfrak{F} = (W, C, F, I)$ *an arrow frame and* v *is a valuation, i.e. a function mapping propositional*

[1]There are other notations in the literature as well, e.g. $\bullet$ or ; for $\circ$, $\varphi^{\smile}$ for $\otimes\varphi$, and *id* or δ for $\iota\delta$.

variables to subsets of W. Truth of the modal formulas is inductively defined as follows:

$$\begin{array}{lll}\mathfrak{M}, a \Vdash \varphi \circ \psi & \textit{iff} & \textit{there are } b, c \textit{ with } Cabc,\ \mathfrak{M}, b \Vdash \varphi \textit{ and } \mathfrak{M}, c \Vdash \psi \\ \mathfrak{M}, a \Vdash \otimes\varphi & \textit{iff} & \textit{there is a } b \textit{ with } Fab \textit{ and } \mathfrak{M}, b \Vdash \varphi \\ \mathfrak{M}, a \Vdash \iota\delta & \textit{iff} & Ia.\end{array}$$

Note that on pair frames, this definition gives the intended meaning to the modalities. Pair frames can be depicted as graphs, but we can also draw them in a way that reflects their two-dimensional aspects more clearly, as in Figure 1.

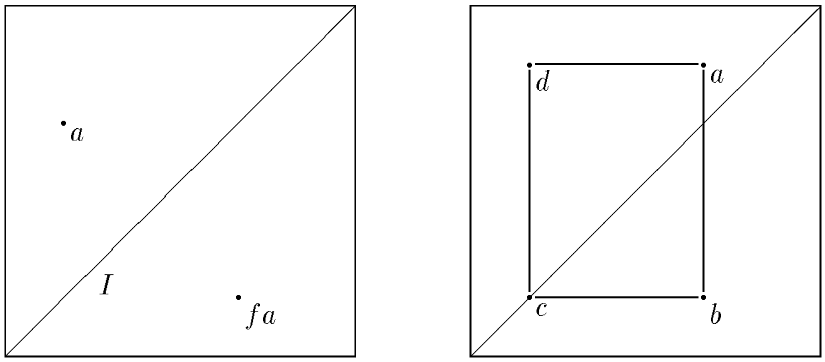

Figure 1: Two–dimensional aspects of arrow logic.

The identity relation I consists of the *diagonal* elements of the universe; f is the operation of *mirroring* in the diagonal. C consists of all triples (a, b, d) with the property that we can draw a *rectangle abcd* such that: b lies on the vertical line through a; d lies on the horizontal line through a; and c (which does not have to belong to the universe of the model) lies on the diagonal.

We will study the following classes of pair frames.

Definition 1.7 *Let R stand for reflexive, S for symmetric, and T for transitive. Let H be a subset of $\{R, S, T\}$, and V a binary relation. We use "V is an H-relation" to abbreviate that V has the properties mentioned in H.* PF_H *denotes the class of all pair frames whose universe is an H-relation.* SQ *denotes the class of pair frames in which the universe is a full Cartesian product $U \times U$, and* LSQ *denotes the class* PF_{RS}*. Members of* SQ *and* LSQ *will be called squares and local squares, respectively.*

1.2.3 Extra expressive power

We will now consider several modalities with a "global character": the universal modality, the difference operator, and the graded modalities.

Universal modality. The *universal modality* $\Diamond\!\!\!+$ is defined on any modal model $\mathfrak{M} = (W, R_i, \mathsf{v})_{i \in I}$ as

$$\mathfrak{M}, w \Vdash \Diamond\!\!\!+\, \varphi \text{ iff there exists a } w' \in W \text{ such that } \mathfrak{M}, w' \Vdash \varphi. \tag{8}$$

As usual, its dual $\neg \Diamond\!\!\!+ \neg$ is denoted by $\boxplus$. We define the global consequence relation for a class of frames K as follows:

$$\varphi_1, \ldots, \varphi_n \models^*_{\mathsf{K}} \varphi_0 \text{ iff for all } \mathfrak{M} \in K,\ \mathfrak{M} \models \varphi_1, \ldots, \varphi_n \text{ implies } \mathfrak{M} \models \varphi_0. \tag{9}$$

In this case, we also say that $\varphi_1, \ldots, \varphi_n / \varphi_0$ is a K-*admissible or valid rule*. We can also define the notion of an *admissible or valid sequent*: $\varphi_1, \ldots \varphi_n \Rightarrow \psi_1, \ldots, \psi_k$ is a K-admissible sequent, if for all $\mathfrak{M} \in K$,

$$\mathfrak{M} \models \varphi_1, \ldots, \varphi_n \text{ implies that } \mathfrak{M} \models \psi_i \text{ for some } 1 \leq i \leq k. \tag{10}$$

Once a modal logic has a (term-definable) *universal modality*, the complexity classes of validities, valid rules and valid sequents collapse.

Theorem 1.8 *Let L be the modal logic of a class K. Assume that on K a universal modality is term-definable in L, and assume that the validity problem of L is in complexity class $\mathcal{C}$. Then the valid sequent problem (whence also the valid rule problem and the global consequence problem) of L is also in $\mathcal{C}$.*

PROOF. Let $\Diamond\!\!\!+$ denote the (term definition of the) universal modality. It is easy to show that $\varphi_1, \ldots \varphi_n \Rightarrow \psi_1, \ldots, \psi_k$ is a K-valid sequent iff $\boxplus(\varphi_1 \wedge \ldots \wedge \varphi_n) \to (\bigvee_{1 \leq i \leq k} \boxplus\, \psi_i)$ is valid in K. So we found an easy (in particular LOGSPACE) reduction to a problem in $\mathcal{C}$. QED

Remark 1.9 There are plenty examples of modal logics whose validity problem is in a lower complexity class than the valid-rule problem (e.g., the modal logic **K**, with PSPACE and EXPTIME). Are there modal logics where the three problems are complete for three "different" complexity classes (assuming P≠NP)?

Difference operator. The *difference operator* D is defined on any modal model $\mathfrak{M} = (W, R_i, \mathsf{v})_{i \in I}$ as

$$\mathfrak{M}, w \Vdash \mathsf{D}\varphi \text{ iff there exists a } w' \in W \text{ such that } w' \neq w \text{ and } \mathfrak{M}, w' \Vdash \varphi. \tag{11}$$

The difference operator provides the ability to define the universal modality, to *name* states, and to count up to two. The universal modality $\Diamond\!\!\!+\,\varphi$ – which means "there is at least one φ-state"– can be defined as $\varphi \vee \mathsf{D}\varphi$. A proposition letter p functions as a name or *nominal* for a world if that world satisfies $p \wedge \neg \mathsf{D}p$. Finally we can say that there are at least two φ-worlds by $\mathsf{D}(\varphi \wedge \mathsf{D}\varphi)$. It is easy to see that this is all the counting power the difference operator can provide.

Graded modalities. The graded modalities are defined as follows: for any modal model $\mathfrak{M} = (W, R_i, \mathsf{v})_{i \in I}$ and $w \in W$,

$$\mathfrak{M}, w \Vdash [k]\varphi \text{ iff there are at most } k \text{ many } w' \in W \text{ such that } \mathfrak{M}, w' \Vdash \varphi. \tag{12}$$

We will denote the "at least" version by $\langle k \rangle$, it can be defined by $\langle k \rangle = \neg[k-1]\varphi$. Note that $[k]$ is not a modality in the strict sense since it doesn't distribute.

Remark 1.10 Let L be any modal logic whose satisfaction problem is in complexity class $\mathcal{C}$. Suppose the universal modality is term-definable in L, and in L we can also express the existence of *nominals*, i.e., we can define an L-formula φ with one variable p such that in any L-model $\mathfrak{M} = (W, R_i, \mathsf{v})_{i \in I}$ we have

$$\mathfrak{M}, w \Vdash \varphi \iff w \text{ is the only world in } W \text{ where } p \text{ holds.}$$

We show that if we expand L with the graded modalities $[k]$ for every $k \in \omega$, then the satisfaction problem of this expansion is also in complexity class $\mathcal{C}$. Let φ be a formula in the larger language. In two steps we construct a formula $\varphi' \wedge \theta'$ without graded modalities such that φ is satisfiable iff $\varphi' \wedge \theta'$ is. $\varphi' \wedge \theta$ is constructed inductively with respect to the number of graded modalities in φ. If φ doesn't contain any we are done. Otherwise, consider a subformula of the form $[k]\psi$ with ψ without graded modalities. Let φ' be φ with this subformula substituted by a new variable p and let $\theta_0 = \boxplus(p \leftrightarrow [k]\psi)$. Claerly φ is satisfiable iff $\varphi' \wedge \theta_0$ is. An inductive application of this procedure eventually yields a formula φ' without graded modalities and a conjunction θ of formulas of the form of θ_0 with ψ without graded modalities.

Now we "axiomatise" every $[k]\psi$ using nominals as follows. For every conjunct θ_i we create a formula θ'_i as the conjunction of

$$\boxplus\, p \vee \boxplus\, \neg p$$
$$p \rightarrow \boxplus(\psi \rightarrow \textstyle\bigvee_{1\leq i\leq k} n_i)$$
$$\neg p \rightarrow \Diamond\!\!\!+(\psi \wedge \textstyle\bigwedge_{1\leq i\leq k} \neg n_i) \wedge \textstyle\bigwedge_{1\leq i\leq k} \Diamond\!\!\!+(\psi \wedge n_i \wedge \textstyle\bigwedge_{j\neq i} \neg n_j).$$

It is clear that $\varphi' \wedge \theta$ is satisfiable iff $\varphi' \wedge \theta'$ is, and θ' contains no graded modalities. *Assuming that k is given in unary* (which is very reasonable), this is a PTIME-reduction, so we obtain the claimed result.

These results go through if we "relativise" all notions to some accessibility relation. E.g., in MLR_n models we can define *coordinate wise nominals* and coordinate wise graded modalities. The coordinate wise nominals are then *functions*.

1.3 Overview of results

As mentioned in the introduction one of our aims is to provide price tags in terms of complexity for the various choices of modelling modal logics of relations. We will give an overview of our results ordered along three of these choices: number of variables involved, conditions on the universes of our models, and the similarity types used.

Number of variables. In cylindric modal logic with the classical full semantics the number 3 forms a border between decidable and undecidable systems, cf Table 1.

MLR_1	MLR_2	MLR_n, $n \geq 3$
NP-complete	NEXPTIME-complete	undecidable
(Ladner, 1977)	Theorem 4.12	Tarski 40's

Table 1: Complexity classes of MLR_n languages on classical cubic models.

There is also an axiomatic analogon of this table. Let $S5^n$ stand for the modal logic of n diamonds defined by the standard **K**-axioms plus axioms making each diamond an $S5$-modality plus the commutativity axioms $\Diamond_i\Diamond_j \leftrightarrow \Diamond_j\Diamond_i$. Then we have the following situation.

We note that Maddux' Theorem states something stronger: namely every axiomatic system extending $S5^n$ with cubic-valid axioms is undecidable as

$S5^1$	$S5^2$	$S5^n$, $n \geq 3$
NP-complete	NEXPTIME-complete	undecidable
(Ladner, 1977)	Corollary 4.14	(Maddux, 1980)

Table 2: Complexity classes of $S5^n$.

well. If we delete the commutativity axioms, ending up with a multi $S5$ system, then the satisfaction problem for all $n \geq 2$ is PSPACE-*complete* (Halpern and Moses, 1992).

For arrow logic we have a similar picture, namely every arrow logic defined by the **K**-axioms plus associativity of composition plus any extension with square-valid axioms is undecidable (Andréka et al., 1996) (for related results see section 5). When interpreted on square models arrow logic is precisely as expressive as L_3^2 —first order logic with 3 variables and only binary relation symbols (Tarski and Givant, 1987).

Conditions on the universes. The idea behind the "relativisation" or "taming logic" approach is to weaken the conditions imposed on the universes of the models. In the classical full models, the universes are full Cartesian products, at the other extreme we have models whose universe is an arbitrary subset of a Cartesian product. In between we have a whole scala of possibilities. We will indicate what kind of conditions form a border between decidability and undecidability.

Let K_n be any class of n-assignment frames between the class R_n of all assignment frames and the class LC_n of local cubes. We have the following results on the complexity of the satisfaction problem of MLR_n formulas interpreted on K_n-frames,

MLR_2 on K_2	MLR_n on K_n for $n \geq 3$
PSPACE-complete	EXPTIME-complete
Corollary 4.7	Corollary 3.3

So if we do not require any condition on our assignment frames, the complexity drops from NEXPTIME to PSPACE in the case of K_2 and from undecidable to EXPTIME in the higher cases. The difference between the class of all assignment frames and its subclasses containing the cubes, is that on these subclasses we ask for the *existence* of certain sequences given that other sequences are already there. The outstanding feature of these existence claims in the decidable K_n above is that they are local, in the sense that they can all be written as

If $s \in V$, then $t \in V$, where t is a sequence built up using only elements occurring in s.

(E.g., the locally cubic condition says that if $s \in V$, then also every $s \circ \sigma \in V$ for any $\sigma \in {}^n n$.)

The condition corresponding to the commutativity axiom has a different flavour. Namely an n-assignment frame with universe V validates $\Diamond_i \Diamond_j p \to \Diamond_j \Diamond_i p$ iff V satisfies

if $s, r, t \in V$ and $s \equiv_i r \equiv_j t$, then there exists an $r' \in V$ such that $s \equiv_j r' \equiv_i t$.

The sequence r' is constructed using elements occurring *both* in s and t. The condition corresponding to the associativity axiom in arrow logic has a similar *global* nature. For arrow logic we have the same result as for MLR_n: the satisfaction problem for arrow logic interpreted on any class of pair-frames satisfying only local conditions is EXPTIME-complete.

So the trend seems to be –at least for these similarity types– that local conditions are harmless for decidability, while some global ones cause undecidability.

Different similarity types. Even if we keep the number of variables and the restrictions on the universes fixed we see quite some variation in complexity depending on the chosen similarity type.

First we discuss cylindric modal logic. Let $n \geq 3$. Above we saw that MLR_n interpreted on the class of all assignment frames is EXPTIME-complete. We still have this result if we take the $\Diamond_i, \mathsf{s}^i_j$-reduct of this language. But if we consider the $\Diamond_i$-reduct the satisfaction problem becomes PSPACE-complete (Theorem 2.12).

On the other hand it seems that once we are in EXPTIME, we can add quite some powerful operations without getting out of it. E.g., the expansion of MLR_n with all graded modalities (so including the universal modality and the difference operator) is EXPTIME-complete when interpreted on the class R_n of all assignment frames (Theorem 3.11). So the ability to state information which must hold everywhere in the model (using the universal modality) and to count globally does not affect our complexity results. A slight variation on the ability to count, namely to count "coordinate wise" using a counting quantifier defined as

$$\mathfrak{M}, s \Vdash [k]_i \varphi \text{ iff there are at most } k \text{ many } t \equiv_i s \text{ such that } \mathfrak{M}, t \Vdash \varphi, \quad (13)$$

can create undecidability: the expansion of CML_3 with only the operator $[1]_1$ is already undecidable.

Some of our results concerning arrow logic are summarized in Table 3.

	Similarity type			
	$\{\circ\}$	$\{\circ, \otimes, \iota\delta\}$	$\{\circ, \otimes, \iota\delta, [1]_1\}$	$\{\circ, [1]_1\}$
all pair frames	PSPACE	EXPTIME	undecidable	?
reflexive and symmetric pair-frames	EXPTIME	EXPTIME	undecidable	undecidable

Table 3: Complexity of reducts of arrow logic.

2 Lower bounds

In ultimo a lower bound for satisfiability (*sat*) is obtained by showing that a logic can encode an accepting computation of a Turing machine in some complexity class. This gives us a direct reduction from a computational problem (is X an accepting computation on a Turing machine?) to a combinatorial problem (is φ satisfiable?). (van Emde Boas, 1996) calls such a reduction where "there is an intended directly visible correspondence between records of accepting computations from the computational problem and solutions of the combinatorial problem under consideration" a *master reduction*. The prime example of such a master reduction is the traditional proof of the Cook–Levin Theorem establishing NP-completeness of satisfiability for propositional logic.

On first sight a master reduction is more appealing than an arbitrary reduction, but on the other hand there are many peculiarities of Turing machines which "pollute" the encoding with emphasis on rather unimportant details. We will use reductions to *tiling problems* instead. (van Emde Boas, 1996) contains a modern introduction to tiling problems and gives examples of problems which are complete for the most common complexity classes arising in logic. The reduction of a Turing machine computation to a tiling problem is so direct that we can always take the composition of the two reductions to be a master reduction.

Finally if we use a problem A (which is not the Turing machine problem) with a known complexity $\mathcal{C}$ to show that problem B is $\mathcal{C}$-hard, we might do far more work than strictly needed. Let f be the reduction from A to B, and g from the Turing machine problem to A. Then g most probably covers only a subset of all A-instances, so in order to get our result f need only work on this subset. (van Emde Boas, 1996) calls this a violation of the *principle of orthogonality for reductions*, as he says

> "the correctness of the reduction f should be based on the combinatorial structure of the problems A and B and nothing else."

As it turns out, we won't have to violate this principle in what comes below.

Tiling problems. We just list two tiling problems to get the flavour. More can be found in the literature (e.g., (van Emde Boas, 1996)).

$n \times n$ **tiling**: Given a finite set $\mathcal{T}$ of tiles, can $\mathcal{T}$ tile $n \times n$?.

That is, does there exists a function t from $n \times n$ to $\mathcal{T}$ such that:

$$\begin{array}{lcl} \mathrm{right}(t(n,m)) & = & \mathrm{left}(t(n+1,m)), \text{ and} \\ \mathrm{up}(t(n,m)) & = & \mathrm{down}(t(n,m+1))? \end{array}$$

Here "right, left, up and down" are functions determining the four "colors" of a tile. Tiles are assumed to be fixed in orientation. So the last condition says that colors of adjacent tiles match.

Theorem 2.1 *(i) $n \times n$-tiling (with n given in unary) is* NP*-complete.*
(ii) $n \times n$-tiling (with n given in binary) is NEXPTIME*-complete.*
(iii) $\mathbb{N} \times \mathbb{N}$-tiling is undecidable, in fact it is co-r.e. (recursively enumerable) complete.

Standard translation. While the tiling problems are already much simpler than Turing machines, we would be stupid not to exploit the fact that the modal logics we are investigating are *modal logics of relations.* As explained in the introduction, the modal logic MLR_n –when interpreted on standard models– is precisely as expressive as first order logic with n variables. The well-known standard translation (van Benthem, 1983) of traditional modal logic into first order logic brings us to a fragment of first order logic with just two variables. So even if we abandon the standard models, the modal language MLR_n (with $n \geq 2$) is rich enough to serve as the range of the standard translation.

We will use (slight adaptations of) the standard translation to obtain lower bounds for our logics. We will now present the logics we will interpret into MLR_n and arrow logic.

Theorem 2.2 *(i) [(Ladner, 1977)] The validity (and the satisfiability-) problem of modal logic **K** is* PSPACE*-complete.*
*(ii) The valid-rule/global consequence problem of modal logic **K** is* EXPTIME*-complete.*
*(iii) [(Fisher and Ladner, 1979; Spaan, 1993)] The validity (and the satisfiability-) problem of modal logic **K** plus the universal modality is* EXPTIME*-complete. (Whence, by Theorem 1.8 also the valid rule and the valid sequent problems are* EXPTIME*-complete.)*

Proof. For (i), Ladner gives a direct reduction to a Turing machine computation. (iii) is shown in (Spaan, 1993) p.17. Fisher–Ladner provide a lower bound for Propositional Dynamic Logic. The crucial observation is that the formula encoding an accepting computation has the form

$$\psi \wedge [a^*]\varphi, \tag{14}$$

where φ and ψ are build up using only the booleans and $[a]$. Now $[a]$ is just a **K**-modality and $[a^*]$ is the reflexive transitive closure of $[a]$, but it is easy to see that we do not need such a restricted modality. Any modality $\Box$ which is reflexive, transitive and extends $[a]$, would do the job, so certainly the universal modality works.

(van Emde Boas, 1996) provides a (in our opinion nicer and simpler) reduction to a two-person corridor tiling game. His formula has the same structure as (14) and even ψ is just a conjunction of propositional variables. Using this reduction we can also obtain (ii). Let us view the formula (14) as $\boxplus\,\varphi \wedge \psi$ where $\boxplus$ is the universal modality and φ is a formula in the language of modal logic **K**. Then, obviously $\boxplus\,\varphi \wedge \psi$ is satisfiable in modal logic **K** plus the universal modality if and only if $\varphi \not\models^* \neg\psi$ in modal logic **K**. This gives us the lower bound, the upper bound is immediately obtained from (iii).

One can also use the van Emde boas version of formula (14) to show PSPACE-hardness of modal logic **K**, by reducing it to the two person $n \times n$ tiling game. For the proof, use the same formula but substitute $\bigwedge_{0 \le i \le n} \Box^i$ for $\boxplus$.

(Whether we violated the principle of orthogonality here is a matter of taste, which we leave to the reader to decide.) QED

We will also need an undecidable logic. Following (Spaan, 1993), we will tailor-make a modal logic to the $\mathbb{N} \times \mathbb{N}$-tiling problem.

Define Tile as the modal logic with two unary modalities $\langle \mathsf{r} \rangle$ and $\langle \mathsf{u} \rangle$ interpreted on the class of frames $(W, R_{\langle \mathsf{r} \rangle}, R_{\langle \mathsf{u} \rangle})$ where $R_{\langle \mathsf{r} \rangle}$ and $R_{\langle \mathsf{u} \rangle}$ are commuting total functions. We usually denote such a frame by (W, r, u) where r and u are the functions giving meaning to $\langle \mathsf{r} \rangle$ and $\langle \mathsf{u} \rangle$, respectively. Tile consists of all formulas valid in this class of frames.

Theorem 2.3 *The satisfaction problem of* Tile *is* NP*-complete, but the global satisfaction problem is undecidable.*

PROOF. The NP-completeness can easily be shown using a selective generation argument, which shows that Tile has the poly-size model property. (Cf. (Blackburn et al., 1997); the argument works for any modal logic where the accessibility relations are partial functions.)

For the undecidability we of course encode $\mathbb{N} \times \mathbb{N}$-tiling. Let $\mathcal{T} = \{T_1, \ldots, T_k\}$ be a set of tiles. Let $t_1, \ldots, t_k$ be propositional variables. Define $\varphi_{\mathcal{T}}$ as the conjunction of the following formulas:

$AX1 \quad \bigvee_{1\leq i\leq k} t_i$
$AX2_i \quad t_i \to \bigwedge_{i\neq j} \neg t_j \qquad (1 \leq i \leq k)$
$AX3_i \quad t_i \to [\mathsf{r}] \bigvee\{t_j \mid \mathrm{right}(T_i) = \mathrm{left}(T_j)\} \quad (1 \leq i \leq k)$
$AX4_i \quad t_i \to [\mathsf{u}] \bigvee\{t_j \mid \mathrm{up}(T_i) = \mathrm{down}(T_j)\} \quad (1 \leq i \leq k).$

It is almost immediate that $\mathcal{T}$ tiles $\mathbb{N} \times \mathbb{N}$ if and only if there exists a Tile model where $\varphi_{\mathcal{T}}$ holds everywhere. QED

Corollary 2.4 *(i) The satisfaction problem of* Tile *plus the universal modality is undecidable.*
(ii) [(Gurevich, 1976)] The satisfaction problem of first order logic with one variable, two binary functions and countably many unary predicates is undecidable.

PROOF. (i) is immediate by the previous theorem. (ii) follows using the (universal closure of the) standard translation of Tile-formulas Then φ is globally satisfied on a Tile-model iff $\forall x[u(r(x)) = r(u(x))] \wedge \forall x \mathrm{ST}(\varphi)(x)$ is first order satisfiable. Note that if the first order language didn't have function symbols, we would need *three* variables to state that the binary relations are commuting total functions. QED

Now we have all the ingredients for our lower bound proofs. We will interpret modal logic **K** plus universal modality into cylindric modal logic of dimension 3 and higher and into arrow logic. The modal logic Tile will be used to show that decidability is lost whenever a modal logic of relations is able to express that a (binary) relation is in fact a function. We hope that we made the reduction chain from these logics to Tiling problems (and whence to accepting Turing machine computations) sufficiently clear that the reader is convinced that these can indeed be seen as master reductions.

In the last part of this introduction to lower bounds we make a slight detour, providing a variation on a theme from (Spaan, 1993). It is not needed to understand the rest of this section, but it shows some light on the power of the relativisation approach.

2.1 Undecidable universal Horn definable modal logics

A beautifully simple result by Edith Hemaspaandra is the following

Theorem 2.5 ((Spaan, 1993)) *There exists a uni-modal logic $ML(\mathsf{K})$ such that*

1. K *is first order definable by a universal sentence*

2. *there is a finite strongly complete Sahlqvist axiomatisation for $ML(\mathsf{K})$*
3. *the satisfaction problem for $ML(\mathsf{K})$ is decidable (in fact it is in* NEXPTIME*)*
4. *the global consequence problem of $ML(\mathsf{K})$ is undecidable.*

The undecidability proof uses a reduction to $\mathbb{N} \times \mathbb{N}$-tiling. The first order definition of K says that every world has at most two one-step successors and at most three two-step successors (a universal sentence with 7 variables and identity). Her theorem is stated about the expansion with the universal modality but using the argument given in the proof of Theorem 2.2.(ii) we can also obtain the result in part 4.

We will show here that if we allow more modalities, we can put more stringent conditions on the first order sentence defining K, finally arriving at a universal Horn sentence in three variables without identity. In number of variables this is the limit, because first order logic with two variables is decidable, so using the standard translation (which goes to first order logic with two variables) we could decide the modal validities.

(Kracht, 1995) also contains undecidable universal Horn definable modal logics, but has no restriction on the number of variables or the use of identity. Also his logics are more abstract, since they arise as "translations" of Thue-systems.

We will define three tiling logics, each one bringing us a bit closer to our final goal. First we slightly adapt the logic Tile, in order to get rid of the existential quantifiers in the first order definition of its frame class.

Let Tile_1 be the logic with the same two modalities $\langle \mathsf{r} \rangle$ and $\langle \mathsf{u} \rangle$ as Tile but interpreted on the class of all frames (W, R, U) satisfying

$T1$ U and R are partial functions
$T2$ $\forall xyzw(xRy \wedge yUz \wedge xUw \rightarrow wRz)$.

Clearly Tile_1 can be axiomatized by the Sahlqvist axioms

$T1^*$ $\langle \mathsf{r} \rangle p \rightarrow [\mathsf{r}]p$ and $\langle \mathsf{u} \rangle p \rightarrow [\mathsf{u}]p$
$T2^*$ $\langle \mathsf{r} \rangle \langle \mathsf{u} \rangle p \rightarrow [\mathsf{u}]\langle \mathsf{r} \rangle p$.

Theorem 2.6 *The satisfaction problem of* Tile_1 *is* NP*-complete, but the global consequence problem is undecidable.*

PROOF. Since U and R are partial functions, the selective generation argument from the proof of Theorem 2.3 still works. For the undecidability result add the following axiom to the "tiling formula" φ_T:

$AX0 \quad \langle \mathsf{r} \rangle \top \wedge \langle \mathsf{u} \rangle \top$.

If this extended $\varphi_{\mathcal{T}}$ is true everywhere in a Tile_1-model, then U and R are commuting total functions, so our previous argument still works. QED

Note that the first order definition of the class of frames of Tile_1 is a conjunction of universal Horn sentences in 4 variables and identity. The price we paid is that it contains an extra relation symbol. Also note that $T2$ is equivalent to

$T2' \quad \forall xyz((xRy \wedge yUz) \rightarrow \forall y(xUy \rightarrow yRz))$,

a formula in just three variables.

In the next step we delete the identity symbol in the definition of K. Let Tile_2 be the logic of the class of all frames (W, U, R) satisfying $T2$.

Theorem 2.7 *The global consequence problem of* Tile_2 *is undecidable, even for questions of the form* $\varphi \models^* \bot$.

PROOF. We show that for any set of tiles $\mathcal{T}$, there exists a Tile_1 model where $\varphi_{\mathcal{T}} \wedge AX0$ is true everywhere iff there exists such a Tile_2 model. Then the result follows because this question is undecidable for Tile_1. From left to right is trivial since every Tile_1 model is a Tile_2 model. For the other side suppose $\mathfrak{M} = (W, R, U, \mathsf{v}) \models \varphi_{\mathcal{T}} \wedge AX0$ and $\mathfrak{M}$ is a Tile_2 model. Because $AX0$ holds every world has both R and U successors. Choose a *substructure* of (W, R, U), where R and U are commuting total functions (this can easily be shown to exist). Because $T2$ is a universal formula, this substructure is a Tile_2-frame. Let $\mathfrak{M}^*$ be the submodel of $\mathfrak{M}$ whose frame is this substructure. By the special form of $\varphi_{\mathcal{T}}$, truth of $\varphi_{\mathcal{T}}$ is preserved under taking substructures, whence also $\mathfrak{M}^* \models \varphi_{\mathcal{T}}$. QED

Finally we delete one more variable, but we have to add a new diamond and a new relation symbol. Let Tile_3 be the modal logic with three modalities $\langle \mathsf{r} \rangle, \langle \mathsf{u} \rangle$ and $\langle \mathsf{diag} \rangle$ of the class of all frames (W, R, U, D) satisfying

$T3 \quad \forall xyz((xRy \wedge yUz) \rightarrow xDz)$

$T4 \quad \forall xyz((xUy \wedge xDz) \rightarrow yRz)$.

These conditions can be captured by the Sahlqvist axioms $\langle \mathsf{r} \rangle \langle \mathsf{u} \rangle p \rightarrow \langle \mathsf{diag} \rangle p$ and $\langle \mathsf{diag} \rangle p \rightarrow [\mathsf{u}] \langle \mathsf{r} \rangle p$.

Theorem 2.8 *The global consequence problem of* Tile_3 *is undecidable, even for questions of the form* $\varphi \models^* \bot$.

PROOF. It is trivial that for any Tile_2 formula φ, there exists a Tile_2 model validating φ iff there exists a Tile_3 model validating φ. Then the result follows from the previous theorem. QED

Remark 2.9 If we want to create universal Horn definable logics with an undecidable set of *validities*, we can add another modality for which we postulate that its accessibility relation is transitive and extends all other accessibility relations. Then the logic obtains a deduction theorem for the global consequence relation (by $\varphi \models^* \psi$ iff $\models (\varphi \wedge \Box\varphi) \rightarrow \psi$, where $\Box$ is the added modality; a simple generated subframe argument suffices). Note that these postulates can also be written using three variables and without identity. So we obtain

Theorem 2.10 *Let T be the expansion of any of the* Tile *logics with a modality $\Diamond$ satisfying*

$$\Diamond\Diamond p \rightarrow \Diamond p$$
$$\langle x\rangle p \rightarrow \Diamond p \text{ for any } \mathsf{Tile} \text{ modality } \langle x\rangle.$$

Then the satisfaction problem of T is undecidable.
In particular there exists an undecidable 4-modal logic which is Sahlqvist axiomatisable with axioms which correspond to universal Horn sentences in three variables without identity.

We note that the expansion of Tile_3 with such a diamond is an instance of a PTC (pseudo-transitive) logic, defined in (Gabbay and Shehtman, 1997).Definition 5.2. PTC logics have several nice properties, e.g., they are complete with respect to unravelled models. An obvious but intriguing question is

Problem 2.11 Find a PTC logic with less than 4 modalities which is complete for a class of frames axiomatisable by universal Horn formulas in three variables, but still undecidable. Or is perhaps the converse true for such PTC logics containing 1,2 or 3 diamonds?

Situations where relativisation does not work

As explained in section 1, the relativisation approach takes an undecidable modal logic of a class K of frames, and hopes to turn it into a decidable logic by considering the class of all *substructures of* K. For the modal logics Tile_1 and Tile_2, clearly this strategy does not work since their frame classes are already closed under taking substructures. Moreover the class of Tile_1 frames is precisely the class of all substructures of Tile-frames. This follows because the Tile-frames can be axiomatized by $T1, T2$ and

$T3 \quad \forall x \exists y(xRy) \wedge \forall x \exists y(xUy)$.

The same observations go through for the expansions of these logics with a master-modality as mentioned in Remark 2.9. So we even have examples of untameable logics with an undecidable set of validities.

2.2 Cylindric modal logic of finite dimension higher than two

In (van Benthem, 1996) it is shown that the reduct of CML_n without the diagonal constants interpreted on the class R_n of all assignment frames is axiomatisable by making every $\Diamond_i$ an S5 modality. By (Halpern and Moses, 1992) (Theorem 6.16 and 6.17), the satisfiability problem for a multi-S5 logic (with at least *two* S5 modalities) is PSPACE-complete. Combining these two results we get

Theorem 2.12 *Let $n > 1$. The satisfiability problem for the reduct of CML_n without the diagonal constants interpreted on the class R_n of all assignment frames is* PSPACE*-complete.*

Below we will see that the restriction to the class R_n is crucial. If we only ask that the frames satisfy the condition "if $s \in V$, then also $s^1_2 \in V$", then the satisfaction problem becomes EXPTIME-hard.

We will first use the standard translation to show that the satisfaction problem of CML_n is EXPTIME-hard. After that we give a slightly different translation which gives us some stronger results. We will formulate our theorems in the most general way. This usually leads to the inclusion of undecidable logics as well (e.g. when interpreted on the cubes).

Theorem 2.13 *Let $n > 2$. For every class K_n between the class R_n of all assignment frames and class C_n of cubes, the satisfaction-problem for CML_n interpreted on K_n is* EXPTIME*-hard.*

PROOF. We will interpret the EXPTIME-complete problem: *satisfaction for the modal logic $\boldsymbol{K}$ + universal diamond.* We call this logic $\mathbf{K}+\diamondplus$, we use $\Diamond$ for the $\mathbf{K}$-modality and $\diamondplus$ for the universal one.

Define the following translation from $\mathbf{K}+\diamondplus$-formulas to CML_n formulas. Here r is a new propositional variable not occurring in φ.

$$\begin{array}{lcl}
p^t & = & p \\
(\neg\varphi)^t & = & \neg\varphi^t \\
(\varphi \wedge \psi)^t & = & \varphi^t \wedge \psi^t \\
(\Diamond\varphi)^t & = & \Diamond_1(r \wedge \mathsf{s}^0_1\mathsf{s}^1_2\varphi^t) \\
(\diamondplus\varphi)^t & = & \Diamond_0\varphi^t.
\end{array}$$

Note the very close resemblance with the *standard translation*, i.e. $(\Diamond p)^t$ in first order logic is $\exists v_1(Rv_0v_1v_2 \wedge Pv_1v_2v_2)$.

Clearly $(\cdot)^t$ is a LOGSPACE-reduction. So we are finished if we can show that any **K**+ ✧ formula φ is **K**+ ✧-satisfiable if and only if $\varphi^t \wedge \iota\delta_{12}$ is satisfiable in the class K_n. This follows from the following two claims.

Claim 1 For any **K**+ ✧-formula φ, φ is **K**+ ✧-satisfiable only if $\varphi^t \wedge \iota\delta_{12}$ is satisfiable in an n-cube.

PROOF OF CLAIM. We prove the case for $n = 3$, in higher dimensions we would work with sequences which all have the same "w-tail".

Assume the antecedent. Let φ be satisfied in model $\mathfrak{M} = (W, R, \mathsf{v})$ at world w. Let $\mathfrak{M}^*$ be the cubic CML_3 model with base W and the following valuation of the variables:

$$\begin{array}{rcl} \mathsf{v}^*(r) & = & \{(x,y,w) \mid xRy\} \\ \mathsf{v}^*(p) & = & \{(x,w,w) \mid x \in \mathsf{v}(p)\}. \end{array}$$

We now show that for any formula ψ, and for any world $x \in W$, $\mathfrak{M}, x \Vdash \psi$ iff $\mathfrak{M}^*, (x,w,w) \Vdash \psi^t$, which would prove the claim. We proceed by induction on the structure of the formula. For the variables the claim holds by definition. The boolean cases and the ✧-case are trivial. We prove the **K**-diamond case:

$$\begin{array}{ll} \mathfrak{M}, x \Vdash \Diamond\varphi & \iff \\ \text{exists } y \in W \text{ such that } xRy \text{ and } \mathfrak{M}, y \Vdash \varphi & \iff \\ \text{exists a } y \text{ such that } \mathfrak{M}^*, (x,y,w) \Vdash r \text{ and } \mathfrak{M}^*, (y,w,w) \Vdash \varphi^t & \iff \\ \mathfrak{M}^*(x,w,w) \Vdash \Diamond_1(r \wedge \mathsf{s}_1^0\mathsf{s}_2^1\varphi^t) & \iff \\ \mathfrak{M}^*(x,w,w) \Vdash (\Diamond\varphi)^t. & \end{array}$$

◀

Claim 2 For any formula φ, $\varphi^t \wedge \iota\delta_{12}$ is satisfiable in an n-assignment frame only if φ is **K**+ ✧-satisfiable.

PROOF OF CLAIM. Again for convenience we do the case $n = 3$. In higher dimensions we work only with sequences which have the same "tail" as (a,b,c) below. Let $\mathfrak{M} = (V, \mathsf{v})$ be a model with V some subset of a cube and let $\mathfrak{M}, (a,b,c) \Vdash \varphi^t \wedge \iota\delta_{12}$. Thus $b = c$. We define a Kripke model $\mathfrak{M}^* = (W, R, \mathsf{v}^*)$ as follows:

$$\begin{array}{lcl} W & = & \{(x,b,b) \in V\} \\ (x,b,b)R(y,b,b) & \iff & \text{there exists a } (x,y,b) \in V \text{ such that } \mathfrak{M}, (x,y,b) \Vdash r \wedge \mathsf{s}_1^0\top \\ \mathsf{v}^*(p) & = & \{s \in W \mid s \in \mathsf{v}(p)\}. \end{array}$$

Again we claim, for all $s \in W$ and all $\mathbf{K}+$ ⊕-formulas ψ, $\mathfrak{M}^*, s \Vdash \psi$ iff $\mathfrak{M}, s \Vdash \psi^t$.

The only interesting case is again the $\mathbf{K}$-diamond.

$$\begin{array}{ll}
\mathfrak{M}^*, (x,b,b) \Vdash \Diamond\varphi & \Longleftrightarrow \\
\text{there exists a } (x,y,b) \in V \text{ such that } \mathfrak{M}, (x,y,b) \Vdash r \wedge \mathsf{s}_1^0\top \text{ and } \mathfrak{M}^*, (y,b,b) \Vdash \varphi & \Longleftrightarrow \\
\text{there exists a } (x,y,b) \in V \text{ such that } \mathfrak{M}, (x,y,b) \Vdash r \wedge \mathsf{s}_1^0\top \text{ and } \mathfrak{M}, (y,b,b) \Vdash \varphi^t & \Longleftrightarrow \\
\mathfrak{M}, (x,b,b) \Vdash \Diamond_1(r \wedge \mathsf{s}_1^0\mathsf{s}_2^1\varphi^t) & \Longleftrightarrow \\
\mathfrak{M}, (x,b,b) \Vdash (\Diamond\varphi)^t. &
\end{array}$$

◀

With these two claims we finish the proof of the theorem. QED

We only used $\iota\delta$ to express s_j^i, so it seems that the difference in complexity we found in the last two theorems is due to availability of these s_j^i-operators. We will now show that this is indeed the case, and strengthen the above results by using a somewhat less intuitive translation.

In the language CML_n it is trivial to mark the difference between the classes R_n and D_n by a formula: i.e. $\mathsf{D}_n \models \mathsf{s}_1^0\top$, while this formula fails in R_n. Szabolcs Mikulás found a formula using only $\Diamond_i$ modalities marking this difference:

Theorem 2.14 *Let $n \geq 3$. The formula $\Diamond_1\Diamond_0\Diamond_1\Diamond_0\Diamond_1 p \to \Diamond_1\Diamond_0\Diamond_1 p$ is valid on any class of n-assignment frames validating $\mathsf{s}_2^1\top$, but fails on R_n.*

PROOF. Left to the reader. QED

The intuition behind this formula led to the next theorem which is a joint result with Mikulás.

Theorem 2.15 *Let $n \geq 3$. (i) The satisfaction problem of the $\{\Diamond_0, \Diamond_1, \mathsf{s}_2^1\}$-reduct of MLR_n interpreted on any class between R_n and C_n is* EXPTIME-*hard.*
(ii) Let D_n^* be $\{\mathfrak{F} \in \mathsf{R}_n \mid \mathfrak{F} \models \mathsf{s}_2^1\top\}$. The satisfaction problem of the $\{\Diamond_0, \Diamond_1\}$-reduct of CML_n interpreted on any class between D_n^* and C_n is EXPTIME-*hard.*

So the theorem tells us that in order to interpret an EXPTIME problem in CML_n we only need the diamonds $\Diamond_0$ and $\Diamond_1$ for the translation, and we need the operator s_2^1 to ensure that the universes of the models obey certain restrictions.

PROOF. We interpret $\mathbf{K}+\oplus$, using the following translation. Here r and d are new propositional variables not occurring in φ.

$$\begin{array}{lcl} p^t & = & p \\ (\neg\varphi)^t & = & \neg\varphi^t \\ (\varphi\wedge\psi)^t & = & \varphi^t \wedge \psi^t \\ (\Diamond\varphi)^t & = & \Diamond_1(r \wedge \Diamond_0(d \wedge \mathsf{s}_2^1\top \wedge \varphi^t)) \\ (\oplus\,\varphi)^t & = & \Diamond_1\Diamond_0\Diamond_1(d \wedge \mathsf{s}_2^1\top \wedge \varphi^t). \end{array}$$

We claim that for any $\mathbf{K}+\oplus$-formula φ, φ is $\mathbf{K}+\oplus$-satisfiable iff $\varphi^t \wedge d \wedge \mathsf{s}_2^1\top$ is R_n-satisfiable. This immediately implies (i) of the theorem. (ii) also follows because on D_n^*-models $\mathsf{s}_2^1\top$ always holds, so we could dispense with the $\mathsf{s}_2^1\top$-part in the translation.

To prove the claim, first suppose φ is satisfied on a $\mathbf{K}+\oplus$-model $\mathfrak{M} = (W, R, \mathsf{v})$ at world w. Define a cubic model $\mathfrak{M}^* = ({}^nW, \mathsf{v}^*)$ where

$$\begin{array}{lcl} \mathsf{v}^*(d) & = & \{(x, x, \vec{w}) \mid x \in W\} \\ \mathsf{v}^*(r) & = & \{(x, y, \vec{w}) \mid xRy\} \\ \mathsf{v}^*(p) & = & \{(x, x, \vec{w}) \mid x \in \mathsf{v}(p)\}. \end{array}$$

We leave it to the reader to show that for any $x \in W$, and any formula ψ,

$$\mathfrak{M}, x \Vdash \psi \iff \mathfrak{M}^*, (x, x, \vec{w}) \Vdash \psi^t.$$

For the other side suppose that $\varphi^t \wedge d \wedge \mathsf{s}_2^1\top$ is satisfiable on a R_n model $\mathfrak{M} = (V, \mathsf{v})$ at the sequence (a, b, s') for s' some $n-2$-sequence. Define a Kripke model $\mathfrak{M}^* = (W, R, \mathsf{v}^*)$ as

$$\begin{array}{lcl} W & = & \{(x, y, s') \in V \mid \mathfrak{M}^*, (x, y, s') \Vdash d \wedge \mathsf{s}_2^1\top\} \\ (x, y, s')R(x', y', s') & \iff & \mathfrak{M}, (x, y', s') \Vdash r \\ \mathsf{v}^* & = & \mathsf{v}\restriction_W . \end{array}$$

We claim that for all $t \in W$, and for all ψ,

$$\mathfrak{M}, t \Vdash \psi^t \iff \mathfrak{M}^*, t \Vdash \psi,$$

which has a simple inductive proof. The extra $d \wedge \mathsf{s}_2^1\top$ in the translation is used in the case for $\oplus$. QED

We summarize the lower bound results on cylindric modal logic in Table 4.

2.3 Arrow logic

We obtain very similar results for arrow logic, in basically the same way.

Theorem 2.16 *Let C be any class of pair frames between the class of all pair frames and the class of squares. The satisfaction problem for the arrow logic of C is* EXPTIME*-hard.*

$n \geq 3$	*reduct*	
Class of assignment frames	$\{\Diamond_0, \Diamond_1\}$	$\{\Diamond_0, \Diamond_1, \mathsf{s}^1_2\}$
R_n	PSPACE	EXPTIME
Any K_n s.t. $\mathsf{C}_n \subseteq \mathsf{K}_n \subseteq \mathsf{R}_n$ and $\mathsf{K}_n \models \mathsf{s}^1_2\top$	EXPTIME	EXPTIME

Table 4: Lower bounds for (reducts of) cylindric modal logics, with $n \geq 3$.

Proof. We will translate the same problem –satisfiability of $\mathbf{K}+\,\diamondplus$-formulas– as with CML_n, but now into arrow logic. Let $\Diamond_0\varphi$ abbreviate $\top \circ (\top \circ \varphi)$. The translation is very similar (again r is a new variable):

$$
\begin{array}{lcl}
p^t & = & p \\
(\neg\varphi)^t & = & \neg\varphi^t \\
(\varphi \wedge \psi)^t & = & \varphi^t \wedge \psi^t \\
(\Diamond\varphi)^t & = & r \circ (\varphi^t \wedge \otimes\top) \\
(\diamondplus\,\varphi)^t & = & \Diamond_0(\varphi^t \wedge \otimes\top).
\end{array}
$$

Claim 1 For any formula φ, φ is $\mathbf{K}+\,\diamondplus$-satisfiable only if $\varphi^t \wedge \otimes\top \wedge \top \circ \iota\delta$ is satisfiable in a square.

Proof of Claim. Let φ be satisfied in model $\mathfrak{M} = (W, R, \mathsf{v})$ at world w. Let $\mathfrak{M}^*$ be the square arrow logical model with base W and the following valuation of the variables:

$$
\begin{array}{lcl}
\mathsf{v}^*(r) & = & R \\
\mathsf{v}^*(p) & = & \{(x, w) \mid x \in \mathsf{v}(p)\}.
\end{array}
$$

We claim that for any formula ψ, and for any world x, $\mathfrak{M}, x \Vdash \psi$ iff $\mathfrak{M}^*, (x, w) \Vdash \psi^t$. For the diamond:

$$
\begin{array}{ll}
\mathfrak{M}, x \Vdash \Diamond\varphi & \iff \\
\text{exists } y \in W \text{ such that } xRy \text{ and } \mathfrak{M}, y \Vdash \varphi & \iff \\
\text{exists a } y \text{ such that } \mathfrak{M}^*, (x, y) \Vdash r \text{ and } \mathfrak{M}^*, (y, w) \Vdash \varphi^t & \iff \\
\mathfrak{M}^*(x, w) \Vdash r \circ (\otimes\top \wedge \varphi^t) & \iff \\
\mathfrak{M}^*(x, w)) \Vdash (\Diamond\varphi)^t. &
\end{array}
$$

◀

Claim 2 For any formula φ, $\varphi^t \wedge \otimes\top \wedge \top \circ \iota\delta$ is satisfiable in a pair frame only if φ is $\mathbf{K}+\,\diamondplus$-satisfiable.

Proof of Claim. Let $\mathfrak{M} = (V, \mathsf{v})$ be a model with V any binary relation. Suppose $\mathfrak{M}, (a, b) \Vdash \varphi^t \wedge \otimes\top \wedge \top \circ \iota\delta$. Whence $\{(b, b), (b, a)\} \subseteq V$. We define a Kripke model $\mathfrak{M}^* = (W, R, \mathsf{v}^*)$ as follows:

$$\begin{array}{lcl} W & = & \{(x, b) \in V \mid (b, x) \in V\} \\ (x, b)R(y, b) & \iff & (x, y) \in V \text{ and } \mathfrak{M}, (x, y) \Vdash r \\ \mathsf{v}^*(p) & = & \{s \in W \mid s \in \mathsf{v}(p)\}. \end{array}$$

Note that $(a, b) \in W$. Again we claim, for all $s \in W$ and all $\mathbf{K}+\diamondplus$-formulas ψ, $\mathfrak{M}^*, s \Vdash \psi$ iff $\mathfrak{M}, s \Vdash \psi^t$. We now do the proof for the universal diamond.

$$\begin{array}{ll} \mathfrak{M}^*, (x, b) \Vdash \diamondplus\varphi & \iff \\ \text{there exists a } y\text{: } \{(y, b), (b, y)\} \subseteq V \text{ and } \mathfrak{M}^*, (y, b) \Vdash \varphi & \iff \\ \text{there exists a } y\text{: } \{(y, b), (b, y)\} \subseteq V \text{ and } \mathfrak{M}, (y, b) \Vdash \varphi^t & \iff \\ \mathfrak{M}, (x, b) \Vdash \top \circ (\top \circ (\otimes\top \wedge \varphi^t)) & \iff \\ \mathfrak{M}, (x, b) \Vdash (\diamondplus\varphi)^t. & \end{array}$$

In the fourth step we use that $(b, b) \in V$ when going down. ◀

We finished the proof of the theorem. QED

We can get some more results out of the proof, and out of the translation. Because we only use $\iota\delta$ and $\otimes$ in the translation to guarantee the existence of certain pairs in the domain, we do not need them when these pairs are already in the domain (because we look at restricted classes). Also we only use $\iota\delta$ to express s_1^0.

We only used the "extra part" $\top \circ \iota\delta \wedge \otimes\top$ in the proof for the universal modality. So if we forget the universal diamond, we do not need $\iota\delta$ nor $\otimes$ in the translation. Because the satisfaction problem for a $\mathbf{K}$-modality is PSPACE-complete (Ladner, 1977), this means that every arrow logic containing at least the composition operator is pspace-hard.

The results on reducts are summarized in the next table.

Theorem 2.17 *All arrow logics in Table 5 labelled with* E *have an* EXPTIME-*hard satisfaction problem. All the other ones have* PSPACE-*hard satisfaction problems.*

2.4 Undecidability by expressing functionality

We will show that once we have the possibility to express that a relation is a partial function in arrow logic, the satisfaction problem becomes undecidable. A very similar result holds for cylindric modal logic.

(Andréka et al., 1996) shows that the arrow logics of the previous section all have the *finite base property.* This means that none of these logics is

		Restrictions on universe			
		none	reflexive	symmetric	refl & symm
	$\circ$	P	P	P	E
similarity	$\circ, \otimes$	P	E	P	E
types	$\circ, \iota\delta$	P	P	E	E
	$\circ, \mathsf{s}_1^0$	P	P	E	E
	$\circ, \mathsf{s}_1^0, \otimes$	E	E	E	E

Table 5: Hardness results for reducts of arrow logic

strong enough to express that the *base* of a model is infinite. Square arrow logic has the ability to force infinite models: e.g., we can express that a relation p is an injective function with a start point by saying

$$\Diamond\!\!\!\!+ \neg(\top \circ p) \wedge \boxplus[p \circ \top \wedge (p \rightarrow \neg(p \circ \neg\iota\delta)) \wedge (\otimes p \rightarrow \neg((\otimes p) \circ \neg\iota\delta))],$$

where $\Diamond\!\!\!\!+ \varphi$ abbreviates $\top \circ \varphi \circ \top$, and indeed on squares behaves as the universal modality. The "unary modality" $\varphi \circ \neg\iota\delta$ behaves on squares as a *coordinate-wise difference operator*, that is

$$\mathfrak{M}, (x, y) \Vdash \varphi \circ \neg\iota\delta \text{ iff there exists a } z \neq y \text{ such that } \mathfrak{M}, (x, z) \Vdash \varphi.$$

It is precisely this operation which we will add as a primitive to the similarity type of arrow logic (and later a related one to cylindric modal logic), and show that undecidability is then inevitable.

Definition 2.18 *Let $AL{+}\mathsf{d}_1$ denote the expansion of arrow logic with a unary modality d_1 with the following interpretation on a model $\mathfrak{M} = (V, \mathsf{v})$. For $(x, y) \in V$,*

$$\mathfrak{M}, (x, y) \Vdash \mathsf{d}_1\varphi \textit{ iff there exists a } (x, z) \in V \textit{ and } z \neq y \textit{ and } \mathfrak{M}, (x, z) \Vdash \varphi.$$

‘

Theorem 2.19 *The satisfaction problem of $AL{+}\mathsf{d}_1$ interpreted on any class of frames between all pair-frames and the squares is undecidable.*

Proof. We will interpret the logic Tile plus universal modality using the translation from the proof of Theorem 2.16. This logic is undecidable by

Corollary 2.4. We keep the same translation except for the two new diamonds:

$$\begin{array}{rcl} (\langle \mathsf{u} \rangle \varphi)^t & = & u \circ (\otimes \top \wedge \varphi^t) \\ (\langle \mathsf{r} \rangle \varphi)^t & = & r \circ (\otimes \top \wedge \varphi^t). \end{array}$$

We must now also encode that u and r are commuting total functions. Define

$$\begin{array}{lll} *_1 & u \circ \top \wedge r \circ \top & \text{(totality)} \\ *_2 & \neg[(u \wedge \mathsf{d}_1 u) \circ \top] \wedge \neg[(r \wedge \mathsf{d}_1 r) \circ \top] & \text{(functionality)} \\ *_3 & [u \circ r \wedge r \circ u] \circ \top & \text{(commutativity).} \end{array}$$

Recall that $\Diamond_0 \varphi$ abbreviates $\top \circ (\top \circ \varphi)$. $\Box_0$ abbreviates –as usual– $\neg \Diamond_0 \neg$. We claim that for any $\mathsf{Tile} + \oplus$ formula φ,

φ is $\mathsf{Tile} + \oplus$-satisfiable iff $\varphi^t \wedge \Box_0(*_1 \wedge *_2 \wedge *_3) \wedge \top \circ \iota\delta \wedge \otimes \top$ is satisfiable on a pair frame.

Now we just copy the proof of Theorem 2.16. In the first claim we turned a Kripke model into an arrow logical model. We do that here in the same way with a Tile model. The inductive proof is the same, and clearly the square arrow logical model $\mathfrak{M}^*$ validates $*_1 \wedge *_2 \wedge *_3$.

In the second claim we turn an arrow logical model into an Kripke model. The inductive proof of the "truth-preserving claim" is the same, but now we also have to show that the defined model is a model of the logic Tile, i.e., that the two relations are commuting total functions. Totality is immediate by $*_1$. For functionality suppose that u is not a function. So we have, for some x, y, y' and $y \neq y'$,

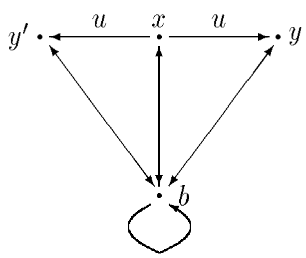

But then $\mathfrak{M}, (x, y) \Vdash u \wedge \mathsf{d}_1 u$. So $\mathfrak{M}, (x, b) \Vdash (u \wedge \mathsf{d}_1 u) \circ \top$, whence $\mathfrak{M}, (a, b) \Vdash \Diamond_0[(u \wedge \mathsf{d}_1 u) \circ \top]$ which contradicts the fact that $\mathfrak{M}, (a, b) \Vdash \Box_0(*_2)$. So u is a total function, and similarly for r.

For commutativity, suppose $(x, b) \in W$ and

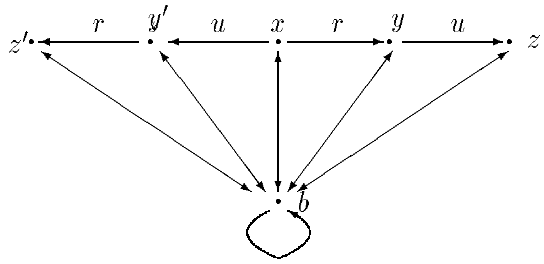

But since $\mathfrak{M}, (a, b) \Vdash \Box_0 *_3$, $\mathfrak{M}, (x, b) \Vdash *_3$, so using the fact that u and r are functions, we must have the following situation, which shows that $z = z'$.

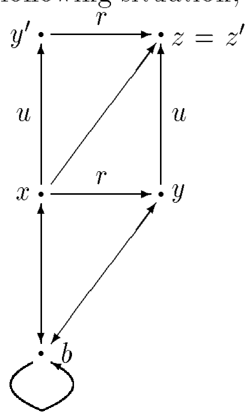

This finishes the proof of the theorem. QED

Again we get the same extra results out of our proof as before since we only used $\iota\delta$ and $\otimes$ to ensure the existence of certain reflexive and symmetric arrows.

Theorem 2.20 *Let X be any arrow logic from Table 5 labelled with E. Then the expansion of X with the vertical difference operator d_1 has an undecidable satisfaction problem.*

PROOF. By the same argument as was given just above Theorem 2.17. QED

The same result can be obtained by adding the "counting quantifier" $[1]_1$, defined as

$(x, y) \Vdash [1]_1\varphi$ iff there is at most one z such that $(x, z) \in V$ and $(x, z) \Vdash \varphi$.

Then we could write axiom $*_2$ as $\neg[(u \wedge \neg[1]_1 u) \circ \top]$ and the proof that $*_2$ forces functionality of u goes through. So we obtain

Corollary 2.21 *Let X be any arrow logic from Table 5 labelled with E. Then the expansion of X with the counting quantifier $[1]_1$ has an undecidable satisfaction problem.*

Cylindric modal logic with functions

In relativised cylindric modal logic we also cannot express that a relation is a (partial) function. If we add the "difference versions" of the cylindric modalities to the similarity type we obtain this ability. Just as in arrow logic, this increase in expressive power leads to undecidability.

The coordinate-wise difference operators d_i are defined as follows. Let $\mathfrak{M} = (V, \mathsf{v})$ be a CML_n-model. For $s \in V$ and $i < n$, we define

$$\mathfrak{M}, s \Vdash \mathsf{d}_i\varphi \text{ iff there exists a } t \in V \text{ such that } s \equiv_i t,\ t \neq s \text{ and } \mathfrak{M}, t \Vdash \varphi.$$

Using the idea behind the proof in arrow logic given above, (Mikulás, to appear) shows (in the equivalent setting of cylindric relativised set algebras),

Theorem 2.22 *Let $n \geq 3$. The expansion of CML_n with the operator d_1 has an undecidable satisfaction problem for any class of frames between the class of all assignment frames and the n-cubes.*

An immediate consequence of the theorem is that this expansion cannot have the finite model or finite base property. A recent result by (Andréka et al., 1996) shows that the expansion of MLR_n with the difference operator interpreted on the class of all assignment frames has the finite base property.

Just as we obtained Corollary 2.21, we find here

Corollary 2.23 *Let $n \geq 3$. The expansion of CML_n with the counting quantifier $[1]_1$ has an undecidable satisfaction problem for any class between R_n and C_n.*

This finishes our section on lower bounds. After we summarise our results we turn to find matching upper bounds.

Summary

We saw that in general the modal logics of relations we are studying are EXPTIME-hard. If we consider reducts, this might go down to PSPACE. By (Andréka et al., 1996) we know that both MLR_n and arrow logic have the finite base property when interpreted on the class of all assignment and arrow frames, respectively. On the other hand, expanding the language in order to be able to express functionality of a relation leads not only to infinity axioms, but even to undecidability. At present we do not have examples —except in the case of MLR_2 expanded with coordinate wise counting quantifiers, cf., section 4— of decidable modal logics of relations without the finite base property. The proof in (Andréka et al., 1996) of this property uses in a crucial way the fact that we can amalgamate finite structures within the class of all assignment frames. In the next section we will see that also in our upper complexity bound proof this is a crucial feature. Is there a deeper connection? Is it the case that in modal logics of relations of sufficient strength (e.g., which are undecidable when interpreted on classical models) decidability and finite base property come and go together?

An at present unexplored terrain is that of the actual size of the finite models. By (Ladner, 1977), MLR_1 —which is just **S5**— has the poly-size model property. (Halpern and Moses, 1992) contains a multi-**S5** formula which forces models to contain a number of worlds exponential in the length of the formula. The formula describes a binary tree of depth n with a formula polynomial in n. Clearly this formula works for MLR_2 and higher, and can easily be changed to work for arrow logic as well. The proofs in (Andréka et al., 1996) provide an upper bound on the size of the models satisfying a formula, but there is still a gap between these bounds and the single exponential lower bound obtained by describing a binary tree.

3 Upper bounds

We will provide matching upper bounds for the EXPTIME lower bounds we found in the previous section. These results can be obtained from the literature by fine-tuning the decidability proofs for CML_n, MLR_n and arrow logic. We even find stronger results about expansions of the languages with the universal modality, the difference operators and the graded modalities.

The proof idea and the main notion of a *saturated set of mosaics* goes back to (Németi, 1986). There it is shown that the validity problem of the logic CML_n interpreted on the classes $\mathsf{R}_n, \mathsf{D}_n$ and LC_n is decidable. A more accessible version for modal logicians of this purely algebraic proof for the language MLR_n interpreted on the class LC_n can be found in (Venema and Marx, 1995). We also refer to this paper for an intuitive explanation of the so called *mosaic-method* (cf., also (Marx, 1997; Reynolds, 1996) for applications of the mosaic-idea in different situations, e.g., in temporal logic).

Instead of referring to existing proofs in the literature and pointing out where changes need to be made to get an EXPTIME decision procedure, we just redo the whole proof. This makes the paper self-contained and hardly uses more space. We will redo the mosaic-proof given in (Venema and Marx, 1995) in a more general setting, following ideas from (Mikulás, 1995).

Before we start we provide the main line of the argument, using the example of the MLR_n-language interpreted on the class of local cubes from (Venema and Marx, 1995). Readers unfamiliar with the notion of a mosaic might first want to read the definition given below.

Exptime upper bound proof strategy using mosaics

The proof uses the same idea as the EXPTIME-completeness proof of Propositional dynamic logic due to (Pratt, 1979), now with mosaics instead of worlds. Here are the main steps.

1. satisfiability of φ is equivalent to the existence of a saturated set of mosaics containing a φ-mosaic –in short a φ-SSM;

2. each mosaic is a first order structure with at most n-elements where every n-tuple is labelled by a set of subformulas of φ;

3. so up to isomorphism the number of mosaics is bounded by an exponential in $|\varphi|$;

4. the "shrinking"-algorithm will start with the set of all mosaics labelled with subformulas of φ;

5. in the first step it will delete all mosaics which are not *coherent*;

6. in each next step it either stops because there is no more mosaic with a sequence with φ in its label. Then it outputs *unsatisfiable*. Or it tries to delete mosaics with a defect $(s, \Diamond_i\psi)$ which cannot be solved in the present set of mosaics. If it cannot delete anything, it outputs *satisfiable*, if it deletes some mosaic it goes to the next round with the smaller set of mosaics;

7. Clearly the algorithm stops after at most exponentially many steps;

8. Each step takes at most exponential time, so the whole algorithm runs in EXPTIME;

9. the algorithm is correct because of the first point in this list.

The remainder of this section is organised as follows. We first show that the satisfaction problem of an expansion of MLR_n interpreted on the local cubes is in EXPTIME. We will strengthen this result to some other classes, most notably the packed-dense local cubes. We use this result to show the relation between relativisation and the loosely guarded fragment of first order logic, and obtain a complexity bound for that fragment as a corollary. As a bonus we obtain that also the satisfaction problem of arrow logic is in EXPTIME. Finally we look at expansions with the graded modalities and meet some NP-complete reducts.

3.1 Modal logic of relations with generalised cylindrifications

We add the generalised cylindrifications to the language MLR_n. They are $\Diamond_{(\Gamma)}$ for every $\Gamma \subseteq n$. On n-assignment models $\mathfrak{M} = (V, \mathsf{v})$ they are defined for any $s \in V$,

$$\mathfrak{M}, s \Vdash \Diamond_{(\Gamma)}\varphi \text{ iff there exists a } t \in V \text{ such that } s \equiv_\Gamma t \text{ and } \mathfrak{M}, t \Vdash \varphi,$$

where $s \equiv_\Gamma t$ holds iff for all $i \notin \Gamma$, $s(i) = t(i)$. So in first order logic, $\Diamond_{\{0,1\}}$ corresponds to the generalised quantifier $\exists v_0 v_1$, which is *with the classical semantics* definable as $\exists v_0 \exists v_1$. In the relativised cases we lose this ability, as is easy to see. We call the expansion of MLR_n with $\Diamond_{(\Gamma)}$ for every non-empty $\Gamma \subseteq n$, MLR_n^+. In order to avoid confusion between $\Diamond_i$ and $\Diamond_\Gamma$ where Γ is a set and i is an ordinal, we denote cylindrification over sets Γ by $\Diamond_{(\Gamma)}$. We reserve the symbol $\diamondplus$ for $\Diamond_{\{0,\ldots,n-1\}}$ (which we abbreviate also as $\Diamond_{(n)}$), because it behaves as the universal modality. From now on we will reserve

$\Diamond_{(\Gamma)}$ for cases when $\Gamma \neq n$ and use $\diamond$ for $\Diamond_{(n)}$. We need some new notation for the new diamonds. Let f be a surjective function from Γ to some set Δ. Then $s^{\Gamma}_{f\Gamma}$ is defined as

$$s^{\Gamma}_{f\Gamma}(i) = \begin{cases} s(i) & \text{if } i \notin \Gamma \\ f(i) & \text{otherwise.} \end{cases}$$

Notation: if the function $f : \Gamma \longrightarrow \Delta$ is clear from context, or it doesn't matter which f is used, we merely write s^{Γ}_{Δ}.

From now on let n be arbitrary but fixed, and ξ an arbitrary but fixed MLR^{+}_{n}-formula. We start with the definition of a saturated set of mosaics.

Mosaics. For a formula ξ, we define $For(\xi)$ as the set of subformulas of ξ plus their single negations together with the set $\{\iota\delta_{ij} \mid i, j < n\}$ Note that $For(\xi)$ is polynomial in $|\xi|$, since n is a fixed constant.

A ξ-**mosaic** is a structure $\mu = (X, \ell)$ such that X is a set of size at most n and ℓ labels elements of ${}^{n}X$ with subsets of $For(\xi)$. We call X the *base* of μ.

Two mosaics $\mu = (X, \ell)$ and $\mu' = (X', \ell')$ are *isomorphic*, notation: $\mu \simeq \mu'$, if there is a bijection $f : X \to X'$ such that $\ell(s) = \ell'(f^{+} \circ s)$, where $f^{+} : {}^{n}X \to {}^{n}X'$ is the bijection induced by f.

$\mu' = (X', \ell')$ is a *submosaic* of $\mu = (X, \ell)$ if $X' \subseteq X$, and ℓ' is the restriction of ℓ to ${}^{n}X'$

A ξ-mosaic is **coherent** if it satisfies the following conditions, for all $s, t \in {}^{n}X$ and $\varphi, \psi \in For(\xi)$:

(FC1)	$\iota\delta_{ij} \in \ell(s)$	iff	$s_i = s_j$,	
(FC2)	$\neg\varphi \in \ell(s)$	iff	$\varphi \notin \ell(s)$	provided $\neg\varphi \in For(\xi)$
(FC3)	$\varphi \wedge \psi \in \ell(s)$	iff	$\varphi \in \ell(s)$ & $\psi \in \ell(s)$	provided $\varphi \wedge \psi \in For(\xi)$
(FC4)	$\odot_{\sigma}\varphi \in \ell(s)$	iff	$\varphi \in \ell(s \circ \sigma)$	provided $\odot_{\sigma}\varphi \in For(\xi)$
(FC5)	$\varphi \in \ell(s)$	implies	$\Diamond_{(\Gamma)}\varphi \in \ell(s)$	provided $\Diamond_{(\Gamma)}\varphi \in For(\xi)$, including $\Gamma = n$,
(FC6)	$\Diamond_{(\Gamma)}\varphi \in \ell(s)$	iff	$\Diamond_{(\Gamma)}\varphi \in \ell(t)$	provided $s \equiv_{\Gamma} t$, including $\Gamma = n$.

Fix a set Y. A **saturated set of ξ-mosaics** (short: SSM) M over Y is a set of ξ-mosaics with each base set contained in Y and satisfying the following two conditions:

Sat$\diamond$ If $\mu = (X, \ell) \in M$, $s \in {}^{n}X$ and $\diamond\varphi \in \ell(s)$, then there exists a mosaic $\mu' = (X', \ell')$ in M and a $t \in {}^{n}X'$ such that $\varphi \in \ell'(t)$, and for all $\diamond\psi \in For(\xi)$: $\diamond\psi \in l(s)$ iff $\diamond\psi \in \ell'(t)$;

Sat$\Diamond_{(\Gamma)}$ If $\mu = (X, \ell) \in M$, $s \in {}^nX$ and $\Diamond_{(\Gamma)}\varphi \in \ell(s)$, then either $\varphi \in \ell(t)$ for some $t \in {}^nX$ with $t \equiv_\Gamma s$, or there exists a set $\Delta \subset (Y \setminus X)$, a surjective function $f : \Gamma \longrightarrow \Delta$, and a mosaic $\mu' = (X', \ell')$ such that

(S1)[2] $X' = \mathcal{R}(s^\Gamma_{f\Gamma})$

(S2) $\varphi \in \ell'(s^\Gamma_{f\Gamma})$

(S3) $\ell(t) = \ell'(t)$ for every $t \in {}^nX' \cap {}^nX$.

For a formula ξ, M is ξ-SSM if M is a saturated set of coherent ξ-mosaics over some set Y with $|Y| \leq 2n$ and M contains a mosaic (X, l) and an $s \in {}^nX$ such that $\xi \in l(s)$.

The next lemma makes the crucial connection between satisfiability and the existence of SSM's.

Lemma 3.1 *Let ξ be any MLR_n^+- formula. Then ξ is satisfiable in a generalised assignment model if and only if there exists a ξ-SSM.*

Proof. We start with the easy left to right direction. Let $\mathfrak{M} = (W, v)$ satisfy ξ. We will "cut out" mosaics from W. For any $s \in W$ define $\mu_s = (X_s, \ell_s)$ as follows:

$$\begin{array}{lcl} X_s & = & \mathcal{R}(s) \\ \ell_s(t) & = & \{\psi \in For(\xi) \mid \mathfrak{M}, t \Vdash \psi\}. \end{array}$$

Let $2n$ denote the set $\{0, 1, \ldots, 2n-1\}$. Define $M = \{\mu = (X, \ell) \mid X \subseteq 2n$ and $\mu \simeq \mu_s$ for some $s \in W\}$. It is not hard to see that M is an ξ-SSM over $2n$ because all mosaics come from a model.

Now we prove the right to left direction of the lemma. Let M be a ξ-SSM. Without loss of generality we may assume it is closed under submosaics. This direction can be proved in several ways. The most elegant way –as described in (Hirsch et al., to appear; Andréka et al., 1996; Marx, 1997)– is to use Herwig's Theorem. This construction has the advantage that the final model has a finite base and no step-by-step argument is needed.

Here we will use a step-by-step construction. (The thus constructed model has some nice features which will be used for proving the complexity of the loosely guarded fragment below.) In ω many steps we will construct a labelled n-dimensional graph which we can easily turn into a model for which we can prove a truth lemma stating that truth at a sequence coincides with membership in the label of that sequence.

A *labelled ξ-graph* (short *LG*) is a structure (U, V, ℓ) where U is a set, V is a local cube over U, and ℓ is a function labelling elements of V with subsets of $For(\xi)$. We assume that U is the smallest set such that $V \subseteq {}^nU$. An LG (U, V, ℓ) is called *good* if it satisfies the following conditions:

1. each of its restrictions to cubicles in V is isomorphic to a mosaic in M. Formally:

$$\text{for all } s \in V \text{ there is a } \mu \in M \text{ such that } (\mathcal{R}(s), \ell\restriction_{\mathcal{C}(s)}) \simeq \mu. \quad (15)$$

2. there is a sequence $s \in V$ with $\xi \in \ell(s)$

3. for all $s, t \in V$, for all $\diamondplus\varphi \in For(\xi)$ it holds that $\varphi \in \ell(s)$ implies $\diamondplus\varphi \in \ell(t)$

4. if $s \in V$ and $\diamondplus\varphi \in \ell(s)$, then there exists a $t \in V$ such that $\varphi \in \ell(t)$.

We will now describe the construction. In the first step we will pick out a minimal set $M' \subseteq M$ satisfying

- there is a mosaic (X, ℓ) in M' with a sequence $s \in {}^nX$ containing ξ in its label

- for any $\diamondplus\varphi \in For(\xi)$, if $\diamondplus\varphi \in \ell(s)$, there is a mosaic $(X', \ell') \in M'$ and a $t \in {}^nX'$ with $\varphi \in \ell'(t)$, and for all $\diamondplus\psi \in For(\xi)$: $\diamondplus\psi \in l(s)$ iff $\diamondplus\psi \in \ell'(t)$.

Clearly such a set exists. Form LG_1 by taking the disjoint union of all mosaics in M'. Trivially LG_1 is good and finite (use FC5, FC6 and Sat $\diamondplus$).

The following claim will be the drive behind our construction:

Claim 1 (Defect repair lemma) Let $\mathfrak{G} = (U, V, \ell)$ be a finite good LG. Let s be any sequence in V and $\Diamond_{(\Gamma)}\varphi \in \ell(s)$. Then there exist a good LG $\mathfrak{G}^* = (U^*, V^*, \ell^*)$ such that $\mathfrak{G} \subseteq \mathfrak{G}^*$, the size of U^* is at most $|U| + n - 1$, and V^* contains a sequence $t \equiv_\Gamma s$ with $\varphi \in \ell^*(t)$.

Proof of Claim. Suppose $\Diamond_{(\Gamma)}\varphi \in \ell(s)$ while there is no $t \in V$ such that $s \equiv_\Gamma t$ and $\varphi \in \ell(t)$. (If there is, then we just take $\mathfrak{G}$ itself.)

Let μ be a mosaic in M which is isomorphic to $(\mathcal{R}(s), \ell\restriction_{\mathcal{C}(s)})$. By (15) it exists. For the sake of a smooth presentation, we simply pretend that $(\mathcal{R}(s), \ell\restriction_{\mathcal{C}(s)})$ is *itself* a mosaic of M. Since M is saturated, there are $\Delta \subseteq (Y \setminus \mathcal{R}(s))$ and a mosaic $\mu' = (\mathcal{R}(s^\Gamma_\Delta), \ell')$ satisfying the saturation conditions

(S1)–(S3). Now let Δ' be a set such that $|\Delta| = |\Delta'|$ and $\Delta' \cap U = \emptyset$. We define the new LG $\mathfrak{G}^* = (U^*, V^*, \ell^*)$ as follows:

$$\begin{aligned} U^* &= U \cup \Delta', \\ V^* &= V \cup \mathcal{C}(s^{\Gamma}_{\Delta'}), \\ \ell^*(t) &= \begin{cases} \ell(t) & \text{if } t \in V \\ \ell'(s^{\Gamma}_{\Delta} \circ \sigma) & \text{if } t \in \mathcal{C}(s^{\Gamma}_{\Delta'}),\ t = s^{\Gamma}_{\Delta'} \circ \sigma. \end{cases} \end{aligned}$$

In other words, $\mathfrak{G}^*$ consists of gluing an isomorphic copy of μ' to $\mathfrak{G}$. Note that the domains X and $\mathcal{C}(s^{\Gamma}_{\Delta'})$ of the two constituting parts are not disjoint; fortunately, by (S3) this will cause no problems for the well-definedness of ℓ^*.

It is not difficult to see that $\mathfrak{G}^*$ contains a witness for the "defect" $\Diamond_{(\Gamma)}\varphi \in \ell(s)$: by (S2), $\varphi \in \ell^*(s^{\Gamma}_{\Delta'})$. Note that the new LG satisfies condition (15), because we assumed that M is closed under submosaics. The other good-LG conditions are easy to check, so we have found the desired extension of $\mathfrak{G}$. ◀

Using the defect repair lemma and standard step-by-step techniques we can make sure that in the limit we have found an LG (U, V, ℓ) satisfying for all $s \in V$, and all formulas $\varphi, \psi \in For(\xi)$:

LG1 the conditions FC1–FC4,
LG2 $\Diamond\!\!\!\!+ \varphi \in \ell(s)$ iff there exists a $t \in V$ such that $\varphi \in \ell(t)$,
LG3 $\varphi \in \ell(s)$ implies $\Diamond_{(\Gamma)}\varphi \in \ell(t)$, provided $s \equiv_{\Gamma} t$ and $\Diamond_{(\Gamma)}\varphi \in For(\xi)$,
LG4 $\Diamond_{(\Gamma)}\varphi \in \ell(s)$ implies there exists a $t \in V$ such that $s \equiv_{\Gamma} t$ and $\varphi \in \ell(t)$.

Only the third condition is not trivially true. So assume $s, t \in V$, $s \equiv_{\Gamma} t$ and $\varphi \in \ell(s)$. Then by (15) and FC5, $\Diamond_{(\Gamma)}\varphi \in \ell(s)$. Take an index $j \notin \Gamma$. It exists because $\Gamma \neq n$. Form s' by substituting s_j for every coordinate in Γ in s. Then $s' \in \mathcal{C}(s) \cap \mathcal{C}(t)$ and $s \equiv_{\Gamma} s' \equiv_{\Gamma} t$. s' is in V because V is a local cube. By (15), the mosaics $(\mathcal{R}(s), \ell \restriction_{\mathcal{C}(s)})$ and $(\mathcal{R}(t), \ell \restriction_{\mathcal{C}(t)})$ are both (isomorphic to mosaics) in M. But then by FC6, $\Diamond_{(\Gamma)}\varphi \in \ell(t)$.

Finally, defining a valuation v by $\mathsf{v}(p) = \{s \in V \mid p \in \ell(s)\}$, we obtain a model $\mathfrak{M} = (V, \mathsf{v})$ for which we can easily –by the above LG-conditions– prove the truth lemma

$$\text{For all } \varphi \in For(\xi), \text{ for all } s \in V, \ (\mathfrak{M}, s \Vdash \varphi \text{ iff } \varphi \in \ell(s)).$$

Since in the first step we included a mosaic with a sequence having ξ in its label, $\mathfrak{M}$ satisfies ξ. QED

Theorem 3.2 *The satisfaction problem of the logic MLR_n^+ interpreted on the class of local cubes is in* EXPTIME.

Decidability of this language on this class of frames is due to (Mikulás, 1995).

PROOF. Fix an n. Let ξ be an arbitrary MLR_n^+ formula. We will describe an EXPTIME algorithm which decides whether ξ can be satisfied on a local cube. By Lemma 3.1 this is equivalent to deciding whether there exists a ξ-SSM. This last thing is what we will do.

Let M be the set of all ξ-mosaics whose base is a subset of $2n$. Our algorithm will construct a chain $M \supseteq M_1 \supseteq M_2 \ldots M_k$, of decreasing sets of mosaics which will answer "ξ is satisfiable" if and only if it finds a ξ-SSM.

The algorithm is defined as follows:

begin
1) input ξ;
2) create M;
3) delete all incoherent mosaics from M, let this be M_1;
4) call function **mosaic**(M_1, ξ);
5) if **mosaic**(M_1, ξ) returns true, output "ξ is satisfiable"
    else output "ξ is not satisfiable"
end

function **mosaic**(M, ξ) returns boolean
begin
    * check occurrence of a ξ-mosaic *
    if there is no mosaic $(X, l) \in M$ with an $s \in {}^nX$ and $\xi \in l(s)$,
      then return false

    delete mosaics with unsolvable defects
    for each formula $\Diamond_{(\Gamma)}\varphi \in For(\xi)$,
        for each mosaic $\mu = (X, \ell) \in M$,
            for each $s \in {}^nX$ such that $\Diamond_{(\Gamma)}\varphi \in \ell(s)$, do
            if there is no mosaic $\mu' = (X', \ell') \in M$ such that $\varphi \in t$ for some $t \in {}^nX'$ with $s \equiv_\Gamma t$,
                call **mosaic**$(M \setminus \{\mu\}, \xi)$;
    if **mosaic**$(M \setminus \{\mu\}, \xi)$ returns false
        then return false
    return true
end

First we show that the algorithm is correct. Suppose it outputs that ξ is satisfiable. Then it has found a ξ-SSM. But then, by Lemma 3.1, ξ is indeed satisfiable on a local cube.

For the other side, suppose ξ is satisfiable on a local cube. Then by lemma 3.1, there exists a ξ-SSM M. M is a subset of all ξ-mosaics. The algorithm will never delete any mosaic from M because it can always find witnessing mosaics for defects. So after some time it will call **mosaic**(M', ξ) where either $M' = M$ or M' is a larger ξ-SSM, which returns true, so it outputs "ξ is satisfiable".

Now we show that the complexity of the algorithm is in EXPTIME.

(1) The set of all ξ-mosaics M with base contained in $2n$ can be created in exponential time.

(2) To check whether a mosaic is coherent, or it contains a ξ-sequence takes polynomial time.

(3) So part (1)–(3) of the algorithm runs in exponential time.

(4) Each round of calling the function **mosaic** takes exponential time. There can be at most exponentially many rounds since the size of M is exponential in ξ.

(5) So also part (4) of the algorithm takes exponential time, whence the whole algorithm runs in EXPTIME. QED

Because MLR_n^+ contains the universal modality, we obtain the following

Corollary 3.3 *The validity problem, the valid rule/global consequence problem and the valid sequent problem of MLR_n^+ interpreted on the local cubes are all in* EXPTIME.

PROOF. By Theorem 3.2, the fact that MLR_n^+ contains the universal modality, Theorem 1.8 and the fact that EXPTIME is a deterministic complexity class, so the validity problem is in it iff the satisfaction problem is in it. QED

A simple relativisation argument also provides these results for any class between the class R_n of all assignment frames and the local cubes.

Corollary 3.4 *Let $\mathsf{LC}_n \subseteq \mathsf{K}_n \subseteq \mathsf{R}_n$. The validity problem, the valid rule/global consequence problem and the valid sequent problem of MLR_n^+ interpreted on the class K_n are all in* EXPTIME.

PROOF. Fix a class K_n between R_n and LC_n. Using the formulas $\{\odot_\sigma \top \mid \sigma \in {}^n n\}$ we can express every condition on the "local cube" part of the universe of a relativised model $\mathfrak{M} = (V, \mathsf{v})$. For example, $\mathfrak{M} \models \{\odot_\sigma \top \mid \sigma \in {}^n n\}$ iff V is a local cube iff $s \in V$ implies $\mathcal{C}(s) \subseteq V$. Let Σ be such a set of formulas such that $\mathfrak{M} \models \Sigma$ iff $\mathfrak{M} \in \mathsf{K}_n$.

Let $(\cdot)^t$ be the following translation from MLR_n^+ to MLR_n^+ formulas, where v is a new variable:

$$\begin{array}{rcl@{\qquad}rcl} p^t &=& p & (\odot_\sigma \varphi)^t &=& \odot_\sigma(v \wedge \varphi^t) \\ \iota\delta_{ij}^t &=& \iota\delta_{ij} & (\Diamond_{(\Gamma)}\varphi)^t &=& \Diamond_{(\Gamma)}(v \wedge \varphi^t) \\ (\neg\varphi)^t &=& \neg\varphi^t & (\diamondplus \varphi)^t &=& \diamondplus(v \wedge \varphi^t) \\ (\varphi \wedge \psi)^t &=& \varphi^t \wedge \psi^t. & & & \end{array}$$

We claim that $\varphi^t \wedge v \wedge \boxplus(v \to \Sigma[\top/v])$ is LC_n-satisfiable iff φ is K_n satisfiable, which is proved by a straightforward relativisation argument. Then the result is immediate by the previous corrolary. QED

3.2 Packed-dense local cubes and the loosely guarded fragment

Theorem 3.2 will also provide us complexity results for arrow logic and the loosely guarded fragment. To obtain them we have to consider a subclass of the local cubes, called the packed-dense local cubes.

Definition 3.5 *Let $V \subseteq {}^n U$ be a local cube and $X \subseteq U$. We say that X is* packed in V *if for all $x, y \in X$, we have ${}^n\{x, y\} \subseteq V$. We say that X is* live in V *if ${}^n X \subseteq V$ (or alternatively, if there is a $s \in V$ such that for all $x \in X$, $x = s(i)$, for some i.). If V is such that for every set X with $|X| \leq n$, X is packed in V only if X is live in V, then we call V* packed-dense. *If $\mathfrak{M} = (V, \mathsf{v})$ is a model and V is packed-dense, we also call $\mathfrak{M}$ packed-dense.*

If V is clear from the context we usually say just packed and live. If we assume that U is chosen minimally, as we always do, then every one-element set is both packed and live, because V is a local cube. Every two-element packed set is clearly live. So every 2-dimensional local cube is packed-dense. The set $\{a, b, c\}$ is packed in the local cube ${}^3\{a, b\} \cup {}^3\{b, c\} \cup {}^3\{a, c\}$ but not live. The packed dense condition can be written as a first order formula; in the case of dimension 3 for instance as

$$\forall xyz((Vxyx \wedge Vxzx \wedge Vyzy) \to Vxyz). \tag{16}$$

Below we will prove that the class of local n-cubes and the class of packed-dense local n-cubes have the same MLR_n^+-validities. This is all the more surprising because just a slight variation of (16) leads to undecidability.

Consider the following condition on V where we have just deleted one of the conjuncts in the antecedent of (16),

$$\forall xyz((Vxyx \wedge Vxzx) \rightarrow Vxyz). \tag{17}$$

It is easy to show that every local 3-cube satisfying (17) is in fact a disjoint union of cubes. Which implies that the MLR_3-theory of the class of local cubes satisfying this condition is undecidable, by Tarski's result mentioned in the introduction.

The syntactic difference between the two formulas is that in the first one, every variable which is quantified over co-occurs with every other variable in an atom in the antecedent, while this does not hold in (17) for the variables z and y. This syntactic feature is the crucial ingredient of the loosely guarded fragment.

Theorem 3.6 *For any n, the class of local n-cubes and the class of packed-dense local n-cubes have the same set of MLR_n^+-validities.*

Proof. One direction is obvious. For the other, let a formula be satisfied on a local cube. By Lemma 3.1, every local cube-satisfiable formula is satisfiable on a model as constructed in the proof of that lemma. We show that that model is packed-dense by an induction on the construction. In the base step, the property trivially holds, since LG_1 is a disjoint union of cubes. For the induction step of the construction, assume that in the labelled graph $\mathfrak{G} = (U, V, l)$, V is packed-dense. Let the extended graph be $\mathfrak{G}^* = (U^*, V^*, l^*)$ as in the proof of the lemma. Let $X \subseteq U^*$ which is packed in V^*. If $X \subseteq U$, then the claim holds by the inductive hypothesis. Otherwise, there is an $x \in X \cap \Delta'$. Because X is packed, for every $y \in X$, ${}^n\{x, y\} \subseteq V^*$. But all sequences in V^* containing members from Δ' are in $\mathcal{C}(s^{\Gamma}_{\Delta'})$, whence ${}^nX \subseteq \mathcal{C}(s^{\Gamma}_{\Delta'}) \subseteq V^*$. Whence the constructed model is packed-dense in every stage of the construction, so it is packed-dense itself. QED

In the introduction we saw that MLR_n interpreted on cubes is just a syntactic variant of first order logic restricted to n variables , L_n. In this modal disguise, we tamed first order logic by changing its semantics. There is also a syntactic approach to taming: keep the classical semantics but consider only certain decidable subsets of all first order formulas. This syntactic approach leads to several bounded or restricted fragments. As we see below, these two approaches are really two sides of the same coin with guarded formulas being classically valid if and only if they are locally cubic valid.

We now define the syntactic counterpart of restricting the set of assignments. Let $\varphi \in L_n^+$. Here L_n^+ denotes L_n expanded with generalised quantifiers $\exists\vec{v}$ for every $\vec{v} \subseteq \{v_0, \ldots, v_{n-1}\}$. Its modal counterpart is unmistakingly MLR_n^+. We define $Rel_n(V, \varphi)$ as the formula $V(v_0, \ldots, v_{n-1}) \wedge \varphi^V$ where φ^V is φ with every subformula of the form $\exists\vec{v}\psi$ substituted by $\exists\vec{v}(V(v_0, \ldots, v_{n-1}) \wedge \psi)$. We assume that V does not occur in φ. As the n-ary predicate symbol V will only occur in the form $V(v_0, \ldots, v_{n-1})$ we often just write V.

Let $(\cdot)^t$ be the formula isomorphism between MLR_n^+ formulas (with the application of the $\odot_\sigma$ restricted to variables) and L_n^+ from the introduction. Then clearly, for every MLR_n^+ formula φ,

φ is valid on the class LC_n iff "V is a local cube" $\models Rel_n(V, \varphi)$ is classically valid.

So, by Theorem 3.2, the set of all L_n^+ formulas where every occurrence of a quantifier is relativised by one and the same locally cubic predicate V is decidable in EXPTIME.

A generalisation of this phenomenon is due to Andréka, van Benthem and Németi, and called the *loosely guarded fragment* of first order logic. For its definition, we expand first order logic (with equality but without function symbols and constants) with generalised quantifiers $\exists\vec{v}$. A formula of this language is called loosely guarded if it is constructed from atoms using the booleans and "loosely guarded quantification":

$$\exists\vec{v}(G\vec{v}\vec{x} \wedge \varphi(\vec{v}\vec{x})),$$

is a loosely guarded formula if G is a conjunction of atomic formulas with variables from $\vec{v}\vec{x}$ and moreover every variable $v \in \vec{v}$ co-occurs with every other variable in $\vec{v}\vec{x}$ in at least one of the conjuncts, and φ is a loosely guarded formula. The variables $\vec{v}$ and $\vec{x}$ may occur in any order and any multiplicity in both G and φ, though they are the only variables which occur there free. All these formulas together form the loosely guarded fragment. We abbreviate this fragment by $\mathcal{LGF}$. The fragment where all predicates are at most n-ary and all variables are from $\{v_0, \ldots, v_{n-1}\}$ is abbreviated by $\mathcal{LGF}_n$. The fragment where the guard is just one predicate symbol is called the *guarded fragment*, abbreviated by $\mathcal{GF}$. (van Benthem, 1997) shows decidability of $\mathcal{LGF}$ and (Grädel, 1997) that $\mathcal{LGF}$ and $\mathcal{GF}$ have the same double-exponential time complexity. Both proofs use an adaptation of Németi's mosaic technique which was used in the previous subsection, but now in the more complicated setting of the $\mathcal{LGF}$. Here we will reduce $\mathcal{LGF}$-validity to MLR_n^+-validity on local cubes.

There is a certain resemblance between the $Rel_n(V,\varphi)$ formulas and the loosely guarded ones, in that in both cases quantification is relativised or guarded by a conjunction of atoms. To appreciate the close connection we generalise the notion of being packed in V.

Definition 3.7 *Let $\mathfrak{M} = (D, \mathsf{v})$ be a first order model and $X \subseteq D$. We say that X is* packed, *if for any $x, y \in X$ there are m-ary R and elements $a_1, \ldots, a_m \in D$ with $\{x, y\} \subseteq \{a_1, \ldots, a_m\}$ and $\mathfrak{M} \models R(a_1, \ldots, a_m)$. For a an assignment, we say that a is packed if the range of a is.*

Now let $\varphi(\vec{x})$ be in $\mathcal{LGF}$. Suppose $\mathfrak{M} \models \varphi(\vec{a})$, for some assignment $\vec{a}$ which is packed. Then the evaluation of $\varphi(\vec{a})$ only uses assignments which are packed as well. To see this let $\mathfrak{M} \models \exists\vec{v}(G(\vec{a},\vec{v}) \wedge \varphi(\vec{a},\vec{v}))$ for an arbitrary $\mathcal{LGF}$ and arbitrary packed $\vec{a}$. Then there exists a $\vec{b}$ such that $\mathfrak{M} \models G(\vec{a},\vec{b}) \wedge \varphi(\vec{a},\vec{b})$. Now $\vec{a}\vec{b}$ is packed because $\vec{a}$ is packed and G is a loose guard.

Now the following theorem should not come as a surprise. We formulate it for a slightly different loosely guarded fragment, $\mathcal{VLGF}_n$, the vacuous loosely guarded fragment with n variables. Here we allow only loosely guarded quantification

$$\exists\vec{v}_{\cup}\vec{z}(G\vec{v}\vec{x} \wedge \varphi(\vec{v}\vec{x})),$$

where $\vec{z}$ contains all variables from $\{v_0, \ldots, v_{n-1}\}$ not in $\vec{v}\vec{x}$ and $\exists\vec{v}(G\vec{v}\vec{x} \wedge \varphi(\vec{v}\vec{x})) \in \mathcal{LGF}_n$. Clearly every $\mathcal{LGF}_n$ φ is equivalent to the $\mathcal{VLGF}_n$ φ' obtained by adding all these vacuous variables into all quantifications in φ'.

Theorem 3.8 *For every $\mathcal{VLGF}_n$ ξ, ξ is satisfiable iff $Rel_n(V,\xi)$ is satisfiable on a model where V is interpreted as a packed-dense local cube. Moreover, ξ and $Rel_n(V,\xi)$ can be satisfied on the same domains.*

Proof. From left to right is obvious: we just interpret V as ${}^nD$ for D the domain of the models where φ is satisfied. For the other side, let $\mathfrak{M} = (D, I) \models Rel_n(V,\xi)[s]$. Because $s \in I(V)$ (by definition of $Rel_n(V,\xi)$) and every occurrence of a quantifier in $Rel_n(V,\xi)$ is relativised by V, we may as well assume that the interpretation of the other predicates is restricted to $I(V)$, as follows:

$$\text{for all predicates } P, \text{ if } s \in I(P), \text{ then } {}^n\mathcal{R}(s) \subseteq I(V), \tag{18}$$

without losing satisfiability of $Rel_n(V,\xi)$ at s.

The n variables in $\mathcal{VLGF}_n$ are $\{v_0, \ldots, v_{n-1}\}$. Because we have only n variables we can equate assignments s with n-tuples such that $s(v_i) = s(i) =$ the value of s on the i-th coordinate. We claim that for all $\mathcal{VLGF}_n$ φ and for all assignments $s \in I(V)$,

$$\mathfrak{M} \models \varphi[s] \text{ iff } \mathfrak{M} \models Rel_n(V, \varphi)[s]. \tag{19}$$

The proof is by induction on the construction of the formula. All cases are trivial except the left to right direction of the guarded quantifier. So suppose $\varphi = \exists \vec{v}_{\cup}\vec{z}(G(\vec{v}, \vec{x}) \wedge \psi(\vec{v}, \vec{x}))$, $s \in I(V)$ and $\mathfrak{M} \models \exists \vec{v}_{\cup}\vec{z}(G(\vec{v}, \vec{x}) \wedge \psi(\vec{v}, \vec{x}))$ $[s]$. Then there exists a $t \equiv_{\vec{v}\cup\vec{z}} s$ such that $\mathfrak{M} \models G(\vec{v}, \vec{x}) \wedge \psi(\vec{v}, \vec{x})$ $[t]$. By the argument given just below Definition 3.7, the set $\{t(i) \mid v_i \in \vec{v} \cup \vec{x}\}$ is packed, so by (18) it is also packed in V. But then , because V is packed-dense, the set is live in V. Let t' be any sequence in ${}^n\{t(i) \mid v_i \in \vec{v} \cup \vec{x}\}$ such that $t'(i) = t(i)$ for $v_i \in \vec{v} \cup \vec{x}$. Then $t' \in I(V)$ because $\{t(i) \mid v_i \in \vec{v} \cup \vec{x}\}$ is live in V. Since t and t' agree on all free variables in $G(\vec{v}, \vec{x}) \wedge \psi(\vec{v}, \vec{x})$, also $\mathfrak{M} \models G(\vec{v}, \vec{x}) \wedge \psi(\vec{v}, \vec{x})$ $[t']$. By the inductive hypothesis $\mathfrak{M} \models V \wedge G(\vec{v}, \vec{y}) \wedge \psi^V(\vec{v}, \vec{y})$ $[t']$. Then, as $t' \equiv_{\vec{v}\cup\vec{z}} s$, we have $\mathfrak{M} \models \exists \vec{v}_{\cup}\vec{z}(V \wedge G(\vec{v}, \vec{x}) \wedge \psi^V(\vec{v}, \vec{x}))$ $[s]$. This proves (19), whence the theorem. QED

E. Grädel (Grädel, 1997) showed that the full (loosely) guarded fragment is double-exponential time complete, and that the finite variable fragments are in EXPTIME. We obtain these upper bounds as corollaries of Theorem 3.2.

Corollary 3.9 *(i) For every n, $\mathcal{LGF}_n$ can be decided in* EXPTIME.
(ii) The full loosely guarded fragment can be decided in 2EXPTIME.
(iii) $\mathcal{GF}$ has the strong finite model property.

PROOF. Given an $\mathcal{LGF}_n$ formula φ, we first change it into a $\mathcal{VLGF}_n$ φ^* and then into a L_n^+ formula φ', by substituting every occurrence of a $n-k$-ary predicate $P(v_{\sigma(0)}, \ldots, v_{\sigma(n-k)})$ by a n-ary predicate $P'(v_{\sigma(0)}, \ldots, v_{\sigma(n-k)}, v_{\sigma(0)}, \ldots, v_{\sigma(0)})$. Clearly this can be done in polynomial time and without loss of validity. φ' is still in $\mathcal{VLGF}_n$, but now it is a syntactic variant of a MLR_n^+ formula $(\varphi')^t$. By the last theorem φ' is valid iff $Rel_n(V, \varphi')$ is valid on each model where V is interpreted as a packed-dense local cube. But $Rel_n(V, \varphi')$ is valid on these models iff $(\varphi')^t$ is valid on the class of packed-dense local cubes. By Theorem 3.6, this class has the same validities as the local cubes, which are decidable in EXPTIME by Theorem 3.2.
For (ii) observe that every $\mathcal{LGF}$ formula lives in some $\mathcal{LGF}_n$-fragment. Only now n is not fixed, so depends only on $|\varphi|$. An inspection of the proof of

Theorem 3.2 shows that this n occurs as an extra exponent in the complexity of the decision procedure, whence we obtain only a double exponential bound.
(iii) For guarded formulas, Theorem 3.8 is true for models where V is just a local cube. By (Andréka et al., 1996), every local cube satisfiable MLR_n^+ formula φ can be satisfied on a finite (locally cubic) model whose size is some function of $|\varphi|$. QED

Our results on lower bounds imply that the guarded fragment with 3 variables and without identity and generalised quantifications is already EXPTIME-hard. In the next section we also see some two-variable EXPTIME-hard guarded fragments. There are however quite interesting guarded fragments with 2 variables which are PSPACE-complete. We will meet them in the next section where we discuss logics of dimension two.

Using $\mathcal{LGF}_3$ we also obtain an upper bound for arrow logic.

Theorem 3.10 *(i) The validity problem of arrow logic interpreted on the classes of pair frames PF_H for $H \subseteq \{R, S\}$ is in* EXPTIME.
(ii) The valid rule/global consequence problem and the valid sequent problem of the above arrow logics are also in EXPTIME.

PROOF. Let φ^t be the following translation from arrow logical formulas into $\mathcal{LGF}_3$:

$$\begin{array}{lcl}
p^t(x,y) & = & P(x,y)\\
(\neg\varphi)^t(x,y) & = & \neg\varphi^t(x,y)\\
(\varphi\wedge\psi)^t(x,y) & = & \varphi^t(x,y)\wedge\psi^t(x,y)\\
\iota\delta^t(x,y) & = & x=y\\
(\otimes\varphi^t)(x,y) & = & V(y,x)\wedge\varphi^t(y,x)\\
(\varphi\circ\psi)^t(x,y) & = & \exists z(V(x,z)\wedge V(z,y)\wedge\varphi^t(x,z)\wedge\psi^t(z,y)).
\end{array}$$

Let Σ be a subset of the $\mathcal{LGF}_2$ sentences $\{\forall xy(V(x,y)\to V(x,x)\wedge V(y,y)), \forall xy(V(x,y)\to V(y,x))\}$. Then the sequent $\varphi_1,\ldots\varphi_m \Rightarrow \psi_1\ldots\psi_k$ is valid on Σ-models iff $(\Sigma\wedge\forall xy[V(x,y)\to\varphi_1^t(x,y)\wedge\ldots\wedge\varphi_m^t(x,y)])\to(\forall xy[V(x,y)\to\psi_1^t(x,y)]\vee\ldots\vee\forall xy[V(x,y)\to\psi_k^t(x,y)])$ is valid. The last formula is clearly in $\mathcal{LGF}_3$ and can be obtained in polynomial time from the original sequent. Whence the result by Corollary 3.9. QED

This might be a good point to try to answer why relativisation in modal logics of relations leads quite often to decidability. We proved our strongest result, for the loosely guarded fragment, by reducing it to the decidability of MLR_n^+ over local cubes —which is essentially the first decidability result by relativisation, due to (Németi, 1986)— and the new observation that mosaics can be amalgamated without losing the packed dense property. In essence it is precisely because we can preserve packed-density, that so many adaptations of Németi's result work. If we allow this as an explanation, this covers all cases where

- the definitions of the modalities involved have the form of a loosely guarded formula (in other words, the "standard translation" to first order logic of the language falls in the loosely guarded fragment), and
- the conditions we impose on the universe have the form of loosely guarded sentences (e.g., the locally cubic conditions, but also the packed-dense condition).

Then we can decide the relativised modal logic inside the decidable loosely guarded fragment.

From a related perspective, (Andréka et al., 1996) reaches almost the same conclusion. That article contains a syntactic characterisation of a set of modal operators which when interpreted on relativised n-dimensional structures lead to a decidable logic with the finite model property. In essence they allow all k-adic diamond-type modalities defined on n-dimensional models $\mathfrak{M} = (V, \mathsf{v})$ as follows

$$\mathfrak{M}, s \Vdash \Diamond(\varphi_1, \ldots, \varphi_k) \text{ iff } s \in V \text{ and there exist } s_1, \ldots, s_k \in V \text{ such that } \bigwedge \mathfrak{M}, s_i \Vdash \varphi_i.$$

Such a modality $\Diamond$ is called *existentially packed* (EP for short) if, assuming $s \in V$, the set $\bigcup_{1 \leq i \leq k} \mathcal{R}(s_i)$ is packed in V. Clearly every unary modality with such a definition is EP, and also the composition operator from arrow logic is EP. The following generalisation of composition (due to Vera Stebletsova) however is not EP.

$$\begin{aligned} \mathfrak{M}, (x, y) \Vdash \Diamond(p, q, r) \quad \Longleftrightarrow \quad & (x, y) \in V \text{ and there exist } \{(x, z_1), (z_1, z_2), (z_2, y)\} \subseteq V \\ & \& \ \mathfrak{M}, (x, z_1) \Vdash p \ \& \ \mathfrak{M}, (z_1, z_2) \Vdash q \ \& \ \mathfrak{M}, (z_2, y) \Vdash r. \end{aligned}$$

Note that the (coordinate-wise) difference operator is also not EP.

In algebraic disguise (and in a bit more general formulation), Theorem 23 in (Andréka et al., 1996) states that every modal logic containing only n-dimensional EP-modalities interpreted on the class of all n-dimensional frames is decidable and has the strong finite model property (any satisfiable formula can be satisfied on a finite model whose size is computable from the sentence).

In dimension two, packed-density becomes a vacuous notion. Is there a connection with the fact that we obtain much better complexity and decidability results there?

3.3 Expansions with the counting modalities

(Mikulás, 1995) shows that the expansion of MLR_n^+ with all graded modalities $[k]$ is still decidable when interpreted on the relativised classes of assignment frames. By (Andréka et al., 1996) the logic even has the finite base property. Similar results hold for arrow logic as well (cf., (Mikulás, 1995) for decidability and (Andréka et al., 1996) for the finite base property).

Here we will indicate how to change the given proof for MLR_n in order to obtain the same EXPTIME upper bounds for this expansion.

Theorem 3.11 *The validity problem of MLR_n^+ expanded with all graded modalities interpreted on any class between LC_n and R_n is in* EXPTIME*. The same holds for the valid rule and the valid sequent problem.*

Proof. As explained in Remark 1.10 it is sufficient to compute the complexity of the expansion with *nominals*, i.e, special variables which can only be true at at most one element of the universe of a model. We will change the proof of Lemma 3.1 along the lines of (Mikulás, 1995). The changes we make are based on the fact that states where a nominal is true are unique. In the proof of Claim 1 in Lemma 3.1 we used mosaics over and over again. It is here that we have to ensure that we do not duplicate mosaics with a nominal in their label. We will do this by enlarging our mosaics as follows: every new mosaic consists of a mosaic in the old sense plus a distinguished part which contains all sequences with a nominal in their label and their local cubes. We will ensure that this distinguished part is the same for all mosaics in an SSM. In this way we can still amalgamate mosaics in the manner of the defect repair lemma. Technically these changes are implemented as follows: Again let a formula ξ be fixed. We use k_i to denote nominals.

A ξ-mosaic now is a structure $\mu = (X, \ell)$ such that X is a set of size at most $n + |\xi| \cdot n$, and the other notions stay the same. Define the named part P_μ of a mosaic as

$$P_\mu = \{s \in {}^nX \mid \text{there exists a nominal } k \in For(\xi) \text{ such that } k \in \ell(s)\}.$$

Define the *distinguished part* D_μ as $\{s \in {}^nX \mid \exists t \in P_\mu \text{ and } s \in \mathcal{C}(t)\}$. We change the notion of a submosaic to reflect the fact that the distinguished part should remain the same. That is, we add a condition saying that $D_\mu = D_{\mu'}$.

We add the following two clauses to the definition of a coherent mosaic. For all $s, t \in {}^nX$:

FC7	$k \in \ell(s)$	implies	$k \notin \ell(t)$, for k a nominal in $For(\xi)$, and $s \neq t$
FC8	$\Diamond_{(\Gamma)} k \in \ell(s)$	implies	there exists a $t \in P_\mu$ such that $s \equiv_\Gamma t$ and $k \in \ell(t)$.

We add to the definition of an SSM M a clause saying that the distinguished part of all mosaics in M is the same:

$$\text{for all } \mu, \mu' \in M,\ D_\mu = D_{\mu'} \text{ and } \ell \restriction_{D_\mu} = \ell' \restriction_{D_{\mu'}},$$

and we change the Sat-conditions accordingly.

Finally we change the definition of a ξ-SSM only in the size of the base set Y. It should reflect the extra space required by the nominals, so $Y \leq 2n + |\xi| \cdot n$. ($|\xi| \cdot n$ denotes the maximum number of elements required in order to make all nominals in ξ true at precisely one n-tuple.)

Clearly we can still decide in EXPTIME whether or not a ξ-SSM exists for arbitrary ξ, so we are left to prove that satisfiability of ξ is indeed equivalent to the existence of a ξ-SSM. The proof of this is very similar to the proof of Lemma 3.1, the only difference having to do with the distinguished part of the mosaics. We leave these details to the reader, who can consult (Mikulás, 1995) for the details. QED

3.4 NP-complete reducts

Even when interpreted on the classical cubic frames, MLR_n has some interesting NP-complete reducts. MLR_1, where the $\iota\delta_{11}$ and the $\odot_{11}$ do not do anything is NP-complete. Modally, MLR_1 is just the logic of the universal modality, thus the modal logic **S5**. NP-completeness now follows from (Ladner, 1977). As a first order logic, MLR_1 is monadic first order logic with one variable.

When we allow more variables (i.e., $n > 1$) the complexity of the full language goes up, but this is solely due to the $\Diamond_i$.

Theorem 3.12 *The boolean plus $\{\iota\delta_{ij}, \odot_\sigma\}_{i,j<n,\, \sigma \in {}^nn}$ reduct of MLR_n interpreted on the class C_n of cubic frames has an NP-complete satisfaction problem. In fact, every formula can be satisfied on a cubic frame whose base set U has size at most n.*

PROOF. NP-hardness follows because propositional logic is NP-complete. For the upper bound let a formula in this reduct be satisfied in a model at a sequence s. Since the subframe generated from s by $\bowtie_\sigma$ is just the cubicle of s we obtain the n-size model property. Since n is fixed the reduct is in NP. QED

So the modal logic of permutations and substitutions (Venema, 1995) has a nice complexity. For completeness we note that MLR_n is just as expressive as its boolean plus $\{\iota\delta_{ij}, \odot_\sigma, \Diamond_0\}$-reduct, when interpreted on any class of frames contained in the local cubes. This follows because on the local cubes the following formula is valid

$$\Diamond_i\varphi \leftrightarrow \otimes_{i0}\Diamond_0\otimes_{i0}\varphi. \tag{20}$$

We will use this fact in the next section to simplify our proofs considerably.

4 Cylindric Modal logic of dimension two

We will show that the complexity of systems in two dimensions is better behaved than in dimension three or higher. In particular we show that for any language L between the $\{\Diamond_0, \Diamond_1\}$-reduct of MLR_2 and MLR_2,

- the satisfiability problem of L interpreted on local squares is PSPACE-complete (subsection 4.1). For MLR_3 this is EXPTIME-complete on the local cubes.
- the satisfiability problem of L interpreted on squares is NEXPTIME-complete (subsection 4.2). For MLR_3 on cubes the problem is undecidable.
- the satisfiability problem of MLR_2 expanded with coordinate wise counting modalities interpreted on squares is decidable (subsection 4.3). For MLR_3 on the class of all assignment frames these operators lead to undecidability.

In a similar way as we obtained exptime upper bounds for the loosely guarded fragment in the previous section, we will obtain pspace upper bounds for certain 2-variable guarded fragments here in section 4.1, which also contains the corresponding pspace upper bounds for the reducts of arrow logic from Table 5.

Instead of $\Diamond_0, \Diamond_1$ we will use the notation $\diamondminus$, $\diamondvert$, respectively, emphasising the two dimensional nature of these operators. The converse modality $\otimes_{01}$ will be abbreviated as $\otimes$, and instead of $\iota\delta_{01}$ we merely write $\iota\delta$.

4.1 Local squares are pspace complete

We show that the satisfaction problem of MLR_2 formulas over local squares is pspace-complete. Pspace hardness can be shown by interpreting the modal logic K. The translation $(\Diamond\varphi)^t = \diamondvert(r \wedge \diamondminus\varphi^t)$ can easily be shown to work. To show that the problem is in pspace requires some more effort. Our technique uses the equivalence between $\forall\exists$-games and alternating Turing machines (cf. (Chandra et al., 1981; Chlebus, 1986)). An alternating Turing machine solves a problem in PTIME iff a standard Turing machine solves it in PSPACE. A $\forall\exists$-game is a two player game with players $\forall$ ("he") and $\exists$ ("she"). Our proof is build on the following steps: for every formula ξ we define a $\forall\exists$-game called ξ-game, which always terminates after at most $|\xi|$ rounds. This means that we can decide whether $\exists$ has a winning strategy in ξ-game using a PTIME algorithm which runs on an alternating Turing machine. Finally we show that winning strategy for $\exists$ and satisfiability of ξ are equivalent.

Remark 4.1 Instead of deciding arbitrary MLR_2-formulas we decide MLR_2 formulas in which no $\diamondminus$ operators occur. This is sufficient by the locally square valid theorem

$$\diamondminus\varphi \leftrightarrow \otimes \diamondvert \otimes\varphi, \tag{21}$$

where $\otimes$ is the permutation modality, which permutes the coordinates 0 and 1. Also we will only decide formulas of the form $\psi \wedge \iota\delta_{01}$. Clearly this is no restriction since a formula φ is local square satisfiable iff $\iota\delta_{01} \wedge \diamondvert\varphi$ is local square satisfiable.

Preliminaries. First, the **cylindric depth** $d(\xi)$ of a formula ξ is defined by the following induction:

$$\begin{array}{rcl}
d(p) & = & 0 \\
d(\iota\delta) & = & 0 \\
d(\neg\varphi) & = & d(\varphi) \\
d(\varphi \wedge \psi) & = & \max(d(\varphi), d(\psi)) \\
d(\otimes\varphi) & = & d(\varphi) \\
d(\diamondvert\varphi) & = & d(\varphi) + 1.
\end{array}$$

The depth of a set of formulas is just the maximal depth of its members. As before, for a formula ξ, $For(\xi)$ is the set of all subformulas of ξ plus single negations plus $\iota\delta$.

Mosaics. Let ξ be a fixed MLR_n-formula. A **mosaic** is a structure $\mu = (X, \ell_r, l, d)$ with X a set of size at most 2, $d \in X$, $\ell_r : {}^2X \longrightarrow Pow(For(\xi))$, and l is a function sending elements $s \in {}^2X$ to a subset of $\ell_r(s)$. Intuitively, $\ell_r(s)$ denotes the set of formulas which are "relevant" at s, and $l(s)$ those formulas in $\ell_r(s)$ which are true at s.

A mosaic $\mu = (\{d', d\}, \ell_r, l, d)$ is **smart** if the cylindric depth of $\ell_r(d, d)$ is strictly smaller than that of $\ell_r(d', d')$.

A mosaic is **coherent** if it satisfies the following conditions for all $s, t \in {}^2X$, $x, y \in X$,

(FC1)	$\iota\delta \in \ell(s)$	iff	$s_0 = s_1$,	
(FC2)	$\neg\varphi \in \ell(s)$	iff	$\varphi \notin \ell(s)$	provided $\neg\varphi \in \ell_r(s)$
(FC3)	$\varphi \wedge \psi \in \ell(s)$	iff	$\varphi \in \ell(s)$ & $\psi \in \ell(s)$	provided $\varphi \wedge \psi \in \ell_r(s)$
(FC4)	$\otimes\varphi \in \ell(x,y)$	iff	$\varphi \in \ell(y,x)$	provided $\otimes\varphi \in \ell_r(x,y)$
(FC5)	$\neg \Diamond \varphi \in \ell(x,x)$	implies	$\neg\varphi \in \ell(x,y)$	
(FC6)	$\neg \Diamond \varphi \in \ell(x,y)$	implies	$\neg \Diamond \varphi \in \ell(x,x)$.	

The ξ-game.
Start $\exists$ plays a mosaic $\mu = (\{a\}, \ell_r, l, a)$ with $\xi \in \ell(a,a)$.
$\forall$-move Let $\mu = (X, \ell_r, l, d)$ be the previous move of $\exists$. $\forall$ wins if μ is not smart and coherent. Otherwise $\forall$ chooses a formula $\Diamond \psi \in \ell(d,d)$ such that $\psi \notin \ell(d,x)$, for any $x \in X$.
$\exists$-move $\exists$ responds with a mosaic $\mu' = (X', \ell'_r, \ell', d')$ for some $d' \in X' \setminus X$ such that

(S1) $X' = \{d, d'\}$
(S2) $\varphi \in \ell'(d,d')$
(S3) $\ell(d,d) = \ell'(d,d)$ and $\ell_r(d,d) = \ell'_r(d,d)$.

$\exists$ wins the game if she can always respond with a smart and coherent mosaic. Otherwise $\forall$ wins.

Proposition 4.2 *ξ-game stops after at most $d(\xi)$ moves.*

Proof. After every move of $\exists$, either $\forall$ wins because the mosaic μ played is not smart and coherent, or otherwise it is smart and $\forall$ chooses a formula $\Diamond \psi \in \ell(d,d)$ where d is the new point in the last played mosaic. Since the mosaic is smart, the label of (d,d) must be of strictly smaller degree than the label of the pair picked by $\forall$ in the previous round. So after $d(\xi)$ rounds the degree of the label of the new sequence becomes 0, and there is no more $\Diamond$-formula for $\forall$ to choose. QED

Proposition 4.3 *If ξ is satisfiable, then $\exists$ has a winning strategy for ξ-game.*

Proof. Assume ξ is satisfied in model $\mathfrak{M}$ at sequence (a,b). By remark 4.1 we may assume that $a = b$. In her first move $\exists$ plays the mosaic $(\{a\}, \ell_r, \ell, a)$ where $\ell_r(a,a) = For(\xi)$ and $\ell(a,a) = \{\psi \in \ell_r(a,a) \mid \mathfrak{M}, (a,a) \Vdash \psi\}$.

Clearly $\exists$ can start the game with this mosaic, and it is coherent since $\mathfrak{M}$ is a model.

To prove that $\exists$ can win any game starting with this mosaic, we do an induction on the length of the game. So assume that up to round k $\exists$ could always respond with a smart and coherent mosaic, which was "cut out from the model $\mathfrak{M}$". Let $\mu = (X, \ell_r, \ell, d)$ be the k-th mosaic $\exists$ played. Assume $\forall$ plays $\diamondsuit\,\psi$. It is almost cheating what $\exists$ does. She looks at the model $\mathfrak{M}$ and finds the pair (d,d). By assumption $\mathfrak{M},(d,d) \Vdash \diamondsuit\,\psi$, so there exists an x such that $\mathfrak{M},(d,x) \Vdash \psi$. x cannot be in X since we assumed that the previous mosaic came from the model.

She just plays the mosaic $\mu' = (\{x,d\}, \ell'_r, \ell', x)$ where

$$
\begin{array}{lcl}
\ell'_r(d,d) & = & \ell_r(d,d)\\
\ell'(d,d) & = & \ell(d,d)\\
\ell'_r(d,x) & = & For(\{\psi\} \cup \{\neg\varphi \mid \neg\diamondsuit\,\varphi \in \ell(d,d)\} \cup \{\diamondsuit\,\varphi \in \ell(d,d)\})\\
\ell'(d,x) & = & \ell'_r(d,x) \cap \{\varphi \mid \mathfrak{M},(d,x) \Vdash \varphi\}\\
\ell'_r(x,d) & = & For(\{\varphi \mid \otimes\varphi \in \ell'(d,x)\})\\
\ell'(x,d) & = & \ell'_r(x,d) \cap \{\varphi \mid \mathfrak{M},(x,d) \Vdash \varphi\}\\
\ell'_r(x,x) & = & For(\{\neg\varphi \mid \neg\diamondsuit\,\varphi \in \ell(x,d)\} \cup \{\diamondsuit\,\varphi \in \ell(x,d)\})\\
\ell'(x,x) & = & \ell'_r(x,x) \cap \{\varphi \mid \mathfrak{M},(x,x) \Vdash \varphi\}.
\end{array}
$$

Clearly μ' satisfies S1–S3, so this is a legal move for $\exists$. By definition of ℓ'_r, the mosaic is smart. Coherence is easy to check since the mosaic is cut out from a model. QED

Proposition 4.4 *If $\exists$ has a winning strategy for ξ-game, then ξ is satisfiable in a finite model whose size of the base set is exponential in $|\xi|$.*

Proof. Assume $\exists$ has a winning strategy for ξ-game. Fix a start mosaic of $\exists$. We will now represent a part of all possible ξ-games which $\exists$ wins starting from the fixed mosaic as a labelled tree as follows.

The nodes in the tree are the mosaics played by $\exists$, the labels on the edges are all the possible moves of $\forall$, one for each $\diamondsuit$-formula in $\ell(d,d)$. Cf. Figure 2.

By Proposition 4.2 the depth of the tree is at most $d(\xi)$. The branching factor is less than $|\xi|$. So the number of nodes in the tree is less than $(|\xi|)^{(|\xi|+1)}$.

In every move of $\exists$, the domain of the mosaic μ changes one element. We may assume that in the tree all these elements are pairwise disjoint. Let l and ℓ_r be the union of all labelling functions of all mosaics in the tree. By S3, l and ℓ_r are functions.

Now we turn this tree into a model. Let $\mathfrak{M} = (W, \mathsf{v})$ be the locally square model defined by

$$
\begin{array}{lcl}
W & = & \bigcup\{{}^2X \mid (X,\ell_r,\ell,d) \text{ is a mosaic in the tree}\}\\
s \in \mathsf{v}(p) & \text{iff} & p \in \ell(s).
\end{array}
$$

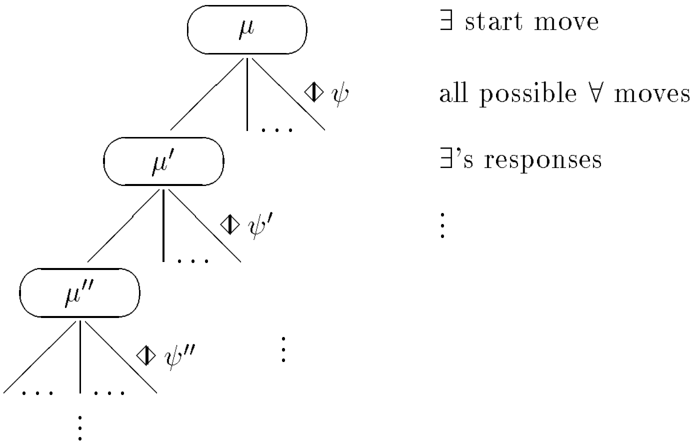

Figure 2: Tree of ξ-games.

Since every node in the mosaic corresponds to an element of the base set U over which W is a binary relation, and the observation about the size of the tree made above, we can conclude that the model is exponential in the length of ξ.

Since ξ is in the label of the pair in the start-mosaic, we are finished when we have proved the truth-lemma:

$$\text{for all } s \in W, \text{ for all } \psi \in \ell_r(s), \mathfrak{M}, s \Vdash \psi \text{ iff } \psi \in \ell(s). \tag{22}$$

We prove this by induction on the construction of ψ. The game structure appears in the diamond-case: to show that $\Diamond\varphi \in l(s)$ implies $s \Vdash \Diamond\varphi$ we use that $\forall$ forces $\exists$ to create a "φ-witness. The other side is proved using the coherence conditions, namely suppose $\Diamond\varphi \in \ell_r(x, y)$, then

$$\neg\Diamond\varphi \in \ell(x, y) \overset{\text{FC6}}{\Rightarrow} \neg\Diamond\varphi \in \ell(x, x) \overset{\text{FC5}}{\Rightarrow} \neg\varphi \in \ell(x, z).$$

QED

Proposition 4.5 *Local square satisfiability is in* PSPACE.

Proof. By the previous two propositions it is sufficient to check whether $\exists$ has a winning strategy in ξ-game. This can be easily implemented on an alternating Turing machine, which runs in ptime because the maximum length of the game is just $d(\xi)$ by the first proposition. Alternatively we can implement this problem on a standard Turing machine. What it will do is to construct the tree given in the figure. Clearly it only has to remember at most one complete branch, which takes space just polynomial in ξ. QED

We conclude with the main result of this section

Theorem 4.6 *For any language L between the $\{\Diamond_0, \Diamond_1\}$-reduct of MLR_2 and MLR_2, the satisfiability problem of L interpreted on local squares is* PSPACE*-complete. Moreover, every satisfiable φ can be satisfied on a model whose base set has size exponential in $|\varphi|$.*

PROOF. That the problem is in PSPACE follows from the previous proposition. That the problem is PSPACE-hard can easily be shown by interpreting modal logic $\mathbf{K}$ using the translating $(\Diamond\varphi)^t = \Diamond_1(r \wedge \Diamond_0\varphi^t)$. The bound on the size of the models is immediate by Propositions 4.3 and 4.4. QED

Corollary 4.7 *For any language L between the $\{\Diamond_0, \Diamond_1\}$-reduct of MLR_2 and MLR_2, the satisfiability problem of L interpreted on any class between R_n and LC_n is* PSPACE*-complete.*

PROOF. Hardness follows as before. The upper bound follows by the relativisation argument of Corollary 3.3. QED

4.1.1 A pspace-complete guarded fragment

The (loosely) guarded fragment can be used to show decidability of many (multi-)modal logics. For an interpretation to work we only need that 1) the modalities involved "standard translate" to a loosely guarded formula, and 2) all conditions on the accessibility relations can also be formulated in $\mathcal{LGF}$. Very often we can even use a fixed variable $\mathcal{LGF}$, so then we obtain even an EXPTIME upper complexity bound. Still there are many modal logics which are pspace-complete, so it is natural to see if we can find guarded fragments which can be decided in pspace. The following fragment is an example. Let $\mathcal{PGF}_2$ denote the sublanguage of $\mathcal{LGF}_2$ where

$$\text{no generalised quantifiers occur in the formulas, and} \tag{23}$$

$$\text{every guard contains an atom of the form } R(v_0, v_1). \tag{24}$$

Theorem 4.8 *$\mathcal{PGF}_2$ is* PSPACE*-complete and every satisfiable $\mathcal{PGF}_2$ formula φ can be satisfied on a model whose size is exponential in $|\varphi|$.*

PROOF. The lower bound follows because the standard translation of modal logic **K** goes to $\mathcal{PGF}_2$. By (24) every $\mathcal{PGF}_2$ is in $\mathcal{VLGF}_2$ defined just above Theorem 3.8. Let $\varphi \in \mathcal{PGF}_2$. First change it into a L_2 formula by stretching out unary predicates, just as in Corollary 3.9. Then by (23) it is the syntactic equivalent of a MLR_2-formula. Now the result follows from Theorems 3.8 and 4.6. QED

$\mathcal{PGF}_2$ is optimal in the sense that if we delete any of the conditions, then the complexity becomes EXPTIME-hard. If we allow generalised quantifiers $\exists v_0 v_1$, we can interpret modal logic **K** plus universal diamond $\oplus$ as follows:

$$\begin{array}{lcl} p^t & = & P(v_0, v_1) \\ (\Diamond\varphi)^t & = & \exists v_1(R(v_0, v_1) \wedge \exists v_0(W(v_0, v_1) \wedge \varphi^t)) \\ (\oplus\varphi)^t & = & \exists v_0 v_1(W(v_0, v_1) \wedge \varphi^t). \end{array}$$

If we don't ask for condition (24), we can interpret as follows:

$$\begin{array}{lcl} p^t(v_0) & = & P(v_0) \\ (\Diamond\varphi)^t(v_0) & = & \exists v_1(R(v_0, v_1) \wedge (W(v_1) \wedge \varphi^t(v_1))) \\ (\oplus\varphi)^t(v_0) & = & \exists v_1(W(v_1) \wedge \varphi^t(v_1)). \end{array}$$

If we allow ternary predicates and 3 variables, but only quantification over two variables, we obtain EXPTIME-hardness by interpreting the $\{\Diamond_0, \Diamond_1, \mathsf{s}_2^1\}$-reduct of MLR_3 evaluated on the class R_3 (cf. Theorem 2.15).

The last theorem accounts for the pspace upper bounds of quite a lot of modal logics. Call a (unary) modality a $\mathcal{PGF}_2$-modality if its standard translation is a $\mathcal{PGF}_2$ formula. We obtain

Theorem 4.9 *Let L be a basic multi-modal logic containing only $\mathcal{PGF}_2$ modalities. Then the validity problem of L is in* PSPACE.

This theorem covers many well known modalities, and rightly excludes some others. We list some examples.

Modalities which are not guarded are the universal modality (no guard at all), the difference operator ($x \neq y$ is not a guard) and also the inaccessibility operator from (Humberstone, 1983):

$$\mathfrak{M}, w \Vdash \triangledown\varphi \text{ iff there exists a } w' \text{ such that } \neg Rww' \text{ and } \mathfrak{M}, w' \Vdash \varphi.$$

In all three cases we obtain a logic which is potentially EXPTIME-hard (in the first two cases we only have to have another **K** modality, in the Humberstone case we would need the complement modality of $\triangledown$, then their disjunction would be a universal modality, cf. Theorem 2.2.(iii)).

Examples of guarded modalities are the tense operators F and P. It also covers funny non-bisimulation invariant modalities as

$$\mathfrak{M}, x \Vdash \Diamond\varphi \quad \Longleftrightarrow \quad \exists y : Rxy \wedge Ryx \ \& \ \mathfrak{M}, y \Vdash \varphi.$$

We could use these modalities to show PSPACE upper bounds for some sublogics of **K**, for instance **T** (accessibility relation is reflexive) and **B** (symmetric). For **B** we easily show that a modal formula is satisfiable on a symmetric model iff the same formula is satisfiable in the logic where $\Diamond$ is defined as above.

Remark 4.10 We would also like to be able to account for the PSPACE completeness of **S4** and possibly other logics with a *transitive* accessibility relation. We could add to MLR_2 an operator *transitive*$(\cdot)$ which can only be applied to variables. (Grädel, 1997) shows that $\mathcal{GF}_3$ plus this operator is already undecidable, so we have to be careful. If we do this without care then even guarded MLR_2 becomes EXPTIME-hard. E.g. we can translate the EXPTIME-hard bimodal logic defined by the extra axioms $\langle 1\rangle p \rightarrow \langle 2\rangle p$ and $\langle 2\rangle\langle 2\rangle p \rightarrow \langle 2\rangle p$ to guarded MLR_2 by

$$\begin{array}{lcl} (\langle 2\rangle\varphi)^t & = & \Diamond_1(r_2 \wedge \mathsf{s}_0^1\varphi^t) \\ (\langle 1\rangle\varphi)^t & = & \Diamond_1(r_1 \wedge r_2 \wedge \mathsf{s}_0^1\varphi^t). \end{array}$$

It is not hard to show that φ is satisfiable iff $\varphi^t \wedge \iota\delta \wedge \mathit{transitive}(r_2)$ is cubic satisfiable.

Is there an elegant way of adding this new operator to a fragment of MLR_2 such that the result is still in PSPACE? The resulting fragment should be expressive enough to account for the PSPACE upper bounds of at least multi- **K4**, **S4** and **S5**. A fragment of guarded MLR_2 which is expressive enough to capture these examples is to allow only guarded formulas of the form

$$\Diamond_i(\bigwedge G \wedge \mathsf{s}_j^i\varphi),$$

where $i \neq j$ and G is a subset of $\{p, \otimes p, \mathsf{s}_1^0 p, \mathsf{s}_0^1 p\}$ containing at least one of $\{p, \otimes p\}$, for any variable p. We conjecture that the extension of this fragment with transitivity statements is in PSPACE when interpreted on square models. (The standard **S4** tableau technique from (Ladner, 1977) seems to work fine.) The crucial point is that we can not express interaction between the different modalities.

4.1.2 Pspace in arrow logic

The proof technique with mosaics having shrinking labels used above can also be applied to show that the arrow logics in Table 5 which are PSPACE-hard are also PSPACE-complete.

Theorem 4.11 *All arrow logics in Table 5 labelled by* P *have a* PSPACE-*complete satisfaction problem.*

PROOF. We have shown hardness already. For the upper bound we sketch how to adjust the proof given for the local square logic above. First we define a suitable notion of degree. Inductively we define the **composition degree** $d(\varphi)$ of a formula φ as follows:

$$\begin{array}{lcl}
d(\iota\delta) & = & d(p) = d(\top) = d(\top \circ \iota\delta) = d(\iota\delta \circ \top) = 0 \\
d(\neg\varphi) & = & d(\varphi) \\
d(\otimes\varphi) & = & d(\varphi) \\
d(\varphi \wedge \psi) & = & max(d(\varphi), d(\psi)) \\
d(\varphi \circ \psi) & = & max(d(\varphi), d(\psi)) + 1 \text{ whenever } \varphi \circ \psi \text{ is not } \iota\delta \circ \top \text{ or } \top \circ \iota\delta.
\end{array}$$

We adjust the definition of a closure set also a bit: for ξ a formula, $For(\xi)$ is the smallest set of formulas containing the subformulas of ξ plus single negations plus

$$\begin{array}{ll}
\{\otimes\top, \top\} & \text{if } \otimes \text{ is in the similarity type} \\
\{\iota\delta, \top, \iota\delta \circ \top, \top \circ \iota\delta\} & \text{if } \iota\delta \text{ is in the similarity type.}
\end{array}$$

From now on let a formula ξ be fixed. We define mosaics as follows: A **mosaic** is a structure $\mu = (V, \ell_r, \ell, s)$ where

$V \subseteq U \times U$, for U a set such that $1 \leq |U| \leq 3$
$\ell_r : V \longrightarrow Pow(For(\xi))$
ℓ is a function choosing for every $s \in V$ a subset of $\ell_r(s)$
$s \in V$ or $s = \emptyset$.

A mosaic (V, ℓ_r, ℓ, s) is **smart** if for any $t \in V$ which has at least one coordinate not occurring in s, the composition degree of $\ell_r(t)$ is strictly less than that of $\ell_r(s)$.

A mosaic (V, ℓ_r, ℓ, s) is **coherent** if it satisfies, for all $(x, y) \in V$:

AC0	The usual boolean conditions and the one for $\iota\delta$		
AC1	$\iota\delta \circ \top \in \ell(x, y)$	$\Rightarrow$	$(x, x) \in V$
AC2	$\top \circ \iota\delta \in \ell(x, y)$	$\Rightarrow$	$(y, y) \in V$
AC3	$\otimes\varphi \in \ell(x, y)$	$\Rightarrow$	$(y, x) \in V$ and $\varphi \in \ell(y, x)$
AC4	$\neg\otimes\varphi \in \ell(x, y)$ and $(y, x) \in V$	$\Rightarrow$	$\neg\varphi \in \ell(y, x)$
AC5	$\neg(\varphi \circ \psi) \in \ell(x, y)$ and $\exists z : (\{(x, z), (z, y)\} \subseteq V)$	$\Rightarrow$	$\neg\varphi \in \ell(x, z)$ or $\neg\psi \in \ell(z, y)$.

Now we can define the game:

ξ-game. **Start.** $\exists$ plays a mosaic (V, ℓ_r, ℓ, s) with $\xi \in \ell(t)$ for some $t \in V$.

$\forall-$**move.** $\forall$ wins if the last mosaic (V, ℓ_r, ℓ, s) played by $\exists$ is not coherent and smart. Otherwise he chooses a $t \in V \setminus \mathcal{C}(s)$ and a formula of the form $\varphi \circ \psi \in \ell(t)$ such that in V there are no witnesses for decomposition of t into a φ and a ψ-arrow.

$\exists$-**move.** $\exists$ responds with a mosaic (V', ℓ'_r, ℓ', t) where
$(\mathcal{C}(t) \cap V) \subseteq V'$ and the restrictions of ℓ' and ℓ'_r to $(\mathcal{C}(t) \cap V)$ coincide with ℓ and ℓ_r, resp.
$\{(t_0, z), (z, t_1)\} \subseteq V'$ for some z not in the base of V
$\varphi \in \ell'(t_0, z)$ and $\psi \in \ell'(z, t_1)$.
$\exists$ wins if $\forall$ cannot move anymore. Otherwise $\forall$ wins.

Claim 1 ξ-game stops after at most $d(\xi)$ moves.

PROOF OF CLAIM. As is Proposition 4.2. ◀

Claim 2 If $\exists$ has a winning strategy in ξ-game, then ξ is satisfiable in a finite model whose size is exponential in ξ.

PROOF OF CLAIM. As in Proposition 4.3. ◀

The converse of the last claim is not always true! It holds only for the specific similarity types and model classes mentioned in the Theorem.

Claim 3 (i) Let ξ be in the language containing only $\circ$ and $\otimes$. If ξ is satisfiable (on a symmetric model), then $\exists$ has a winning strategy for ξ-game (where $\otimes\top$ is in the label of every mosaic $\exists$ plays).
(ii) Let ξ be in the language containing only $\circ$ and $\iota\delta$. If ξ is satisfiable (on a reflexive model), then $\exists$ has a winning strategy for ξ-game (where $\iota\delta \circ \top$ and $\top \circ \iota\delta$ are in the label of every mosaic $\exists$ plays).

PROOF OF CLAIM. We prove (i), show what goes wrong with the full language, and then indicate how to adapt the proof to show (ii). So let a $\{\circ, \otimes\}$-formula ξ be satisfiable. Using the representation theorem for this reduct in (Marx, 1995) (Section 4.3.1), we may assume that ξ is satisfied in a model $\mathfrak{M} = (W, \mathsf{v})$ where W is an *irreflexive* relation (this can also be shown by a simple unravelling argument). Let $\mathfrak{M}, (a, b) \Vdash \xi$. Then $\exists$ starts with playing the mosaic $({}^2\{a, b\} \cap W, \ell_r, \ell, \emptyset)$ where $\ell_r(s) = For(\xi)$ for all s, and $\ell(s) = \{\varphi \in For(\xi) \mid \mathfrak{M}, s \Vdash \varphi\}$. Clearly $\exists$ doesn't loose immediately.

We now prove by induction that $\exists$ can respond to any move of $\forall$ with a coherent and smart mosaic which she "cuts out" of the model. So let (V, ℓ_r, ℓ, s) be the mosaic played in the k-th round. Let $\forall$ choose a $t \in V$ which is not s nor the converse of s and a formula $\varphi \circ \psi \in \ell(t)$ for which there are no φ and ψ "witnesses" in V. $\exists$ looks at the model and plays the mosaic (V', ℓ'_r, ℓ', t) where

$V' = {}^2\{t_0, t_1, z\} \cap W$ for some z not in the base of V

$\mathfrak{M}, (t_0, z) \Vdash \varphi$ and $\mathfrak{M}, (z, t_1) \Vdash \psi$

$\ell'_r \restriction_{V' \cap V} = \ell_r \restriction_{V' \cap V}$ and $\ell' \restriction_{V' \cap V} = \ell \restriction_{V' \cap V}$

for $s \in V' \setminus V$, $\ell'_r(s) = \{\psi \in For(\xi) \mid d(\psi) < d(\ell_r(t))\}$ and $\ell'(s) = \ell'_r(s) \cap \{\psi \mid \mathfrak{M}, s \Vdash \psi\}$.

By definition the mosaic she plays is smart, and it satisfies the conditions of the game. Coherence is easy to check, precisely because we assumed that the model had an irreflexive universe. So she can always respond to the $\forall$-moves. *This proves (i).*

Now we describe a situation in which the requirements of coherence and smartness might clash:

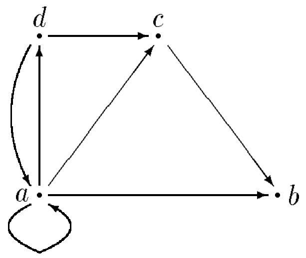

Suppose the first mosaic played was the one with base $\{a, b, c\}$, and the special sequence was (a, b). Also suppose that $d(\ell(a, b)) = d(\ell(a, a)) = 2$. Then by smartness $d(\ell(a, c)) < 2$. Suppose then the mosaic with base (a, d, c) was played because $\ell(a, c)$ had some defect of the form $\varphi \circ \psi$. Then, by smartness, $d(\ell(a, d)) < d(\ell(a, c)) < 2$. But now suppose that $\neg((p \circ q) \circ \top) \in \ell(a, a)$. Then by coherence we must have that $\neg(p \circ q) \in \ell(a, d)$, *because* (a, a) *can be decomposed into* (a, d) *and* (d, a). But that is not possible because the degree of (a, d) is too small.

The proof of (ii) is similar to the one above. What we need is to ensure that ξ is satisfied in a model where we do not have situations as in the picture above. Again by the representation of this reduct in (Marx, 1995) (Section 4.3.1) we can be sure of this. In Figure 3 we indicate how to unravel a model having such situations.

◀

Now we have all the ingredients to argue that the satisfaction problem of the logics under consideration are in PSPACE, just as in Proposition 4.5. QED

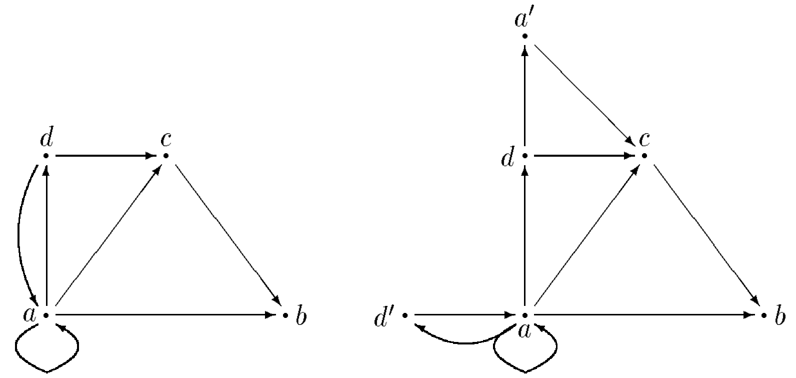

Figure 3: Unravelling the unwanted situation.

4.2 Squares are nexptime complete

Theorem 4.12 *Let σ be any similarity type between $\{\Diamond_0, \Diamond_1\}$ and $\{\Diamond_0, \Diamond_1, \iota\delta, \otimes, \mathsf{s}_1^0, \mathsf{s}_0^1, [k]\}_{k\in\omega}$. Then the satisfaction problem for σ-formulas over the class of square models is* NEXPTIME*-complete.*

Because on square models we can term-define the universal modality as $\Diamond_0\Diamond_1$, we obtain by Theorem 1.8 the following

Corollary 4.13 *Let σ be as in the above Theorem. Then the valid rule/global consequence and the valid sequent problem for σ formulas interpreted on square models is* CO-NEXPTIME*-complete.*

We also obtain the following "axiomatic" result. (It seems unlikely that this was not known but we haven't found it in the literature.)

Corollary 4.14 *The modal logic with two commuting S5 operators has a* NEXPTIME*-complete satisfaction problem.*

PROOF. Because this logic axiomatises the $\{\Diamond_0, \Diamond_1\}$-reduct on the squares, cf. (Marx and Venema, 1997): Theorem 2.2.3. QED

PROOF OF THEOREM 4.12. The NEXPTIME-upper bound for the logic containing the operators $\{\Diamond_0, \Diamond_1, \iota\delta_{01}, \otimes\}$ was shown in (Marx, 1997). An equivalent result in first order logic with two variables is in (Grädel et al., to appear). They also show that every formula φ can be satisfied on a model whose size is exponential in the length of φ. The operators s_j^i are term-definable using $\iota\delta$ and the cylindric diamonds, so they do not cause extra complexity.

The counting operators create extra expressive power. But with $\iota\delta$ and the diamonds we can term-define the difference operator as, (cf. (Venema, 1992))

$$\mathsf{D}\varphi = \Diamond_0(\iota\delta \wedge \Diamond_1(\neg\iota\delta \wedge \Diamond_0\varphi)) \vee \Diamond_1(\iota\delta \wedge \Diamond_0(\neg\iota\delta \wedge \Diamond_1\varphi)).$$

So we have the possibility to count to two using the difference operator. But then –by remark 1.10– adding further counting power doesn't add to the complexity. *This concludes the upper bound part.*

Lower bound. We prove the lower bound for the smallest similarity type. If we have either $\iota\delta$ or $\otimes$ in the language too, the result would follow from (Lewis, 1980), Theorem 4.1: he provides a 2-variable first order formula with only monadic predicates which describes the run of a NEXPTIME Turing machine. We will give a simple proof for the similarity type $\{\Diamond_0, \Diamond_1\}$ using the tiling problem instead of a direct encoding of Turing-machines. We use the following NEXPTIME-complete problem (cf. section 2):

$k{\times}k$ – tiling : given a finite set of tiles $\mathfrak{T}$ and k given in *binary*, can $\mathfrak{T}$ tile the $k{\times}k$ square?

Given an instance $(k, \mathfrak{T})$ of the tiling problem, we will provide a formula in the language with only $\{\Diamond_0, \Diamond_1\}$ which is satisfiable iff $(k, \mathfrak{T})$ is a yes-instance. Without loss we may assume that $k = 2^n$ for some n. From now on let an instance $(2^n, \mathfrak{T})$, where $\mathfrak{T} = \{T_1, \ldots, T_t\}$ be fixed.

We will now define the formula which describes this instance. The "difficult" part of the proof is of course to make sure that we have a PTIME-reduction, i.e., that our formula doesn't get too long.

We first describe a binary tree of depth $2n$, using $2n$ propositional variables $p_1, \ldots p_{2n}$. This will provide us with $(2^n)^2$ leaves each encoding an element in the $2^n{\times}2^n$ grid. Our formula is similar to the one given in (Halpern and Moses, 1992): Proposition 6.5.

Let $\Diamond\varphi$ be an abbreviation for $\Diamond\!\!\!\mid(r \wedge \Diamond\!\!\!\!-(p \wedge \varphi))$, where r and p are two new propositional variables. We use $\Diamond$ as an ordinary **K**-modality.

$$
\begin{aligned}
p \wedge \Box^{2n}\Box\bot \quad \textstyle\bigwedge_{k<2n} \Box^k[\quad & (\Diamond p_{k+1} \wedge \Diamond\neg p_{k+1}) \ \wedge \\
& \bigwedge_{i \leq k} ((p_i \rightarrow \Box p_i) \wedge (\neg p_i \rightarrow \Box\neg p_i))\,]. \qquad (25)
\end{aligned}
$$

Any square model which forces (25) will have $2^{2n} = (2^n)^2$ different pairs –the leaves of the binary tree. These pairs code up elements in the grid $2^n \times 2^n$: the valuation of the first n p_i provides the x-coordinate, and of the next n p_i the y-coordinate. We use the convention as in the following example where $n = 3$: the formula $p_1 \wedge \neg p_2 \wedge p_3$ codes the binary number 101 which equals 5. So the formula $p_1 \wedge \neg p_2 \wedge p_3 \wedge \neg p_4 \wedge \neg p_5 \wedge p_6$ codes the element $(5, 1)$ in the grid to be tiled.

We can recognize the leaves because they force $p \wedge \Box\bot$.

Our final formula will have the form $(25) \wedge \Box_0\Box_1\psi$, where ψ is a conjunction of formulas. $\Box_0\Box_1$ functions as the universal modality, so ψ wants to be true in the whole model.

We have variables $t_1, \ldots, t_t$, one for each tile. The next two formulas say that on any leaf, precisely one tile variable t holds.

$$p \wedge \Box\bot \quad \leftrightarrow \quad \bigvee_{1\leq i\leq t} t_i \tag{26}$$

$$\bigwedge_{1\leq i\leq t} (t_i \quad \rightarrow \quad \bigwedge_{j\neq i} \neg t_j). \tag{27}$$

The hard part is to ensure that adjacent colors match without writing too many formulas. Here we use the two-dimensional nature of our models plus the fact that they are square.

We introduce $4n$ new variables for counting $p^h_1, \ldots p^h_{2n}$, and $p^v_1, \ldots p^v_{2n}$ and $2t$ new tile-variables t^h_i and t^v_i. We use them to compare two leaves in one state of our model. We use *grid* as an abbreviation of $\Diamond_v(p\wedge\Box\bot)\wedge\Diamond_h(p\wedge\Box\bot)$, so *grid* holds at a pair iff it is "in the intersection of two leaves", as x in the picture below.

First we write formulas ensuring inheriting of information:

$$(grid \wedge \Diamond_v(p \wedge \Box\bot \wedge p_i)) \leftrightarrow p^v_i \tag{28}$$

$$(grid \wedge \Diamond_h(p \wedge \Box\bot \wedge p_i)) \leftrightarrow p^h_i \tag{29}$$

$$(grid \wedge \Diamond_v(p \wedge \Box\bot \wedge t_i)) \leftrightarrow t^v_i \tag{30}$$

$$(grid \wedge \Diamond_h(p \wedge \Box\bot \wedge t_i)) \leftrightarrow t^h_i. \tag{31}$$

So if in a model we have the following situation:

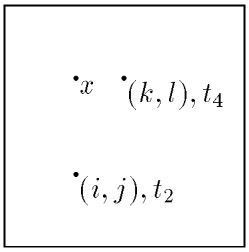

then at x we have $(i^v, j^v), (k^h, l^h), t^v_2$ and t^h_4.

Now we can express that colors match:

$$grid \wedge x_h = x_v \wedge y_h = y_v + 1 \wedge t^v_i \quad \rightarrow \quad \bigvee\{t^h_j \mid \mathrm{top}(T_i) = \mathrm{bottom}(T_j)\}. \tag{32}$$

$$grid \wedge y_h = y_v \wedge x_h = x_v + 1 \wedge t^v_i \quad \rightarrow \quad \bigvee\{t^h_j \mid \mathrm{left}(T_i) = \mathrm{right}(T_j)\}. \tag{33}$$

Here we use the following abbreviations:

$x_h = x_v$	abbreviates	$\bigwedge_{i\leq n}(p^h_i \leftrightarrow p^v_i)$
$x_h = x_v + 1$	abbreviates	$\bigvee_{i\leq n} [\bigwedge_{l<i}(p^h_l \leftrightarrow p^v_l) \wedge p^h_i \wedge \neg p^v_i \wedge \bigwedge_{i<l\leq n}(\neg p^h_l \wedge p^v_l)]$,

and similar for the y coordinates where we use the p_i's between p_{n+1} and p_{2n}.

Finally let $\varphi_{(2^n,\mathfrak{T})}$ be $25 \wedge \boxminus \boxbar(26 \wedge 27 \wedge \bigwedge_{1\leq i\leq 2n}(28 \wedge 29) \wedge \bigwedge_{1\leq i\leq t}(30 \wedge 31 \wedge 32 \wedge 33))$.

Clearly this formula is PTIME-computable in the length of the instance $(2^n, \mathfrak{T})$. That the reduction is correct is easy to show and left to the reader.

QED

4.3 Expansion with coordinate wise counting

If we expand MLR_2 with a vertical difference operator d_1 we obtain infinity axioms. The following formula describes an injective function with a start point and can only be true at a model $({}^2U, \mathsf{v})$ if U is infinite:

$$\iota\delta \wedge \boxminus \neg f \wedge \boxbar \boxminus(\iota\delta \to f) \wedge \boxbar \boxminus(f \to \neg \mathsf{d}_1 f) \wedge \boxbar \boxminus(\otimes f \to \neg \mathsf{d}_1 \otimes f). \qquad (34)$$

In section 2 we saw that adding d_1 to arrow logic or to MLR_n $(n \geq 3)$ led to undecidability even when interpreted on classes of relativised frames. Here we will show that with MLR_2, this expansion does not lead to undecidability. In fact we will show that the expansion with all coordinate wise counting modalities $[k]_i$ is still decidable when interpreted on the classical square models. This result was first obtained, in a first order logic setting, in (Grädel et al., 1996). Both that proof and the one we provide here only establish that the set of validities is recursive. Recently, (Pacholski, 1997) announced a NEXPTIME upper bound for the satisfaction problem. Our proof uses again mosaics and is an adaptation of the NEXPTIME upper bound proof for MLR_2 in (Marx, 1997). So we see that the mosaic-technique is not limited to relativised cases, but can also be applied in classical cases, even when the logic has infinity axioms. It remains an open problem whether we can fine-tune this mosaic-proof in order to get a better upper-bound.

Theorem 4.15 *The expansion of MLR_2 with the coordinate wise counting modalities $\{[k]_i \mid k < \omega,\ i \in \{0,1\}\}$ interpreted on square models is decidable.*

Instead of giving a proof which concerns the whole language we consider an expansion with "vertical nominals", i.e., special variables which can only be true at one place in a column. These nominals then act as unary functions symbols:

$$\text{if } (x,y) \Vdash f \text{ and } (x,z) \Vdash f \text{ then } y = z.$$

From Remark 1.10 it follows that decidability of this expansion leads to decidability of the expansion with all vertical counting quantifiers. The horizontal ones can be term-defined using $\otimes$, by $[k]_0\varphi = \otimes[k]_1\otimes\varphi$. So we will concentrate on proving

Theorem 4.16 *The validity problem of the logic MLR_2 expanded with unary function symbols interpreted on square models is decidable.*

We start with the usual mosaic preliminaries and show that satisfiability is equivalent to the existence of a "linked set of mosaics" (an LSM). As before the theorem can then be easily proved.

Preliminaries. Let ξ be a fixed $CMML_2$-formula without occurrences of $\Diamond\!\!\!\!-$. As explained in Remark 4.1 it is sufficient to decide this reduct. Let Fun be the set containing all function symbols occurring in ξ. We will use f as a metavariable for these function symbols. $For(\xi)$ is defined as follows:

$$\begin{array}{lcl} For_0(\xi) & = & \{\iota\delta\} \cup \{\text{subformulas of } \xi\} \cup \{\Diamond\!\!\!\mid f \mid f \in F\} \\ For_1(\xi) & = & For_0(\xi) \cup \bigcup_{\Delta \subseteq Fun} \{\Diamond\!\!\!\mid(\bigwedge\{f \in \Delta\} \wedge \psi) \mid \Diamond\!\!\!\mid \psi \in For_0(\xi)\} \\ For_2(\xi) & = & For_1(\xi) \text{ plus subformulas} \\ For(\xi) & = & For_2(\xi) \cup \{\otimes\psi \mid \psi \in For_2(\xi) \text{ and } \psi \not\approx \otimes\varphi\}. \end{array}$$

Clearly $For(\xi)$ is finite. Let $L_{For(\xi)}$ be the first order language with binary predicates $\overline{\psi}$ for every $\psi \in For(\xi)$, and two unary predicates F and B. An element x is called *fixed*, if Fx, and *bounded* if Bx holds. We will use the following abbreviations:

$$\begin{array}{lcl} Nofunc(x,y) & =:: & \bigwedge_{f \in For(\xi)} \neg\overline{f}(x,y) \\ Free(x) & =:: & \neg Fx \wedge \neg Bx \\ (x,y) \cong (x',y') & =:: & \bigwedge_{\psi \in For(\xi)} (\overline{\psi}(x,y) \leftrightarrow \overline{\psi}(x',y')) \\ x \simeq x' & =:: & (x,x) \cong (x',x') \wedge (Fx \leftrightarrow Fx') \wedge (Bx \leftrightarrow Bx') \wedge \\ & & \forall z(Fz \rightarrow ((x,z) \cong (x',z))). \end{array}$$

If $a \simeq b$ holds for some a, b, we say that a and b are of the same type.

An $L_{For(\xi)}$ structure $\mathfrak{K}$ is called a *coherent* **linked set of mosaics** (short: LSM) if $\mathfrak{K}$ satisfies the following conditions:

$$\begin{array}{lll} B & \forall x(\neg Bx \rightarrow \forall z(Fz \rightarrow Nofunc(z,x))) & \\ C^{\neg} & \forall xy(\overline{\neg\psi}(x,y) \leftrightarrow \neg\overline{\psi}(x,y)) & \text{for } \neg\psi \in For(\xi); \\ C^{\wedge} & \forall xy(\overline{\varphi \wedge \psi}(x,y) \leftrightarrow \overline{\varphi}(x,y) \wedge \overline{\psi}(x,y)) & \text{for } \varphi \wedge \psi \in For(\xi); \\ C^{\iota\delta} & \forall xy(\overline{\iota\delta}(x,y) \leftrightarrow x = y); & \\ C^{\otimes} & \forall xy(\overline{\otimes\psi}(x,y) \leftrightarrow \overline{\psi}(y,x)) & \text{for } \otimes\psi \in For(\xi); \\ C^{\Diamond\!\!\!\mid\, 1} & \forall xy(\overline{\psi}(x,y) \rightarrow \overline{\Diamond\!\!\!\mid \psi}(x,y)) & \text{for } \Diamond\!\!\!\mid \psi \in For(\xi); \\ C^{\Diamond\!\!\!\mid\, 2} & \forall xyz(\overline{\Diamond\!\!\!\mid \psi}(x,y) \leftrightarrow \overline{\Diamond\!\!\!\mid \psi}(x,z)) & \text{for } \Diamond\!\!\!\mid \psi \in For(\xi); \\ C^{f1} & \forall xyz((\overline{f}(x,y) \wedge \overline{f}(x,z)) \rightarrow y = z) & \text{for } f \in For(\xi); \\ C^{f2} & \forall x \overline{\Diamond\!\!\!\mid f}(x,x) & \text{for } f \in For(\xi); \\ C^{f3} & \forall xy((\overline{f}(x,y) \wedge \overline{\Diamond\!\!\!\mid(f \wedge \psi)}(x,y)) \rightarrow \overline{\psi}(x,y) & \text{for } f, \Diamond\!\!\!\mid \psi \in For(\xi). \end{array}$$

An $L_{For(\xi)}$ structure $\mathfrak{K}$ is called a *saturated* LSM if $\mathfrak{K}$ satisfies the following conditions:

$$\begin{array}{llcl} \textbf{Wit} & \forall xy(\neg Fx \wedge \neg Fy & \rightarrow & \exists x'y'(x' \neq y' \wedge (x,x) \cong (x',x') \wedge (y,y) \cong (y',y') \\ & & & \wedge Nofunc(x',y') \wedge Nofunc(y',x'))) \\ \textbf{Sat}^{\Diamond\!\!\!\mid} & \forall x(Free(x) \wedge \overline{\Diamond\!\!\!\mid \psi}(x,x) & \rightarrow & \exists z(Fz \wedge \overline{\psi}(x,z)) \vee \\ & & & \exists x'y(x \simeq x' \wedge \overline{\psi}(x',y) \wedge Free(y))) \\ \textbf{Sat}^{\Diamond\!\!\!\mid\, 2} & \forall x((\neg Free(x) \wedge \overline{\Diamond\!\!\!\mid \psi}(x,x)) & \rightarrow & \exists y(\overline{\psi}(x,y))). \end{array}$$

A $L_{For(\xi)}$ structure $\mathfrak{K}$ is called a ξ-LSM if $\mathfrak{K}$ is finite, coherent, saturated, and $\mathfrak{K} \models \exists xy\overline{\xi}(x,y)$.

Remark 4.17 The intuitive ideas behind mosaics are rather hidden in the above definition. The mosaics in an LSM $\mathfrak{K}$ are –up to isomorphism– all the substructures of $\mathfrak{K}$ with a domain of just 2 elements. Instead of a labelling function we use the first order valuation and the "predicates as codes"-idea. The coherence conditions for Φ can be stated in this way since it is an S5 modality. Saturation works as follows. Suppose in some mosaic with domain $\{x,z\}$, the formula $\overline{\Phi\,\psi}(x,z)$ holds. Then by $C^{\Phi\,2}$, also $\overline{\Phi\,\psi}(x,x)$, whence by **Sat**$^{\Phi(2)}$, either there is a witness in the structure, or there is a mosaic (with domain $\{x',y\}$) which has a witness (the part $\overline{\psi}(x',y)$) and which can be coherently glued to the $\{x,z\}$-mosaic (the rest of the consequent). The next proposition explains **Sat**$^{\Phi}$ more precisely.

The notion of F (for fixed) elements is needed because the MLR_2 language is strong enough to express that a state is unique (using the definable difference operator). In the construction of a model from an LSM we will ensure that we do not copy those fixed elements. The notion of *bounded* elements is needed to ensure that nominals are not copied in one column.

Proposition 4.18 *Let an $L_{For(\xi)}$-structure $\mathfrak{K}$ be a saturated LSM. Suppose $\overline{\Phi\,\psi}(x,x) \wedge Free(x)$ holds. Then there exists a partial isomorphism $f : K \longrightarrow K$ such that*

- *$fz = z$ for all z such that Fz;*
- *$fx' = x$, for some (x',y) such that $\overline{\psi}(x',y)$ holds;*
- *The domain of f consists of $\{x'\} \cup \{z \in K \mid Fz\}$;*
- *either f is the identity (and Fy holds (i.e., y is fixed)) or $Free(y)$ holds.*

Proof. Obvious, since $(x,y) \cong (x',y')$ means that (x,y) and (x',y') satisfy the same $L_{For(\xi)}$ relation symbols. Because $For(\xi)$ is closed under single $\otimes$, $(x,z) \simeq (x',z)$ implies $(z,x) \simeq (z,x')$. QED

The last proposition explains the meaning of the **Sat**$^{\Phi}$- conditions: they provide us with witnesses for diamond-type formulas. The condition **Wit** has a different flavour: it provides us, for $\neg F$-elements x, y, with a set of formulas which can be coherently made true at the pair (x,y). Checking coherence is greatly facilitated by the fact that this set does not contain any function symbols, nor does its converse. On any coherent LSM, satisfying

Wit we can define a function *Joker* from pairs of non-fixed points to subsets of $For(\xi)$, such that

> $Joker(x,y) = \{\overline{\psi} \mid \mathfrak{K} \models \overline{\psi}(x',y')\}$ for some $x',y' \in K$ such that $x' \neq y'$, $(x,x) \cong (x',x')$, $(y,y) \cong (y',y')$ & $Nofunc(x',y')$ & $Nofunc(y',x')$.

In general we can define several of these functions. We just fix one which has the property that if $Joker(x,y)$ is defined from the pair (x',y'), then $Joker(y,x)$ is defined from (y',x'). In particular these functions obey the following constraints

$\overline{\iota\delta} \notin Joker(x,y)$
$Joker(x,y)$ does not contain any f or $\otimes f$ symbols
$\overline{\otimes\psi} \in Joker(x,y) \iff \overline{\psi} \in Joker(y,x)$.
$\overline{\Phi\,\psi} \in Joker(x,y) \iff \mathfrak{K} \models \overline{\Phi\,\psi}(x,x)$.

These sets function as a kind of "jokers" when we extend structures: because they do not contain function symbols it is easy to show that extensions with them are still coherent. We wil have to extend $For(\xi)$-structures several times. We now define how we do that.

Let $\mathfrak{K}$ be a finite $For(\xi)$-structure, x,y free elements in K and let $\Sigma \subseteq For(\xi)$. We define $extension(\mathfrak{K},y,x,\Sigma)$ as the structure $\mathfrak{K}'$ with domain $K \cup \{y'\}$, for some $y' \notin K$ and the following valuation:

$\mathfrak{K}' \models \overline{\psi}(u,v)$	iff	$\mathfrak{K} \models \overline{\psi}(u,v)$	for all $u,v \in K$
$\mathfrak{K}' \models \neg Fy' \wedge \neg By'$			
$\mathfrak{K}' \models \overline{\psi}(y',y')$	iff	$\mathfrak{K} \models \overline{\psi}(y,y)$	
$\mathfrak{K}' \models \overline{\psi}(z,y')$	iff	$\mathfrak{K} \models \overline{\psi}(z,y)$	for every fixed element z
$\mathfrak{K}' \models \overline{\psi}(y',z)$	iff	$\mathfrak{K} \models \overline{\psi}(y,z)$	for every fixed element z
$\mathfrak{K}' \models \overline{\psi}(x,y')$	iff	$\psi \in \Sigma$	
$\mathfrak{K}' \models \overline{\psi}(y',x)$	iff	$\psi \in \otimes\Sigma$.	

Here $\otimes\Sigma$ abbreviates $\{\otimes\psi \in For(\xi) \mid \psi \in \Sigma\} \cup \{\psi \in For(\xi) \mid \otimes\psi \in \Sigma\}$. Note that because $For(\xi)$ is closed under single $\otimes$, for any coherent $\mathfrak{K}$ and $x,y \in K$, $\otimes\{\overline{\psi} \mid \mathfrak{K} \models \overline{\psi}(x,y)\} = \{\overline{\psi} \mid \mathfrak{K} \models \overline{\psi}(y,x)\}$. Figure 4 shows which pairs in the extension obtained a valuation until now (the ones indicated by a small circle). All other elements obtain a valuation using our jokers.

$$\mathfrak{K}' \models \overline{\psi}(w,y') \text{ iff } \overline{\psi} \in Joker(w,y).$$

Define the set P of *potential defects* as $\omega \times \{\Phi\,\psi \in For(\xi)\}$. For $\mathfrak{K}$ a coherent LSM and $(n,\Phi\,\psi)$ a potential defect, we say that $(n,\Phi\,\psi)$ is a *defect of* $\mathfrak{K}$ if $n \in K$, $\mathfrak{K} \models \overline{\Phi\,\psi}(n,n)$, but there is no $y \in K$ such that $\mathfrak{K} \models \overline{\psi}(n,y)$.

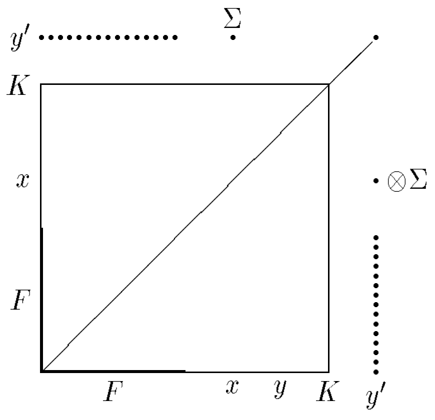

Figure 4: Extension of a $For(\xi)$-structure.

Lemma 4.19 (Extension Lemma) *Let $\mathfrak{K}$ be a coherent finite LSM satisfying **Wit** with $x \in K$ and x is not fixed.*
(i) If $(x, \Phi\,\psi)$ is a defect of $\mathfrak{K}$, and $\exists x'y \in K(x \simeq x' \wedge \neg Fx' \wedge \overline{\psi}(x', y) \wedge Free(y))$, then $extension(\mathfrak{K}, y, x, \{\theta \in For(\xi) \mid \mathfrak{K} \models \overline{\theta}(x', y)\})$ is coherent.
(ii) $extension(\mathfrak{K}, x, x, Joker(x, x))$ is coherent.

Proof. (i) All the coherence conditions, except C^{f1} follow easily from the fact that $\mathfrak{K}$ was coherent (to show that $\mathfrak{K} \models \overline{\neg\iota\delta}(x', y)$, we use that $\Phi\,\psi$ is a defect of x, whence it cannot be that $\mathfrak{K} \models \overline{\psi}(x, x)$). Checking C^{f1} has 3 cases: Let $w \in K'$.

Case 1. $w \neq x$ and $w \neq y'$. If Fw, then $Nofunc(w, y)$ since y is free and coherence condition B, so also $Nofunc(w, y')$. If w is not fixed, then $Joker(w, y)$ does not contain function symbols either.

Case 2. $w = x$. Here we use our extended closure set $For(\xi)$. Suppose $\overline{f}(x, y')$ for some function-symbol f and also $\overline{f}(x, v)$ for some $v \neq y'$. Then we have $\overline{f}(x', y)$. But then also –by $\overline{\psi}(x', y)$ and $C^{\Phi 1}$ and the definition of $For(\xi)$– $\overline{\Phi(f \wedge \psi)}(x', y)$, so also $\overline{\Phi(f \wedge \psi)}(x', x')$ and $\overline{\Phi(f \wedge \psi)}(x, x)$. But then also $\overline{\Phi(f \wedge \psi)}(x, v)$, so by C^{f3}, $\overline{\psi}(x, v)$, which is a contradiction because we assumed that $(x, \Phi\,\psi)$ was a defect.

Case 3. $w = y'$. This is easy, all pairs (y', v) where $v \neq x$ and not Fv don't force function symbols. The valuation of the other elements is a copy of the valuation in $\mathfrak{K}$ of the y-column, which was coherent by assumption.

The proof of (ii) follows the same argument. QED

Lemma 4.20 *Every ξ-LSM can be turned into a model satisfying ξ.*

PROOF. We first show that any defect of a coherent and saturated LSM can be solved.

Claim 1 Let $\mathfrak{K}$ be a ξ-LSM and $(n, \Phi\,\psi)$ a defect of $\mathfrak{K}$. Then there exists a ξ-LSM $\mathfrak{K}' \supseteq \mathfrak{K}$, with size $|K'| = |K| + 1$ in which $(n, \Phi\,\psi)$ is no defect.

PROOF OF CLAIM. Let $\mathfrak{K}$ be a ξ-LSM and $(n, \Phi\,\psi)$ a defect of $\mathfrak{K}$. Let $K = k \times k$ for some natural number k. By $\mathbf{Sat}^{\Phi\,2}$, we know that *Free*(n), because otherwise $(n, \Phi\,\psi)$ was no defect. Also there is no $x \in K$ such that Fx and $\overline{\psi}(n, x)$. So, by $\mathbf{Sat}^{\Phi}$, we know that there are $n', y \in K$, y is free, $\overline{\psi}(n', y)$ holds, and a partial isomorphism f which sends n' to n and leaves the fixed elements fixed. Let $\mathfrak{K}' = extension(\mathfrak{K}, y, n, \{\theta \in For(\xi) \mid \mathfrak{K} \models \overline{\theta}(n', y)\})$.

We claim that $\mathfrak{K}'$ satisfies our purposes. Clearly we have that $\mathfrak{K}' \supseteq \mathfrak{K}$, $|K'| = |K|+1$, and the defect is solved. $\mathfrak{K}'$ is coherent by lemma 4.19.(i). We show that $\mathfrak{K}'$ is *saturated.* We only have to show that $\mathbf{Sat}^{\Phi}$ holds for (k, k), where k is the added element. So assume $\overline{\Phi\,\psi}(k, k)$. Then also $\overline{\Phi\,\psi}(y, y)$. If $\overline{\psi}(y, z)$ for some fixed z then also $\overline{\psi}(k, z)$ holds. Otherwise, there exists a partial isomorphism f "solving the defect". But then that same f works for k as well, because $y \simeq k$. ◀

From here on the proof follows familiar lines. Let $\mathfrak{M}$ be the (countable) structure obtained by an infinite process of defect-solving starting from the ξ-LSM $\mathfrak{K}$. Clearly, $\mathfrak{M}$ is coherent, it satisfies $\overline{\xi}$ somewhere, and every defect is solved.

Now it is easy to prove a truth-lemma, which finishes the proof of the lemma. QED

Lemma 4.21 *For every satisfiable ξ, there exists a ξ-LSM.*

PROOF. Let $\mathfrak{M}$ be a model satisfying ξ with base U. Let *For*(ξ) be the set obtained from ξ as in the preliminaries. Define

$$\Sigma_{(x,y)} \stackrel{\text{def}}{=} \{\psi \in For(\xi) \mid \mathfrak{M}, (x, y) \Vdash \psi\}.$$

Let C be the set of all "colors" $\Sigma_{(x,y)}$. Define $CC(x)$ as the set of all colors occuring in the column of x: $CC(x) = \{\Sigma_{(x,y)} \mid y \in U\}$. Call an element $z \in U$ *fixed* if $\{y \in U \mid CC(y) = CC(z)\}$ is finite. Define for $x, x' \in U$, $x \simeq_2 x'$ if $CC(x) = CC(x')$ and for every fixed z, $\Sigma_{(z,x)} = \Sigma_{(z,x')}$. Call an element $x \in U$ *bounded* if $\{x' \in U \mid x' \simeq_2 x\}$ is finite. Call an element $x \in U$ *free* if it is not fixed and not bounded.

First a claim involving counting.

$$\text{There are finitely many fixed, and finitely many bounded elements in } U. \tag{35}$$

$$\text{If } x \text{ and } y \text{ are not fixed, there are } x', y' \in U \text{ such that } \Sigma_{(x,x)} = \Sigma_{(x',x')},\ \Sigma_{(y,y)} = \Sigma_{(y',y')}, \tag{36}$$
$$x' \neq y', \text{ and } \Sigma_{(x',y')} \text{ does not contain any } f, \text{ nor } \otimes f \text{ symbols.}$$

$$\text{If } x \text{ is not fixed, there is a } y \in U \text{ such that } \Sigma_{(x,x)} = \Sigma_{(y,y)} \text{ and } y \text{ is free.} \tag{37}$$

$$\text{If } f \in \Sigma_{(z,x)} \text{ for some fixed } z \text{ and some function symbol } f, \text{ then } x \text{ is bounded.} \tag{38}$$

(35) is obvious since there are finitely many colors, whence only finitely many $\simeq_2$ equivalence classes.
For (36), note that there are infinitely many x' such that $\Sigma_{(x,x)} = \Sigma_{(x',x')}$ and similar for y, and just finitely many function symbols in $For(\xi)$.
For (37), if x is not fixed, there are infinitely many x' such that $\Sigma_{(x,x)} = \Sigma_{(x',x')}$. Because there are only finitely many $\simeq_2$ equivalence classes, there must be some infinite $\simeq_2$ equivalence class containing an x' with $\Sigma_{(x,x)} = \Sigma_{(x',x')}$.
If the antecedent of (38) holds, then $\{y \in U \mid y \simeq_2 x\} = \{x\}$, because function symbols can occur only once in a column.

We define a set W as a minimal subset of U satisfying

1. there are $x, y \in W$ such that $\xi \in \Sigma_{(x,y)}$;

2. every fixed or bounded element in U is in W;

3. for any fixed or bounded z in W and for any $\Phi\,\psi \in \Sigma_{(z,z)}$, there is $v \in W$ such that $\psi \in \Sigma_{(z,v)}$;

4. if $x \in W$ is free, $\Phi\,\psi \in \Sigma_{(x,x)}$ and there are $x', y \in U$ such that $x' \simeq_2 x$, and y is free, and $\psi \in \Sigma_{(x',y)}$, then $\{x'', y'\} \subseteq W$, for some $x'', y' \in U$ such that $x'' \simeq_2 x$, y' is free and $\psi \in \Sigma_{(x'',y')}$;

5. if $x, y \in U$ are not fixed, then $\{x', y'\} \subseteq W$, for some $x' \neq y'$ such that x' and y' are both not fixed, and $\Sigma_{(x,x)} = \Sigma_{(x',x')}$, $\Sigma_{(y,y)} = \Sigma_{(y',y')}$ and $\Sigma_{(x',y')}$ and $\Sigma_{(y',x')}$ both do not contain any function-symbol.

6. If $y \in U$ is not fixed, then there is a $y' \in W$ such that $\Sigma_{(y,y)} = \Sigma_{(y',y')}$ and y' is free.

By (36) and (37) such W exist, and $W \neq \emptyset$ because $\mathfrak{M}$ satisfies ξ. We are ready to define the $L_{For(\xi)}$ structure $\mathfrak{K}$, which is almost a ξ-LSM:

$K = W$;
$\mathfrak{K} \models Fx$ iff x is fixed;
$\mathfrak{K} \models Bx$ iff x is bounded;
For $\psi \in For(\xi)$, $\mathfrak{K} \models \overline{\psi}(x, y)$ iff $\mathfrak{M}, (x, y) \Vdash \psi$.

Note that $\mathfrak{K} \models \overline{\psi}(x, y)$ iff $\psi \in \Sigma_{(x,y)}$.

Claim 2 $\mathfrak{K}$ is finite, coherent, it satisfies $\overline{\xi}(x, y)$ and **Wit** and $\mathbf{Sat}^{\Phi\, 2}$ hold.

In other words, the only LSM-condition missing is $\mathbf{Sat}^{\Phi}$. We will ensure later that this holds as well.

PROOF OF CLAIM. By clause 1, $\mathfrak{K} \models \exists xy\overline{\xi}(x, y)$. $\mathfrak{K}$ is finite because of the following. By (35), the first 3 clauses only force finitely many elements in W. Clause 4 only brings in finitely many new elements because there are just finitely many $\simeq_2$ equivalence classes. The last two clauses also only bring in finitely many new elements.

$\mathfrak{K}$ is coherent, because $\mathfrak{M}$ is a model, and (38). By clause 5, **Wit** holds, and by clause 3 $\mathbf{Sat}^{\Phi\, 2}$ holds as well. ◀

It might be that there are free elements x and a formula $\Phi\, \psi$ which holds at (x, x) but there is no ψ-witness which is good in the sense of $\mathbf{Sat}^{\Phi}$. Because K contains all fixed elements of $\mathfrak{M}$ and by (4) in the definition of W, the only case in which this can occur is when for all $x' \simeq_2 x$, ψ holds in $\mathfrak{M}$ at (x', y) for some *bounded* y. This is a problem, because $\mathbf{Sat}^{\Phi}$ wants y to be *free*. For these elements we need to do extra work. Let $x \in K$ be free. We say that $(x, \Phi\, \psi)$ is a **Sat**-*defect* of $\mathfrak{K}$ if $\mathbf{Sat}^{\Phi}$ fails for this x and this formula. Since K is finite, there are just finitely many **Sat**-defects. We now show how to solve every one of them, by creating a finite extension of $\mathfrak{K}$.

So let $(x, \Phi\, \psi)$ be a **Sat**-defect. We claim that there exist $x^*, y \in U$ such that $x^* \simeq_2 x$, $x^* \neq y$, y is bounded but not fixed, $\psi \in \Sigma_{(x^*,y)}$ and for no function symbol $f \in For(\xi)$, $f \in \Sigma_{(y,x^*)}$.

Because $(x, \Phi\, \psi)$ is a **Sat**-defect, for all $x^* \simeq_2 x, \mathfrak{M}, (x^*, x^*) \Vdash \Phi\, \psi$, and for all $y \in U$ such that $\mathfrak{M}, (x^*, y) \Vdash \psi$, $x^* \neq y$ and y is bounded but not fixed. Because x is free, there are infinitely many $x^* \simeq_2 x$. Because there are just finitely many bounded elements and finitely many function symbols in $For(\xi)$, we find such x^*, y such that for all $f \in For(\xi)$, $\mathfrak{M}, (y, x^*) \nVdash f$.

Now we are able to solve the **Sat**-defect $(x, \Phi\, \psi)$. First we create a copy of x: let $\mathfrak{K}' = extension(\mathfrak{K}, x, x, Joker(x, x))$. By (ii) of the extension lemma,

$\mathfrak{K}'$ is finite, coherent, and $\mathfrak{K} \subseteq \mathfrak{K}'$. Let x' be the copy of x in $\mathfrak{K}'$. Then $x \simeq x'$ and the only pairs in the column x' which might validate a function symbol are (x', z) for $z = x'$ or z is fixed. This will provide us with the necessary freedom to solve the **Sat**-defect of x above x'.

Let x^*, y be as in the last claim. By clause 6 in the definition of W, there exists a $y^* \in K$ such that $\Sigma_{(y,y)} = \Sigma_{(y^*,y^*)}$ and y^* is free. Define the extension $\mathfrak{K}_{(x,\Phi\,\psi)}$ in which the **Sat**-defect is solved as *extension*$(\mathfrak{K}', y^*, x', \Sigma_{(x^*,y)})$.

Claim 3 $\mathfrak{K}_{(x,\Phi\,\psi)}$ is finite, coherent, $\mathfrak{K}_{(x,\Phi\,\psi)} \supseteq \mathfrak{K}$ and $\mathfrak{K}_{(x,\Phi\,\psi)}$ has strictly less **Sat**-defects than $\mathfrak{K}$.

Proof of Claim. Clearly, for every $x' \simeq x$, $\mathbf{Sat}^{\Phi}$ holds for $\Phi\,\psi$. So this **Sat**-defect is solved. We didn't create any new ones since in both extensions we only added a element which was $\simeq$ to some element in $\mathfrak{K}$. Obviously $\mathfrak{K}_{(x,\Phi\,\psi)}$ is finite and extends $\mathfrak{K}$. Coherence can be shown similarly as in the proof of the extension lemma, except for C^{f1} in columns x' and y' (where y' is the new copy of y^*).

C^{f1} holds in column x' because of the following. We know that $\Phi\,\theta \in \Sigma_{(x^*,y)} \iff \mathfrak{K}' \models \overline{\Phi\,\theta}(x', x')$. Because we created x' deleting as many functions in its column as possible, we also know that if $\mathfrak{K}' \models \overline{f}(x', z)$, then $z = x'$ or z is fixed. Now suppose that $f \in \Sigma_{(x^*,y)}$ and $\mathfrak{K}' \models \overline{f}(x', z)$, for some $z \neq y'$. Then, by $C^{\Phi 1}$, and the fact that $\Phi(f \wedge \psi) \in For(\xi)$, $\Phi(f \wedge \psi) \in \Sigma_{(x^*,y)}$, whence also $\mathfrak{K}' \models \overline{\Phi(f \wedge \psi)}(x', x')$. But then, by $C^{\Phi 2}$, $\mathfrak{K}' \models \overline{\Phi(f \wedge \psi)}(x', z)$. But then by C^{f3}, $\mathfrak{K}' \models \overline{\psi}(x', z)$. But, because $(x, \Phi\,\psi)$ was a **Sat**-defect, $\mathfrak{K}' \not\models \overline{\psi}(x', z)$, for all fixed z and for $z = x'$. A contradiction.

All colors in the column of y' are without function symbols, except possibly the ones at pairs (y', y') or (y', z) for fixed z precisely because $\otimes\Sigma_{(x^*,y)} = \Sigma_{(y,x^*)}$ does not contain any function symbols. So then C^{f1} holds because $y' \simeq y^*$. ◀

Finally, let $\mathfrak{K}^{\#}$ be the result of extending $\mathfrak{K}$ in this way for every **Sat**-defect of $\mathfrak{K}$.

Claim 4 $\mathfrak{K}^{\#}$ is a ξ-LSM.

Proof of Claim. Clearly $\mathfrak{K}^{\#}$ is finite, coherent, and $\mathfrak{K}$ is a substructure of $\mathfrak{K}^{\#}$. So it is left to show that $\mathbf{Sat}^{\Phi}$ holds for every free element and every $\Phi\,\psi \in For(\xi)$. For the elements in $\mathfrak{K}$ this is immediate by the previous claim and the fact that we have solved all **Sat**-defects of $\mathfrak{K}$. But in every extension we have only added elements which are free and of the same type as some free element in K. But then $\mathbf{Sat}^{\Phi}$ holds for these new elements as well. ◀

This finishes the proof of the lemma. QED

PROOF OF THEOREM 4.16. The set of validities is recursively enumerable, because this language is just a syntactic variant of first order logic with 2 variables, and the completeness theorem of first order logic.

The complement is also recursively enumerable, by lemmas 4.20 and 4.21, which say that satisfiability of ξ is equivalent to the existence of a linked set of mosaics for ξ. Though ξ is a formula without any occurrences of $\Diamond_0$, the lemmas still cover the whole language by Remark 4.1. A ξ-LSM is a finite first order structure with finite signature satisfying some finite number of first order conditions. Thus we can enumerate all ξ-LSM's, whence we can enumerate all non-valid formulas. QED

Summary

We showed that if we only have 2 cylindric diamonds, the logic is much more permissive with respect to complexity and decidability: on relativized models the logic becomes PSPACE-complete, and neither classical full Cartesian models, nor expansion with coordinate wise counting lead to undecidability, as it did before with arrow logic and MLR_3 and higher.

We have shown all our decidability results using the mosaic-based technique. The last proof shows that this technique can rather easily be adapted to quite complex classical cases.

5 Axiomatic undecidability

In this section we look at axiomatic approximations of the modal theory of the cubes and squares, and see when we meet undecidability there. We destroy the validity of the following laws when we consider relativised classes of models:

for MLR_n :	$\Diamond_i\Diamond_j p \leftrightarrow \Diamond_j\Diamond_i p$	(commutativity of the quantifiers)
for arrow logic:	$p \circ (q \circ r) \leftrightarrow (p \circ q) \circ r$	(associativity of composition).

We will show that also from an axiomatic point of view these laws can be seen as the culprits. We recall several undecidability results from the literature and provide a new stronger theorem for arrow logic. The relevant results from the algebraic logic literature are all based on the undecidability of the quasi-equational theory of semigroups —also called the *word problem.* These proofs have to establish two things:

A there is some kind of associative binary (composition-like) operator (definable) in the language;

B for the relevant part of the language, a deduction theorem holds.

The first is to mimic the semigroup operation, the second is to go from quasi-equations to equations. In (van Benthem, 1996) it was asked whether the tiling problem would lead to different undecidability results in arrow logic. Here we will show that it does. We start the section with cylindric modal logic.

5.1 Cylindric modal logic

In the introduction we already quoted Maddux' result that the modal logic of three commuting S5 diamonds is undecidable. His result is really a result in cylindric modal logic (disguised as cylindric algebras). We just translate his theorem to our formalism. Let $S5^n$ denote the modal logic of n commuting S5 modalities.

Theorem 5.1 ((Maddux, 1980)) *Let $n \geq 3$. Any extension of $S5^n$ with axioms which are valid on n-cubes has an undecidable validity problem.*

Maddux' proof is surprisingly simple. For A above he shows that the standard way of defining composition in MLR_3 gives in $S5^3$ an operator which is almost associative, namely it satisfies the associative law if one of the arguments is $\top$. For B, one can use the term-definable universal modality $\Diamond\!\!\!\!+ \varphi = \Diamond_0\Diamond_1\Diamond_2\varphi$.

As a contrast we quote the following theorem.

Theorem 5.2
(i)((Halpern and Moses, 1992): Theorems 6.16 and 6.17) Let $n \geq 2$. The validity problem of the modal logic of n S5-modalities is PSPACE*-complete.*
(ii) ((Gabbay and Shehtman, 1997): Theorem 10.15) Let $n \geq 2$. The modal logic of n commuting modalities has the finite model property. As a consequence, its set of validities is decidable.

The second part of the theorem is based on a result in (Fine, 1975), about *uniform* formulas ($\rightarrow$-formulas whose antecedent and consequent have the same modal depth).

The last two theorems suggest that it is really the combination of S5-modalities and their commutativity which causes undecidability. In (Marx, 1997) it is shown how to change Maddux' argument to obtain undecidability for 3 commuting **K45** (transitive and non-branching) modalities. This naturally gives rise to some challenging questions about the decidability of other combinations.

Problem 5.3 Precisely which combinations of the following axioms lead to undecidability with 3 modalities?

$\Diamond_i\Diamond_j p \leftrightarrow \Diamond_j\Diamond_i p$	(commutativity)
$p \rightarrow \Diamond_i p$	(reflexivity)
$p \rightarrow \Box_i\Diamond_i p$	(symmetry)
$\Diamond_i\Diamond_i p \rightarrow \Diamond_i p$	(transitivity).

Is the tiling problem better suited than the wordproblem to obtain results in these cases? Is it possible to prove Theorem 5.1 in a simple way using the tiling problem (or one of our tiling logics)?

A related but more semantical problem is whether the logic of the *product* (in the sense of (Gabbay and Shehtman, 1997)) of three **K** logics is decidable. These logics satisfy at least commutativity of the diamonds, and also "quantifier exchange" $\Diamond_i\Box_j p \rightarrow \Box_j\Diamond_i p$. If this is all, then the logic is decidable, by the same "uniform formula" argument mentioned above. But an axiomatisation of this ternary product is not known, and it might not even be finitely axiomatisable.

5.2 Arrow logic

Tarski proposed the following axiomatic approximation of the arrow logic of the squares. (Lyndon, 1950) showed that no finite Hilbert-style axiomatization is possible. Tarski showed, using the word-problem, that the arrow logic given by these axioms is undecidable. Of course, these results were obtained in the framework of relation algebras. We call this arrow logic RA. On top of the standard K-axioms, it contains

RA_1 $p \circ (q \circ r) \leftrightarrow (p \circ q) \circ r$
RA_2 $p \circ \iota\delta \leftrightarrow p$
RA_3 $\otimes\otimes p \leftrightarrow p$
RA_4 $\otimes(p \circ q) \leftrightarrow \otimes q \circ \otimes p$
RA_5 $\otimes p \circ \neg(p \circ q) \rightarrow \neg q$.

From (Lyndon, 1950) we can deduce that this logic is canonical and complete with respect to a certain elementary class of arrow frames (this also follows by a straightforward Sahlqvist argument).

Maddux invented the logic SA with the same axioms as RA, but $p \circ (\top \circ \top) \leftrightarrow (p \circ \top) \circ \top$ instead of the associative law RA_1. He showed that also this logic is undecidable ((Maddux, 1994).Corollary 16). As in his result on cylindric modal logic, SA is strong enough to show that composition is associative in every formula $p_1 \circ p_2 \circ \ldots \circ p_n$, provided that one of the arguments is $\top$. The defined operator $\top \circ \varphi \circ \top$ works on generated frames as a universal modality.

(Andréka et al., 1995) showed that the arrow logic defined by only the associative law is also undecidable. Here one cannot term-define a universal modality, but $\varphi \vee \top \circ \varphi \vee \varphi \circ \top \vee \top \circ \varphi \circ \top$ is an S4-modality which extends $\circ$. In this way, they obtain a deduction theorem for the relevant part of the language.

All these results go through if one ads square-valid axioms to these logics. These are just two examples of results by Maddux and the Hungarian group. For more we refer to the mentioned articles.

To this list of arrow logical undecidability results we will add one which is quite different in flavour. The difference is probably due to 1) the fact that we use the tiling problem, and 2) our more modal (in this case PDL (propositional dynamic logic)) view on arrow logic.

Define *VWAL* as the arrow logic of the class of all arrow frames satisfying the following (Sahlqvist) axioms

A_1 $p \circ \iota\delta \rightarrow p$ and $\iota\delta \circ p \rightarrow p$
A_2 $(p \circ (q \wedge \iota\delta)) \circ r \leftrightarrow p \circ ((q \wedge \iota\delta) \circ r)$
A_3 $(p \circ \top) \circ \top \rightarrow p \circ \top$ and $\top \circ (\top \circ p) \rightarrow \top \circ p$.

Below we show that these axioms are sufficient to guarantee undecidability. First we have a closer look at them. They are Sahlqvist axioms and correspond on arrow frames to

A_1^c $\forall xyz(Cxyz \wedge Iz \rightarrow x = y)$ and $\forall xyz(Cxyz \wedge Iy \rightarrow x = z)$
A_2^c $\forall xyzv(Cxyz \wedge Cyyv \wedge Iv \leftrightarrow Cxyz \wedge Czvz \wedge Iv)$
A_3^c $\forall xyzuv(Cxyz \wedge Cyuv \rightarrow \exists wCxuw)$ and $\forall xyzuv(Cxyz \wedge Czuv \rightarrow \exists wCxwv)$.

Note that axiom A_2 is rather weak and does not have the existential power of the associative law (we could rewrite the "real" Sahlqvist correspondent of A_2 to the equivalent version A_2^c because we also have axiom A_1). We also remark that the arrow logic WA which is defined by all RA-axioms, but with A_2 instead of the associative law, has a decidable validity problem (Németi, 1987). Actually, since this axiomatization is complete for the class of local squares (Maddux, 1982), we obtain that the validity problem is EXPTIME-complete by Theorem 3.10. So the danger seems to be in the last axiom. Using a few abbreviations, the axioms become familiar from cylindric modal logic. We will use the following:

$$
\begin{aligned}
\Phi\,\varphi &= \varphi \circ \top\\
\ominus\!\!\!\diamond\,\varphi &= \top \circ \varphi\\
\ominus\,\varphi &= \top \circ (\varphi \wedge \iota\delta)\\
\mathsf{d}_1\varphi &= \varphi \circ \neg\iota\delta\\
\langle\varphi\rangle\psi &= (\varphi \circ (\psi \wedge \iota\delta)) \circ \top.
\end{aligned}
$$

The first three we saw in section 4, d_1 is the vertical difference operator from Definition 2.18. $(r \circ (\varphi \wedge \iota\delta)) \circ \top$ is (when interpreted on the diagonal) really the program modality $\langle r\rangle\varphi$ from Propositional Dynamic Logic, also known under the name of *Peirce product*. On square models, the defined operators behave just as expected, as is easy to see. With these definitions, axioms A_2 and A_3 express familiar laws for the cylindric diamonds:

A_2^* $\Phi\,\ominus\!\!\!\diamond(p \wedge \iota\delta) \leftrightarrow \ominus\!\!\!\diamond\,\Phi(p \wedge \iota\delta)$
A_3 $\Phi\,\Phi\, p \rightarrow \Phi\, p$ and $\ominus\!\!\!\diamond\,\ominus\!\!\!\diamond\, p \rightarrow \ominus\!\!\!\diamond\, p$.

Theorem 5.4 *Any extension of VWAL with square valid axioms is undecidable.*

PROOF. We will show that *VWAL* is strong enough to interpret the undecidable global consequence relation of the logic Tile from section 2, using the arrow logical version of the standard translation.

Instead of working with Tile we will work with the expansion of Tile with a transitive modality which extends the other two. Call this logic $\mathsf{Tile}^{\diamond}$. Then we obtain a deduction theorem, whence the following undecidable problem (cf. Remark 2.9),

$$\text{for } \varphi \text{ and } \psi\ \mathsf{Tile}\text{-formulas: is } (\varphi \wedge \Box\varphi) \rightarrow \psi\ \mathsf{Tile}^{\diamond}\text{-valid?}$$

Let $(\cdot)^t$ be the following translation from $\mathsf{Tile}^{\Diamond}$-formulas to arrow logical formulas:

$$
\begin{array}{lclcl}
p^t & = & p & & \\
(\neg\varphi)^t & = & \neg\varphi^t & & \\
(\varphi \wedge \psi)^t & = & \varphi^t \wedge \psi^t & & \\
(\langle \mathsf{r} \rangle \varphi)^t & = & (r \circ (\varphi^t \wedge \iota\delta)) \circ \top & = & \langle r \rangle \varphi^t \\
(\langle \mathsf{u} \rangle \varphi)^t & = & (u \circ (\varphi^t \wedge \iota\delta)) \circ \top & = & \langle u \rangle \varphi^t \\
(\Diamond\varphi)^t & = & (\top \circ (\varphi^t \wedge \iota\delta)) \circ \top & = & \langle \top \rangle \varphi^t.
\end{array}
$$

Let Ax denote the conjunction of the following formulas

$$
\begin{array}{ll}
*^1 & \langle r \rangle \top \wedge \langle u \rangle \top \\
*^2 & \boxbar(r \to \neg \mathsf{d}_1 r) \wedge \boxbar(u \to \neg \mathsf{d}_1 u) \\
*^3 & \boxbar(r \circ u \leftrightarrow u \circ r) \\
*^4 & \boxbar(\iota\delta \circ \top) \\
*^5 & \boxbar(r \circ u \to \neg \mathsf{d}_1(r \circ u)) \wedge \boxbar(u \circ r \to \neg \mathsf{d}_1(u \circ r)) \\
*^6 & \boxbar \diamondbar(r \circ (u \wedge \neg \mathsf{d}_1 u)) \wedge \boxbar \diamondbar(u \circ (r \wedge \neg \mathsf{d}_1 r)).
\end{array}
$$

On square models, the first 3 formulas forces r and u to be commuting total functions, and the last 3 formulas follow. Since our logic is so much weaker we also need these last 3 "helper-formulas".

We claim that, for φ, ψ Tile-formulas, we have

$$\mathsf{Tile}^{\Diamond} \models (\varphi \wedge \Box\varphi) \to \psi \text{ iff } VWAL \models Ax \to (\iota\delta \to ((\varphi \wedge \Box\varphi) \to \psi)^t). \quad (39)$$

The direction from right to left is straightforward. Let $\not\models_{\mathsf{Tile}^{\Diamond}} (\varphi \wedge \Box\varphi) \to \psi$ for φ and ψ Tile-formulas. Then there is a $\mathsf{Tile}^{\Diamond}$-model $\mathfrak{M} = (W, R, U, R_{\Diamond}, \mathsf{v})$, generated by a state w and $\mathfrak{M}, w \Vdash \varphi \wedge \Box\varphi \wedge \neg\psi$. Because $\mathfrak{M}$ is generated by w, $\mathfrak{M} \models \varphi$. Now create a square arrow logical model $\mathfrak{M}^* = (W \times W, \mathsf{v}^*)$ and set $\mathsf{v}^*(r) = R$, $\mathsf{v}^*(u) = U$ and $\mathsf{v}^*(p) = \{(x, x) \mid x \in \mathsf{v}(p)\}$. It is easy to show that i) $\mathfrak{M}^* \models Ax$, and ii) for any Tile-formula θ, for any $w \in W$, $\mathfrak{M}, x \Vdash \theta \iff \mathfrak{M}^*, (x, x) \Vdash \theta^t$. But then $\mathfrak{M}^*, (w, w) \Vdash (\varphi \wedge \Box\varphi)^t \wedge \neg\psi^t$, as desired.

We will show that the left to right direction holds for any $\mathsf{Tile}^{\Diamond}$-formula, using the complete axiomatization of $\mathsf{Tile}^{\Diamond}$, which is on top of the standard **K**-axioms

$$
\begin{array}{ll}
Ti_1 & \langle \mathsf{r} \rangle \top \text{ and } \langle \mathsf{u} \rangle \top \\
Ti_2 & \langle \mathsf{r} \rangle p \wedge \langle \mathsf{r} \rangle q \to \langle \mathsf{r} \rangle (p \wedge q) \text{ and } \langle \mathsf{u} \rangle p \wedge \langle \mathsf{u} \rangle q \to \langle \mathsf{u} \rangle (p \wedge q) \\
Ti_3 & \langle \mathsf{r} \rangle \langle \mathsf{u} \rangle p \leftrightarrow \langle \mathsf{u} \rangle \langle \mathsf{r} \rangle p \\
Ti_4 & \Diamond\Diamond p \to \Diamond p \\
Ti_5 & \langle \mathsf{r} \rangle p \to \Diamond p \text{ and } \langle \mathsf{u} \rangle p \to \Diamond p.
\end{array}
$$

We claim that

$$VWAL \models (Ax \wedge \iota\delta) \rightarrow (Ti_1 \wedge \ldots \wedge Ti_5)^t. \tag{40}$$

Then we can transform a $\mathsf{Tile}^{\Diamond}$ proof of φ to a $VWAL$ proof of $(Ax \wedge \iota\delta) \rightarrow \varphi^t$, as desired. So to finish the proof of (39) we show (40).

Claim 5 $VWAL \models (Ti_4)^t$ and $VWAL \models (Ti_5)^t$.

PROOF OF CLAIM. For Ti_4 we reason as follows,

$$\begin{array}{lcl}
(\Diamond\Diamond p)^t & = & \text{(by definition)} \\
\Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.65em{-}\kern0.2em(\iota\delta \wedge \Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.65em{-}\kern0.2em(p \wedge \iota\delta)) & \rightarrow & \text{(by } A_2^*) \\
\Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.65em{-}\kern0.2em \Diamond\kern-0.65em{-}\kern0.2em \Diamond\kern-0.6em\mid\kern0.3em(p \wedge \iota\delta) & \rightarrow & \text{(by } A_3) \\
\Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.65em{-}\kern0.2em \Diamond\kern-0.6em\mid\kern0.3em(p \wedge \iota\delta) & \leftrightarrow & \text{(by } A_2^*) \\
\Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.65em{-}\kern0.2em(p \wedge \iota\delta) & \rightarrow & \text{(by } A_3) \\
\Diamond\kern-0.6em\mid\kern0.3em \Diamond\kern-0.65em{-}\kern0.2em(p \wedge \iota\delta) & = & \\
(\Diamond p)^t. & &
\end{array}$$

Ti_5 follows immediately using monotonicity (i.e., if $\models \varphi \rightarrow \psi$, then $\models \Diamond\varphi \rightarrow \Diamond\psi$, for $\Diamond$ any modality). ◀

To show that the other axioms hold, we need some extra $VWAL$-theorems, which we prove first.

Claim 6 The following formulas are all $VWAL$-valid:

$$\begin{array}{llcl}
(i) & p \wedge (\iota\delta \wedge q) \circ \top & \leftrightarrow & (\iota\delta \wedge q) \circ p \text{ or equivalently } p \wedge \Diamond\kern-0.6em\mid\kern0.3em(\iota\delta \wedge q) \leftrightarrow (\iota\delta \wedge q) \circ p \\
(ii) & p \wedge \top \circ (\iota\delta \wedge q) & \leftrightarrow & p \circ (\iota\delta \wedge q) \\
(iii) & \ominus p \wedge \ominus q & \leftrightarrow & \ominus(p \wedge q) \\
(iv) & \neg \mathsf{d}_1 p \wedge \Diamond\kern-0.6em\mid\kern0.3em(p \wedge q) & \rightarrow & (p \wedge q) \\
(v) & \iota\delta \wedge p \wedge \Diamond\kern-0.6em\mid\kern0.3em(\iota\delta \circ q) & \rightarrow & \Diamond\kern-0.6em\mid\kern0.3em((\iota\delta \wedge p) \circ q) \\
(vi) & \iota\delta \wedge \Box\kern-0.5em\mid\kern0.3em(\iota\delta \circ \top) \wedge \Diamond\kern-0.6em\mid\kern0.3em p \wedge \Diamond\kern-0.6em\mid\kern0.3em q & \rightarrow & \Diamond\kern-0.6em\mid\kern0.3em(p \wedge \Diamond\kern-0.6em\mid\kern0.3em q) \\
(vii) & \neg \mathsf{d}_1 p \wedge q \wedge \Diamond\kern-0.6em\mid\kern0.3em(p \circ r) & \rightarrow & \Diamond\kern-0.6em\mid\kern0.3em((p \wedge q) \circ r).
\end{array}$$

PROOF OF CLAIM. The $\leftarrow$ side of (i) follows from A_1 and monotonicity. For the other side,

$$\begin{array}{lcl}
p \wedge (\iota\delta \wedge q) \circ \top & \leftrightarrow & \text{(prop logic)} \\
p \wedge (\iota\delta \wedge q) \circ (p \vee \neg p) & \leftrightarrow & \text{(distribution)} \\
[p \wedge (\iota\delta \wedge q) \circ p] \vee [p \wedge (\iota\delta \wedge q) \circ \neg p] & \rightarrow & \text{(by } A_1) \\
(\iota\delta \wedge q) \circ p. & &
\end{array}$$

(ii) has the same proof. The $\leftarrow$ side of (iii) is immediate. For the other side,

$$\begin{array}{lll} \top \circ (\iota\delta \wedge p) \wedge \top \circ (\iota\delta \wedge q) & \leftrightarrow & \text{(by } (ii)) \\ [\top \circ (\iota\delta \wedge p)] \circ (\iota\delta \wedge q) & \leftrightarrow & \text{(by } A_2) \\ \top \circ [(\iota\delta \wedge p) \circ (\iota\delta \wedge q)] & \rightarrow & \text{(by } A_1) \\ \top \circ (\iota\delta \wedge p \wedge q) & = & \\ \ominus(p \wedge q). & & \end{array}$$

For (iv),

$$\begin{array}{lll} \neg \mathsf{d}_1 p \wedge \Diamond(p \wedge q) & \leftrightarrow & \\ \neg \mathsf{d}_1 p \wedge (p \wedge q) \circ (\iota\delta \vee \neg\iota\delta) & \leftrightarrow & \\ [\neg \mathsf{d}_1 p \wedge (p \wedge q) \circ \iota\delta] \vee [\neg \mathsf{d}_1 p \wedge (p \wedge q) \circ \neg\iota\delta] & \rightarrow & \text{(because } \mathsf{d}_1 p = p \circ \neg\iota\delta \text{ and } A_1) \\ p \wedge q. & & \end{array}$$

For (v),

$$\iota\delta \wedge p \wedge \Diamond(\iota\delta \circ q) \quad \leftrightarrow \quad [\iota\delta \wedge p \wedge \Diamond((\iota\delta \wedge p) \circ q)] \vee [\iota\delta \wedge p \wedge \Diamond((\iota\delta \wedge \neg p) \circ q)].$$

But the second disjunct leads to the absurdum because,

$$\begin{array}{lll} \iota\delta \wedge p \wedge \Diamond((\iota\delta \wedge \neg p) \circ q) & \rightarrow & \text{(by monotonicity)} \\ \iota\delta \wedge p \wedge ((\iota\delta \wedge \neg p) \circ \top) \circ \top & \rightarrow & \text{(by } A_3) \\ \iota\delta \wedge p \wedge (\iota\delta \wedge \neg p) \circ \top & \leftrightarrow & \text{(by } (i)) \\ (\iota\delta \wedge p) \circ (\iota\delta \wedge \neg p) & \rightarrow & \text{(by } A_1) \\ p \wedge \neg p. & & \end{array}$$

(vi) now is easy with the help of (v):

$$\begin{array}{lll} \iota\delta \wedge \boxminus(\iota\delta \circ \top) \wedge \Diamond p \wedge \Diamond q & \rightarrow & \text{(modal logic and } (i)) \\ \iota\delta \wedge \Diamond(\iota\delta \circ p) \wedge \Diamond q & \leftrightarrow & \text{(by } (v)) \\ \iota\delta \wedge \Diamond((\iota\delta \wedge \Diamond q) \circ p & \leftrightarrow & \text{(by } (i)) \\ \iota\delta \wedge \Diamond(p \wedge \Diamond(\iota\delta \wedge \Diamond q)) & \rightarrow & \text{(by } A_3) \\ \Diamond(p \wedge \Diamond q). & & \end{array}$$

(vii) finally uses a by now familiar argument:

$$\neg \mathsf{d}_1 p \wedge q \wedge \Diamond(p \circ r) \quad \leftrightarrow \quad \neg \mathsf{d}_1 p \wedge q \wedge [\, \Diamond((p \wedge q) \circ r) \vee \Diamond((p \wedge \neg q) \circ r)].$$

But the second disjunct leads to the absurdum because,

$$\begin{array}{lll} \neg \mathsf{d}_1 p \wedge q \wedge \Diamond((p \wedge \neg q) \circ r) & \rightarrow & \\ \neg \mathsf{d}_1 p \wedge q \wedge ((p \wedge \neg q) \circ \top) \circ \top & \rightarrow & \text{(by } A_3) \\ \neg \mathsf{d}_1 p \wedge q \wedge (p \wedge \neg q) \circ \top & \rightarrow & \text{(by } (iv)) \\ q \wedge \neg q. & & \end{array}$$

◀

By (ii), $(\langle \mathsf{r} \rangle p)^t = \langle r \rangle p = (r \circ (p \wedge \iota\delta)) \circ \top = \Diamond(r \circ (p \wedge \iota\delta)) = \Diamond(r \wedge \ominus p)$. We will often use the last notation.

Claim 7 $VWAL \models (*^1 \wedge \iota\delta) \rightarrow (Ti_1)^t$ and $VWAL \models (*^2 \wedge *^4 \wedge \iota\delta) \rightarrow (Ti_2)^t$.

Proof of Claim. The first claim is immediate since $*^1$ is $\langle r \rangle \top \wedge \langle u \rangle \top$. For the second we reason as follows

$$
\begin{array}{lll}
\boxbar(r \to \neg \mathsf{d}_1 r) \wedge \boxbar(\iota\delta \circ \top) \wedge \iota\delta \wedge (\langle \mathsf{r} \rangle p)^t \wedge (\langle \mathsf{r} \rangle q)^t & \leftrightarrow & \\
\boxbar(r \to \neg \mathsf{d}_1 r) \wedge \boxbar(\iota\delta \circ \top) \wedge \iota\delta \wedge \Phi(r \wedge \ominus p) \wedge \Phi(r \wedge \ominus q) & \to & \text{(by } (vi)) \\
\boxbar(r \to \neg \mathsf{d}_1 r) \wedge \Phi(r \wedge \ominus p \wedge \Phi(r \wedge \ominus q)) & \to & \text{(by modal logic)} \\
\Phi(\neg \mathsf{d}_1 r \wedge r \wedge \ominus p \wedge \Phi(r \wedge \ominus q)) & \to & \text{(by } (iv)) \\
\Phi(r \wedge \ominus p \wedge \ominus q) & \to & \text{(by } (iii)) \\
\Phi(r \wedge \ominus(p \wedge q)) & = & \\
(\langle \mathsf{r} \rangle (p \wedge q))^t. & &
\end{array}
$$

◀

Claim 8 $VWAL \models (Ax \wedge \iota\delta) \to (\langle \mathsf{r} \rangle\langle \mathsf{u} \rangle p \leftrightarrow \langle \mathsf{u} \rangle\langle \mathsf{r} \rangle p)^t$.

PROOF OF CLAIM. Since we know already that $\langle r \rangle$ and $\langle u \rangle$ are functional modalities (i.e., they distribute over all the booleans) assuming the axioms, it is sufficient to show that $\langle r \rangle\langle u \rangle p \wedge \langle u \rangle\langle r \rangle q \to \langle r \rangle\langle u \rangle(p \wedge q)$ follows from $(Ax \wedge \iota\delta)$. We will show that

$$VWAL \models (\iota\delta \wedge *^2 \wedge *^6) \to [(\langle r \rangle\langle u \rangle p \to \langle r \circ u \rangle p) \wedge (\langle u \rangle\langle r \rangle p \to \langle u \circ r \rangle p)]. \quad (41)$$

Then, since we have axioms $*^3$ and $*^5$, we obtain the desired result, by the same proof as used for Ti_4. So let us show the first conjunct of (41).

$$
\begin{array}{lll}
\iota\delta \wedge *^3 \wedge *^6 \wedge \langle r \rangle\langle u \rangle p & = & \\
\iota\delta \wedge \boxbar(r \to \neg \mathsf{d}_1 r) \wedge \boxbar\, \Phi(r \circ (u \wedge \neg \mathsf{d}_1 u)) \wedge \Phi(r \wedge \ominus \langle u \rangle p) & \to & \text{(modal logic)} \\
\Phi[\Phi(r \circ (u \wedge \neg \mathsf{d}_1 u)) \wedge \neg \mathsf{d}_1 r \wedge \ominus \langle u \rangle p] & \to & \text{(by } (vii)) \\
\Phi\, \Phi([r \wedge \ominus \langle u \rangle p] \circ (u \wedge \neg \mathsf{d}_1 u)) & \to & \text{(by } A_3) \\
\Phi([r \wedge \ominus \langle u \rangle p] \circ (u \wedge \neg \mathsf{d}_1 u)) & \to & \text{(by monotonicity)} \\
\Phi(r \circ u \wedge \ominus \langle u \rangle p \circ \neg \mathsf{d}_1 u) & = & \\
\Phi(r \circ u \wedge [\top \circ (\iota\delta \wedge \langle u \rangle p)] \circ \neg \mathsf{d}_1 u) & \leftrightarrow & \text{(by } A_2) \\
\Phi(r \circ u \wedge \top \circ [(\iota\delta \wedge \langle u \rangle p) \circ \neg \mathsf{d}_1 u]) & \leftrightarrow & \text{(by } (i)) \\
\Phi(r \circ u \wedge \top \circ [\Phi(\iota\delta \wedge \langle u \rangle p) \wedge \neg \mathsf{d}_1 u]) & \to & \text{(mon \& def. of } \langle u \rangle) \\
\Phi(r \circ u \wedge \top \circ [\Phi\, \Phi(u \wedge \ominus p) \wedge \neg \mathsf{d}_1 u]) & \to & \text{(by } A_3) \\
\Phi(r \circ u \wedge \top \circ [\Phi(u \wedge \ominus p) \wedge \neg \mathsf{d}_1 u]) & \to & \text{(by } (iv)) \\
\Phi(r \circ u \wedge \top \circ \ominus p) & = & \\
\Phi(r \circ u \wedge \top \circ (\top \circ (p \wedge \iota\delta)) & \to & \text{(by } A_3) \\
\Phi(r \circ u \wedge \top \circ (p \wedge \iota\delta)) & = & \\
\langle r \circ u \rangle p. & &
\end{array}
$$

◀

This finishes the proof of (39). In the proof we only used that the our logic contains the squares as models and the *VWAL*-axioms. Since the left hand side of (39) is an undecidable problem, the theorem indeed holds for any square-valid extension of *VWAL*. QED

6 Conclusion

We have provided a fairly complete overview of the complexity of the two most prominent modal logics of relations and several expansions and reducts. We indicated what price tags are connected to changing features of the logic in terms of the similarity type, the dimension of the relations, and the restrictions on the universes of the models. We saw that the EXPTIME-completeness of MLR_n^+ on packed-dense local cubes gave us –via the loosely guarded fragment– complexity results for a lot of other modal logics (of relations) as well. In dimension two the mosaic-technique could be modified to work not only in unrelativised settings but even with the dangerous counting quantifiers in the language. Another challenging direction in which we may try to use the mosaic-techniques is in reducts of the used modal languages where we don't have all the boolean operators, but we keep the classical semantics. Positive results are available here, e.g. the Lambek Calculus (Lambek, 1958) and the negation-free fragment of arrow logic and CML_n (Andréka, 1991).

Acknowledgements. Thanks are due to Johan van Benthem, Patrick Blackburn, Stephane Demri, Edith Hemaspaandra and Valentin Shehtman. I especially thank the members of the London Algebraic Logic Group: Robin Hirsch, Ian Hodkinson, Szabolcs Mikulás and Mark Reynolds, for extensive support and help.

References

Andréka, H. (1991). Complexity of the equations valid in algebras of relations. Thesis for D.Sc. (a post–habilitation degree) with Math. Inst. Hungar. Ac. Sci. Budapest. A slightly updated version will appear in *Annals of Pure and Applied Logic*.

Andréka, H., Hodkinson, I., and Németi, I. (1996). Finite algebras of relations are representable on finite sets. *Journal of Symbolic Logic*. Accepted for publication.

Andréka, H., Kurucz, Á., Németi, I., Sain, I., and Simon, A. (1995). Decidable and undecidable modal logics with a binary modality. *Journal of Logic, Language and Information*, 4:191–206.

Andréka, H., Kurucz, Á., Németi, I., Sain, I., and Simon, A. (1996). Causes and remedies for undecidability in arrow logics and in multi-modal logics. In (Marx et al., 1996), pages 63–100.

Blackburn, P., de Rijke, M., and Venema, Y. (1997). *A course in modal logic*. in preparation.

Buszkowski, W., Marciszewski, W., and van Benthem, J., editors (1988). *Categorial Grammar*. John Benjamins, Amsterdam.

Chandra, A., Kozen, D., and Stockmeyer, L. (1981). Alternation. *J. Assoc. Comput. Mach.*, 28(1):114–133.

Chlebus, B. (1986). Domino-tiling games. *J. Comput. Syst. Sci.*, 32:374–392.

Csirmaz, L., Gabbay, D., and de Rijke, M., editors (1995). *Logic Colloquium '92*, number 1 in Studies in Logic, Language and Information, Stanford. CSLI Publications.

Fine, K. (1975). Normal forms in modal logic. *Notre Dame Journal of Formal Logic*, 16(2):229–237.

Fisher, M. and Ladner, R. (1979). Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.*, 18(2):194–211.

Gabbay, D. and Shehtman, V. (1997). Products of modal logics. manuscript, Imperial College, London, UK. Also to appear as a chapter in the book *Fibring Logics* by D. Gabbay.

Grädel, E. (1997). On the restraining power of guards. manusscript, June.

Grädel, E., Kolaitis, P., and Vardi, M. (to appear). On the decision problem for two–variable first order logic. *Bulletin of symbolic logic*. Available on WWW:
`ftp://ftp.informatik.rwth-aachen.de/pub/FMT/graedel/basl.ps`.

Grädel, E., Otto, M., and Rosen, E. (1996). Two-variable logic with counting is decidable. Preliminary version June 1996. Available via
`ftp://ftp.informatik.rwth-aachen.de/pub/FMT/graedel/gorc2.ps`.

Gurevich, Y. (1976). The decision problem for standard classes. *Journal of Symbolic Logic*, 41:460–464.

Halpern, J. and Moses, Y. (1992). A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379.

Henkin, L., Monk, J. D., and Tarski, A. (1971 & 1985). *Cylindric Algebras, Parts I & II.* North-Holland.

Hirsch, R., Hodkinson, I., Marx, M., Mikulás, S., and Reynolds, M. (to appear). Mosaics and step–by–step. Remarks on "A modal logic of relations". In (Orlowska, 199). Appendix to the paper (Venema and Marx, 1995). Available on WWW:
`http://theory.doc.ic.ac.uk/ mm28/pub/abstracts/sl96.html`.

Humberstone, I. (1983). Inaccessible worlds. *Notre Dame Journal of Formal Logic*, 24:346–352.

Kracht, M. (1995). Highway to the danger zone. *Journal of Logic and Computation*, 5(1):93–109.

Ladner, R. (1977). The computational complexity of provability in systems of modal propositional logic. *SIAM journal of computing*, 6(3):467–480.

Lambek, J. (1958). The mathematics of sentence structure. *American Mathematical Monthly*, 65:154–169. Reprinted in (Buszkowski et al., 1988).

Lewis, H. (1980). Complexity results for classes of quantificational formulas. *Journal of Computer and System Sciences*, 21:317–353.

Lyndon, R. C. (1950). The representation of relation algebras. *Annals of Mathematics*, 51:707–729. Part II appeared in volume 63 (1956) 294–307.

Maddux, R. (1980). The equational theory of CA_3 is undecidable. *Journal of Symbolic Logic*, 45(2):311–316.

Maddux, R. (1982). Some varieties containing relation algebras. *Transactions of the American Mathematical Society*, 272:501–526.

Maddux, R. (1991). Introductory course on relation algebras, finite dimensional cylindric algebras, and their interconnections. In Andréka, H., Monk, J. D., and Németi, I., editors, *Algebraic Logic (Proc. Conf. Budapest 1988)*, pages 361–392, Amsterdam. Colloq. Math. Soc. J. Bolyai, North-Holland.

Maddux, R. (1994). Undecidable semi-associative relation algebras. *Journal of Symbolic Logic*, 59(2).

Marx, M. (1995). *Algebraic Relativization and Arrow Logic.* PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, University of Amsterdam. ILLC Dissertation Series 1995–3.

Marx, M. (1997). Mosaics and cylindric modal logic of dimension two. In Kracht, M., de Rijke, M., and Wansing, H., editors, *Advances in Modal*

Logic 1997, number ? in Studies in Logic, Language and Information, Stanford. CSLI Publications. Available on WWW:
`http://theory.doc.ic.ac.uk/ mm28/pub/abstracts/aiml96.html`.

Marx, M., Pólos, L., and Masuch, M., editors (1996). *Arrow Logic and Multimodal Logics*, Studies in Logic, Language and Information, Stanford. CSLI Publications.

Marx, M. and Venema, Y. (1997). *Multi-dimensional Modal Logic*. Applied Logic Series. Kluwer Academic Publishers.

Mikulás, S. (1995). Taming first-order logic. manuscript, Department of Mathematics and Computer Science, University of Amsterdam. To appear in: IGPL. Available on WWW:
`http://www.dcs.kcl.ac.uk/staff/szabolcs/tamefol.ps`.

Mikulás, S. (1995). *Taming Logics*. PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, University of Amsterdam. ILLC Dissertation Series 95-12.

Mikulás, S. (to appear). Undecidable relativization of algebras of relations. *Journal of Symbolic Logic*. Accepted for publication. Available on WWW:
`http://www.dcs.kcl.ac.uk/staff/szabolcs/reltile2.ps`.

Németi, I. (1986). Free algebras and decidability in algebraic logic. Thesis for D.Sc. (a post–habilitation degree) with Math. Inst. Hungar. Ac. Sci. Budapest. In Hungarian, the English version is (Németi, 1995a).

Németi, I. (1987). Decidability of relation algebras with weakened associativity. *Proc. Amer. Math. Soc.*, 100, 2:340–344.

Németi, I. (1991). Algebraizations of quantifier logics: an overview. *Studia Logica*, L(3/4):485–569.

Németi, I. (1995a). Decidability of weakened versions of first–order logic. In (Csirmaz et al., 1995), pages 177–242.

Németi, I. (1995b). A fine-structure analysis of fisrt-order logic. In (Marx et al., 1996), pages 221–247.

Orlowska, E., editor (199?). *Memorial Volume for Elena Rasiowa*. Studia Logica Library. Kluwer Academic Publishers, Dordrecht.

Pacholski, L. (1997). Complexity of two-variable logic with counting. Paper to be presented at Twelfth Annual IEEE Symposium June 29 - July 2, 1997, Warsaw, Poland. Instytut Informatyki, Wroclaw, Poland.

Papadimitriou, C. (1994). *Computational Complexity*. Addison–Wesley.

Pratt, V. (1979). Models of program logics. In *Proceedings of the 20th IEEE symposium on Foundations of Computer Science*, pages 115–122.

Reynolds, M. (1996). A decidable logic of parallelism. Submitted. Available on WWW:

http://www.dcs.kcl.ac.uk/staff/markr/paruni.ps.Z.

Spaan, E. (1993). *Complexity of modal logics.* PhD thesis, University of Amsterdam, Institute for Logic, Language and Computation.

Tarski, A. and Givant, S. (1987). *A Formalization of Set Theory without Variables*, volume 41. AMS Colloquium publications, Providence, Rhode Island.

van Benthem, J. (1983). *Modal Logic and Classical Logic.* Bibliopolis, Naples.

van Benthem, J. (1996). *Exploring Logical Dynamics.* Studies in Logic, Language and Information. CSLI Publications, Stanford.

van Benthem, J. (1997). Dynamic bits and pieces. Technical Report LP-97-01, Institute for Logic, Language and Computation, University of Amsterdam.

van Emde Boas, P. (1996). The convenience of tilings. Technical Report CT-96-01, Institute for Logic, Language and Computation, University of Amsterdam.

Venema, Y. (1989). Two-dimensional modal logics for relational algebras and temporal logic of intervals. ITLI-prepublication series LP–89–03, University of Amsterdam.

Venema, Y. (1992). *Many-Dimensional Modal Logic.* PhD thesis, Institute for Logic, Language and Computation, University of Amsterdam, Universiteit van Amsterdam.

Venema, Y. (1995). A modal logic of quantification and substitution. In (Csirmaz et al., 1995), pages 293–309.

Venema, Y. and Marx, M. (1995). A modal logic of relations. In (Orlowska, 199). Available on WWW:
http://theory.doc.ic.ac.uk/ mm28/pub/abstracts/sl96.html.