

Decidability of Quasi-Dense Modal Logics

Piotr Ostropolski-Nalewaja
Technische Universität Dresden
Dresden, Germany
University of Wrocław
Wrocław, Poland
postropolski@cs.uni.wroc.pl

Tim S. Lyon
Technische Universität Dresden
Dresden, Germany
timothy_stephen.lyon@tu-dresden.de

ABSTRACT

The decidability of axiomatic extensions of the modal logic K with *modal reduction principles*, i.e. axioms of the form $\diamond^k p \rightarrow \diamond^n p$, has remained a long-standing open problem. In this paper, we make significant progress toward solving this problem and show that decidability holds for a large subclass of these logics, namely, for *quasi-dense logics*. Such logics are extensions of K with modal reduction axioms such that $0 < k < n$ (dubbed *quasi-density axioms*). To prove decidability, we define novel proof systems for quasi-dense logics consisting of disjunctive existential rules, which are first-order formulae typically used to specify ontologies in the context of database theory. We show that such proof systems can be used to generate proofs and models of modal formulae, and provide an intricate model-theoretic argument showing that such generated models can be encoded as finite objects called *templates*. By enumerating templates of bound size, we obtain an EXPSPACE decision procedure as a consequence.

CCS CONCEPTS

• **Theory of computation** → **Modal and temporal logics; Automated reasoning**; Constructive mathematics; Proof theory.

KEYWORDS

Chase, Decidability, Existential rule, Kripke model, Modal logic, Modal reduction principle, Model theory, Quasi-density axiom

ACM Reference Format:

Piotr Ostropolski-Nalewaja and Tim S. Lyon. 2024. Decidability of Quasi-Dense Modal Logics. In *39th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS '24)*, July 8–11, 2024, Tallinn, Estonia. ACM, New York, NY, USA, 8 pages. <https://doi.org/10.1145/3661814.3662111>

1 INTRODUCTION

Modal logic is a rich sub-discipline of mathematical logic, having important applications in numerous areas from computer science to legal theory. For example, temporal logics have been used in the verification of programs [2], epistemic logics have found use in distributed systems [11], and multi-agent logics have been used to analyze legal concepts [5]. Modal logics are typically obtained by extending propositional logic with a set of modalities, which prefix

formulae and qualify the truth of a proposition. In this paper, we focus on a class of normal modal logics [4], dubbed *quasi-dense (modal) logics*, whose language extends propositional classical logic with the possibility modality \diamond and the necessity modality \square .

Quasi-dense logics extend the minimal normal modal logic K (cf. Blackburn et al. [4]) with a finite set of *quasi-dense axioms* of the form $\diamond^k p \rightarrow \diamond^n p$ such that $0 < k < n$.¹ This class of axioms forms a prominent subclass of the so-called *modal reduction principles* [14], which are modal axioms of the same form, though without the restriction that $k < n$. Logics extending K with modal reduction principles were studied at least as far back as 1976 by van Bentham [14]. The decidability of such logics has remained a longstanding open problem (see [15] for a discussion), though partial solutions have been obtained. For example, it is well-known that the decidability of T , $K4$, $S4$, and $K \oplus \diamond p \rightarrow \diamond \diamond p$ can be obtained via filtration techniques [4] and each logic $K \oplus \diamond^n p \rightarrow \diamond p$ is known to be decidable as well (cf. [13]). Notable however, is the result of Zakharyashev [16], who showed that any extension of $K4$ with modal reduction principles is decidable. In this paper, we obtain a new substantial result, showing all quasi-dense logics decidable via a novel model-theoretic argument. Thus, we provide additional tools for approaching the general problem and make further headway toward a complete solution.

Our approach to decidability is unique to the literature on modal logics in that we rely on tools from database theory to make our argument. In particular, we rely on *disjunctive existential rules*—normally used in ontology specification—which are first-order formulae of the form $\forall \bar{x}, \bar{y} \beta(\bar{x}, \bar{y}) \rightarrow \bigvee_{1 \leq i \leq n} \exists \bar{z}_i \alpha_i(\bar{y}_i, \bar{z}_i)$ such that $\beta(\bar{x}, \bar{y})$ and each $\alpha_i(\bar{y}_i, \bar{z}_i)$ are conjunctions of atomic formulae over constants and the variables \bar{x}, \bar{y} and \bar{y}_i, \bar{z}_i , respectively [7]. In the context of database theory, a finite collection \mathcal{R} of such rules constitutes an *ontology*, while a finite collection \mathcal{D} of first-order atoms over constants constitutes a *database*. Together, $(\mathcal{D}, \mathcal{R})$ form a *knowledge base*, which may be queried to extract implicit information from the explicitly specified knowledge. The primary tool in querying a knowledge base $(\mathcal{D}, \mathcal{R})$ is the so-called (*disjunctive*) *chase*, which is an algorithm that exhaustively applies rules from \mathcal{R} to \mathcal{D} , ultimately resulting in a selection of models of the knowledge base to which queries can be mapped [6, 7].

In this paper, we rely on the above tools, representing a modal formula φ as a database \mathcal{D}_φ and encoding a tableau calculus with a

¹The name ‘quasi-dense’ is obtained from the fact that such axioms are canonical for a generalization of the density property. In particular, the axiom $\diamond^k p \rightarrow \diamond^n p$ is canonical for the following property: if a directed R path of length k exists from a world w to u in a Kripke model $M = (W, R, V)$, then a directed R path of length n exists from w to u as well. When $0 < k < n$, this implies the existence of $n - 1$ points ‘between’ w and u when an R path of length k exists, thus generalizing the usual density property.

Publication rights licensed to ACM. ACM acknowledges that this contribution was authored or co-authored by an employee, contractor or affiliate of a national government. As such, the Government retains a nonexclusive, royalty-free right to publish or reproduce this article, or to allow others to do so, for Government purposes only. Request permissions from owner/author(s).

LICS '24, July 8–11, 2024, Tallinn, Estonia

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 979-8-4007-0660-8/24/07

<https://doi.org/10.1145/3661814.3662111>

set \mathcal{R} of disjunctive existential rules. Running the disjunctive chase over $(\mathcal{D}_\varphi, \mathcal{R})$ simulates a tableau algorithm, effectively yielding a tableau proof or model of φ . We find this approach to be preferable to utilizing a standard tableau approach for a couple of reasons: first, this saves us from introducing the various notions required to specify a tableau calculus, which only forms a minor part of our overall decidability argument. Second, the tools used in the context of (disjunctive) existential rules are already perfectly suited for the intricate model-theoretic arguments we employ. To put it succinctly, this approach is more parsimonious than it would otherwise be.

We note that the product of the disjunctive chase, which we denote by $Ch(\mathcal{D}_\varphi, \mathcal{R})$, consists of a potentially infinite set of *instances*, i.e. potentially infinite sets of atomic first-order formulae, which intuitively correspond to maximal branches in a tableau. When all such instances contain a contradictory formula *contr*, $Ch(\mathcal{D}_\varphi, \mathcal{R})$ witnesses the unsatisfiability of φ , and when at least one such instance omits *contr*, that instance can be transformed into a Kripke model witnessing the satisfiability of φ . Since $Ch(\mathcal{D}_\varphi, \mathcal{R})$ can be infinitely large, the central issue that needs to be overcome to obtain decidability is the ‘finitization’ of $Ch(\mathcal{D}_\varphi, \mathcal{R})$. Indeed, we show that an instance $C \in Ch(\mathcal{D}_\varphi, \mathcal{R})$ is free of *contr* iff a finite encoding of the model C exists, which we refer to as a *template*. As only a finite number of possible templates exist, we obtain an EXPSPACE decision procedure for each quasi-dense logic as a corollary.

This paper is organized as follows: In Section 2, we define the language and semantics of quasi-dense logics, preliminary notions for disjunctive existential rules, and the various model-theoretic tools required for our decidability argument. In Section 3, we show that the disjunctive chase constitutes a sound and complete ‘proof system’ for quasi-dense logics. In Section 4, we employ an intricate model-theoretic argument showing a correspondence between models generated by the disjunctive chase and finite encodings of such models (i.e. templates), yielding an EXPSPACE decision procedure for quasi-dense logics. Last, in Section 5, we share some concluding remarks and discuss future work.

2 PRELIMINARIES

2.1 Modal Logic

We let $\text{Prop} := \{p, q, r, \dots\}$ be a set of *propositional atoms* (which may be annotated), and define the *modal language* \mathcal{L}_M to be the smallest set of *modal formulae* generated by the following grammar:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \diamond \varphi \mid \square \varphi$$

where p ranges over Prop . The negation $\neg \varphi$ of a formula φ is defined recursively as shown below:

- $\neg p := \neg p$;
- $\neg \neg p := p$;
- $\neg(\psi \vee \chi) := \neg \psi \wedge \neg \chi$;
- $\neg(\psi \wedge \chi) := \neg \psi \vee \neg \chi$;
- $\neg \diamond \psi := \square \neg \psi$;
- $\neg \square \psi := \diamond \neg \psi$.

We define the *modal depth* $d(\varphi)$ of a modal formula φ recursively as follows: (1) $d(p) = d(\neg p) = 0$, (2) $d(\psi \circ \chi) = \max\{d(\psi), d(\chi)\}$ for $\circ \in \{\vee, \wedge\}$, and (3) $d(\nabla \psi) = d(\psi) + 1$ for $\nabla \in \{\diamond, \square\}$. Given a modal formula φ , define the set $ms_i(\varphi)$ of φ ’s *depth i modal subformulae* as follows: $\psi \in ms_i(\varphi)$ iff ψ is a sub-formula of φ , ψ is of the form $\nabla \chi$ with $\nabla \in \{\diamond, \square\}$, and $d(\varphi) - d(\psi) = i$.

Modal formulae are interpreted over special kinds of Kripke models, which we call \mathcal{P} -models, defined below. We note that \mathcal{P} is

taken to be a finite set of first-order formulae called *quasi-density properties*. A *quasi-density property* (or, *QDP* for short) is a first-order formula $p = \forall x, y (R^k(x, y) \rightarrow R^n(x, y))$ such that $0 < k < n$ and where $R^\ell(x, y) := (x = y)$ if $\ell = 0$ and $R^{\ell+1} := \exists z (R^\ell(x, z) \wedge R(z, y))$. We also use $k \rightarrow k_+$ to denote a QDP of the above form, where $n = k_+$ suggests that k_+ is larger than k . This notation will be used when k should be made explicit and is relevant to the discussion.

Definition 1 (\mathcal{P} -Model). Let \mathcal{P} be a finite set of QDPs. A \mathcal{P} -frame is defined to be an ordered pair $F(\mathcal{P}) := (W, R)$ such that W is a non-empty set of points, called *worlds*, and the *accessibility relation* $R \subseteq W \times W$ satisfies every $p \in \mathcal{P}$. A \mathcal{P} -model is defined to be a tuple $M(\mathcal{P}) = (F(\mathcal{P}), V)$ such that F is a \mathcal{P} -frame and $V : \text{Prop} \rightarrow 2^W$ is a *valuation function* mapping propositions to sets of worlds.

Definition 2 (Semantic Clauses). Let $M(\mathcal{P}) = (W, R, V)$ be a \mathcal{P} -model. We define a forcing relation $\Vdash_{\mathcal{P}}$ such that

- $M(\mathcal{P}), w \Vdash_{\mathcal{P}} p$ iff $w \in V(p)$;
- $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \neg p$ iff $w \notin V(p)$;
- $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \varphi \vee \psi$ iff $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \varphi$ or $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \psi$;
- $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \varphi \wedge \psi$ iff $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \varphi$ and $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \psi$;
- $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \diamond \varphi$ iff $\exists u \in W, R(w, u)$ and $M(\mathcal{P}), u \Vdash_{\mathcal{P}} \varphi$;
- $M(\mathcal{P}), w \Vdash_{\mathcal{P}} \square \varphi$ iff $\forall u \in W$, if $R(w, u)$, then $M(\mathcal{P}), u \Vdash_{\mathcal{P}} \varphi$;
- $M(\mathcal{P}) \Vdash_{\mathcal{P}} \varphi$ iff $\forall w \in W, M(\mathcal{P}), w \Vdash_{\mathcal{P}} \varphi$.

A modal formulae $\varphi \in \mathcal{L}_M$ is \mathcal{P} -valid, written $\Vdash_{\mathcal{P}} \varphi$, iff for all \mathcal{P} -models $M(\mathcal{P}), M(\mathcal{P}) \Vdash_{\mathcal{P}} \varphi$. We define the modal logic $L(\mathcal{P}) \subseteq \mathcal{L}_M$ to be the smallest set of \mathcal{P} -valid formulae, and note that $L(\emptyset)$ is the well-known minimal, normal modal logic K.

In this paper, we prove that each logic $L(\mathcal{P})$ is decidable, that is, we provide a computable function f such that for any quasi-dense logic $L(\mathcal{P})$ and modal formula $\varphi \in \mathcal{L}_M$, $f(L(\mathcal{P}), \varphi) = 1$ if $\varphi \in L(\mathcal{P})$ and $f(L(\mathcal{P}), \varphi) = 0$ if $\varphi \notin L(\mathcal{P})$. We note that this provides a partial answer to a decades old problem, which is discussed in more detail by Wolter and Zakharyashev [15, Problem 8]. Therefore, the main theorem we demonstrate is the following:

THEOREM 3. *Given a logic $L(\mathcal{P})$ and a modal formula φ , it is decidable to check if $\varphi \in L(\mathcal{P})$.*

2.2 Existential Rules and Model-Theoretic Notions

Formulae and Syntax. We let C and V be two disjoint, denumerable sets of *constants* and *variables*. We use a, b, c, \dots to denote constants, and use x, y, z, \dots to denote variables (potentially annotated). We define the set of *terms* T to be the set of words over $C \cup V$ including the empty word ε ;² we denote terms by t, s, \dots (potentially annotated). We use \cdot to denote the concatenation operation over terms.

Moreover, we let $\text{Pred} := \{\text{contr}\} \cup \{A, B, C, \dots\}$ be a denumerable set of *predicates* such that *contr* is a predicate of arity 0 and the remaining predicates are of arity 1 and 2. The nullary predicate *contr* serves a special purpose in our work and will be used to encode contradictory information (see Section 3). We use $ar(A) = n$ to denote that $A \in \text{Pred}$ is of arity $n \in \{0, 1, 2\}$. An *atom* is a formula of the form *contr*, $A(t)$, or $B(t, t')$ such that $t, t' \in T$, $ar(A) = 1$, and $ar(B) = 2$. We will often write atomic formulae as $A(\bar{t})$ with \bar{t} denoting a single term t or a pair of terms t, t' .

²Although this notion of term is non-standard, it is tailored to ease our proofs later on (see Definition 13 in particular).

An instance \mathcal{I} is defined to be a (potentially infinite) set of atoms. We use $\mathcal{I}, \mathcal{J}, \dots$ (potentially annotated) to denote instances. We define $A \in \text{labels}_{\mathcal{I}}(t)$ iff $A(t) \in \mathcal{I}$. When \mathcal{I} is clear from the context, we simply write $\text{labels}(t)$ omitting the subscript \mathcal{I} . We define a *path* in an instance \mathcal{I} from a term t to a term t' to be a sequence of binary atoms $P(t, t') = A_1(s_0, s_1), \dots, A_n(s_{n-1}, s_n) \in \mathcal{I}$ for $n > 0$, where $t = s_0$ and $t' = s_n$. We define the *length* of a path to be equal its cardinality. Given a path $P(t, t')$ of the above form, we define the *word of the path* to be $\text{word}(P(t, t')) = s_0 \cdots s_{n-1}$; note that we omit the end term of the path in its word. An instance \mathcal{I} is a *directed acyclic graph* (DAG) iff it is free of cycles, i.e. there exists no path from a term t in \mathcal{I} to itself. A DAG \mathcal{I} is *rooted* at a term t iff for every other term t' of \mathcal{I} there exists a path from t to t' . A *tree* is a rooted DAG such that each path from its root to any other term is unique. We define a *multi-tree* to be like a tree, but allowing for multiple edges per pair of terms. We remark that every multi-tree is a rooted DAG. We use \mathcal{T} and annotated versions thereof to denote (multi-)trees, and given a tree \mathcal{T} and a term t therein, we let \mathcal{T}_t denote the sub-tree of \mathcal{T} rooted at t . Moreover, we will employ standard graph-theoretic notions; e.g. the notion of a *descendant* and the *depth* of a tree. Given a rooted DAG \mathcal{I} with source t containing a term s , we define the *depth* of s in \mathcal{I} , denoted $\text{dep}_{\mathcal{I}}(s)$, to be equal to the minimal length among all paths from t to s in \mathcal{I} .

Substitutions and Homomorphisms. We define a *substitution* to be a partial function over T . A *homomorphism* from an instance \mathcal{I} to an instance \mathcal{J} is a substitution h from the terms of \mathcal{I} to the terms of \mathcal{J} such that (1) if $A(t_1, \dots, t_n) \in \mathcal{I}$, then $A(h(t_1), \dots, h(t_n)) \in \mathcal{J}$, and (2) $h(a) = a$, for each constant $a \in \mathsf{C}$. A homomorphism h from an instance \mathcal{I} to \mathcal{J} is an *isomorphism* iff it is bijective and h^{-1} is also a homomorphism. A homomorphism $h : \mathcal{I} \rightarrow \mathcal{J}$ is a *partial isomorphism* iff it is an isomorphism from \mathcal{I} to a sub-instance of \mathcal{J} induced by the terms of $h(\mathcal{I})$.

We define the *core* (up to isomorphism) of a finite instance \mathcal{I} to be the instance $\mathcal{J} = \text{core}(\mathcal{I})$ such that $\mathcal{J} \subseteq \mathcal{I}$, \mathcal{I} can be homomorphically mapped to \mathcal{J} , and if there exists a homomorphism h from \mathcal{J} to itself then h is an isomorphism. We also define the core of a tree \mathcal{T} of finite depth in the same manner.³

Existential Rules. A *disjunctive existential rule* is a first-order formula of the form:

$$\rho = \forall \bar{x}, \bar{y} \left(\beta(\bar{x}, \bar{y}) \rightarrow \bigvee_{1 \leq i \leq n} \exists \bar{z}_i \alpha_i(\bar{y}_i, \bar{z}_i) \right)$$

such that $\bar{y} = \bar{y}_1, \dots, \bar{y}_n$ and $\beta(\bar{x}, \bar{y}) = \text{body}(\rho)$ (called the *body*) is a conjunction of atoms over constants and the variables \bar{x}, \bar{y} , and $\bigvee_{1 \leq i \leq n} \alpha_i(\bar{y}_i, \bar{z}_i) = \text{head}(\rho)$ (called the *head*) is a disjunction such that each $\alpha_i(\bar{y}_i, \bar{z}_i)$ is a conjunction of atoms over constants and the variables \bar{y}_i, \bar{z}_i . As usual, we will often treat $\beta(\bar{x}, \bar{y})$ and $\alpha_i(\bar{y}_i, \bar{z}_i)$ as sets of atoms, rather than conjunctions of atoms; cf. [3]. We call a finite set \mathcal{R} of disjunctive existential rules a *rule set* and refer to disjunctive existential rules as rules for simplicity. A *disjunctive*

datalog rule is a rule without existential quantifiers. A rule is *non-disjunctive* iff its head contains a single disjunct.⁴

Triggers and Rule Applications. Given an instance \mathcal{I} , a rule ρ , and a homomorphism h from $\text{body}(\rho)$ to \mathcal{I} we call the pair $\langle \rho, h \rangle$ a *trigger* in \mathcal{I} . A trigger is *active* if there exists no disjunct γ of $\text{head}(\rho)$ such that h can be extended to map $\text{body}(\rho) \wedge \gamma$ to \mathcal{I} .

Given an active trigger $\pi = \langle \rho, h \rangle$ over some instance \mathcal{I} , we define the *application* of the trigger π to the instance \mathcal{I} as the following set of instances:

$$\{ \mathcal{I} \cup \alpha(h(\bar{y}), \bar{t}) \mid \exists \bar{z} \alpha(\bar{y}, \bar{z}) \text{ is a disjunct of } \text{head}(\rho) \}$$

where each \bar{t} is a tuple of fresh terms. We denote an application of a trigger to an instance by $\text{apply}(\mathcal{I}, \pi)$.

Disjunctive Chase. We use a variant of the restricted disjunctive chase presented by Carral et al. [6]. Given an instance \mathcal{I} and a set of rules \mathcal{R} we define a *fair derivation sequence* of \mathcal{I} and \mathcal{R} to be any (potentially infinite) sequence $\{\mathcal{I}_i\}_{i \in \mathbb{N}}$ satisfying the following:

- $\mathcal{I}_0 = \mathcal{I}$;
- there exists a trigger π_i such that $\mathcal{I}_{i+1} \in \text{apply}(\mathcal{I}_i, \pi_i)$, for every $i \in \mathbb{N}$;
- there exists no trigger π that is active indefinitely.

Given an instance \mathcal{I} and a set of rules \mathcal{R} we define the *chase* $Ch(\mathcal{I}, \mathcal{R})$ as the set containing $\bigcup_{i \in \mathbb{N}} \mathcal{I}_i$ for every fair derivation sequence $\{\mathcal{I}_i\}_{i \in \mathbb{N}}$ of \mathcal{I} and \mathcal{R} .

Formulae and Semantics. We define a *formula* X to be an expression of the following form: $X := \text{contr} \mid A(\bar{t}) \mid X \vee X \mid X \wedge X$. The *complexity* $|X|$ of a formula X is recursively defined as follows: $|\text{contr}| = |A(\bar{t})| = 0$ and $|X \circ Y| = |X| + |Y| + 1$ for $\circ \in \{\vee, \wedge\}$. We let \models denote the typical first-order entailment relation, which is defined between instances, rule sets, and formulae. Note that $\mathcal{I} \models \mathcal{R}$ iff there are no active \mathcal{R} -triggers in \mathcal{I} , and $Ch(\mathcal{I}, \mathcal{R}) \models X$ iff for each $\mathcal{I} \in Ch(\mathcal{I}, \mathcal{R})$, $\mathcal{I} \models X$.

OBSERVATION 4. For any two instances \mathcal{I}, \mathcal{J} and a set \mathcal{R} of non-disjunctive existential rules, if \mathcal{I} homomorphically maps to \mathcal{J} , then $Ch(\mathcal{I}, \mathcal{R})$ homomorphically maps to $Ch(\mathcal{J}, \mathcal{R})$.

3 FROM MODAL LOGIC TO RULES

For the remainder of the paper, let us fix a set \mathcal{P} of quasi-density properties and a modal formula φ . We let the signature Σ be the set $\{P_\psi \mid \psi \text{ is a sub-formula of } \varphi\} \cup \{R, E, \text{contr}\}$ such that each predicate P_ψ is unary, and both R and E are binary. We will often refer to unary atoms as *labels* of terms. Below, we define a set of disjunctive existential rules that effectively function like a tableau system for φ .⁵ This set is comprised of two distinct parts: the QDPs in \mathcal{P} and a set \mathcal{R}_φ that encodes the semantic clauses from Definition 2. By standard first-order equivalences, every QDP of the form $\forall x, y (R^k(x, y) \rightarrow R^n(x, y))$ is equivalent to a non-disjunctive

³We remark that cores of infinite instances do not exist in general [9], despite existing for infinite trees of finite depth.

⁴Non-disjunctive rules are also referred to as *tuple-generating dependencies* (TGDs) [1], *conceptual graph rules* [12], *Datalog[±]* [8], and *$\forall\exists$ -rules* [3] in the literature, though the name *existential rule* is most commonly used.

⁵The relationship between existential rules and sequent systems (which are well-known to be equivalent to tableaux) has been studied previously [10].

existential rule of the form:

$$\forall x, \bar{y}, yR(x, y_1) \wedge \dots \wedge R(y_{k-1}, y) \rightarrow \exists \bar{z}R(x, z_1) \wedge \dots \wedge R(z_{n-1}, y)$$

where $\bar{y} = y_1, \dots, y_k$ and $\bar{z} = z_1, \dots, z_n$. Therefore, we are permitted to treat \mathcal{P} as a rule set. The second rule set \mathcal{R}_φ , which encodes the modal semantics, is defined as follows:

Definition 5. Let \mathcal{R}_φ be the rule set consisting of the following:

$$\begin{aligned} P_{\psi \wedge \chi}(x) &\rightarrow P_\psi(x) \wedge P_\chi(x) & \text{(i)} \\ P_{\psi \vee \chi}(x) &\rightarrow P_\psi(x) \vee P_\chi(x) & \text{(ii)} \\ P_{\Box\psi}(x) \wedge R(x, y) &\rightarrow P_\psi(y) & \text{(iii)} \\ P_{\Diamond\psi}(x) &\rightarrow \exists y R(x, y) \wedge E(x, y) \wedge P_\psi(y) & \text{(iv)} \\ P_\psi(x) \wedge P_{\neg\psi}(x) &\rightarrow \text{contr} & \text{(v)} \end{aligned}$$

for ψ, χ ranging over sub-formulae of φ .

The above rule set encodes the the semantic clauses from Definition 2 in the following way: Rules (i) and (ii) encode the fact that if $\psi \wedge \chi$ or $\psi \vee \chi$ is satisfied at a world x in a Kripke model, then φ and (or) χ is satisfied at x . Rule (iii) encodes the fact that if $\Box\psi$ is satisfied at a world x , which relates to a world y via the accessibility relation R , in a Kripke model, then ψ holds at y . Rule (iv) captures the fact that if $\Diamond\psi$ is satisfied at x in a Kripke model, then there exists a world y such that $R(x, y)$ with ψ holding at y . Observe that the rule (iv) includes the additional binary predicate E . This predicate has a technical purpose and is used to mark which binary predicates R are introduced by rules of the form (iv) (as opposed to rules from \mathcal{P}) during the chase. The usefulness of the predicate E will become apparent in the subsequent section. Last, rule (v) expresses that the satisfaction of contradictory formulae implies a contradiction, which we denote by the special nullary predicate contr . We interpret bodies and heads of rules from $\mathcal{R}_\varphi \cup \mathcal{P}$ accordingly:

Definition 6. Let $M(\mathcal{P}) = (W, R, V)$ be a \mathcal{P} -model with $\iota : \mathsf{T} \rightarrow W$ a partial function called an *assignment*, and X, Y formulae over the signature Σ . Then,

- $M(\mathcal{P}), \iota \not\models \text{contr}$;
- $M(\mathcal{P}), \iota \models P_\psi(t)$ iff $M(\mathcal{P}), \iota(t) \models \mathcal{P} \psi$;
- $M(\mathcal{P}), \iota \models F(t, s)$ iff $(\iota(t), \iota(s)) \in R$ for $F \in \{E, R\}$;
- $M(\mathcal{P}), \iota \models X \wedge Y$ iff $M(\mathcal{P}), \iota \models X$ and $M(\mathcal{P}), \iota \models Y$;
- $M(\mathcal{P}), \iota \models X \vee Y$ iff $M(\mathcal{P}), \iota \models X$ or $M(\mathcal{P}), \iota \models Y$.

If there exists an assignment ι such that $M(\mathcal{P}), \iota \models X$, then we write $M(\mathcal{P}) \models X$.

We now demonstrate that when the chase is run with $\mathcal{R}_\varphi \cup \mathcal{P}$ over the database $\mathcal{D}_\varphi = \{P_\varphi(a)\}$ with a a fixed constant, the resulting set $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ witnesses the unsatisfiability of φ if for each $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$, $\text{contr} \in C$, and witnesses the satisfiability of φ otherwise. This soundness (Theorem 8) and completeness (Theorem 9) result is proven below.

LEMMA 7. Let $M(\mathcal{P}) = (W, R, V)$ be a \mathcal{P} -model and $\rho \in \mathcal{P} \cup \mathcal{R}_\varphi$. If $M(\mathcal{P}) \models \text{body}(\rho)$, then $M(\mathcal{P}) \models \text{head}(\rho)$.

PROOF. We prove the result by a case distinction on $\rho \in \mathcal{P} \cup \mathcal{R}_\varphi$. We consider the cases where ρ is of form (iv) in Definition 5 and the case where $\rho \in \mathcal{P}$, and note that the remaining cases are similar.

First, let $\rho = P_{\Diamond\psi}(x) \rightarrow \exists y R(x, y) \wedge E(x, y) \wedge P_\psi(y)$ and suppose $M(\mathcal{P}), \iota \models P_{\Diamond\psi}(x)$ with ι an assignment. Then, $M(\mathcal{P}), \iota(x) \models \mathcal{P} \Diamond\psi$,

implying the existence of a world $u \in W$ such that $(\iota(x), u) \in R$ and $M(\mathcal{P}), u \models \mathcal{P} \psi$. Let us define ι' such that $\iota'(t) = \iota(t)$ if $t \neq y$ and $\iota'(y) = u$. Then, it follows that $M(\mathcal{P}), \iota' \models R(x, y) \wedge E(x, y) \wedge P_\psi(y)$.

For the second case, suppose $\rho = \forall x, y(R^k(x, y) \rightarrow R^n(x, y))$ with $M(\mathcal{P}), \iota \models R^k(x, y)$ with ι an assignment. Since $M(\mathcal{P})$ is a \mathcal{P} -model, it immediately follows that there exist worlds $u_1, \dots, u_{k-1} \in W$ such that $(\iota(x), u_1), \dots, (u_{k-1}, \iota(y)) \in R$. Let z_1, \dots, z_{n-1} be the existential variables in $R^n(x, y)$, and let us define $\iota'(t) = \iota(t)$ if $t \notin \{u_1, \dots, u_{n-1}\}$ with $\iota'(z_i) = u_i$ otherwise. Then, it is clear that $M(\mathcal{P}), \iota' \models R^n(x, y)$. \square

THEOREM 8 (SOUNDNESS). If for each $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$, $\text{contr} \in C$, then $\not\models \mathcal{P} \dot{\neg} \varphi$.

PROOF. Suppose for a contradiction that $\text{contr} \in C$ for each $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$, but $\not\models \mathcal{P} \dot{\neg} \varphi$. Then, there exists a \mathcal{P} -model $M(\mathcal{P})$ such that $M(\mathcal{P}) \models P_\varphi(a)$. By Lemma 7 and the fact that each step of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ only applies rules from $\mathcal{P} \cup \mathcal{R}_\varphi$, we know that $M(\mathcal{P}) \models \text{contr}$, giving a contradiction, meaning $\not\models \mathcal{P} \dot{\neg} \varphi$. \square

THEOREM 9 (COMPLETENESS). If $\not\models \mathcal{P} \dot{\neg} \varphi$, then for each $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$, $\text{contr} \in C$.

PROOF. We prove the result by contraposition and assume that there exists a $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{contr} \notin C$. We now define a \mathcal{P} -model $M(\mathcal{P}) = (W, R, V)$ from C such that W contains all terms in C , $(t, s) \in R$ iff $R(t, s) \in C$, and $t \in V(p)$ iff $P_p(t) \in C$. Let us now verify that $M(\mathcal{P})$ is indeed a \mathcal{P} -model. First, observe that $W \neq \emptyset$ as it contains the constant a . Second, observe that by the definition of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$, we have that R satisfies each QDP in \mathcal{P} . Last, we note that V is well-defined, i.e. it cannot be the case that $t \in V(p)$ and $t \notin V(p)$ since then both $P_p(t), P_{\neg p}(t) \in C$, meaning rule $P_\psi(x) \wedge P_{\neg\psi}(x) \rightarrow \text{contr}$ would be applied at some step of the chase, ensuring that $\text{contr} \in C$.

We now show by induction on the complexity of ψ that if $P_\psi(t) \in C$, then $M(\mathcal{P}), t \models \psi$. We argue the case where $\psi = \Diamond\chi$ as the remaining cases are similar. If $P_{\Diamond\chi}(t) \in C$, then at some step of the chase, $P_{\Diamond\chi}(x) \rightarrow \exists y R(x, y) \wedge E(x, y) \wedge P_\chi(y)$ will be applied, ensuring that $R(t, y), P_\chi(y) \in C$ with y fresh. By definition, $(t, y) \in R$, and by IH, $M(\mathcal{P}), y \models \mathcal{P} \chi$, implying $M(\mathcal{P}), t \models \Diamond\chi$.

Last, as $P_\varphi(a) \in C$ by definition, it follows that $M(\mathcal{P}), a \models \mathcal{P} \varphi$, showing that $\not\models \mathcal{P} \dot{\neg} \varphi$. \square

Last, observe that every step of a fair derivation sequence of \mathcal{D}_φ and $\mathcal{R}_\varphi \cup \mathcal{P}$ starts with \mathcal{D}_φ and successively ‘grows’ a DAG rooted at a . As a consequence, the following observations hold, which is used in the next section.

OBSERVATION 10. Every element of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ is a DAG rooted at a .

OBSERVATION 11. The maximal length of E paths in any element of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ is bound by $d(\varphi)$.

PROOF. Note that creating E paths during the chase is possible only using the (iv) rules and that their usage decreases the modal depth of the subscript modal formula by one. In addition, when a (iii) rule is applied the modal depth of the ‘propagated’ formula ψ is one less than $\Diamond\psi$. As the (i) and (ii) rules only decompose conjunctions and disjunctions, the (v) rules only introduce the nullary predicate




Figure 1: Example unfolding of a rooted DAG I . The dashed blue edges are E , while solid gray edges are R . Note how terms in the unfolded instance are uniquely identified by paths from the root of I .

contr, and the \mathcal{P} rules only introduce R paths (leaving E atoms unaffected), it follows that E paths must be bounded by $n = d(\varphi)$. \square

4 DECIDABILITY VIA TEMPLATES

The goal of this section is to establish Lemma 12 (shown below), which, in conjunction with Theorems 8 and 9, implies our main decidability result (Theorem 3).

LEMMA 12. *Given a set \mathcal{P} of QDPs and a modal formula φ , it is decidable to check if there exists a $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{contr} \notin C$.*

To demonstrate the above lemma, we first show that every element of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ can be unfolded into a multi-tree (see Definition 13), trimmed to a certain depth (see Definition 15), and simplified via the core operation (see Section 2), yielding a finite instance called a template (Definition 25, and Observation 26). Second, we show that the existence of a template without contr is sufficient to conclude the existence of a $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{contr} \notin C$ (Lemma 27). In other words, we establish a correspondence between templates (of which there can be only finitely many) and members of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ (Corollary 28), culminating in a proof of the above lemma. As the size of templates depends on the modal depth of φ and elements of \mathcal{P} , we fix a few parameters that will be instrumental to conclude our proofs. In particular, we let $n = d(\varphi)$, and $N := 4n + K$ such that $K := \max\{k_+ \mid k \rightarrow k_+ \in \mathcal{P}\}$ for the remainder of the section.

4.1 Proof of Lemma 12

Let us first define three crucial operations – unfold , trim , and unravel – which will be used in our proofs. To increase comprehensibility, we provide examples of various operations we consider, which are included in Figs. 1 to 3.

Definition 13 (unfold). Given an instance I , one of its terms t and a term $s \in \mathbb{T}$, we inductively define $\text{unfold}(I, t, s)$ as:

$$\text{labels}(t)[t \mapsto s \cdot t] \cup \bigcup_{S(t,r) \in I} S(s \cdot t, s \cdot t \cdot r) \cup \bigcup_{r \text{ is a successor of } t} \text{unfold}(I, r, s \cdot t)$$

for $S \in \{R, E\}$ and where the operation $\mathcal{J}[t \mapsto s]$ replaces each occurrence of t in an instance \mathcal{J} with s . We will write $\text{unfold}(I)$ instead of $\text{unfold}(I, r, \varepsilon)$ if I is a DAG rooted at r .

An example of the unfold operation can be found in Fig. 1.

OBSERVATION 14. *Given two rooted DAGs I and J , if there exists a homomorphism h from I to J then there exists a homomorphism h' from $\text{unfold}(I)$ to $\text{unfold}(J)$*

PROOF. Let $h'(w) := h(t_1) \cdot h(t_2) \cdots h(t_n)$ for each term $w = t_1 \cdot t_2 \cdots t_n$ in $\text{unfold}(I)$. We know that $h'(w)$ exists since in I there exists a path $P(t_1, t_n)$ such that $\text{word}(P(t_1, t_n)) \cdot t_n = w$ (recall that word omits the last element of its path). The path $P(t_1, t_n)$ has to be mapped to the path $h(P(t_1, t_n))$ as h is a homomorphism from I to J . Thus, $h(P(t_1, t_n))$ witnesses the existence of the term $h'(w)$ in J . Next, let us show that h' is indeed a homomorphism. First, observe that if $A(w \cdot t) \in \text{unfold}(I)$, then $A(t) \in I$, meaning $A(h(t)) \in J$, and so, $A(h'(w \cdot t)) \in \text{unfold}(J)$. Second, let $S(w, w \cdot t)$ be an atom of $\text{unfold}(I)$. Then, $S(s, t)$ is an atom of I where $w = w' \cdot s$. Since $h(S(s, t))$ is an atom of J , $S(h'(w), h'(w \cdot t)) = S(h'(w') \cdot h(s), h'(w') \cdot h(s) \cdot h(t))$ will be an atom of $\text{unfold}(J)$. \square

Definition 15 (trim). Given a multi-tree \mathcal{T} over $\{R, E\}$ and a natural number i , we define $\text{trim}_i(\mathcal{T})$ as the multi-tree \mathcal{T} with all the elements past depth i removed, except for the elements reachable with E -paths from elements at depths no larger than i .

An example of how trim works can be found in Fig. 2. The following observation is straightforward to prove using the above definition.

OBSERVATION 16. *Given two multi-trees \mathcal{T} and \mathcal{T}' , if there exists a homomorphism h from \mathcal{T} to \mathcal{T}' then there exists a homomorphism h' from $\text{trim}_i(\mathcal{T})$ to $\text{trim}_j(\mathcal{T}')$ for any $i, j \in \mathbb{N}$ such that $i \leq j$.*

From Observation 10 and Observation 11, we get the following corollary:

COROLLARY 17. *The depth of $\text{trim}_i \circ \text{unfold}(C)$ for any element C of $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ is bounded by $i + n$. Thus, $\text{core} \circ \text{trim}_i \circ \text{unfold}(C)$ is well-defined for any $i \in \mathbb{N}$.*

Definition 18 (unravel). We define $\text{unravel}_i := \text{core} \circ \text{trim}_i \circ \text{unfold}$. Note that unravel_i is idempotent for any $i \in \mathbb{N}$, that is, for any instance I , $(\text{unravel}_i \circ \text{unravel}_i)(I)$ is isomorphic to $\text{unravel}_i(I)$.

Remark 19. We assume that the unfold , trim , and unravel operations preserve the nullary atom contr if it appears in the input instance.

OBSERVATION 20. *For any instance I and $i \in \mathbb{N}$, if $I \models \mathcal{R}_\varphi$, then $\text{unravel}_i(I) \models \mathcal{R}_\varphi$.*

PROOF. The observation holds for rules of the form (i), (ii), (iv), and (v) as every operation composed into unravel preserves the satisfaction of such rules. For a rule ρ of the form (iii), unfold trivially preserves the satisfaction of ρ , trim never removes witnesses necessary for the satisfaction of the rule, and core also trivially




Figure 2: Example trimming of a multi-tree to the depth of two. The dashed blue edges are E , while solid gray edges are R . Note how E paths are preserved if they start below depth of three.

preserves the satisfaction of the rule. Hence, such a rule will be satisfied as well. \square

OBSERVATION 21. *Given two rooted DAGs \mathcal{I} and \mathcal{J} , if there exists a homomorphism h from \mathcal{I} to \mathcal{J} , then there exists a homomorphism h' from $\text{unravel}_i(\mathcal{I})$ to $\text{unravel}_j(\mathcal{J})$ for $i, j \in \mathbb{N}$ such that $i \leq j$.*

PROOF. From Observation 14, and Observation 16 we have that there exists homomorphism h'' from $\text{trim}_i \circ \text{unfold}(\mathcal{I})$ to $\text{trim}_j \circ \text{unfold}(\mathcal{J})$. Note that $\text{core} \circ \text{trim}_i \circ \text{unfold}(\mathcal{I})$ is a sub-instance of $\text{trim}_j \circ \text{unfold}(\mathcal{I})$. Thus, there exists a homomorphism f from the former to $\text{trim}_j \circ \text{unfold}(\mathcal{J})$ that is a restriction of h'' . Moreover, by definition, there exists a homomorphism g from $\text{trim}_j \circ \text{unfold}(\mathcal{J})$ to $\text{core} \circ \text{trim}_j \circ \text{unfold}(\mathcal{J})$. Therefore, $g \circ f$ is the required homomorphism. \square

Definition 22 (embed). Given a multi-tree \mathcal{T} and two of its terms t' and t_+ such that there is a partial isomorphism f from \mathcal{T}_{t_+} to $\mathcal{T}_{t'}$ we write $\text{embed}(\mathcal{T}, t', t_+, f)$ to indicate an instance $\tilde{f}(\mathcal{T})$ where \tilde{f} is an extension of f over the set of all terms that is the identity function outside the domain of f . If such f exists we say that t_+ is *embeddable* in t' . We say that embedding is *proper* iff $\text{unravel}_\ell \circ \text{embed}(\mathcal{T}, t', t_+, f)$ is isomorphic to \mathcal{T} , where ℓ is the maximal length of a path in \mathcal{T} .

An example of how embed works can be found in Fig. 3.

Definition 23 (compliance). Given a QDP $k \rightarrow k_+$, a multi-tree \mathcal{T} and one of its terms t , we say that t is $(k \rightarrow k_+)$ -compliant iff for every descendant term t' of t at a distance k there exists a descendant term t_+ of t at a distance k_+ such that t_+ is properly embeddable in t' through the partial isomorphism f . We call every such $\langle t', t_+, f \rangle$ a *compliance witness* for t . We extend this notion to sets of QDPs in the natural way.

Observe that the right side of Fig. 3 shows that s' and s_+ , along with the partial isomorphism depicted as the purple arrows, serves as a $(2 \rightarrow 3)$ -compliance witness for s and s' .

OBSERVATION 24. *Given a rooted DAG \mathcal{I} that is a model of \mathcal{P} we have that every term t of $\text{unravel}_N(\mathcal{I})$ at a depth no greater than $2n$ is \mathcal{P} -compliant.*

PROOF. We prove that $\text{unravel}_N(\mathcal{I}) = \text{core} \circ \text{trim}_N \circ \text{unfold}$ preserves the compliance requirement by arguing that each of the composed operations preserves it. First, $\text{unfold}(\mathcal{I})$ trivially satisfies the compliance requirement. Second, assume toward a contradiction that the result $\mathcal{J} = (\text{trim}_N \circ \text{unfold})(\mathcal{I})$ does not satisfy that

requirement. Let t and t' be two terms of \mathcal{J} that are missing a compliance witness for $k \rightarrow k_+ \in \mathcal{P}$. Let $\langle t', t_+, f \rangle$ be the witness for t and t' from $\text{unfold}(\mathcal{I})$. Observe that $\langle t', t_+, f \rangle$ cannot be used as a compliance witness for t in \mathcal{J} since there exists a term s_+ that is descendant t_+ in \mathcal{J} such that $f(s_+)$ is not a term of \mathcal{J} . This, however, is a contradiction for the following three reasons: (1) Both t' and t_+ are terms of \mathcal{J} as the depth of t is $\leq 2n$, and thus, the depths of t' and t_+ do not exceed $2n + K \leq N$. (2) The paths in $\text{unfold}(\mathcal{I})$ from t_+ to s_+ and from t' to $f(s_+)$ are isomorphic. (3) The depth of s_+ in $\text{unfold}(\mathcal{I})$ is $k_+ - k$ greater than that of $f(s_+)$.

Note that if s_+ is reachable by an E path from t_+ , then the same can be said about the path from t' to $f(s_+)$. Therefore, trim_N should preserve $f(s_+)$ as t' is at depth smaller than N in $\text{unfold}(\mathcal{I})$. If s_+ is not reachable from t_+ by an E -path, it means the same for the path from t' to $f(s_+)$; however, as the depth of s_+ is greater than that of $f(s_+)$ in $\text{unfold}(\mathcal{I})$ we know that s_+ should be removed by trim_N as $f(s_+)$ is. Finally, observe that $\text{core}(\mathcal{J})$ trivially satisfies the compliance requirement. \square

Enough groundwork has been laid to define the notion of a template. We encourage the reader to briefly skim through the rest of this section to understand how templates will be used to prove Lemma 12 as this justifies the imposition of the four properties mentioned in the definition below.

Definition 25. A *template* \mathcal{N} is a multi-tree that satisfies the following properties:

- (1) $\text{unravel}_N(\mathcal{N})$ is isomorphic to \mathcal{N} .
- (2) Every term of \mathcal{N} at a depth no greater than $2n$ is \mathcal{P} -compliant.
- (3) \mathcal{N} is a model of \mathcal{R}_φ .
- (4) The root of \mathcal{N} is labeled with P_φ .

The core idea behind a template is that it serves as both a relaxation of the chase (as witnessed by Observation 26 below), yet is sufficiently strong enough to encode at least one element of the chase (see Lemma 27 below).

OBSERVATION 26. *For any $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$, $\text{unravel}_N(C)$ is a template. Moreover, if $\text{contr} \notin C$, then $\text{contr} \notin \text{unravel}_N(C)$.*

PROOF. We argue that $\text{unravel}_N(C)$ satisfies properties (1)–(4) of Definition 25. Observe that (1) follows from the idempotence of unravel_N , (2) follows from Observation 24, (3) follows from Observation 20, and (4) is trivial. As unravel does not introduce new nullary atoms, the last assertion is satisfied as well. \square

In essence, the above observation tells us that if there exists a $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{contr} \notin C$, then there exists a template witnessing this fact. To complete our decidability argument however, we will also need to establish the converse of this, i.e. the lemma below. Importantly, the definition of a template has been tailored so that the proof of this more difficult direction goes through.

LEMMA 27. *If \mathcal{N} is a template such that $\text{contr} \notin \mathcal{N}$, then there exists a $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{contr} \notin C$.*

We shall prove the above in the next section. Before that, we show how the lemma fits into the scheme of our decidability proof. From Observation 26 and Lemma 27 above we have:

COROLLARY 28. *There exists a template \mathcal{N} such that $\text{contr} \notin \mathcal{N}$ iff there exists a $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{contr} \notin C$.*

One can leverage the above to decide whether $\text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P}) \models \text{contr}$ by enumerating all the templates (of which there are only a finite number to consider) and checking whether there exists one that does not contain contr . From this we get Lemma 12, and thus, the decidability result (Theorem 3). The complexity of this decision procedure is straightforward to determine:

- (1) Checking whether a multi-tree is a template is trivially in PSPACE with respect to the size of the input multi-tree.
- (2) The size of a template is at most exponential in size of the modal formula φ and \mathcal{P} as template branching is bounded by 2^n and its depth is at most $N + n$.

Therefore, one can iterate in EXPSpace through all multi-trees of exponential size and check whether they are templates or not.

4.2 Proof of Lemma 27

Definition 29. Given a DAG \mathcal{I} rooted at a and a template \mathcal{N} we say that \mathcal{I} follows \mathcal{N} through homomorphism h iff the following hold:

- (1) h is a homomorphism from $\text{unravel}_n(\mathcal{I})$ to \mathcal{N} .
- (2) For every term t of \mathcal{I} and all words w, w' , if wt and $w't$ are terms of $\text{unravel}_n(\mathcal{I})$, then $\text{labels}(h(wt)) = \text{labels}(h(w't))$.

Note that in the above definition we use unravel_n as opposed to unravel_N . The following is the main result of this section:

LEMMA 30. *For every instance \mathcal{I} , $(\mathcal{R}_\varphi \cup \mathcal{P})$ -trigger π in \mathcal{I} , and template \mathcal{N} , if \mathcal{I} follows \mathcal{N} , then there exists an $\mathcal{I}' \in \text{apply}(\mathcal{I}, \pi)$ such that it follows \mathcal{N} .*

PROOF. Let $\pi = \langle \rho, f \rangle$, and h be the homomorphism from $\text{unravel}_n(\mathcal{I})$ to \mathcal{N} . We make a case distinction, assuming $\rho \in \mathcal{R}_\varphi$ first, and assuming $\rho \in \mathcal{P}$ second.

Let $\rho \in \mathcal{R}_\varphi$ and suppose ρ is non-disjunctive. As the template \mathcal{N} is a model of \mathcal{R}_φ there are no active triggers in it for any rule of \mathcal{R}_φ . Thus, $\text{Ch}(\mathcal{N}, \mathcal{R}_\varphi) = \{\mathcal{N}\}$. The rest follows from Observation 4, as the instance in the singleton $\text{apply}(\mathcal{I}, \pi)$ can follow the template \mathcal{N} simply through h . If $\rho \in \mathcal{R}_\varphi$ is disjunctive, then ρ is of the form (ii) in Definition 5. Let $t = f(x)$ for x as in (ii). From the fact that \mathcal{I} follows the template \mathcal{N} we know that all the copies of t that were not removed by trim or core are mapped to the set U of terms of \mathcal{N} sharing the same set of labels (by Definition 29). As \mathcal{N} is a model of \mathcal{R}_φ we can conclude that there exists a symbol S from Σ such that $S(u) \in \mathcal{N}$ for each element u of U and that $S(x)$ is a disjunct

head(ρ). From this we can set \mathcal{I}' to be $\mathcal{I} \cup \{S(t)\}$. We then have that \mathcal{I}' follows \mathcal{N} through h .

Let us now assume that $\rho = k \rightarrow k_+ \in \mathcal{P}$. Let \mathcal{I}' be the result of applying π to \mathcal{I} , $P(t, t')$ be the path from some t to t' witnessing that π was not satisfied in \mathcal{I} and $P_+(t, t')$ be the freshly created path in \mathcal{I}' . Let s denote term wt of $\text{unravel}_n(\mathcal{I})$ for some word w . Let $s' = w \cdot \text{word}(P) \cdot t'$ and $s_+ = w \cdot \text{word}(P_+(t, t')) \cdot t'$. We consider two cases depending on whether $s' \in \text{unravel}_n(\mathcal{I}')$.

Assume $s' \notin \text{unravel}_n(\mathcal{I}')$. This case is trivial. Note that elements created by unfold that appear on the copy of P_+ can be homomorphically mapped to respective elements of the copy of P . Note that the remaining part of the copy of P_+ is no longer than that of P . Therefore, those elements can be “removed” by the core operation and will not have to be homomorphically mapped to \mathcal{N} .

Assume $s' \in \text{unravel}_n(\mathcal{I}')$. Note that from Corollary 17 we have $|w| + |P| < 2n$. Thus we can use the Property 2 of template \mathcal{N} for its nodes $u = h(s)$ and $u' = h(s')$. Let $\langle u', u_+, g \rangle$ be the $k \rightarrow k_+$ compliance witness for u . We can simply extend h to map the sub-tree \mathcal{T}_{s_+} of s_+ to the sub-tree \mathcal{T}_{u_+} of u_+ and the path from s to s_+ to map to the path from u to u_+ . Note that the path from s to s_+ contains only terms that have no labels (disregarding s and s_+). Observe that the sub-tree \mathcal{T}_{s_+} is by definition equal to $\text{core} \circ \text{trim}_{n-(k_+-k)}(\mathcal{T}_{s'})$ where $\mathcal{T}_{s'}$ is the sub-tree rooted at s' , and that \mathcal{T}_{u_+} is simply equal to $\text{core} \circ \text{trim}_{N-(k_+-k)}(\mathcal{T}_{u'})$ where $\mathcal{T}_{u'}$ is the sub-tree rooted at u' . This follows from the fact that the embedding from u_+ to u' is proper. As there exists a homomorphism from $\mathcal{T}_{s'}$ to $\mathcal{T}_{u'}$ by Observation 21 we know there exists a homomorphism from \mathcal{T}_{s_+} to \mathcal{T}_{u_+} since $N > n$.

All that remains is to show that \mathcal{I}' and such extended h adheres to the second point of Definition 29. Let S be the set of all copies of t appearing in $\text{unravel}_n(\mathcal{I}')$ that satisfy $s \cdot \text{word}(P) \in \text{unravel}_n(\mathcal{I}')$ for each $s \in S$, let S' be the set $\{s \cdot \text{word}(P) \mid s \in S\}$, and S_+ be the set $\{s \cdot \text{word}(P_+) \mid s \in S\}$. By assumption, elements of $h(S')$ share the same labels, and thus, $h(S_+)$ share same set of labels as $h(S')$ due to how h was extended using Property 2 of the template. This reasoning trivially extends to sub-trees of elements in S' and S_+ \square

By inductively applying Observation 10 and Lemma 30 we obtain:

LEMMA 31. *Given a template \mathcal{N} there exists an instance $C \in \text{Ch}(\mathcal{D}_\varphi, \mathcal{R}_\varphi \cup \mathcal{P})$ such that $\text{unravel}_n(C)$ homomorphically maps to \mathcal{N} .*

Finally, as unravel preserves the nullary atom contr , we obtain Lemma 27 as a corollary of the above.

5 CONCLUDING REMARKS

In this paper, we provided a novel algorithm that uniformly decides quasi-dense logics – a substantial subclass of ‘modal reduction logics’ (i.e. logics axiomatized by extending K with modal reduction principles). Our method relied on an established correspondence between members of the disjunctive chase and finite instances, called templates, which encode Kripke models witnessing the satisfiability of a modal formula. In the process, we provided a toolkit of model-theoretic operations which we aim to extend and adapt as required to approach the general decidability problem of modal reduction logics. We remark that our result is orthogonal to the other, notable class of modal reduction logics shown decidable by Zakharyashev [16],

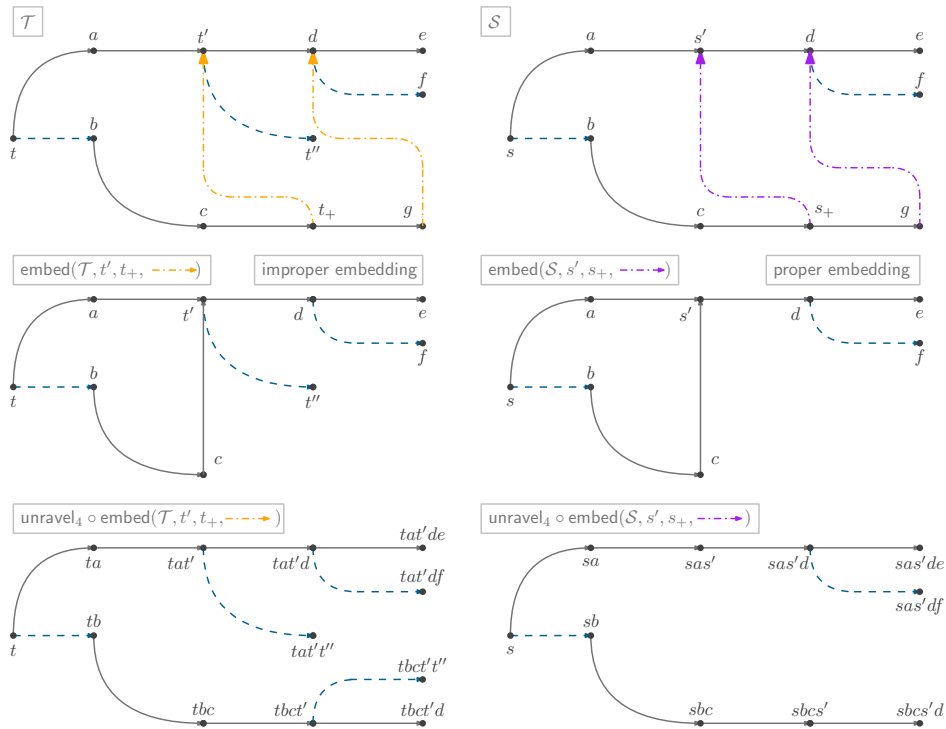


Figure 3: Two embeddings along with their unravelings where dash-dotted lines represent partial isomorphisms. The left embedding is improper as witnessed by an additional outgoing E-atom of t_+ . The right embedding is proper, as witnessed by its unraveling being isomorphic to \mathcal{S} . Note that in both cases t_+ and s_+ are deeper than the counterparts t' and s' , and thus the sub-trees of $tbct'$ and $sbcs'$ are only partially isomorphic to the respective sub-trees tat' and sas' .

and thus we have taken a meaningful step forward and improved our understanding of this longstanding open problem.

ACKNOWLEDGMENTS

Work supported by the European Research Council (ERC) Consolidator Grant 771779 (DeciGUT).

REFERENCES

- [1] S. Abiteboul, R. Hull, and V. Vianu. *Foundations of Databases*. Addison-Wesley, 1995.
- [2] B. Alpern and F. B. Schneider. Defining liveness. *Information Processing Letters*, 21(4):181–185, 1985.
- [3] J.-F. Baget, M. Leclère, M.-L. Mugnier, and E. Salvat. On rules with existential variables: Walking the decidability line. *Artificial Intelligence*, 175(9):1620–1654, 2011.
- [4] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.
- [5] J. Broersen. Deontic epistemic stit logic distinguishing modes of mens rea. *Journal of Applied Logic*, 9(2):137–152, 2011. Special Issue: Selected and revised papers from the Ninth International Conference on Deontic Logic in Computer Science (DEON 2008).
- [6] D. Carral, I. Dragoste, and M. Krötzsch. Restricted chase (non)termination for existential rules with disjunctions. 2017.
- [7] A. Deutsch and V. Tannen. Reformulation of xml queries and constraints. *Lecture Notes in Computer Science (including subseries Lecture Notes in Artificial Intelligence and Lecture Notes in Bioinformatics)*, 2572, 2003.
- [8] G. Gottlob. Datalog+/-: A unified approach to ontologies and integrity constraints. In V. D. Antonellis, S. Castano, B. Catania, and G. Guerrini, editors, *Proceedings of the 17th Italian Symposium on Advanced Database Systems, (SEBD'09)*, pages 5–6. Edizioni Seneca, 2009.
- [9] P. Hell and J. Nešetřil. The core of a graph. *Discrete Mathematics*, 109(1):117–126, 1992.
- [10] T. S. Lyon and P. Ostropolski-Nalewaja. Connecting proof theory and knowledge representation: Sequent calculi and the chase with existential rules. In P. Marquis, T. C. Son, and G. Kern-Isberner, editors, *Proceedings of the 20th International Conference on Principles of Knowledge Representation and Reasoning (KR 2023)*, pages 769–773, 2023.
- [11] J.-J. C. Meyer, J.-J. C. Meyer, and W. van der Hoek. *Epistemic logic for AI and computer science*. Number 41. Cambridge University Press, 2004.
- [12] E. Salvat and M.-L. Mugnier. Sound and complete forward and backward chainings of graph rules. In P. W. Eklund, G. Ellis, and G. Mann, editors, *Proceedings of the 4th International Conference on Conceptual Structures (ICCS'96)*, volume 1115 of *LNCS*, pages 248–262. Springer, 1996.
- [13] A. Tiu, E. Ianovski, and R. Goré. Grammar logics in nested sequent calculus: Proof theory and decision procedures. In T. Bolander, T. Bräuner, S. Ghilardi, and L. S. Moss, editors, *Advances in Modal Logic 9, papers from the ninth conference on "Advances in Modal Logic"*, pages 516–537. College Publications, 2012.
- [14] J. van Benthem. Modal reduction principles. *The Journal of Symbolic Logic*, 41(2):301–312, 1976.
- [15] F. Wolter and M. Zakharyashev. Modal decision problems. In P. Blackburn, J. Van Benthem, and F. Wolter, editors, *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 427–489. Elsevier, 2007.
- [16] M. Zakharyashev. Canonical formulas for k4. part i: Basic results. *Journal of Symbolic Logic*, 57(4):1377–1402, 1992.

Received 20 February 2007; revised 12 March 2009; accepted 5 June 2009