Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning
Main Track

Putting Perspective into OWL [sic]: Complexity-Neutral Standpoint Reasoning for Ontology Languages via Monodic S5 over Counting Two-Variable First-Order Logic

Lucía Gómez Álvarez[1] , **Sebastian Rudolph**[2]
[1]INRIA, Université Grenoble Alpes
[2]TU Dresden and Center for Scalable Data Analytics and Artificial Intelligence Dresden/Leipzig
lucia.gomez-alvarez@inria.fr, sebastian.rudolph@tu-dresden.de

Abstract

Standpoint extensions of KR formalisms have been recently introduced to incorporate multi-perspective modelling and reasoning capabilities. In such modal extensions, the integration of conceptual modelling and perspective annotations can be more or less tight, with monodic standpoint extensions striking a good balance as they enable advanced modelling while preserving good reasoning complexities.

We consider the extension of $\mathcal{C}^2$– the counting two-variable fragment of first-order logic – by monodic standpoints. At the core of our treatise is a polytime translation of formulae in said formalism into standpoint-free $\mathcal{C}^2$, requiring elaborate model-theoretic arguments. By virtue of this translation, the NExpTime-complete complexity of checking satisfiability in $\mathcal{C}^2$ carries over to our formalism. As our formalism subsumes monodic S5 over $\mathcal{C}^2$, our result also significantly advances the state of the art in research on first-order modal logics.

As a practical consequence, the very expressive description logics $\mathcal{SHOIQB}_s$ and $\mathcal{SROIQB}_s$ which subsume the popular W3C-standardized OWL 1 and OWL 2 ontology languages, are shown to allow for monodic standpoint extensions without any increase of standard reasoning complexity.

We prove that NExpTime-hardness already occurs in much less expressive DLs as long as they feature both nominals and monodic standpoints. We also show that, with inverses, functionality, and nominals present, minimally lifting the monodicity restriction leads to undecidability.

1 Introduction

Integrating knowledge from diverse, independently developed sources is a central problem in knowledge representation, particularly given the proliferation of available ontologies and other knowledge sources. Many of these ontologies – often expressed in W3C-standardized dialects of the Web Ontology Language (OWL) (Bao et al. 2009) – cover overlapping domains but embody varying conceptual frameworks and modelling choices. As an example scenario, imagine that some biomedical ontology ($\mathcal{O}_{\mathsf{Process}}$) might define `Tumour` as a dynamic biological process, whereas another ($\mathcal{O}_{\mathsf{Tissue}}$) might view it as a static abnormal tissue structure. While the description logics (DLs) (Baader et al. 2017; Rudolph 2011) underpinning OWL are well-suited to coherently model a domain, they lack mechanisms for managing heterogeneous or conflicting perspectives, leading to notorious challenges whenever such sources are to be integrated.

Standpoint logic (SL) (Gómez Álvarez and Rudolph 2021) is a recently proposed modal logic framework for multi-perspective reasoning and ontology integration. In a similar vein to epistemic logic, propositions with labelled modal operators $\Box_{\mathsf{s}}\,\phi$ and $\Diamond_{\mathsf{s}}\,\phi$ express information relative to the *standpoint* s and read, respectively: "according to s, it is *unequivocal/conceivable* that ϕ". For instance, the formula $\Box_{\mathsf{Process}}[\Diamond_{\mathsf{Tissue}}[\mathtt{Tumour}] \sqsubseteq {=}1\mathtt{TriggeredBy.Tumour}]$ expresses that, according to the $\mathsf{Process}$ standpoint, it is unequivocal that everything that is conceivably a `Tumour` from the Tissue standpoint has been triggered by exactly one `Tumour` (process). Similarly, $\Box_{\mathsf{Tissue}}[\{\mathtt{patient1}\} \sqsubseteq \exists\mathtt{HasBodyPart.}(\mathtt{Tumour} \sqcap \{\mathtt{t1}\})]$ states that according to the Tissue standpoint, it is unequivocal that `patient1` has the `Tumour` `t1` as a body part. From both, we infer that according to the $\mathsf{Process}$ standpoint, `t1` was triggered by one `Tumour`. Natural reasoning tasks over multi-standpoint specifications include gathering undisputed knowledge, determining knowledge that is relative to certain standpoints, and contrasting the knowledge from different standpoints.

The SL framework has promising applications in ontology integration, particularly in facilitating the interoperability of ontologies developed in isolation. For this reason, recent work has explored how it can be combined with logic-based formalisms underpinning the OWL family – most notably with the DLs $\mathcal{EL}$ (Gómez Álvarez, Rudolph, and Strass 2023b), $\mathcal{EL}+$ (Gómez Álvarez, Rudolph, and Strass 2023a) and $\mathcal{SHIQ}$ (Gómez Álvarez and Rudolph 2024). It has been shown that monodic extensions[1] of these languages with SL preserve the complexity of the standpoint-free DL, showing that joint reasoning over the integrated combination of possibly many ontologies is not fundamentally harder than reasoning with the ontologies in separation.

Hitherto, an open question has been whether the same holds for the very expressive side of modelling languages, in particular DLs that would fully cover high-end contemporary ontology languages such as OWL 2 DL. The results obtained so far for such languages only considered *sentential fragments* (Gómez Álvarez, Rudolph, and Strass 2022), which is an easier but much more restricted case with no interplay

[1]Monodic extensions of first-order modal logic allow for one free variable in the scope of the modal operator, and for modalised axioms and concept expressions in the case of modal DLs.

between quantification and modal operators (e.g., in DLs, the modal operators can only occur on the axiom-level).

In this paper, we address this open question by considering the extension of $\mathcal{C}^2$– the counting two-variable fragment of first-order logic, which in fact has already gained some popularity for serving as a logic to embed very expressive DLs into – by monodic standpoints. After the preliminaries (in Section 2), we provide, in Section 3, a polytime translation of formulae in said formalism into plain $\mathcal{C}^2$, using elaborate model-theoretic arguments. From this, we establish that the NExpTime-completeness of checking (finite) satisfiability in $\mathcal{C}^2$ carries over to *monodic standpoint* $\mathcal{C}^2$. As our formalism subsumes monodic S5 over $\mathcal{C}^2$, our result also significantly advances the state of the art in first-order modal logic.

Section 4 exposes how, as a consequence, the very expressive DLs $\mathcal{SHOIQB}_s$ and $\mathcal{SROIQB}_s$ which subsume the OWL 1 and OWL 2 ontology languages, also allow for monodic standpoint extensions without any increase of standard reasoning complexity. Moreover, in Section 5 we prove that NExpTime-hardness already occurs in much less expressive DLs as long as they feature both nominals and monodic standpoints. What is more, we show that, with inverses, functionality, and nominals present, minimally lifting the monodicity restriction by allowing for one distinguished rigid binary predicate leads to undecidability. The full proofs for most proof sketches can be found in the extended version of this paper (Gómez Álvarez and Rudolph 2025).

2 Preliminaries

2.1 First-Order Standpoint Logic

We introduce syntax and semantics of first-order standpoint logic (FOSL, see Gómez Álvarez, Rudolph, and Strass 2022).

Definition 1. The syntax of any FOSL formula is based on a set $\mathbf{V}$ of *variables*, typically denoted with $x, y, \ldots$, and a *signature* $\langle \mathbf{P}, \mathbf{C}, \mathbf{S}\rangle$, consisting of *predicates* $\mathbf{P}$ (each associated with an arity $n \in \mathbb{N}$), *constants* $\mathbf{C}$ and *standpoint symbols* $\mathbf{S}$, usually denoted s, s'. In particular, $\mathbf{S}$ also contains $*$, the *universal standpoint*. $\mathbf{V}$, $\mathbf{P}$, $\mathbf{C}$, and $\mathbf{S}$ are assumed to be pairwise disjoint. The set $\mathbf{T}$ of *terms* contains all constants and variables, that is, $\mathbf{T} = \mathbf{C} \cup \mathbf{V}$.

The set $\mathbf{E_S}$ of *standpoint expressions* is defined by

$$\mathsf{e}_1, \mathsf{e}_2 ::= \mathsf{s} \mid \mathsf{e}_1 \cap \mathsf{e}_2 \mid \mathsf{e}_1 \cup \mathsf{e}_2 \mid \mathsf{e}_1 \setminus \mathsf{e}_2,$$

with $\mathsf{s} \in \mathbf{S}$. The set $\mathbb{S}_{\mathrm{FO}}$ of FOSL *formulae* is then given by

$$\phi, \psi ::= \mathsf{P}(t_1, \ldots, t_k) \mid t_1 \doteq t_2 \mid \neg\phi \mid \phi \wedge \psi \mid \exists^{\lhd n} x.\phi \mid \Diamond_{\mathsf{e}}\, \phi,$$

where $\mathsf{P} \in \mathbf{P}$ is a k-ary predicate; $t_1, \ldots, t_k \in \mathbf{T}$ are terms; $\lhd$ is any of $\leq$, $=$, or $\geq$; $n \in \mathbb{N}$; $x \in \mathbf{V}$; and $\mathsf{e} \in \mathbf{E_S}$. ♢

For a formula ϕ, we denote the set of all of its subformulae by $Sub(\phi)$. The *size* of a formula is $|\phi| := |Sub(\phi)|$. The connectives and operators $\mathbf{t}$, $\mathbf{f}$, $\phi \vee \psi$, $\phi \rightarrow \psi$, $\phi \leftrightarrow \psi$, $\forall x.\phi$, and $\Box_{\mathsf{e}}\, \phi$ are introduced as syntactic macros as usual – in particular, $\forall x.\phi$ is used to abbreviate $\exists^{=0} x.\neg\phi$. In line with intuition, we may just write $\exists x.\phi$ instead of $\exists^{\geq 1} x.\phi$. We note that in full first-order logic, the somewhat exotic *counting quantifiers* $\exists^{\lhd n}$ do not add extra expressivity compared to the non-counting ones, but they do make a difference when the number of variables is restricted. As this is where we are heading, it is convenient to start from this syntax definition.

A first-order standpoint logic formula ϕ is called

- *monodic* if in each of its subformulae of the shape $\Diamond_{\mathsf{e}}\, \psi$, the formula ψ has at most one free variable,
- $\mathcal{C}^2$ if it only uses the two variables x and y and predicates of arity ≤ 2, and *plain* $\mathcal{C}^2$ if it is $\mathcal{C}^2$ and does not use $\Diamond$,
- *S5* if the only standpoint expression used is $*$,
- *nullary-free* if it does not use predicates of arity zero,
- *constant-free* if it does not use constants.

Moreover, we will call formulae of the form $\Diamond_*\, \phi$ *monodic modal formulae* if they have exactly one free variable and *sentential modal formulae* if they have no free variables.

Definition 2. Given a signature $\langle \mathbf{P}, \mathbf{C}, \mathbf{S}\rangle$, a *(first-order) standpoint structure* $\mathfrak{M}$ is a tuple $\langle \Delta, \Pi, \sigma, \gamma\rangle$ where:

- Δ is a non-empty set, the *domain* of $\mathfrak{M}$;
- Π is a non-empty set, called *precisifications* or *worlds*;
- σ is a function mapping each standpoint symbol from $\mathbf{S}$ to a set of worlds (i.e., a subset of Π), with $\sigma(*) = \Pi$ fixed;
- γ is a function mapping each precisification from Π to an ordinary first-order structure $\mathcal{I}$ over the domain Δ, whose interpretation function $\cdot^{\mathcal{I}}$ maps:
  - every predicate symbol $\mathsf{P} \in \mathbf{P}$ of arity k to a k-ary relation $\mathsf{P}^{\mathcal{I}} \subseteq \Delta^k$,
  - each constant symbol $\mathsf{a} \in \mathbf{C}$ to a domain element $\mathsf{a}^{\mathcal{I}} \in \Delta$.

For any two $\pi_1, \pi_2 \in \Pi$ and every $\mathsf{a} \in \mathbf{C}$ we require $\mathsf{a}^{\gamma(\pi_1)} = \mathsf{a}^{\gamma(\pi_2)}$ (i.e., we enforce *rigid constants*). ♢

If in $\mathfrak{M}$, some predicate $\mathsf{P} \in \mathbf{P}$ satisfies $\mathsf{P}^{\gamma(\pi_1)} = \mathsf{P}^{\gamma(\pi_2)}$ for every $\pi_1, \pi_2 \in \Pi$, we say that P is *rigid* (in $\mathfrak{M}$) and allow ourselves to write $\mathsf{P}^{\mathfrak{M}}$ instead of $\mathsf{P}^{\gamma(\pi_1)}$.

Definition 3. Let $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma\rangle$ be a first-order standpoint structure for the signature $\langle \mathbf{P}, \mathbf{C}, \mathbf{S}\rangle$ and $\mathbf{V}$ be a set of variables. A *variable assignment* is a function $v : \mathbf{V} \to \Delta$ mapping variables to domain elements. Given a variable assignment v, we denote by $v_{\{x\mapsto\delta\}}$ the function mapping x to $\delta \in \Delta$ and any other variable y to $v(y)$.

An interpretation function $\cdot^{\mathcal{I}}$ together with a variable assignment v specify how to interpret any term t from $\mathbf{T}$ by a domain element from Δ: We let $t^{\mathcal{I},v} = v(x)$ if $t = x \in \mathbf{V}$, and $t^{\mathcal{I},v} = a^{\mathcal{I}}$ if $t = a \in \mathbf{C}$.

To interpret standpoint expressions, we lift σ from $\mathbf{S}$ to all of $\mathbf{E_S}$ via $\sigma(\mathsf{e}_1 \bowtie \mathsf{e}_2) = \sigma(\mathsf{e}_1) \bowtie \sigma(\mathsf{e}_2)$ for $\bowtie \in \{\cup, \cap, \setminus\}$.

The satisfaction relation for formulae is defined in the usual way via structural induction. In what follows, let $\pi \in \Pi$ and let $v : \mathbf{V} \to \Delta$ be a variable assignment; we now establish the definition of the satisfaction relation $\models$ for FOSL using pointed first-order standpoint structures:

$$
\begin{array}{ll}
\mathfrak{M},\pi,v \models \mathsf{P}(t_1, \ldots, t_k) & \text{iff } (t_1^{\gamma(\pi),v}, \ldots, t_k^{\gamma(\pi),v}) \in \mathsf{P}^{\gamma(\pi)} \\
\mathfrak{M},\pi,v \models t_1 \doteq t_2 & \text{iff } t_1^{\gamma(\pi),v} = t_2^{\gamma(\pi),v} \\
\mathfrak{M},\pi,v \models \neg\phi & \text{iff } \mathfrak{M},\pi,v \not\models \phi \\
\mathfrak{M},\pi,v \models \phi \wedge \psi & \text{iff } \mathfrak{M},\pi,v \models \phi \text{ and } \mathfrak{M},\pi,v \models \psi \\
\mathfrak{M},\pi,v \models \exists^{\lhd n} x\phi & \text{iff } |\{\delta \mid \mathfrak{M},\pi,v_{\{x\mapsto\delta\}} \models \phi\}| \lhd n \\
\mathfrak{M},\pi,v \models \Diamond_{\mathsf{e}}\, \phi & \text{iff } \mathfrak{M},\pi',v \models \phi \text{ for some } \pi' \in \sigma(\mathsf{e}) \\
\mathfrak{M},\pi \models \phi & \text{iff } \mathfrak{M},\pi,v \models \phi \text{ for all } v : \mathbf{V} \to \Delta \\
\mathfrak{M} \models \phi & \text{iff } \mathfrak{M},\pi \models \phi \text{ for all } \pi \in \Pi
\end{array}
$$

As usual, $\mathfrak{M}$ is a *model* for a formula ϕ iff $\mathfrak{M} \models \phi$. ♢

Lemma 4. *Let ϕ be an $\mathbb{S}_{\mathrm{FO}}$ sentence and $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma\rangle$ be a model of ϕ. Then, for any $n \geq |\Pi|$, there exists a model $\mathfrak{M}' = \langle \Delta, \Pi', \sigma', \gamma'\rangle$ of ϕ with $|\Pi'| = n$.*

Proof Sketch. We just pick one precisification from $\mathfrak{M}$ and add as many isomorphic copies of it to $\mathfrak{M}$ as required. □

2.2 Transformations

The results obtained in the first part of this paper concern the fragment of all FOSL formulae that are monodic and $\mathcal{C}^2$ – from here on, we will refer to this logical fragment as *monodic standpoint* $\mathcal{C}^2$, short $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$. For technical reasons, we prefer to focus on formulae that additionally are S5, nullary-free, and constant-free; we will call these *frugal* for brevity. This section establishes that any $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formula can be efficiently transformed into an equisatisfiable frugal one.

Theorem 5. *For any FOSL formula ϕ, an equisatisfiable S5 FOSL formula $\mathsf{S5}(\phi)$ can be computed in polynomial time. The transformation preserves $\mathcal{C}^2$-ness and monodicity.*

Proof Sketch. Let ϕ be a FOSL formula based on a signature $\langle \mathbf{P}, \mathbf{C}, \mathbf{S}\rangle$. We show that for any formula ϕ, the formula $\mathrm{trans}(\phi)$, based on the signature $\langle \mathbf{P} \cup \mathbf{S}, \mathbf{C}, \{*\}\rangle$ is equisatisfiable and preserves $\mathcal{C}^2$-ness and monodicity. For instance, the function trans replaces $\Diamond_{\mathsf{s}}\, \psi$ by $\Diamond_*(\mathsf{s} \wedge \psi)$, introducing one nullary predicate for every standpoint symbol $\mathsf{s} \in \mathbf{S}$ and translating set expressions for standpoints into boolean expressions. The function trans is formally defined as follows:

$$\begin{aligned}
\mathrm{trans}(\mathsf{P}(t_1, \ldots, t_k)) &= \mathsf{P}(t_1, \ldots, t_k)\\
\mathrm{trans}(\neg\psi) &= \neg\mathrm{trans}(\psi)\\
\mathrm{trans}(\psi_1 \wedge \psi_2) &= \mathrm{trans}(\psi_1) \wedge \mathrm{trans}(\psi_2)\\
\mathrm{trans}(\forall x \psi) &= \forall x(\mathrm{trans}(\psi))\\
\mathrm{trans}(\Diamond_*\, \psi) &= \Diamond_*(\mathrm{trans}(\psi))\\
\mathrm{trans}(\Diamond_{\mathsf{e}}\, \psi) &= \Diamond_*(\mathrm{trans}_{\mathbf{E}}(\mathsf{e}) \wedge \mathrm{trans}(\psi))
\end{aligned}$$

Therein, $\mathrm{trans}_{\mathbf{E}}$ implements the semantics of standpoint expressions, providing for each expression $\mathsf{e} \in \mathbf{E}_{\mathbf{S}} \setminus \{*\}$ a propositional formula $\mathrm{trans}_{\mathbf{E}}(\mathsf{e})$ as follows

$$\begin{aligned}
\mathrm{trans}_{\mathbf{E}}(\mathsf{s}) &= \mathsf{s}\\
\mathrm{trans}_{\mathbf{E}}(\mathsf{e}_1 \cup \mathsf{e}_2) &= \mathrm{trans}_{\mathbf{E}}(\mathsf{e}_1) \vee \mathrm{trans}_{\mathbf{E}}(\mathsf{e}_2)\\
\mathrm{trans}_{\mathbf{E}}(\mathsf{e}_1 \cap \mathsf{e}_2) &= \mathrm{trans}_{\mathbf{E}}(\mathsf{e}_1) \wedge \mathrm{trans}_{\mathbf{E}}(\mathsf{e}_2)\\
\mathrm{trans}_{\mathbf{E}}(\mathsf{e}_1 \setminus \mathsf{e}_2) &= \mathrm{trans}_{\mathbf{E}}(\mathsf{e}_1) \wedge \neg\mathrm{trans}_{\mathbf{E}}(\mathsf{e}_2)
\end{aligned}$$

Equisatisfiability follows by an easy induction. A routine check of the translation also yields that it preserves $\mathcal{C}^2$-ness and monodicity, and it can be computed in polynomial time. Similar translations have been described before, for instance by Kurucz, Wolter, and Zakharyaschev (2023). □

Theorem 6. *For any FOSL formula ϕ, one can compute an equisatisfiable nullary-free FOSL formula $\mathsf{NF}(\phi)$ in polynomial time. The transformation preserves $\mathcal{C}^2$-ness, S5-ness, and monodicity.*

Proof Sketch. For any nullary predicate $\mathtt{N}$ occurring in ϕ, introduce a fresh unary predicate $\mathtt{P}_{\mathtt{N}}$ and replace any occurrence of $\mathtt{N}$ inside ϕ by $\forall x.(\mathtt{P}_{\mathtt{N}}(x))$. □

Theorem 7. *For any $\mathcal{C}^2$ FOSL formula ϕ, one can compute in polynomial time an equisatisfiable constant-free $\mathcal{C}^2$ FOSL formula $\mathsf{CF}(\phi)$. If ϕ is S5 and nullary-free, then so is $\mathsf{CF}(\phi)$.*

Proof Sketch. For every constant $\mathtt{a}$ occurring in ϕ, introduce a unary predicate $\mathtt{P}_{\mathtt{a}}$. Let ϕ^{consts} be the conjunction over all $\exists^{=1}x.\mathtt{P}_{\mathtt{a}}(x) \wedge \exists^{=1}x.\,\Box_*\,\mathtt{P}_{\mathtt{a}}(x)$ for all such $\mathtt{a}$. Further, obtain ϕ' by replacing every atom using constants $\mathtt{a}$, $\mathtt{b}$ as follows:

$$\begin{array}{rcll}
\mathtt{P}(\mathtt{a}) & \mapsto & \exists x.(\mathtt{P}_{\mathtt{a}}(x) \wedge \mathtt{P}(x)) & \\
\mathtt{P}(\mathtt{a}, x) & \mapsto & \exists y.(\mathtt{P}_{\mathtt{a}}(y) \wedge \mathtt{P}(y, x)) & \\
\mathtt{P}(x, \mathtt{a}) & \mapsto & \exists y.(\mathtt{P}_{\mathtt{a}}(y) \wedge \mathtt{P}(x, y)) & \\
\mathtt{P}(\mathtt{a}, y) & \mapsto & \exists x.(\mathtt{P}_{\mathtt{a}}(x) \wedge \mathtt{P}(x, y)) & \\
\mathtt{P}(y, \mathtt{a}) & \mapsto & \exists x.(\mathtt{P}_{\mathtt{a}}(x) \wedge \mathtt{P}(y, x)) & \\
\mathtt{P}(\mathtt{a}, \mathtt{b}) & \mapsto & \exists x.\exists y.(\mathtt{P}_{\mathtt{a}}(x) \wedge \mathtt{P}_{\mathtt{b}}(y) \wedge \mathtt{P}(x, y)) & \\
x \doteq \mathtt{a} & \mapsto & \mathtt{P}_{\mathtt{a}}(x) & \text{(same for } \mathtt{a} \doteq x\text{)}\\
y \doteq \mathtt{a} & \mapsto & \mathtt{P}_{\mathtt{a}}(y) & \text{(same for } \mathtt{a} \doteq y\text{)}\\
\mathtt{a} \doteq \mathtt{b} & \mapsto & \exists x.(\mathtt{P}_{\mathtt{a}}(x) \wedge \mathtt{P}_{\mathtt{b}}(x)) &
\end{array}$$

Then we let $\mathsf{CF}(\phi) = \phi^{\mathrm{consts}} \wedge \phi'$. □

Thus given an arbitrary $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formula ϕ, consecutively applying the transformations of the above theorems yields the equisatisfiable frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formula $\mathsf{CF}(\mathsf{NF}(\mathsf{S5}(\phi)))$. The transformation is polytime and, in particular, the result is of polynomial size with respect to the input.

3 Satisfiability in Monodic Standpoint $\mathcal{C}^2$

In this section, we study the satisfiability problem of frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ and prove NExpTime-completeness – which, by the previous section's results, carries over to full $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$. This finding constitutes the central result of our paper.

To get started, we provide an overview of the argument used to establish the result. In Section 3.1, we show that the satisfiability of a frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formula ϕ coincides with the existence of a structure $\mathfrak{M}$ having exponentially many precisifications with respect to ϕ's size, from which a specific kind of model – called the $\mathbf{P}_{\mathrm{E}}$-*stable permutational closure* of $\mathfrak{M}$ – can be constructed. In Section 3.2, we introduce *stacked interpretations*: these are plain first-order interpretations that closely reflect the form of standpoint structures for $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$. We also define stacked formulae ϕ^m_{stack}, which enforce models to take the form of stacked interpretations corresponding to standpoint structures with 2^m precisifications. With these components in place, we present in Section 3.3 an equisatisfiable translation from frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formulae into plain $\mathcal{C}^2$, which is polynomial in the size of the input formula.

Throughout the section, we will use a running example to help the reader navigate through the technical details.

Example 8. Consider the monodic standpoint $\mathcal{C}^2$ sentence E in Figure 1(1), expressing that there is exactly one unequivocally good thing (E_0); that everything is either unequivocally good or conceivably the best (somewhere), with no two distinct things being the best simultaneously (E_1); and that it is conceivable that everything is good or the best (E_2).

Figure 1(2) shows a model of E. Notably, in models of E with infinite domains – such as the one in Fig. 1(2) – there must also be infinitely many precisifications. This is because only one element satisfies `Good` everywhere, while every other element must be the `Best` in some precisification, with at most one such element per precisification. ♢

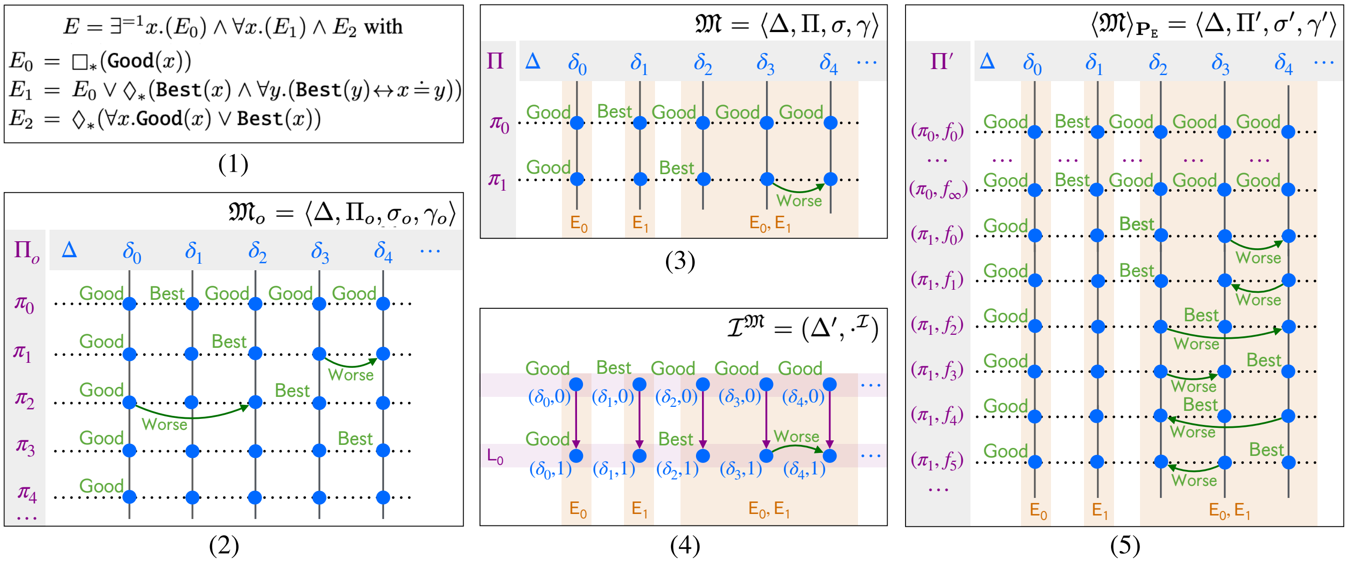

Figure 1: (1) The formula E from Example 8, and illustrations of (2) a model $\mathfrak{M}_o$ of E, (3) an interpretation $\mathfrak{M}$ with the signature $\langle \mathbf{P} \uplus \mathbf{P}_{\mathtt{E}}, \emptyset, \{*\} \rangle$, (4) the stacked interpretation $\mathcal{I}^{\mathfrak{M}}$ of $\mathfrak{M}$, and (5) $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$, the $\mathbf{P}_{\mathtt{E}}$-stable permutational closure of $\mathfrak{M}$. In the graphics, all points within a coloured area labelled with a unary predicate are in the interpretation of that predicate.

3.1 Permutational Representatives

Next, we show that for any satisfiable frugal $\mathbb{S}_{\mathcal{C}^2}^{\mathrm{mon}}$ formula ϕ, there is a structure $\mathfrak{M}$ with only exponentially many precisifications in $|\phi|$ from which a model of ϕ of a specific shape can be created (while $\mathfrak{M}$ may not be a model itself).

Definition 9. Let $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma \rangle$ be a standpoint structure for the signature $\langle \mathbf{P} \uplus \mathbf{P}_{\mathtt{E}}, \emptyset, \{*\} \rangle$, where $\mathbf{P}$ contains only unary and binary predicates, and $\mathbf{P}_{\mathtt{E}} = \{\mathtt{E}_0, \ldots, \mathtt{E}_\ell\}$ is a set of special rigid unary predicates. Let $\mathbb{P}_{\mathtt{E}}$ denote the set of all permutations (i.e., bijective functions) $f : \Delta \to \Delta$ which preserve (non)membership in every $\mathtt{E}_i$, that is, we require $\delta \in \mathtt{E}_i^{\mathfrak{M}} \Leftrightarrow f(\delta) \in \mathtt{E}_i^{\mathfrak{M}}$ for every $i \in \{0, \ldots, \ell\}$ and $\delta \in \Delta$.

Then, the $\mathbf{P}_{\mathtt{E}}$-*stable permutational closure* of $\mathfrak{M}$, denoted $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$ is the standpoint structure $\langle \Delta, \Pi', \sigma', \gamma' \rangle$ defined by

- $\Pi' = \Pi \times \mathbb{P}_{\mathtt{E}}$,
- $\sigma' = \{* \mapsto \Pi'\}$,
- $\mathtt{P}^{\gamma'((\pi,f))} = \{f(\delta) \mid \delta \in \mathtt{P}^{\gamma(\pi)}\}$ for unary predicates $\mathtt{P} \in \mathbf{P}$
- $\mathtt{P}^{\gamma'((\pi,f))} = \{(f(\delta_1), f(\delta_2)) \mid (\delta_1, \delta_2) \in \mathtt{P}^{\gamma(\pi)}\}$ for binary predicates $\mathtt{P} \in \mathbf{P}$ $\Diamond$

As we can see, the structure $\mathfrak{M}$ contains a set of special rigid unary predicates $\mathbf{P}_{\mathtt{E}}$. These predicates induce "E-types", corresponding to the subsets $T \subseteq \mathbf{P}_{\mathtt{E}}$, so a domain element is said to have the E-type T if it belongs to the interpretation of each $\mathtt{E}_i$ in T and to none outside it. We say T is realized (in $\mathfrak{M}$) if at least one domain element has it.

The $\mathbf{P}_{\mathtt{E}}$-stable permutational closure of $\mathfrak{M}$ produces a much larger structure that contains, for each initial precisification in Π, the set of precisifications with all possible permutations between domain elements belonging to the same E-type. Locally, all permuted versions of any world $\pi \in \Pi$ in the closure are isomorphic to each other, they just have their elements "swapped around", preserving the worlds' internal structure. This intuition is materialised in the next lemma.

Lemma 10. *Let ϕ be a frugal $\mathbb{S}_{\mathcal{C}^2}^{\mathrm{mon}}$ formula and let $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$ the $\mathbf{P}_{\mathtt{E}}$-stable permutational closure of some standpoint structure $\mathfrak{M}$. Let (π, f) and (π, f') be precisifications of $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$ and let $v' = f' \circ f^{-1} \circ v$. Then,*

$$\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}, (\pi, f), v \models \phi \Longleftrightarrow \langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}, (\pi, f'), v' \models \phi.$$

At the global level, E-types of domain elements, which are preserved under the permutations (by construction), will be utilized to witness the elements' "membership" in diamond-preceded subformulae in the following sense: we call some $\delta \in \Delta$ a *member* of a formula $\Diamond_* \psi$ with one free variable z if $\mathfrak{M}, \pi, \{z \mapsto \delta\} \models \Diamond_* \psi$ for some/all $\pi \in \Pi$ (note that by the semantics, the choice of π is irrelevant in this case). We will denote the set of members of $\Diamond_* \psi$ in $\mathfrak{M}$ by $(\Diamond_* \psi)^{\mathfrak{M}}$.

Let us now investigate what conditions $\mathfrak{M}$ must meet so that $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$ is a model of ϕ. First, for $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$ to witness membership to the $\Diamond_* \psi$ formulae, the number of E-types must be at least as large as the number of types induced by the monodic modal formulae – these we refer to as $\Diamond$-types. The two sets of types will be aligned: all elements with the same E-type will also have the same $\Diamond$-type. Moreover, for each formula $\Diamond_* \psi$ in a given realised $\Diamond$-type, $\mathfrak{M}$ must include at least one precisification in which some element of that type satisfies ψ. The $\mathbf{P}_{\mathtt{E}}$-stable permutational closure then ensures that every other element of the same E-type also satisfies ψ in some permutation of that precisification.

Theorem 11. *Let ϕ be a satisfiable frugal $\mathbb{S}_{\mathcal{C}^2}^{\mathrm{mon}}$ formula over the signature $\langle \mathbf{P}, \emptyset, \{*\} \rangle$. Let Dia_ϕ denote the diamond subformulae of ϕ and $FreeDia_\phi$ the diamond subformulae with one free variable. Then there is a standpoint structure $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma \rangle$ over $\langle \mathbf{P} \uplus \mathbf{P}_{\mathtt{E}}, \emptyset, \{*\} \rangle$ with*

- $|\mathbf{P}_{\mathtt{E}}| = \ell = |FreeDia_\phi|$
- $|\Pi| \leq |Dia_\phi| \cdot 2^{|Dia_\phi|}$

such that $\langle \mathfrak{M} \rangle_{\mathbf{P}_{\mathtt{E}}}$ is a model of ϕ.

Proof Sketch. To prove Theorem 11, we start from an arbitrary model $\mathfrak{M}'$ of ϕ and let $FreeDia_\phi = \{\Diamond_* \phi_1, \dots \Diamond_* \phi_\ell\}$ be the set of diamond subformulae of ϕ with one free variable. We enrich $\mathfrak{M}'$ by $\mathbf{P}_{\mathtt{E}}$, setting the extension of $\mathtt{E}_i$ to $(\Diamond_* \phi_i)^{\mathfrak{M}'}$ for every $i \in \{1, \dots, \ell\}$. Then we create a new structure $\mathfrak{M}$ by selecting at most exponentially many precisifications from the enriched $\mathfrak{M}'$ and removing the rest. Specifically,

- select an arbitrary π in case there are no diamond subformulae of ϕ at all. Otherwise,
- for each $\Diamond_* \psi$ with no free variables that is satisfied in $\mathfrak{M}'$, select some π with $\mathfrak{M}', \pi \models \psi$ and
- for each realised $\mathtt{E}$-type $T \subseteq \mathbf{P}_{\mathtt{E}}$, pick some δ having T, and select, for every $\mathtt{E}_i \in T$, one π with $\mathfrak{M}', \pi, \{z \mapsto \delta\} \models \phi_i$.

The first point ensures that Π is nonempty. The second adds witnesses for sentential modal formulae. The third provides witnesses of all monodic modal formulae from $FreeDia_\phi$. The construction ensures that at least one domain element witnesses each $\Diamond_* \phi_i$ formula of each $\Diamond$-type occurring in $\mathfrak{M}'$. Once these seed witnesses are in $\mathfrak{M}$, the rest of the elements belonging to that type in $\mathfrak{M}'$ will be witnessed by a permutation in $\langle\mathfrak{M}\rangle_{\mathbf{P}_{\mathtt{E}}}$. One can then show by induction on the structure of ϕ that $\langle\mathfrak{M}\rangle_{\mathbf{P}_{\mathtt{E}}}$ is a model iff $\mathfrak{M}'$ is a model. $\square$

Example 12. Revisiting Example 8, note that E contains two monodic modal subformulae, E_0 and E_1. From the model of E shown in Fig. 1(2), we can extract a structure $\mathfrak{M}$ with $\mathbf{P}_{\mathtt{E}} = \{\mathtt{E}_0, \mathtt{E}_1\}$ (depicted in Fig. 1(3)), such that the corresponding model $\langle\mathfrak{M}\rangle_{\mathbf{P}_{\mathtt{E}}}$ (shown in Fig. 1(5)) also satisfies E.

In constructing $\mathfrak{M}$, we proceed as follows:

- We use π_0 to witness the sentential modal subformula E_2.
- The type $\{\}$ is not realised.
- For type $\{\mathtt{E}_0\}$, we use π_0 as witnes since $\delta_0 \in \mathtt{Good}^{\gamma_o(\pi_0)}$.
- For type $\{\mathtt{E}_1\}$ we use π_0 as witnes since $\delta_1 \in \mathtt{Best}^{\gamma_o(\pi_0)}$.
- For type $\{\mathtt{E}_0, \mathtt{E}_1\}$ we use π_0 and π_1 as witnesses since $\delta_2 \in \mathtt{Good}^{\gamma_o(\pi_0)}$ and $\delta_2 \in \mathtt{Best}^{\gamma_o(\pi_1)}$. $\Diamond$

Notice that the permutations of δ_2 ensure that the membership to the formulae E_0 and E_1 in $\mathfrak{M}$ carries on to $\langle\mathfrak{M}\rangle_{\mathbf{P}_{\mathtt{E}}}$.

3.2 Stacked Interpretations

We now define a specific kind of $\mathcal{C}^2$ interpretation obtained from a given standpoint structure $\mathfrak{M}$ with 2^m precisifications, called the *stacked interpretation* of $\mathfrak{M}$ and denoted $\mathcal{I}^{\mathfrak{M}}$. This structure is designed to closely mirror the shape of $\mathfrak{M}$.

Definition 13. Let $\mathfrak{M} = \langle\Delta, \Pi, \sigma, \gamma\rangle$ with $|\Pi|=2^m$ be a standpoint structure for the signature $\langle\mathbf{P}, \emptyset, \{*\}\rangle$ where $\mathbf{P}$ contains only unary and binary predicates. Assume Π is linearly ordered with elements named $\pi_0, \pi_1, \dots, \pi_{2^m-1}$.

The *stacked interpretation* of $\mathfrak{M}$ is the FO-interpretation $\mathcal{I}^{\mathfrak{M}} = (\Delta', \cdot^{\mathcal{I}})$ with signature $\langle\mathbf{P} \uplus \{\mathtt{F}, \mathtt{L}_0, \dots, \mathtt{L}_{m-1}\}, \emptyset\rangle$, where $\mathtt{F}$ is a fresh binary predicate and $\mathtt{L}_0, \dots, \mathtt{L}_{m-1}$ are fresh unary predicates, such that

(S1) $\Delta' = \Delta \times \{0, \dots, 2^m - 1\}$
(S2) $\mathtt{L}_j^{\mathcal{I}} = \{(\delta, i) \mid \text{the } j^{\text{th}} \text{ bit of } i \text{ in binary encoding is } 1\}$
(S3) $\mathtt{F}^{\mathcal{I}} = \{((\delta, i), (\delta, i+1)) \mid \delta \in \Delta,\ 0 \leq i < 2^m - 1\}$
(S4) $\mathtt{P}^{\mathcal{I}} = \bigcup_{0 \leq i < 2^m} \mathtt{P}^{\gamma(\pi_i)} \times \{i\}$ for all unary $\mathtt{P} \in \mathbf{P}$,
(S5) $\mathtt{P}^{\mathcal{I}} = \{((\delta_1, i), (\delta_2, i)) \mid 0 \leq i < 2^m, (\delta_1, \delta_2) \in \mathtt{P}^{\gamma(\pi_i)}\}$ for all binary $\mathtt{P} \in \mathbf{P}$. $\Diamond$

Our approach constructs a stacked domain by creating one copy of the original domain Δ for each precisification in $\mathfrak{M}$, so that each new element (δ, π) mimics δ at precisification π. A set of new unary predicates $\mathtt{L}_0, \dots, \mathtt{L}_{m-1}$ bit-encodes the index of the associated precisification (each $< 2^m$). Also, a new binary predicate $\mathtt{F}$ links each element (δ, π_i) to its "next-world-twin" (δ, π_{i+1}). Thus, for every original element δ, the stacked interpretation forms an $\mathtt{F}$-chain tracking δ across all precisifications in $\mathfrak{M}$. In Figure 1, (3) depicts the stacked interpretation of (2) with the $\mathtt{F}$-chains in purple.

Definition 14. For a given $m \in \mathbb{N}$, we define ϕ^m_{stack} as the conjunction of the following formulae

(F1) $\forall x.(\bigvee_{0 \leq j < m} \neg\mathtt{L}_j(x)) \to \exists^{=1} y.\mathtt{F}(x, y)$
(F2) $\forall x.(\bigwedge_{0 \leq j < m} \mathtt{L}_j(x)) \to \exists^{=0} y.\mathtt{F}(x, y)$
(F3) $\forall x.(\bigvee_{0 \leq j < m} \mathtt{L}_j(x)) \to \exists^{=1} y.\mathtt{F}(y, x)$
(F4) $\forall x.(\bigwedge_{0 \leq j < m} \neg\mathtt{L}_j(x)) \to \exists^{=0} y.\mathtt{F}(y, x)$
(F5) $\forall xy.\mathtt{F}(x, y) \to \bigwedge_{0 \leq j < m} \Big(\big(\mathtt{L}_j(x) \leftrightarrow \mathtt{L}_j(y)\big) \leftrightarrow \bigvee_{0 \leq j' < j} \neg\mathtt{L}_{j'}(x)\Big)$
(F6) $\forall xy.\mathtt{P}(x, y) \to \bigwedge_{0 \leq j < m} \mathtt{L}_j(x) \leftrightarrow \mathtt{L}_j(y)$ for all binary $\mathtt{P} \in \mathbf{P}$. $\Diamond$

The stacked formula of size m, denoted ϕ^m_{stack}, is used to enforce that models are stacked models. Clause (F1) enforces that all elements except those with the highest index (as determined by the $\mathtt{L}$ predicates) have exactly one $\mathtt{F}$-successor. Conversely, (F2) ensures that elements with the highest index have none. Clauses (F3) and (F4) impose analogous constraints on $\mathtt{F}$-predecessors. Clause (F5) encodes that any two $\mathtt{F}$-connected elements have consecutive indexes, via a binary level-counter using the $\mathtt{L}$ predicates. Lastly, Clause (F6) enforces that all binary predicates (except $\mathtt{F}$) relate only elements with matching indices.

Lemma 15. *Any stacked interpretation $\mathcal{I}^{\mathfrak{M}}$ satisfies ϕ^m_{stack}.*

Proof sketch. We verify that each clause of ϕ^m_{stack} is satisfied by the stacked interpretation $\mathcal{I}^{\mathfrak{M}}$ as defined. In particular, the structure of the domain, and the interpretation of the predicates $\mathtt{F}, \mathtt{L}_0, \dots, \mathtt{L}_{m-1}$, and $\mathtt{P} \in \mathbf{P}$ ensure that all required properties hold. $\square$

Theorem 16. *A first-order interpretation $\mathcal{I}$ over the signature $\langle\mathbf{P} \uplus \{\mathtt{F}, \mathtt{L}_0, \dots, \mathtt{L}_{m-1}\}, \emptyset\rangle$ satisfies ϕ^m_{stack} if and only if it is isomorphic to a stacked interpretation $\mathcal{I}^{\mathfrak{M}}$ of some standpoint structure $\mathfrak{M}$ over signature $\langle\mathbf{P}, \emptyset, \{*\}\rangle$ with 2^m precisifications.*

Proof Sketch. If $\mathcal{I}$ is isomorphic to a stacked model $\mathcal{I}^{\mathfrak{M}}$ then it satisfies ϕ^m_{stack} by Lemma 15. It remains to prove the other direction, i.e., for any $\mathcal{I}$ that satisfies ϕ^m_{stack}, there exists a standpoint structure $\mathfrak{M}$ for which $\mathcal{I}^{\mathfrak{M}}$ is isomorphic to $\mathcal{I}$. We show how to construct $\mathfrak{M} = \langle\Delta, \Pi, \sigma, \gamma\rangle$ given $\mathcal{I}$ with domain Δ': For any $\delta' \in \Delta'$ we let $\mathsf{level}(\delta')$ denote the unique number $i < 2^m$ that satisfies, for every $j < m$, that the j^{th} bit in the binary encoding of i is 1 if and only if $\delta' \in \mathtt{L}_j^{\mathcal{I}}$. Moreover, we let $\approx$ be the smallest equivalence relation containing $\mathtt{F}^{\mathcal{I}}$ and let Δ consist of the $\approx$-equivalence classes of Δ'. As in Definition 13, let $\Pi = \{\pi_0, \dots, \pi_{2^m-1}\}$. Obviously, $\sigma = \{* \mapsto \Pi\}$. Finally, we set $\mathtt{P}^{\gamma(\pi_i)}$ to

- $\{[\delta']_\approx \mid \delta' \in \mathsf{P}^\mathcal{I}, \mathsf{level}(\delta') = i\}$ for unary $\mathsf{P} \in \mathbf{P}$, and
- $\{([\delta'_1]_\approx, [\delta'_2]_\approx) \mid (\delta'_1, \delta'_2) \in \mathsf{P}^\mathcal{I}, \mathsf{level}(\delta'_1) = \mathsf{level}(\delta'_2) = i\}$ for binary $\mathsf{P} \in \mathbf{P}$.

Then, the bijection $\mathsf{stacked} : \Delta' \to \Delta \times \{0, \ldots, 2^m - 1\}$ defined via $\mathsf{stacked}(\delta') = ([\delta']_\approx, \mathsf{level}(\delta'))$ can be shown to constitute an isomorphism from $\mathcal{I}$ to $\mathcal{I}^\mathfrak{M}$. □

3.3 Translating Formulae

So far, we have shown that the satisfiability of a frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formula ϕ coincides with the existence of a structure $\mathfrak{M}$ of size exponential in $|\phi|$ from which a model can be extracted. Furthermore, we demonstrated that such structures can be characterized in plain $\mathcal{C}^2$ through their corresponding stacked interpretations. In this subsection, we leverage these results to define a translation from frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ into plain $\mathcal{C}^2$, such that a frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ formula ϕ is satisfiable if and only if its translation into $\mathcal{C}^2$ is satisfiable. Together with the translation from $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ to frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$, this entails the upper complexity bound for all of $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$.

Definition 17. Given some $m \in \mathbb{N}$, we define the function Trans_m that maps frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ sentences ϕ over the signature $\langle \mathbf{P} \uplus \mathbf{P}_\mathtt{E}, \emptyset, \{*\}\rangle$ with $\mathbf{P}_\mathtt{E} = \{\mathtt{E}_0, \ldots \mathtt{E}_\ell\}$ into $\mathcal{C}^2$ as follows: $\mathsf{Trans}_m(\phi)$ is the sentence $\forall x.\forall y.(x \doteq y \to \mathsf{tr}(\phi))$, where the function tr is recursively defined via

$$\begin{array}{rcl}
\psi & \mapsto & \psi \quad \text{if } \psi \text{ is of the form } \mathsf{P}(z), \mathsf{P}(z,z') \text{ or } z \doteq z' \\
\neg\psi & \mapsto & \neg(\mathsf{tr}(\psi)) \\
\psi \wedge \psi' & \mapsto & \mathsf{tr}(\psi) \wedge \mathsf{tr}(\psi') \\
\exists^{\lhd n} z.\psi & \mapsto & \exists^{\lhd n} z.(\phi^=_\mathtt{L}(x,y) \wedge \mathsf{tr}(\psi)) \\
\Diamond_* \psi & \mapsto & \forall z_{\mathrm{nf}}.x \doteq y \to \exists z_{\mathrm{mf}}.\phi^=_\mathtt{E}(x,y) \wedge \mathsf{tr}(\psi)
\end{array}$$

where $z, z' \in \{x, y\}$ and

- z_{nf} is a variable from $\{x, y\}$ that is not free in ψ and $\{z_{\mathrm{mf}}\} = \{x, y\} \setminus \{z_{\mathrm{nf}}\}$,
- $\phi^=_\mathtt{L}(x,y)$ abbreviates $\bigwedge_{0 \le j < m} \mathtt{L}_j(x) \leftrightarrow \mathtt{L}_j(y)$, and
- $\phi^=_\mathtt{E}(x,y)$ abbreviates $\bigwedge_{0 \le i \le \ell} \mathtt{E}_i(x) \leftrightarrow \mathtt{E}_i(y)$. ◊

The key components of the translation are the handling of counting quantification and modal operators. The translation of a counting existential quantification employs the formula $\phi^=_\mathtt{L}(x,y)$ to ensure that quantification ranges only over elements belonging to the current layer of the stacked interpretation – namely, those whose counterparts correspond to the domain elements at the current precisification. In contrast, the translation of modal subformulae of the form $\Diamond_* \psi$ makes use of $\phi^=_\mathtt{E}(x,y)$ to ensure that quantification ranges over the elements belonging to the current $\mathtt{E}$-type. Recall that $\mathcal{I}^\mathfrak{M}$ is constructed to mirror the structure $\mathfrak{M}$, from which in turn we obtain the model $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}}$. Consequently, if any element of the same $\mathtt{E}$-type satisfies ψ, then there exists some permutation within $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}}$ that satisfies ψ and thus the formula $\Diamond_* \psi$ is satisfied. Notice that when $\Diamond_* \psi$ is sentential, the variable assignment does not make a difference. The following Lemma formally establishes the discussed correspondence.

Lemma 18. *Let ϕ be a frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ sentence over the signature $\langle \mathbf{P}, \emptyset, \{*\}\rangle$. Let $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma\rangle$ be a standpoint structure for the signature $\langle \mathbf{P} \uplus \mathbf{P}_\mathtt{E}, \emptyset, \{*\}\rangle$, with all predicates from $\mathbf{P}_\mathtt{E} = \{\mathtt{E}_0, \ldots, \mathtt{E}_\ell\}$ rigid, and $|\Pi| = 2^m$. Then,*

$$\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}} \models \phi \iff \mathcal{I}_\mathfrak{M} \models \mathsf{Trans}_m(\phi).$$

Proof Sketch. Toward the result, we first prove the claim that, for $\pi_i \in \Pi$, $f \in \mathbb{P}_\mathtt{E}$ and a variable assignment v, we have $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}}, (\pi_i, f), v \models \phi$ iff $\mathcal{I}_\mathfrak{M}, v' \models \mathsf{tr}(\phi)$, where $v'(z) = (f(v(z)), i)$ for $z \in \{x, y\}$. Due to Lemma 10, it suffices to show that $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}}, (\pi_i, f_{id}), v \models \phi$ iff $\mathcal{I}_\mathfrak{M}, v' \models \mathsf{tr}(\phi)$ where $v'(z) = (v(z), i)$; this follows by structural induction on ϕ.

Toward the statement of the Lemma, assume $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}} \models \phi$, thus $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}}, (\pi_i, f), v \models \phi$ holds for all $\pi_i \in \Pi$, $f \in \mathbb{P}_\mathtt{E}$ and assignments v. Then, $\mathcal{I}_\mathfrak{M}, v' \models \mathsf{tr}(\phi)$ where $v'(z) = (v(z), i)$, thus $\mathcal{I}_\mathfrak{M}, v' \models \mathsf{tr}(\phi)$ for all v' where x and y have equal index, thus $\mathcal{I}_\mathfrak{M} \models \mathsf{Trans}_m(\phi)$. The converse direction proceeds in a similar fashion. □

Example 19. Revisiting Example 8, we compute $\mathsf{Trans}_2(E)$. With some simplifications, we obtain the following:

$$\begin{aligned}
\mathsf{Trans}_2(E) &= \forall x.\forall y.x \doteq y \to \mathsf{tr}(E) \\
\mathsf{tr}(E) &= \exists^{=1} x.(\phi^=_\mathtt{L}(x,y) \wedge \mathsf{tr}(E_0)) \\
&\quad \wedge \forall x.\,(\phi^=_\mathtt{L}(x,y) \to \mathsf{tr}(E_1)) \wedge \mathsf{tr}(E_2) \\
\mathsf{tr}(E_0) &= \exists y.\, x \doteq y \wedge \forall x.\, \phi^=_\mathtt{E}(x,y) \to \mathtt{Good}(x) \\
\mathsf{tr}(E_1) &= \mathsf{tr}(E_0) \vee \forall y.x \doteq y \to (\exists x.\, \phi^=_\mathtt{E}(x,y) \wedge \\
&\quad \mathtt{Best}(x) \wedge \forall y.\, \phi^=_\mathtt{L}(x,y) \to (\mathtt{Best}(y) \leftrightarrow x \doteq y)) \\
\mathsf{tr}(E_2) &= \forall x.\, x \doteq y \to \exists y.\big(\phi^=_\mathtt{E}(x,y) \wedge \\
&\quad \forall x.\phi^=_\mathtt{L}(x,y) \to (\mathtt{Good}(x) \vee \mathtt{Best}(x))\big)
\end{aligned}$$

One may verify that the structure in Figure 1(4) indeed satisfies $\mathsf{Trans}_2(E)$. Roughly, there exists x, e.g., $(\delta_0, 0)$, s.t. all elements of its $\mathtt{E}$-Type, i.e. $(\delta_0, 0)$ and $(\delta_0, 1)$, are $\mathtt{Good}$ (E_0), and for all elements, either they satisfy E_0 (like $(\delta_0, 0)$ and $(\delta_0, 1)$) or there is some element of their $\mathtt{E}$-Type (e.g., $(\delta_1, 0)$ for $\{\mathtt{E}_1\}$ and $(\delta_2, 1)$ for $\{\mathtt{E}_0, \mathtt{E}_1\}$) which is the only $\mathtt{Best}$ element on their layer (E_1). Finally there is some element (e.g., $(\delta_0, 0)$) such that all elements on its layer are $\mathtt{Good}$ (E_3). ◊

The last ingredient for our satisfiability translation is to ensure that the predicates in $\mathbf{P}_\mathtt{E}$ are indeed rigid.

Definition 20. We let $\phi^\ell_{rig\mathtt{E}}$ denote the $\mathcal{C}^2$ sentence $\forall x.\forall y.\mathtt{F}(x,y) \to \bigwedge_{0 \le i \le \ell} \mathtt{E}_i(x) \leftrightarrow \mathtt{E}_i(y)$.

Lemma 21. *Let $\mathfrak{M} = \langle \Delta, \Pi, \sigma, \gamma\rangle$ be a first-order standpoint structure for the signature $\langle \mathbf{P} \uplus \mathbf{P}_\mathtt{E}, \emptyset, \{*\}\rangle$. Then, all predicates from $\mathbf{P}_\mathtt{E} = \{\mathtt{E}_0, \ldots, \mathtt{E}_\ell\}$ are rigid iff $\mathcal{I}_\mathfrak{M} \models \phi^\ell_{rig\mathtt{E}}$.*

Theorem 22. *Let ϕ be an arbitrary frugal $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ sentence. Let $\ell = |FreeDia_\phi|$ and $m = \lceil |Dia_\phi| + \log_2(|Dia_\phi|)\rceil$. Then ϕ is satisfiable iff $\phi^m_{\mathrm{stack}} \wedge \phi^\ell_{rig\mathtt{E}} \wedge \mathsf{Trans}_m(\phi)$ is satisfiable.*

Proof. Assume that ϕ is satisfiable. Then by Theorem 11 there is a standpoint structure $\mathfrak{M}$ over the signature $\langle \mathbf{P} \uplus \mathbf{P}_\mathtt{E}, \emptyset, \emptyset\rangle$ with $|\mathbf{P}_\mathtt{E}| = \ell = |FreeDia_\phi|$ and $|\Pi| \le |Dia_\phi| \cdot 2^{|Dia_\phi|}$, such that $\langle\mathfrak{M}\rangle_{\mathbf{P}_\mathtt{E}}$ is a model of ϕ. Then, from Lemma 18, we have that $\mathcal{I}_\mathfrak{M} \models \mathsf{Trans}_m(\phi)$. Moreover, from Lemma 15, $\mathcal{I}_\mathfrak{M} \models \phi^m_{\mathrm{stack}}$. Finally, by Definition 9, all predicates from $\mathbf{P}_\mathtt{E}$ are rigid and thus by Lemma 21 $\mathcal{I}_\mathfrak{M} \models \phi^\ell_{rig\mathtt{E}}$.

For the other direction, assume that there is a model $\mathcal{I}$ over the signature $\langle\{\mathtt{F}, \mathtt{L}_0, \ldots, \mathtt{L}_{m-1}\} \uplus \mathbf{P}_\mathtt{E} \uplus \mathbf{P}, \emptyset\rangle$ such that $\mathcal{I} \models \phi^m_{\mathrm{stack}} \wedge \phi^\ell_{rig\mathtt{E}} \wedge \mathsf{Trans}_m(\phi)$. Then, by Theorem 16, $\mathcal{I}$ is isomorphic to a stacked interpretation $\mathcal{I}^\mathfrak{M}$ of some standpoint structure $\mathfrak{M}$ over $\langle \mathbf{P} \uplus \mathbf{P}_\mathtt{E}, \emptyset, \emptyset\rangle$ with 2^m precisifications. Moreover, since $\mathcal{I} \models \phi^\ell_{rig\mathtt{E}}$ then by Lemma 21 the predicates in $\mathbf{P}_{\mathrm{E}}$ are rigid. And since $\mathcal{I}^{\mathfrak{M}} \models \mathsf{Trans}_m(\phi)$, then by Lemma 18, we have that $\langle\mathfrak{M}\rangle_{\mathbf{P}_{\mathrm{E}}} \models \phi$ as desired. $\square$

Therefore (and taking into account Section 2.2), there is a polytime equisatisfiable translation from $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ to plain $\mathcal{C}^2$. On the other hand, every plain $\mathcal{C}^2$ formula is $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$, thus in view of the known NExpTime completeness of plain $\mathcal{C}^2$ (Pratt-Hartmann 2005), we arrive at the below corollary.

Corollary 23. *Satisfiability in monodic standpoint $\mathcal{C}^2$ is* NExpTime*-complete.*

Our equisatisfiable "standpoint removal" technique turns out to be robust under some variations. Let us call a monodic standpoint $\mathcal{C}^2$ formula ϕ *finitely satisfiable* if it has a model $\mathfrak{M} = \langle\Delta, \Pi, \sigma, \gamma\rangle$ where Δ is finite. It is easy to see that all equisatisfiable transformations in Section 2.2 are also "equi-finitely-satisfiable", because the underlying model transformations do not alter the domain whatsoever; the same holds for the argument behind Theorem 11. Last not least, the domain Δ' of the stacked interpretation $\mathcal{I}^{\mathfrak{M}} = (\Delta', \cdot^{\mathcal{I}})$ corresponding to a structure $\mathfrak{M} = \langle\Delta, \Pi, \sigma, \gamma\rangle$ is finite whenever Δ is (by the construction of Definition 13, we get $|\Delta'| = |\Delta| \cdot 2^m$). Thus the correspondence established in Theorem 22 also holds for finite satisfiability. On the other hand, finite satisfiability of plain $\mathcal{C}^2$ is also known to be NExpTime-complete (Pratt-Hartmann 2005), thus we obtain the following result.

Corollary 24. *Finite satisfiability in monodic standpoint $\mathcal{C}^2$ is* NExpTime*-complete.*

Last not least, more recently Benedikt, Kostylev, and Tan (2020) introduced two-variable FO with a more expressive version of counting quantifiers, denoted $\exists^S$, where S is any semilinear subset of $\mathbb{N} \cup \{\infty\}$. For example, by means of such quantifiers one can express quantities like "evenly many x" or also "infinitely many x", which go beyond what can be stated by the counting quantifiers of $\mathcal{C}^2$. Satisfiability of the ensuing logic, denoted $\mathrm{FO}^2_{\mathrm{Pres}}$ was established to be decidable in N2ExpTime and NExpTime hard. We note that our definitions, constructions, and arguments seamlessly extend from $\mathcal{C}^2$ to this logic, leading to the subsequent corollary.

Corollary 25. *Satisfiability and finite satisfiability in monodic standpoint* $\mathrm{FO}^2_{\mathrm{Pres}}$ *is in* N2ExpTime *and hard for* NExpTime.

We believe that – beyond their applicability to ontology reasoning as demonstrated in the next sections – the results presented here also provide significant novel insights for the area of first-order modal logics (Gabbay et al. 2005). As indicated by our naming, the subcase of $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ where the only standpoint expression used is $*$ coincides with the monodic fragment of modal counting two-variable FO with a S5 modal operator. While it has been observed earlier that restricting to the monodic setting is crucial for maintaining decidability in non-trivial combinations of FO fragments with modalities of varying kinds (Wolter and Zakharyaschev 2001), existing decidability results explicitly exclude FO fragments with equality or function symbols, which are notoriously harder, leaving such cases as an open question. We transcend this boundary, since $\mathcal{C}^2$ supports equality and unary functions (via axiomatising binary predicates as functional), and beyond mere decidability, we establish tight complexity bounds.

4 Application to Ontology Languages

We now show that adding monodic standpoints to popular ontology languages does not negatively affect the computational complexity of standard reasoning tasks. To this end, we begin by adding monodic standpoints to the description logic $\mathcal{ALCOIQB}^{\mathsf{Self}}$ which is very closely related to $\mathcal{C}^2$, and then we show how we can also accommodate role chain axioms, thus obtaining monodic standpoint extensions of $\mathcal{SHOIQB}_s$ and $\mathcal{SROIQB}_s$, which subsume the W3C ontology standards OWL 1 and OWL 2 DL respectively.[2] For the following, some familiarity with description logics (Baader et al. 2017; Rudolph 2011) will be very helpful.

4.1 Monodic Standpoint $\mathcal{ALCOIQB}^{\mathsf{Self}}$

We first introduce $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ obtained by adding monodic standpoints to the description logic $\mathcal{ALCOIQB}^{\mathsf{Self}}$. Just like $\mathcal{C}^2$ FOSL, $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ is based on a signature $\langle\mathbf{P}, \mathbf{C}, \mathbf{S}\rangle$ where $\mathbf{P}$ only contains unary and binary predicates, also referred to as *concept names* and *role names*, respectively. Based on these, we define the set $\mathbf{E}_{\mathrm{rol}}$ of *role expressions*

$$R, R' ::= \mathtt{R} \mid \mathtt{R}^- \mid \neg R \mid R \cap R'$$

with $\mathtt{R} \in \mathbf{P}$ binary, and the set $\mathbf{E}_{\mathrm{con}}$ of *concept expressions*

$$C, D ::= \mathtt{A} \mid \neg C \mid \{o\} \mid C \sqcap D \mid \geqslant nR.C \mid \exists R.\mathsf{Self} \mid \Diamond_{\mathsf{e}} C$$

with $\mathtt{A} \in \mathbf{P}$ unary, $o \in \mathbf{C}$, $n \in \mathbb{N}$, $\mathsf{e} \in \mathbf{E_S}$ (see Definition 1). Finally the set of $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ sentences is defined by

$$\phi, \psi ::= C \sqsubseteq D \mid \neg\phi \mid \phi \wedge \psi \mid \Diamond_{\mathsf{e}}\, \phi.$$

We introduce $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ with a minimalistic syntax, but note that all the usual description logic constructs can be introduced as syntactic sugar. For example, we obtain $\bot$ as $\mathtt{A} \sqcap \neg\mathtt{A}$ and $\top$ as $\neg\bot$; we may write $\exists R.C$ instead of $\geqslant 1R.C$ and also $\forall R.C$ instead of $\neg{\geqslant}1R.\neg C$; last not least we may write $\Box_{\mathsf{e}}\, C$ to denote $\neg \Diamond_{\mathsf{e}} \neg C$. We also remind the reader that other usual axiom types can all be rewritten into statements of the form $C \sqsubseteq D$ (referred to as *general concept inclusions*, short: *GCIs*) in the presence of *nominals* (i.e., expressions of the form $\{o\}$) and role expressions. Following DL naming conventions, a $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ sentence will be referred to as *TBox* if it is a conjunction of GCIs.

For later discussions, we single out some fragments of $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$: We obtain $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIF}}$ by excluding $\neg$ and $\cap$ from role expressions as well as disallowing concept expressions that use Self or $\geqslant k$ for $k \geq 2$, with the notable exception of axioms of the specific form $\top \sqsubseteq \neg{\geqslant}2\mathtt{F}.\top$ stating the functionality for binary predicates $\mathtt{F}$, which are then often abbreviated by $\mathit{func}(\mathtt{F})$. We obtain $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCO}}$ from $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIF}}$ by disallowing role expressions of the form $\mathtt{R}^-$ (known as *inverses*), and functionality axioms.

The semantics of standpoint-enhanced description logics is usually provided in a model-theoretic way using standpoint structures as in Definition 3 (Gómez Álvarez, Rudolph,

[2] The less mainstream letter $\mathcal{B}$ in the DL names refers to *boolean role constructors*, where $\mathcal{B}_s$ denotes boolean role constructors over simple roles only (see, e.g., Rudolph, Krötzsch, and Hitzler 2008). This modelling feature is not available in OWL 1 or OWL 2 DL.

and Strass 2022). For space reasons, we will instead define the semantics by directly providing a translation into $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$. To justify this "shortcut" we point out that said translation truthfully reflects the model-theoretic semantics of all earlier described standpoint-enhanced DLs and that existing translations from standpoint-free DLs to plain $\mathcal{C}^2$ (Kazakov 2008) naturally arise as a special case of ours.

The translation of a $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ sentence ϕ into a $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ sentence is obtained by replacing every GCI $C \sqsubseteq D$ inside ϕ by $\forall x.\big(\mathsf{ctrans}(x,C) \rightarrow \mathsf{ctrans}(x,D)\big)$, where $\mathsf{ctrans} : \{x,y\} \times \mathbf{E}_{\mathrm{con}} \to \mathbb{S}_{\mathcal{C}^2}$ is inductively defined:

$$\begin{aligned}
\mathsf{ctrans}(z,\mathtt{A}) &= \mathtt{A}(z)\\
\mathsf{ctrans}(z,\neg C) &= \neg\mathsf{ctrans}(z,C)\\
\mathsf{ctrans}(z,\{o\}) &= z\dot{=}o\\
\mathsf{ctrans}(z,C\sqcap D) &= \mathsf{ctrans}(z,C)\wedge\mathsf{ctrans}(z,D)\\
\mathsf{ctrans}(x,\geqslant nR.C) &= \exists^{\geq n}y.\mathsf{rtrans}(x,y,R)\wedge\mathsf{ctrans}(y,C)\\
\mathsf{ctrans}(y,\geqslant nR.C) &= \exists^{\geq n}x.\mathsf{rtrans}(y,x,R)\wedge\mathsf{ctrans}(x,C)\\
\mathsf{ctrans}(z,\exists R.\mathsf{Self}) &= \mathsf{rtrans}(z,z,R)\\
\mathsf{ctrans}(z,\Diamond_{\mathsf{e}}C) &= \exists z.\Diamond_{\mathsf{e}}\,\mathsf{ctrans}(z,C),
\end{aligned}$$

using $\mathsf{rtrans} : \{x,y\}\times\{x,y\}\times\mathbf{E}_{\mathrm{rol}} \to \mathbb{S}_{\mathcal{C}^2}$ defined by

$$\begin{aligned}
\mathsf{rtrans}(z,z',\mathtt{R}) &= \mathtt{R}(z,z')\\
\mathsf{rtrans}(z,z',\mathtt{R}^-) &= \mathtt{R}(z',z)\\
\mathsf{rtrans}(z,z',\neg R) &= \neg\mathsf{rtrans}(z,z',R)\\
\mathsf{rtrans}(z,R\cap R') &= \mathsf{rtrans}(z,z',R)\wedge\mathsf{rtrans}(z,z',R).
\end{aligned}$$

It can be readily checked that the translation described is computable in polytime (hence polynomial in output) and indeed yields a $\mathbb{S}^{\mathrm{mon}}_{\mathcal{C}^2}$ sentence. Therefore, and in view of the fact that satisfiability is already NExpTime-hard for the standpoint-free sublogic $\mathcal{ALCOIF}$ (Tobies 2000), we obtain the following tight complexity bounds.

Theorem 26. *Checking satisfiability of* $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$ *sentences is* NExpTime*-complete.*

4.2 Adding Role Chain Axioms

In order to fully cover the web ontology languages OWL 1 and OWL 2 DL, we need to extend our formalism by so-called *role chain axioms*, arriving at the description logics $\mathcal{SHOIQB}_s$ (when allowing just role chain axioms expressing transitivity such as $\mathtt{FriendOf}\circ\mathtt{FriendOf}\sqsubseteq\mathtt{FriendOf}$) or $\mathcal{SROIQB}_s$ (when admitting more complex forms like $\mathtt{FriendOf}\circ\mathtt{EnemyOf}\sqsubseteq\mathtt{EnemyOf}$), respectively. Luckily, by combining known standpoint encoding tricks (Gómez Álvarez, Rudolph, and Strass 2022) and removal techniques for role-chain axioms (Kazakov 2008; Demri and de Nivelle 2005) with some novel ideas, it is possible to translate $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SHOIQB}_s}$ and $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SROIQB}_s}$ sentences back into $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCOIQB}^{\mathsf{Self}}}$. The translation is polynomial for $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SHOIQB}_s}$ and exponential for $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SROIQB}_s}$.

Theorem 27. *Checking satisfiability of* $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SHOIQB}_s}$ *sentences is* NExpTime*-complete. Checking satisfiability of* $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SROIQB}_s}$ *sentences is* N2ExpTime*-complete.*

Therein, the hardness part for $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SROIQB}_s}$ follows from the known N2ExpTime hardness of its fragment $\mathcal{SROIQ}$ (Kazakov 2008). This concludes our argument that adding monodic standpoints to OWL 1 and OWL 2 does not increase complexity of standard reasoning tasks.

5 Nominals Cause Trouble

We finish our considerations by presenting two results that provide some context for our main results and corroborate the intuition (cf. Gómez Álvarez, Rudolph, and Strass 2023b as well as Gómez Álvarez and Rudolph 2024) that the interplay of nominals and standpoint modalities is particularly troublesome for reasoning. To this end, we will use the tiling problem in two variations, which we introduce next.

A *tiling system* $\mathbb{T} = \langle k, H, V\rangle$ consists of a number $k \in \mathbb{N}$ indicating the number of tiles, and horziontal and vertical compatibility relations $H, V \subseteq \{1,\ldots,k\}\times\{1,\ldots,k\}$. For a downward-closed set $S \subseteq \mathbb{N}$ of natural numbers, a $\mathbb{T}$-*tiling* of $S \times S$ with initial condition $\langle t_0,\ldots t_n\rangle \in \{1,\ldots,k\}^n$ for some $n \in S$ is a mapping $\mathsf{tile} : S\times S \to \{1,\ldots,k\}$ such that $\mathsf{tile}(i,0) = t_i$ for $i \in \{1,\ldots,n\}$, and for all $i \in S$ with $i{+}1 \in S$ and all $j \in S$ holds $(\mathsf{tile}(i,j), \mathsf{tile}(i{+}1,j)) \in H$ as well as $(\mathsf{tile}(j,i), \mathsf{tile}(j,i{+}1)) \in V$. We recall the following:

- There is a tiling system $\mathbb{T}_{\mathrm{exp}}$ such that the following problem is NExpTime-hard: Given an initial condition of size n, does there exist a corresponding $\mathbb{T}_{\mathrm{exp}}$-tiling of $\{0,\ldots,2^n-1\}\times\{0,\ldots,2^n-1\}$?
- There is a tiling system $\mathbb{T}_{\mathrm{und}}$ such that the following problem is undecidable: Given an initial condition of size n, does there exist a corresponding $\mathbb{T}_{\mathrm{und}}$-tiling of $\mathbb{N}\times\mathbb{N}$?

5.1 NExpTime Hardness for $\mathcal{ALCO}$ TBoxes

From prior works, it is known that monodic standpoint $\mathcal{SHIQ}$, a sublogic of $\mathbb{S}^{\mathrm{mon}}_{\mathcal{SHOIQB}_s}$ has an ExpTime-complete satisfiability problem (Gómez Álvarez and Rudolph 2024), which means that the complexity of $\mathcal{SHIQ}$ is unaltered if monodic standpoints are added. Two other popular ExpTime-complete sub-DLs of $\mathcal{SHOIQB}_s$ (incomparable to $\mathcal{SHIQ}$) are $\mathcal{SHIO}$ and $\mathcal{SHOQ}$ (Hladik and Model 2004; Glimm, Horrocks, and Sattler 2008). This poses the question if adding monodic standpoints to these DLs preserves ExpTime reasoning, like it does for $\mathcal{SHIQ}$.

Interestingly, we can answer this question in the negative (unless NExpTime = ExpTime), and identify nominals as the common culprit by showing that satisfiability even of monodic standpoint TBoxes in $\mathcal{ALCO}$ (a restricted sublogic of both $\mathcal{SHIO}$ and $\mathcal{SHOQ}$) is already NExpTime-hard.

To this end, we provide a polynomial reduction from the first of the two above tiling problems to the satisfiability problem of a $\mathbb{S}^{\mathrm{mon}}_{\mathcal{ALCO}}$ TBox of size polynomial in n, using just one nominal concept $\{o\}$. We use atomic concepts $\mathtt{T}_1,\ldots,\mathtt{T}_k$ for the k tiles and atomic concepts $\mathtt{X}_1,\ldots,\mathtt{X}_n$ as well as $\mathtt{Y}_1,\ldots,\mathtt{Y}_n$ to encode grid x- and y-coordinates in binary. First, we declare all these concepts as "almost rigid": they hold uniformly across precisifications for all elements but o.

$$\neg\{o\}\sqcap\mathtt{T}_\ell \sqsubseteq \Box_*\mathtt{T}_\ell \quad \neg\{o\}\sqcap\mathtt{X}_i \sqsubseteq \Box_*\mathtt{X}_i \quad \neg\{o\}\sqcap\mathtt{Y}_i \sqsubseteq \Box_*\mathtt{Y}_i$$

Above and below, we let i range from 1 to n and let ℓ range from 1 to k. Next, we ensure that, in every precisification, every non-o element with x-coordinate (y-coordinate) smaller than 2^n-1 has a horizontal (vertical) neighbour with that coordinate incremented and the same y-coordinate (x-coordinate). We let j range from 1 to $i-1$.

$$\neg\{o\} \sqcap \bigsqcup_i \neg\mathtt{X}_i \sqsubseteq \exists\mathtt{H}.\neg\{o\} \qquad \neg\{o\} \sqcap \bigsqcup_i \neg\mathtt{Y}_i \sqsubseteq \exists\mathtt{V}.\neg\{o\}$$
$$\mathtt{X}_i \sqcap \neg\mathtt{X}_j \sqsubseteq \forall\mathtt{H}.\mathtt{X}_i \qquad \mathtt{Y}_i \sqcap \neg\mathtt{Y}_j \sqsubseteq \forall\mathtt{V}.\mathtt{Y}_i$$
$$\neg\mathtt{X}_i \sqcap \neg\mathtt{X}_j \sqsubseteq \forall\mathtt{H}.\neg\mathtt{X}_i \qquad \neg\mathtt{Y}_i \sqcap \neg\mathtt{Y}_j \sqsubseteq \forall\mathtt{V}.\neg\mathtt{Y}_i$$
$$\mathtt{X}_i \sqcap \textstyle\bigsqcap_j \mathtt{X}_j \sqsubseteq \forall\mathtt{H}.\neg\mathtt{X}_i \qquad \mathtt{Y}_i \sqcap \bigsqcap_j \mathtt{Y}_j \sqsubseteq \forall\mathtt{V}.\neg\mathtt{Y}_i$$
$$\neg\mathtt{X}_i \sqcap \textstyle\bigsqcap_j \mathtt{X}_j \sqsubseteq \forall\mathtt{H}.\mathtt{X}_i \qquad \neg\mathtt{Y}_i \sqcap \bigsqcap_j \mathtt{Y}_j \sqsubseteq \forall\mathtt{V}.\mathtt{Y}_i$$

$$\mathtt{Y}_i \sqsubseteq \forall\mathtt{H}.\mathtt{Y}_i \quad \neg\mathtt{Y}_i \sqsubseteq \forall\mathtt{H}.\neg\mathtt{Y}_i \quad \mathtt{X}_i \sqsubseteq \forall\mathtt{V}.\mathtt{X}_i \quad \neg\mathtt{X}_i \sqsubseteq \forall\mathtt{V}.\neg\mathtt{X}_i$$

We next ensure that there exists a non-o element with x and y set to zero, which together with its horizontal neighbours realises the initial condition $\langle t_0, \dots t_n\rangle \in \{1, \dots, k\}^n$.

$$\top \sqsubseteq \exists\mathtt{R}.(\neg\{o\} \sqcap \neg\mathtt{X}_1 \sqcap \dots \sqcap \neg\mathtt{X}_n \sqcap \neg\mathtt{Y}_1 \sqcap \dots \sqcap \neg\mathtt{Y}_n \sqcap$$
$$\mathtt{T}_{t_1} \sqcap \forall\mathtt{H}.(\mathtt{T}_{t_2} \sqcap \forall\mathtt{H}.(\dots(\mathtt{T}_{t_{n-1}} \sqcap \forall\mathtt{H}.\mathtt{T}_{t_n})\dots))$$

For every non-o element, there exists some precisification wherein it is P-linked to o and propagates its x- and y-coordinate as well as its tile assignment via this link to o.

$$\neg\{o\} \sqsubseteq \Diamond_* \exists\mathtt{P}.\{o\} \qquad \mathtt{X}_i \sqsubseteq \forall\mathtt{P}.\mathtt{X}_i \qquad \mathtt{Y}_i \sqsubseteq \forall\mathtt{P}.\mathtt{Y}_i$$
$$\mathtt{T}_i \sqsubseteq \forall\mathtt{P}.\mathtt{T}_i \qquad \neg\mathtt{X}_i \sqsubseteq \forall\mathtt{P}.\neg\mathtt{X}_i \qquad \neg\mathtt{Y}_i \sqsubseteq \forall\mathtt{P}.\neg\mathtt{Y}_i$$

In every precisification, every non-o element is P′-linked to o and, should its assigned x- and y-coordinate coincide with those assigned to o, then its tile-assignment will be made to coincide with the one of o as well.

$$\neg\{o\} \sqsubseteq \exists\mathtt{P}'.\{o\}$$
$$\exists\mathtt{P}'.\mathtt{T}_\ell \sqcap \textstyle\bigsqcap_i \big((\mathtt{X}_i \sqcap \exists\mathtt{P}'.\mathtt{X}_i) \sqcup (\neg\mathtt{X}_i \sqcap \exists\mathtt{P}'.\neg\mathtt{X}_i)\big)$$
$$\sqcap \textstyle\bigsqcap_i \big((\mathtt{Y}_i \sqcap \exists\mathtt{P}'.\mathtt{Y}_i) \sqcup (\neg\mathtt{Y}_i \sqcap \exists\mathtt{P}'.\neg\mathtt{Y}_i)\big) \sqsubseteq \mathtt{T}_\ell$$

Note that this way, the tile assignments will be synchronized between all elements carrying the same coordinates. We finally make sure that in every precisification, every domain element must be assigned a tile. Moreover the H- and V-neighbouring pairs of elements must conform with the horizontal and vertical compatibility relation.

$$\top \sqsubseteq \mathtt{T}_1 \sqcup \dots \sqcup \mathtt{T}_k$$
$$\mathtt{T}_\ell \sqsubseteq \forall\mathtt{H}.\neg\mathtt{T}_{\ell'} \quad \text{for } (\ell,\ell') \in \{1,\dots,k\} \times \{1,\dots,k\} \setminus H$$
$$\mathtt{T}_\ell \sqsubseteq \forall\mathtt{V}.\neg\mathtt{T}_{\ell'} \quad \text{for } (\ell,\ell') \in \{1,\dots,k\} \times \{1,\dots,k\} \setminus V$$

This finishes the description of the TBox (obtained by taking the conjunction of all the introduced GCIs). We note that these axioms do not enforce the H and V relation to form a proper grid (in any precisification). Rather, the axioms ensure that for any two horizontally (vertically) neighbouring coordinate pairs, there exists a H-connected (V-connected) pair of domain elements carrying said coordinates. Since the tile assignments are rigid (except for o) and synchronized over all elements carrrying equal coordinates, this suffices to ensure that satisfiability of our TBox coincides with the existence of a $\mathbb{T}_{\text{exp}}$-tiling, so we obtain the following theorem.

Theorem 28. *In any sublogic of $\mathbb{S}^{\text{mon}}_{\mathcal{SHOIQB}_s}$ that subsumes $\mathbb{S}^{\text{mon}}_{\mathcal{ALCO}}$ TBoxes, satisfiability is* NEXPTIME*-complete.*

5.2 Lifting Monodicity Causes Undecidability

A crucial restriction underlying all logical formalisms that we have considered so far is monodicity requiring that modal operators can only be put in front of subformulae with at most one free variable. The arguably mildest way of lifting monodicity is by imposing that one distinguished binary predicate, say E, must be rigidly interpreted. Note that rigidity of a binary predicate E could be expressed by the FOSL formula $\forall x,y.\big(\mathtt{E}(x,y) \to \Box_* \mathtt{E}(x,y)\big)$, which is not monodic. By a reduction from the second of the above tiling problems, we show that allowing one rigid binary predicate causes undecidability even for a sublogic of $\mathbb{S}^{\text{mon}}_{\mathcal{SHOIQB}_s}$.

Theorem 29. *Satisfiability of $\mathbb{S}^{\text{mon}}_{\mathcal{ALCOIF}}$ TBoxes with one rigid binary predicate is undecidable, even when using just one nominal and one functionality statement.*

For space reasons, we only briefly provide a set of GCIs enforcing an $\mathbb{N} \times \mathbb{N}$ grid, noting that a $\mathbb{T}_{\text{und}}$-tiling on top can be obtained very similarly to the previous case. Let E be the distinguished rigid binary predicate, which we use to represent both horizontal and vertical grid connections (distinguishing them via extra unary rigid "markers" for even/odd grid rows). Let *func*(Point) specify that the binary "pointer predicate" Point is functional and put the following GCIs:

$$\top \sqsubseteq \exists\mathtt{E}.\Box_*\,\mathtt{Even} \sqcap \exists\mathtt{E}.\Box_*\,\mathtt{Odd} \sqcap \Diamond_*\,\mathtt{Pick}$$
$$\mathtt{Pick} \sqsubseteq \forall\mathtt{E}.\big(\neg\mathtt{Even} \sqcup \forall\mathtt{E}.(\neg\mathtt{Odd} \sqcup \exists\mathtt{Point}^-.\{o\})\big)$$
$$\mathtt{Pick} \sqsubseteq \forall\mathtt{E}.\big(\neg\mathtt{Odd} \sqcup \forall\mathtt{E}.(\neg\mathtt{Even} \sqcup \exists\mathtt{Point}^-.\{o\})\big)$$

In a nutshell, these GCIs ensure that every grid element δ will be Picked in some precisification, and in that precisification the upper neighbour of δ's right neighbour is forced to coincide with the right neighbour of δ's upper neighbour, by having both being "functionally Pointed to" from o.

We note that this finding contrasts with a positive result by Artale, Lutz, and Toman (2007), according to which – in our nomenclature – the satisfiability of TBoxes over $\mathcal{ALCIQ}$ with arbitrarily many rigid roles and one S5 modality allowed to occur in front of concept *and role expressions*, is decidable in 2EXPTIME. Once more, this underlines the previous observation that while counting and inverses go reasonably well with standpoint modalities, nominals do not.

6 Conclusions and Future Work

We have shown that monodic standpoints can be added to $\mathcal{C}^2$ without increasing the NEXPTIME reasoning complexity. We obtained this result by establishing a polynomial translation into plain $\mathcal{C}^2$, whose justification required rather elaborate model-theoretic constructions and arguments. On one hand, this finding advances the research into first-order modal logics, since our result subsumes the case of monodic S5 over $\mathcal{C}^2$ and even generalises to logics with more expressive counting. On the other hand, we leveraged the obtained result to prove that very expressive DLs subsuming popular W3C-standardized ontology languages can be endowed with monodic standpoints in a complexity-neutral way. We finally showed that in the presence of nominal concepts, NEXPTIME-hardness already arises for much less expressive DLs, and lifting monodicity even incurs undecidability.

For future work, it would be interesting to investigate the data complexity of our formalism. Also it would be advantageous to find translations from versions of monodic standpoint OWL into plain OWL rather than $\mathcal{C}^2$, since this would allow to deploy existing highly optimized OWL reasoners for standpoint-aware ontological reasoning. While our results show that there are no complexity-theoretic barriers for this, our current translation approach heavily relies on features of $\mathcal{C}^2$ that are beyond the capabilities of plain OWL.

Acknowledgements

We thank the anonymous reviewers for their constructive and useful comments, which helped us to improve the paper. This work is funded by the Agence Nationale de la Recherche, France (ANR) in project ANR-25-CE23-2478-01 (SPaRK), Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) grant 389792660 as part of TRR 248 – CPEC (Center for Perspicuous Systems), as well as the German Bundesministerium für Forschung, Technologie und Raumfahrt (BMFTR, Federal Ministry of Research, Technology and Space) and Deutscher Akademischer Austauschdienst (DAAD, German Academic Exchange Service) in project 57616814, SECAI (School of Embedded and Composite AI).

References

Artale, A.; Lutz, C.; and Toman, D. 2007. A description logic of change. In Veloso, M. M., ed., *Proc. 20th Int. Joint Conf. on Artificial Intelligence*, 218–223. AAAI Press.

Baader, F.; Horrocks, I.; Lutz, C.; and Sattler, U. 2017. *An Introduction to Description Logic*. Cambridge University Press.

Bao, J.; Calvanese, D.; Grau, B. C.; Dzbor, M.; Fokoue, A.; Golbreich, C.; Hawke, S.; Herman, I.; Hoekstra, R.; Horrocks, I.; Kendall, E.; Krötzsch, M.; Lutz, C.; McGuinness, D. L.; Motik, B.; Pan, J.; Parsia, B.; Patel-Schneider, P. F.; Rudolph, S.; Ruttenberg, A.; Sattler, U.; Schneider, M.; Smith, M.; Wallace, E.; Wu, Z.; and Zimmermann, A. 2009. OWL 2 Web Ontology Language: Document Overview. W3C Recommendation. http://www.w3.org/TR/owl2-overview/.

Benedikt, M.; Kostylev, E. V.; and Tan, T. 2020. Two variable logic with ultimately periodic counting. In Czumaj, A.; Dawar, A.; and Merelli, E., eds., *Proc. 47th Int. Colloquium on Automata, Languages and Programming*, volume 168 of *LIPIcs*, 112:1–112:16. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

Demri, S., and de Nivelle, H. 2005. Deciding regular grammar logics with converse through first-order logic. *Journal of Logic Language and Information* 14(3):289–329.

Gabbay, D. M.; Kurucz, A.; Wolter, F.; and Zakharyaschev, M. 2005. Many-dimensional modal logics: Theory and applications. *Studia Logica* 81(1):147–150.

Glimm, B.; Horrocks, I.; and Sattler, U. 2008. Unions of conjunctive queries in SHOQ. In Brewka, G., and Lang, J., eds., *Proc. 11th Int. Conf. on Principles of Knowledge Representation and Reasoning*, 252–262. AAAI Press.

Gómez Álvarez, L., and Rudolph, S. 2021. Standpoint logic: Multi-perspective knowledge representation. In Neuhaus, F., and Brodaric, B., eds., *Proc. 12th Int. Conf. on Formal Ontology in Information Systems*, volume 344 of *FAIA*, 3–17. IOS Press.

Gómez Álvarez, L., and Rudolph, S. 2025. Putting Perspective into OWL [sic]: Complexity-Neutral Standpoint Reasoning for Ontology Languages via Monodic S5 over Counting Two-Variable First-Order Logic (Extended Version with Appendix) .

Gómez Álvarez, L., and Rudolph, S. 2024. Reasoning in SHIQ with axiom- and concept-level standpoint modalities. In Marquis, P.; Ortiz, M.; and Pagnucco, M., eds., *Proc. 21st Int. Conf. on Principles of Knowledge Representation and Reasoning*, 383–393. ijcai.org.

Gómez Álvarez, L.; Rudolph, S.; and Strass, H. 2022. How to agree to disagree - managing ontological perspectives using standpoint logic. In Sattler, U.; Hogan, A.; Keet, C. M.; Presutti, V.; Almeida, J. P. A.; Takeda, H.; Monnin, P.; Pirrò, G.; and d'Amato, C., eds., *Proc. 21st Int. Semantic Web Conf.*, volume 13489 of *LNCS*, 125–141. Springer.

Gómez Álvarez, L.; Rudolph, S.; and Strass, H. 2023a. Pushing the boundaries of tractable multiperspective reasoning: A deduction calculus for standpoint EL+. In Marquis, P.; Son, T. C.; and Kern-Isberner, G., eds., *Proc. 20th Int. Conf. on Principles of Knowledge Representation and Reasoning*, 333–343. ijcai.org.

Gómez Álvarez, L.; Rudolph, S.; and Strass, H. 2023b. Tractable diversity: Scalable multiperspective ontology management via standpoint EL. In Elkind, E., ed., *Proc. 32nd Int. Joint Conf. on Artificial Intelligence*, 3258–3267. ijcai.org.

Hladik, J., and Model, J. 2004. Tableau systems for SHIO and SHIQ. In Haarslev, V., and Möller, R., eds., *Proc. 2004 Int. Workshop on Description Logics*, volume 104 of *CEUR Workshop Proceedings*. CEUR-WS.org.

Kazakov, Y. 2008. RIQ and SROIQ are harder than SHOIQ. In Brewka, G., and Lang, J., eds., *Proc. 11th Int. Conf. on Principles of Knowledge Representation and Reasoning*, 274–284. AAAI Press.

Kurucz, A.; Wolter, F.; and Zakharyaschev, M. 2023. Definitions and (uniform) interpolants in first-order modal logic. In Marquis, P.; Son, T. C.; and Kern-Isberner, G., eds., *Proc. 20th Int. Conf. on Principles of Knowledge Representation and Reasoning*, 417–428. ijcai.org.

Pratt-Hartmann, I. 2005. Complexity of the two-variable fragment with counting quantifiers. *Journal of Logic, Language, and Information* 14(3):369–395.

Rudolph, S.; Krötzsch, M.; and Hitzler, P. 2008. Cheap Boolean role constructors for description logics. In Hölldobler, S.; Lutz, C.; and Wansing, H., eds., *Proc. 11th European Conf. on Logics in Artificial Intelligence*, volume 5293 of *LNCS*, 362–374. Springer.

Rudolph, S. 2011. Foundations of description logics. In Polleres, A.; d'Amato, C.; Arenas, M.; Handschuh, S.; Kroner, P.; Ossowski, S.; and Patel-Schneider, P. F., eds., *Reasoning Web. Semantic Technologies for the Web of Data - 7th International Summer School 2011, Tutorial Lectures*, volume 6848 of *LNCS*, 76–136. Springer.

Tobies, S. 2000. The complexity of reasoning with cardinality restrictions and nominals in expressive description logics. *Journal of Artificial Intelligence Research* 12:199–217.

Wolter, F., and Zakharyaschev, M. 2001. Decidable fragments of first-order modal logics. *Journal of Symbolic Logic* 66(3):1415–1438.