

LTL over Description Logic Axioms

Franz Baader

TU Dresden, Germany

Silvio Ghilardi

Università degli Studi di Milano, Italy

and

Carsten Lutz

University of Bremen, Germany

Most of the research on temporalized Description Logics (DLs) has concentrated on the case where temporal operators can be applied to concepts, and sometimes additionally to TBox axioms and ABox assertions. The aim of this paper is to study temporalized DLs where temporal operators on TBox axioms and ABox assertions are available, but temporal operators on concepts are not. While the main application of existing temporalized DLs is the representation of conceptual models that explicitly incorporate temporal aspects, the family of DLs studied in this paper addresses applications that focus on the temporal evolution of data and of ontologies. Our results show that disallowing temporal operators on concepts can significantly decrease the complexity of reasoning. In particular, reasoning with rigid roles (whose interpretation does not change over time) is typically undecidable without such a syntactic restriction, whereas our logics are decidable in elementary time even in the presence of rigid roles. We analyze the effects on computational complexity of dropping rigid roles, dropping rigid concepts, replacing temporal TBoxes with global ones, and restricting the set of available temporal operators. In this way, we obtain a novel family of temporalized DLs whose complexity ranges from 2-EXPTIME-complete via NEXPTIME-complete to EXPTIME-complete.

Categories and Subject Descriptors: I.2.4 [Artificial Intelligence]: Knowledge Representation Formalisms and Methods; F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic

General Terms: Knowledge Representation, Complexity

Additional Key Words and Phrases: Description Logics, Temporal Extensions

1. INTRODUCTION

Description logics (DLs) [Baader et al. 2003] are a family of logic-based knowledge representation formalisms, which are employed in various application domains, such as natural language processing, configuration, databases, and bio-medical ontologies, but their most notable success so far is the adoption of the DL-based language OWL [Horrocks et al. 2003] as the standard ontology language for the web. In many applications of DLs, such as the use of DLs as ontology languages or conceptual modeling languages, being able to represent dynamic aspects of the application domain would be quite useful. This is, for instance, the case if one wants to use DLs as conceptual modeling languages for temporal databases as proposed in [Artale et al. 2002]. In this area, dynamics is introduced via evolution constraints such as “every ordered item will eventually be a shipped item.” Another example are medical ontologies, where the faithful representation of concepts often requires the description of temporal patterns. As a simple example, consider the concept “Concussion with

no loss of consciousness,” which is modeled as a primitive (i.e., not further defined) concept in the well-known medical ontology SNOMED CT.¹ As argued in [Schulz et al. 2006], a correct representation of this concept should actually say that, after the concussion, the patient remained conscious until the examination.

Since the expressiveness of pure DLs is not sufficient to describe such temporal patterns, a multitude of temporal extensions of DLs have been investigated in the literature.² These include approaches as diverse as the combination of DLs with Halpern and Shoham’s logic of time intervals [Schmiedel 1990], formalisms inspired by action logics [Artale and Franconi 1998], the treatment of time points and intervals as a concrete domains [Lutz 2001], and the combination of standard DLs with standard (propositional) temporal logics into logics with a two-dimensional semantics, where one dimension is for time and the other for the DL domain [Schild 1993; Wolter and Zakharyashev 1999; Gabbay et al. 2003]. In this paper, we follow the last approach, where we use the basic DL \mathcal{ALC} [Schmidt-Schauß and Smolka 1991] in the DL component and linear temporal logic (LTL) [Pnueli 1977] (sometimes also called propositional temporal logic (PTL) [Gabbay et al. 2003]) in the temporal component. However, even after the DL and the temporal logic to be combined have been fixed, there remain several degrees of freedom when defining the resulting temporalized DL.

On the one hand, one must decide to which pieces of syntax temporal operators can be applied. Temporal operators may be allowed to occur within concept descriptions, as required by the above example of a concussion with no loss of consciousness, which could be defined using the until-operator U of LTL as follows:

$$\exists \text{finding.Concussion} \sqcap \text{Conscious} U \exists \text{procedure.Examination}. \quad (1)$$

Alternatively or in addition, temporal operators may be applied to TBox axioms (i.e., general concept inclusions, GCIs) and/or to ABox assertions. For example, the temporalized TBox axiom

$$\diamond \square (\text{USCitizen} \sqsubseteq \exists \text{insured.by.HealthInsurer})$$

says that there is a future time point from which on US citizens will always have health insurance, and the formula ψ :

$$\diamond ((\text{BOB} : \exists \text{finding.Concussion}) \wedge (\text{BOB} : \text{Conscious}) U (\text{BOB} : \exists \text{procedure.Examination})) \quad (2)$$

says that, sometime in the future, Bob will have a concussion with no loss of consciousness between the concussion and the examination.

On the other hand, one must decide whether one wants to have rigid concepts and/or roles, i.e., concepts/roles whose interpretation does not vary over time. For example, the concept `Human` and roles such as `has_bloodtype` and `has_genedefect` should probably be rigid since a human being will stay a human being and have the same blood type and gene defects over his/her life-time, whereas the concept

¹This ontology underlies the standardized medical terminology used in the health-care systems of various countries such as Australia, USA, and UK; see <http://www.ihtsdo.org/our-standards/>

²For a more thorough survey of the literature on temporalized DLs, see the survey papers [Artale and Franconi 2000; 2001; Lutz et al. 2008].

LTL operators on concepts	LTL operators on TBox/ABox axioms	Rigid roles	
X			EXPTIME-complete [Schild 1993]
X	X		EXPSpace-complete [Gabbay et al. 2003]
X		X	undecidable [Gabbay et al. 2003]

Table I. Some known complexity results for combinations of \mathcal{ALC} and LTL

Conscious should be flexible (i.e., not rigid) since someone who is conscious at the moment need not always be conscious. Similarly, `insured_by` should be modeled as a flexible role. Using a logic that cannot enforce rigidity of concepts/roles may result in unintended models, and thus prevent certain useful inferences to be drawn. For example, the concept description $\exists \text{has_bloodtype.AB} \sqcap \diamond (\forall \text{has_bloodtype.A})$ is only unsatisfiable w.r.t. the TBox $\text{AB} \sqsubseteq \neg \text{A}$ if both `has_bloodtype` and (at least one of) `AB` and `A` are rigid.

Related work. The combination of \mathcal{ALC} with LTL was first considered by Schild [Schild 1993] and, since then, has developed into a lively research area recently surveyed in [Lutz et al. 2008]. In the temporalized DL proposed by Schild, temporal operators can be applied to concept descriptions, but not to TBox axioms (and ABoxes are not considered at all). Rigid concepts are definable in this logic, but rigid roles are not. Schild observes that his logic behaves similarly to the so-called *fusion* of \mathcal{ALC} and LTL,³ which shows that the interaction between the \mathcal{ALC} component and the LTL component is limited. This observation forms the basis for Schild’s proof that reasoning in his logic is EXPTIME-complete.

The combination of (extensions of) \mathcal{ALC} and LTL in which temporal operators are applied to concept descriptions, TBox axioms, and ABox assertions has been studied by Wolter, Zakharyashev, and others (see, e.g., [Wolter and Zakharyashev 1999; Gabbay et al. 2003]). In this more general setting, the interaction between the DL component and LTL is stronger, and reasoning is EXPSpace-complete. As in Schild’s logic, rigid concepts can be defined, but rigid roles cannot. In fact, as shown in [Gabbay et al. 2003], the addition of rigid roles causes undecidability. This already holds for concept satisfiability w.r.t. a global TBox (i.e., where the same TBox axioms must hold at all time points), without ABoxes, and with only a single rigid role. Decidability can be regained by dropping TBoxes altogether, but the decision problem is still hard for non-elementary time [Gabbay et al. 2003]. The most relevant results mentioned up to here are summarized in Table I.

Decidable combinations of DLs and temporal logics that allow for rigid roles can be obtained by strongly restricting either the temporal or the DL component. In [Artale et al. 2007], temporal operators can be applied to concept descriptions, TBoxes are global, and there are no ABoxes. The reason for decidability (more precisely, 2-EXPTIME-completeness) also in the presence of rigid roles is that the only available temporal operators are an undirected diamond that says “at some time point” and an undirected box that says “at all time points.” Here, undirected

³Note, however, that Schild’s proof of this fact is incorrect; a correct proof can be found in [Lutz et al. 2008].

means that these operators cannot discriminate between the past, the future, and the current time point. The restriction imposed in [Artale et al. 2007] is different: the temporal component is LTL, but \mathcal{ALC} is replaced with the lightweight DL $\text{DL-Lite}_{\text{bool}}$. Temporal operators can be applied to concept descriptions, TBoxes, and ABoxes. Here, it is the weak expressive power of the DL component that is responsible for decidability (more precisely, $\text{EXPSpace-completeness}$) of reasoning also in the presence of rigid roles. In [Artale et al. 2009], the expressive power is restricted further by (i) admitting only global TBoxes and (ii) restricting LTL to the operators \diamond and \square , with the consequence that the complexity of reasoning drops to NP-complete even in the presence of rigid roles. In [Artale et al. 2007], it is shown that concept subsumption w.r.t. global TBoxes and with rigid roles is undecidable already in the lightweight description logic \mathcal{EL} , which provides only for the constructors conjunction and existential restriction.

Our contribution. In this paper, we study combinations of \mathcal{ALC} with LTL in which temporal operators are allowed to occur only in front of axioms (i.e., ABox assertions and TBox axioms), but not as concept constructors. We show that reasoning becomes simpler in this setting: with rigid roles, satisfiability is decidable (more precisely: 2-EXPTIME-complete); without rigid roles (but with rigid concepts), the complexity decreases to NEXPTIME-complete; and without any rigid symbols, it decreases further to EXPTIME-complete (i.e., the same complexity as reasoning in \mathcal{ALC} alone). We also consider two other ways of decreasing the complexity of satisfiability to EXPTIME. On the one hand, satisfiability without rigid roles (but with rigid concepts) becomes EXPTIME-complete if GCIs can occur only as global axioms that must hold in every temporal world. Note that, in this case, ABox assertions are *not* assumed to be global, i.e., the valid ABox assertions may vary over time. On the other hand, satisfiability with rigid concepts and roles becomes EXPTIME-complete if the temporal component is restricted appropriately by replacing the temporal operators until (U) and next (X) of LTL with diamond (\diamond), which says “sometime in the future.”

The situation we concentrate on in this paper (i.e., where temporal operators are allowed to occur only in front of axioms) has been considered before only for the case where there are no rigid concepts or roles. The combination approach introduced in [Finger and Gabbay 1992] yields a decision procedure for this case, whose worst-case complexity is, however, non-optimal. Our EXPTIME upper bound for this case actually also follows from more general results in [Gabbay et al. 2003] (see the remark following Theorem 14.15 on page 605). However, also in [Gabbay et al. 2003], the setting where temporal operators are allowed to occur only in front of axioms is considered only in the absence of rigid symbols.

Obviously, the temporalized DLs we investigate in this paper cannot be used to define temporal concepts such as (1) for concussion with no loss of consciousness. However, they are very useful in applications that focus on the temporal evolution of data (stored in an ABox) or ontologies (represented by a TBox). As a concrete example, consider an emergency ward where the vital parameters of a patient are monitored in short intervals (not longer than 10 minutes), interpreted qualitatively, and then stored in an ABox. Additional information is manually added by doctors and nurses, and imported from historic patient records for the hospitalized

patient. All medical information in the ABox is represented using concepts defined in the medical ontology SNOMED CT; note that this is a realistic assumption as SNOMED CT codes are adopted by standard formats for patient records such as CDA (the Clinical Document Architecture).⁴ The sequence of ABoxes obtained in this way can be captured by a single temporal ABox in the temporalized DLs defined in this paper. Critical situations, which require the intervention of a doctor and should thus result in an alarm being raised, can then be described by a formula in our temporalized DL, and recognized using the reasoning procedures developed in this paper. For example, given a formula ϕ encoding a sequence of ABoxes that describe the medical status of Bob, starting at some time point t_0 , and the formula ψ defined in (2), we can check whether Bob sometime after t_0 had a concussion with no loss of consciousness by testing $\phi \wedge \neg\psi$ for unsatisfiability. This concrete application and the general ideas behind it are further developed in [Baader et al. 2009]. In particular, the authors develop an approach to runtime verification in the temporalized DLs defined in this paper.

The paper is organized as follows. In Section 2, we introduce \mathcal{ALC} -LTL, TBoxes, ABoxes, and related notions. We define the different versions of the satisfiability problem studied in this paper in Section 3. In the subsequent sections, we analyze the complexity of reasoning with rigid roles and concepts (Section 4), without any rigid symbols (Section 5), and with only rigid concepts (Section 6). The reason for this order of Sections 5 and Section 6 is that the upper bound established in Section 5 can be seen as a warmup exercise for the upper bound proved in Section 6. We then move on to restricting the available temporal operators to \diamond and \square in Section 7. Finally, we briefly summarize the paper and discuss some potential future research issues in Section 8. This paper is an extended version of the conference paper [Baader et al. 2008].

2. BASIC DEFINITIONS

The temporalized DL \mathcal{ALC} -LTL introduced in this paper combines the basic DL \mathcal{ALC} [Schmidt-Schauß and Smolka 1991] with linear temporal logic (LTL) [Pnueli 1977]. We start by recalling the relevant definitions for \mathcal{ALC} .

DEFINITION 2.1. *Let N_C , N_R , and N_I respectively be disjoint sets of concept names, role names, and individual names. The set of \mathcal{ALC} -concept descriptions is the smallest set such that*

- all concept names are \mathcal{ALC} -concept descriptions;
- if C and D are \mathcal{ALC} -concept descriptions, then so are $\neg C$, $C \sqcup D$, and $C \sqcap D$;
- if C is an \mathcal{ALC} -concept description and $r \in N_R$, then $\exists r.C$ and $\forall r.C$ are \mathcal{ALC} -concept descriptions.

A general concept inclusion axiom (GCI) is of the form $C \sqsubseteq D$, where C, D are \mathcal{ALC} -concept descriptions, and an assertion is of the form $a : C$ or $(a, b) : r$ where C is an \mathcal{ALC} -concept description, r is a role name, and a, b are individual names. We call both GCIs and assertions \mathcal{ALC} -axioms. A Boolean combination of \mathcal{ALC} -axioms is called a Boolean \mathcal{ALC} -knowledge base, i.e.,

⁴<http://www.hl7.org/implement/standards/cda.cfm>

- every \mathcal{ALC} -axiom is a Boolean \mathcal{ALC} -knowledge base;
- if \mathcal{B}_1 and \mathcal{B}_2 are Boolean \mathcal{ALC} -knowledge bases, then so are $\mathcal{B}_1 \wedge \mathcal{B}_2$, $\mathcal{B}_1 \vee \mathcal{B}_2$, and $\neg\mathcal{B}_1$.

An \mathcal{ALC} -TBox is a conjunction of GCIs, and an \mathcal{ALC} -ABox is a conjunction of assertions.

According to this definition, TBoxes and ABoxes are special kinds of Boolean knowledge bases. However, note that they are often written as sets of axioms rather than as conjunctions of these axioms.

The semantics of \mathcal{ALC} is defined through the notion of an interpretation.

DEFINITION 2.2. An interpretation is a pair $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ where the domain $\Delta^{\mathcal{I}}$ is a non-empty set, and $\cdot^{\mathcal{I}}$ is a function that assigns to every concept name A a set $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$, to every role name r a binary relation $r^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$, and to every individual name a an element $a^{\mathcal{I}} \in \Delta^{\mathcal{I}}$. This function is extended to \mathcal{ALC} -concept descriptions as follows:

- $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, $(C \sqcup D)^{\mathcal{I}} = C^{\mathcal{I}} \cup D^{\mathcal{I}}$, $(\neg C)^{\mathcal{I}} = \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}}$;
- $(\exists r.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \text{there is a } y \in \Delta^{\mathcal{I}} \text{ with } (x, y) \in r^{\mathcal{I}} \text{ and } y \in C^{\mathcal{I}}\}$;
- $(\forall r.C)^{\mathcal{I}} = \{x \in \Delta^{\mathcal{I}} \mid \text{for all } y \in \Delta^{\mathcal{I}}, (x, y) \in r^{\mathcal{I}} \text{ implies } y \in C^{\mathcal{I}}\}$.

The interpretation \mathcal{I} is a model of the \mathcal{ALC} -axioms $C \sqsubseteq D$, $a : C$, and $(a, b) : r$ iff it respectively satisfies $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, $a^{\mathcal{I}} \in C^{\mathcal{I}}$, and $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in r^{\mathcal{I}}$. The notion of a model is extended to Boolean \mathcal{ALC} -knowledge bases as follows:

- \mathcal{I} is a model of $\mathcal{B}_1 \wedge \mathcal{B}_2$ iff it is a model of both \mathcal{B}_1 and \mathcal{B}_2 ;
- \mathcal{I} is a model of $\mathcal{B}_1 \vee \mathcal{B}_2$ iff it is a model of \mathcal{B}_1 or of \mathcal{B}_2 ;
- \mathcal{I} is a model of $\neg\mathcal{B}_1$ iff it is not a model of \mathcal{B}_1 .

We say that the Boolean \mathcal{ALC} -knowledge base \mathcal{B} is consistent iff it has a model. The concept description C is satisfiable w.r.t. the GCI $D_1 \sqsubseteq D_2$ iff there is a model \mathcal{I} of $D_1 \sqsubseteq D_2$ with $C^{\mathcal{I}} \neq \emptyset$.

As usual, we will use \top as an abbreviation for $A \sqcup \neg A$, i.e., \top denotes the *top concept*, which is always interpreted as the whole interpretation domain.

For LTL, we use the variant with a *non-strict until* (U) and a *next* (X) operator. Instead of first introducing the propositional temporal logic LTL, we directly define our new temporalized DL, called \mathcal{ALC} -LTL. The difference between \mathcal{ALC} -LTL and LTL is that \mathcal{ALC} -axioms replace propositional letters.

DEFINITION 2.3. \mathcal{ALC} -LTL formulae are defined by induction:

- if α is an \mathcal{ALC} -axiom, then α is an \mathcal{ALC} -LTL formula;
- if ϕ, ψ are \mathcal{ALC} -LTL formulae, then so are $\phi \wedge \psi$, $\phi \vee \psi$, $\neg\phi$, $\phi U \psi$, and $X\phi$.

As usual, we use true as an abbreviation for $A(a) \vee \neg A(a)$, $\diamond\phi$ as an abbreviation for $\text{true} U \phi$ (*diamond*, which should be read as “sometime in the future”), and $\square\phi$ as an abbreviation for $\neg \diamond \neg \phi$ (*box*, which should be read as “always in the future”).

The semantics of \mathcal{ALC} -LTL is based on \mathcal{ALC} -LTL structures, which are sequences of \mathcal{ALC} -interpretations over the same non-empty domain Δ . Assuming the same

domain for every time point is usually called the *constant domain assumption*. Other possible choices include expanding domains, decreasing domains, and varying domains [Gabbay et al. 2003]. In contrast to temporalized DLs that allow the application of temporal operators to concepts, there does not appear to be an easy way to simulate these other choices by constant domains in \mathcal{ALC} -LTL. However, constant domains are usually considered the most natural choice. It is outside the scope of this paper to investigate the other choices (see also Section 8, where this is discussed as possible future work).

We also assume that every individual name stands for a unique element of Δ , i.e., the interpretation of individual names does not change over time (rigid individual names). This is a standard assumption in temporalized DLs [Gabbay et al. 2003]. As usual in DLs, we also make the *unique name assumption (UNA)*, i.e., different individual names are interpreted by different elements of Δ .

DEFINITION 2.4. *An \mathcal{ALC} -LTL structure is a sequence $\mathfrak{I} = (\mathcal{I}_i)_{i=0,1,\dots}$ of \mathcal{ALC} -interpretations $\mathcal{I}_i = (\Delta, \cdot^{\mathcal{I}_i})$ obeying the UNA (called worlds) such that $a^{\mathcal{I}_i} = a^{\mathcal{I}_j}$ for all individual names a and all $i, j \in \{0, 1, 2, \dots\}$. Given an \mathcal{ALC} -LTL formula ϕ , an \mathcal{ALC} -LTL structure $\mathfrak{I} = (\mathcal{I}_i)_{i=0,1,\dots}$, and a time point $i \in \{0, 1, 2, \dots\}$, validity of ϕ in \mathfrak{I} at time i (written $\mathfrak{I}, i \models \phi$) is defined inductively:*

$$\begin{aligned}
\mathfrak{I}, i \models C \sqsubseteq D & \quad \text{iff } C^{\mathcal{I}_i} \subseteq D^{\mathcal{I}_i} \\
\mathfrak{I}, i \models a : C & \quad \text{iff } a^{\mathcal{I}_i} \in C^{\mathcal{I}_i} \\
\mathfrak{I}, i \models (a, b) : r & \quad \text{iff } (a^{\mathcal{I}_i}, b^{\mathcal{I}_i}) \in r^{\mathcal{I}_i} \\
\mathfrak{I}, i \models \phi \wedge \psi & \quad \text{iff } \mathfrak{I}, i \models \phi \text{ and } \mathfrak{I}, i \models \psi \\
\mathfrak{I}, i \models \phi \vee \psi & \quad \text{iff } \mathfrak{I}, i \models \phi \text{ or } \mathfrak{I}, i \models \psi \\
\mathfrak{I}, i \models \neg \phi & \quad \text{iff not } \mathfrak{I}, i \models \phi \\
\mathfrak{I}, i \models \mathbf{X}\phi & \quad \text{iff } \mathfrak{I}, i+1 \models \phi \\
\mathfrak{I}, i \models \phi \mathbf{U}\psi & \quad \text{iff there is } k \geq i \text{ such that } \mathfrak{I}, k \models \psi \text{ and} \\
& \quad \mathfrak{I}, j \models \phi \text{ for all } j, i \leq j < k
\end{aligned}$$

As mentioned above, for some concepts and roles, it is not desirable that their interpretation changes over time. Thus, we will sometimes assume that a subset of the set of concept and role names can be designated as being rigid. We will call the elements of this subset *rigid concept names* and *rigid role names*. Concept and role names that do not belong to this subset will be called *flexible*.

DEFINITION 2.5. *We say that the \mathcal{ALC} -LTL structure $\mathfrak{I} = (\mathcal{I}_i)_{i=0,1,\dots}$ respects rigid concept names (role names) iff $A^{\mathcal{I}_i} = A^{\mathcal{I}_j}$ ($r^{\mathcal{I}_i} = r^{\mathcal{I}_j}$) holds for all $i, j \in \{0, 1, 2, \dots\}$ and all rigid concept names A (rigid role names r).*

3. THE SATISFIABILITY PROBLEM IN \mathcal{ALC} -LTL

Depending on whether rigid concept and role names are considered or not, we obtain different variants of the satisfiability problem.

DEFINITION 3.1. *Let ϕ be an \mathcal{ALC} -LTL formula and assume that a subset of the set of concept and role names has been designated as being rigid.*

—We say that ϕ is satisfiable w.r.t. rigid names iff there is an \mathcal{ALC} -LTL structure \mathfrak{I} respecting rigid concept and role names such that $\mathfrak{I}, 0 \models \phi$.

- We say that ϕ is satisfiable w.r.t. rigid concepts iff there is an \mathcal{ALC} -LTL structure \mathfrak{I} respecting rigid concept names such that $\mathfrak{I}, 0 \models \phi$.
- We say that ϕ is satisfiable without rigid names (or simply satisfiable) iff there is an \mathcal{ALC} -LTL structure \mathfrak{I} such that $\mathfrak{I}, 0 \models \phi$.

In this article, we show that the complexity of the satisfiability problem for \mathcal{ALC} -LTL strongly depends on which of the above cases one considers. Note that it does not really make sense to consider satisfiability w.r.t. rigid role names, but without rigid concept names, as a separate case when investigating the complexity of the satisfiability problem. In fact, rigid concepts can be simulated by rigid roles: just introduce a new rigid role name r_A for each rigid concept name A , and then replace A by $\exists r_A.T$.

Another dimension that influences the complexity of the satisfiability problem is whether GCIs occur globally or locally in the formula. Intuitively, a GCI occurs globally if it must hold in every world of the \mathcal{ALC} -LTL structure.

DEFINITION 3.2. *We say that ϕ is an \mathcal{ALC} -LTL formula with global GCIs iff it is of the form $\phi = \Box \mathcal{B} \wedge \xi$ where \mathcal{B} is a conjunction of \mathcal{ALC} -axioms and ξ is an \mathcal{ALC} -LTL formula that does not contain GCIs. We denote the fragment of \mathcal{ALC} -LTL that contains only \mathcal{ALC} -LTL formulae with global GCIs by \mathcal{ALC} -LTL_{gGCI}.*

In the above definition, saying that \mathcal{B} is a conjunction of \mathcal{ALC} -axioms just means that \mathcal{B} is a TBox together with an ABox. Equivalently, we could have restricted \mathcal{B} to being a conjunction of GCIs (i.e., a TBox) since assertions α in \mathcal{B} could be moved as conjuncts $\Box \alpha$ to ξ .⁵ However, it turns out to be more convenient to allow also ABox assertions to occur in the “global part” $\Box \mathcal{B}$ of ϕ .

Note that, semantically, $\Box \mathcal{B}$ does not express the condition that \mathcal{B} is true at every time point because \Box is directed towards the future. However, we are studying satisfiability questions, and satisfiability in time point 0 is equivalent to satisfiability in any time point: due to the lack of past temporal operators, it is always possible to drop the time points that precede the point satisfying the formula. For this reason, using $\Box \mathcal{B}$ is sufficient to assert that \mathcal{B} holds globally.

Instead of restricting to \mathcal{ALC} -LTL formulae with global GCIs, we can also restrict the temporal component, by considering the fragment \mathcal{ALC} -LTL _{\diamond} of \mathcal{ALC} -LTL in which \diamond is the only temporal operator. In this fragment, neither U nor X is definable.

DEFINITION 3.3. *\mathcal{ALC} -LTL _{\diamond} formulae are defined by induction:*

- if α is an \mathcal{ALC} -axiom, then α is an \mathcal{ALC} -LTL _{\diamond} formula;
- if ϕ, ψ are \mathcal{ALC} -LTL _{\diamond} formulae, then so are $\phi \wedge \psi$, $\phi \vee \psi$, $\neg \phi$, and $\diamond \phi$.

The semantics of \mathcal{ALC} -LTL _{\diamond} formulae is defined as in the case of \mathcal{ALC} -LTL. In particular, the interpretation of the diamond operator is defined as

$$\mathfrak{I}, i \models \diamond \phi \text{ iff there is } k \geq i \text{ such that } \mathfrak{I}, k \models \phi.$$

Table II summarizes the results of our investigation of the complexity of the satisfiability problem in \mathcal{ALC} -LTL and its fragments. This table shows that the com-

⁵This is the reason why we talk about \mathcal{ALC} -LTL formulae with global GCIs in this case, rather than about \mathcal{ALC} -LTL formulae with global axioms.

	W.r.t. rigid names	W.r.t. rigid concepts	Without rigid names
$\mathcal{ALC}\text{-LTL}$	2-EXPTIME-complete	NEXPTIME-complete	EXPTIME-complete
$\mathcal{ALC}\text{-LTL} _{gGCI}$	2-EXPTIME-complete	EXPTIME-complete	EXPTIME-complete
$\mathcal{ALC}\text{-LTL} _{\diamond}$	EXPTIME-complete	EXPTIME-complete	EXPTIME-complete

Table II. Complexity of the satisfiability problem in $\mathcal{ALC}\text{-LTL}$ and its fragments.

plexity of the satisfiability problem in $\mathcal{ALC}\text{-LTL}$ increases from EXPTIME (which is also the complexity of the satisfiability problem in \mathcal{ALC}) to NEXPTIME if rigid concept names are available. The additional presence of rigid role names further increases the complexity to 2-EXPTIME. The restriction to $\mathcal{ALC}\text{-LTL}|_{gGCI}$ (i.e., global GCIs) has no effect on the complexity in the presence of rigid role names. However, it decreases the complexity to EXPTIME if only rigid concept names are available. In $\mathcal{ALC}\text{-LTL}|_{\diamond}$, the satisfiability problem is only EXPTIME-complete even w.r.t. rigid names.

4. REASONING WITH RIGID NAMES

In this section, we investigate the complexity of the satisfiability problem in $\mathcal{ALC}\text{-LTL}$ and its fragment $\mathcal{ALC}\text{-LTL}|_{gGCI}$ if rigid concept and role names are available.

THEOREM 4.1. *Satisfiability w.r.t. rigid names is 2-EXPTIME-complete both in $\mathcal{ALC}\text{-LTL}$ and in $\mathcal{ALC}\text{-LTL}|_{gGCI}$.*

4.1 The complexity lower bound

2-EXPTIME-hardness for satisfiability w.r.t. rigid names and with global GCIs (i.e., in $\mathcal{ALC}\text{-LTL}|_{gGCI}$) can be shown by a (quite intricate) reduction of the word problem for exponentially space-bounded alternating Turing machines. Obviously, this also yields 2-EXPTIME-hardness for the more general case with arbitrary GCIs (i.e., in $\mathcal{ALC}\text{-LTL}$).

LEMMA 4.2. *The satisfiability problem in $\mathcal{ALC}\text{-LTL}|_{gGCI}$ w.r.t. rigid names is 2-EXPTIME-hard.*

Proof. The proof is by reduction of the word problem for exponentially space-bounded alternating Turing machines (ATMs). An ATM is of the form $\mathcal{M} = (Q, \Sigma, \Gamma, q_0, \Theta)$, where $Q = Q_{\exists} \uplus Q_{\forall} \uplus \{q_a, q_r\}$ is a finite set of *states*, partitioned into *existential states* from Q_{\exists} , *universal states* from Q_{\forall} , an *accepting state* q_a , and a *rejecting state* q_r ; Σ is the *input alphabet* and $\Gamma \supseteq \Sigma$ the *work alphabet* containing a *blank symbol* $B \notin \Sigma$; $q_0 \in Q_{\exists} \cup Q_{\forall}$ is the *initial state*; and the *transition relation* Θ is of the form $\Theta \subseteq Q \times \Gamma \times Q \times \Gamma \times \{L, R\}$. We write $\Theta(q, a)$ for $\{(q', b, M) \mid (q, a, q', b, M) \in \Theta\}$.

A *configuration* of an ATM is a word uqu' with $u, u' \in \Gamma^*$ and $q \in Q$. The intended meaning is that the (one-sided infinite) tape contains the word uu' with only blanks behind it, the machine is in state q , and the head is on the left-most symbol of u' . The *successor configurations* of a configuration uqu' are defined in the usual way (i.e., as with normal, non-alternating Turing machines) in terms of the transition relation Θ [Chandra et al. 1981]. A *halting configuration* is of the form

uqu' with $q \in \{q_a, q_r\}$. We may assume w.l.o.g. that any configuration other than a halting configuration has at least one successor configuration. A *computation* of an ATM \mathcal{M} on a word w is a (finite or infinite) sequence of successive configurations K_1, K_2, \dots , where $K_1 = q_0w$ is the *initial configuration* for the input w . For the ATMs considered here, we may assume without loss of generality that they have only finite computations on any input. Since this case is simpler than the general one, we define acceptance for ATMs with finite computations and refer to [Chandra et al. 1981] for the full definition. Let \mathcal{M} be such an ATM. A halting configuration is *accepting* iff it is of the form $uq_a u'$. For other configurations $K = uqu'$, the acceptance behavior depends on q : if $q \in Q_\exists$, then K is accepting iff at least one successor configuration is accepting; if $q \in Q_\forall$, then K is accepting iff all successor configurations are accepting. Finally, the ATM \mathcal{M} with initial state q_0 *accepts* the input w iff the initial configuration q_0w is accepting. We use $L(\mathcal{M})$ to denote the language accepted by \mathcal{M} , i.e.,

$$L(\mathcal{M}) = \{w \in \Sigma^* \mid \mathcal{M} \text{ accepts } w\}.$$

The *word problem* for \mathcal{M} is the following decision problem: given a word $w \in \Sigma^*$, does $w \in L(\mathcal{M})$ hold or not?

There exists an exponentially space-bounded ATM $\mathcal{M} = (Q, \Sigma, \Gamma, q_0, \Theta)$ with finite computations only whose word problem is 2-EXPTIME-hard [Chandra et al. 1981]. Our aim is to reduce the word problem for this ATM \mathcal{M} to satisfiability in \mathcal{ALC} -LTL w.r.t. rigid names. We may assume that the length of every computation of \mathcal{M} on $w \in \Sigma^k$ is bounded by 2^{2^k} , and all the configurations uqu' in such computations satisfy $|uu'| \leq 2^k$. We may also assume w.l.o.g. that \mathcal{M} never attempts to move to the left when it is on the left-most tape cell.

Let $w = \sigma_0 \cdots \sigma_{k-1} \in \Sigma^*$ be an input to \mathcal{M} . In the following, we construct an \mathcal{ALC} -LTL_{|gGCI} formula $\phi_{\mathcal{M},w}$ such that $w \in L(\mathcal{M})$ iff $\phi_{\mathcal{M},w}$ is satisfiable w.r.t. rigid names. The formula $\phi_{\mathcal{M},w}$, that will be defined below is actually not of the syntactic form $\Box \mathcal{B} \wedge \xi$ (where \mathcal{B} is a conjunction of \mathcal{ALC} -axioms and ξ is an \mathcal{ALC} -LTL formula that does not contain GCIs) required for \mathcal{ALC} -LTL formulae with global GCIs. Instead, $\phi_{\mathcal{M},w}$ is a conjunction of formulae of the form

- $\Box \alpha$ where α is an \mathcal{ALC} -axiom,
- ψ where ψ is an \mathcal{ALC} -LTL formula not containing GCIs.

Since \Box distributes over conjunction, it is obvious that such a formula is equivalent to an \mathcal{ALC} -LTL formula with global GCIs.

In an \mathcal{ALC} -LTL structure satisfying $\phi_{\mathcal{M},w}$, an accepting computation of \mathcal{M} is encoded as a tree. The root of the tree is identified by an individual name a and its edges are represented using a single *rigid* role r ; thus, we can find the same tree in every world of the structure. Each node in the tree corresponds to a single tape cell of a configuration of \mathcal{M} . Going to an r -successor of such a node either leads to the subsequent tape cell of the same configuration (if the current cell was not the right-most one) or to the left-most cell of a successor configuration (otherwise).

Note that configurations are of length 2^k , and thus each configuration is represented as a sequence of 2^k nodes in the tree. Since the length of each computation is bounded by 2^{2^k} , the depth of the tree is thus bounded by $2^{2^k} \cdot 2^k$, the maximum number of successive configurations multiplied by their length.

Branching occurs whenever the Turing machine is in a universal configuration, i.e., a configuration where the state is a universal one: then, we need to examine (and thus represent) *all* successor configurations instead of just a single one. The actual branching does not occur at the end of the universal configuration, but at the tape cell where the head is located in that configuration. The representation of the remaining part of the configuration (tape content after the head) is then copied in every branch.

To realize this intuition in the formal definition of $\phi_{\mathcal{M},w}$, we will use the following symbols to represent the tree:

- a single individual name a identifies the root of the tree;
- a single *rigid* role name r represents the edges of the tree;
- the elements of Q and Γ are viewed as *rigid* concept names and used to represent the tape content, the current state, and the head position in each configuration in the tree; more precisely, if \mathcal{M} is in state q and the head is on the i -th tape cell, then the concept name q is true at the node that represents this cell; similarly, if σ is the symbol in the i -th tape cell, then the concept name σ is true at the node that represents this cell;
- rigid* concept names A_0, \dots, A_{k-1} are the bits of a binary counter that numbers the tape cells in each configuration;
- rigid* concept name I and H are used as markers: I is true for all nodes representing the initial configuration and H is true for all nodes representing a tape cell that is located (an arbitrary number of steps) to the right of the head in the represented configuration;
- rigid* concept names $T_{q,\sigma,M}$ for all $q \in Q$, $\sigma \in \Gamma$, and $M \in \{L, R\}$ also serve as markers; intuitively, $T_{q,\sigma,M}$ is true at a tape cell if, in the represented configuration, the head is on the left neighboring cell and the machine executes transition (q, σ, M) ;

The main problem to solve when defining the reduction is the synchronization of successor configurations, i.e., to ensure that the content of tape cells under the head changes according to the chosen transition, and that all other cells remain unchanged. Additionally, we have to exchange and move the concept name from Q that marks the current state and head position to reflect the new state and head position. To implement this synchronization, we will use several technical tricks that make use of the temporal dimension (which is not used otherwise in the reduction). To this end, we introduce the following additional symbols:

- flexible* concept names A'_0, \dots, A'_{k-1} realize a counter in the temporal dimension (instead of along r); this counter is dual to the counter A_0, \dots, A_{k-1} : it has a different value in different temporal worlds, but within each single world all elements of the tree agree on the value of the A' -counter; a world where this counter has value i will be used to synchronize the i -th tape cell of any configuration with the i -th tape cell of its successor configurations;
- for each element of Q and Γ , a *flexible* concept name, which is distinguished from its rigid version by a prime; these concept names will be used to ‘memorize’ the state or content of a tape cell in the temporal worlds discussed in the previous item.

In the following, we use $\phi \rightarrow \psi$ as an abbreviation for $\neg\phi \vee \psi$, $C \Rightarrow D$ as an abbreviation for $\neg C \sqcup D$, and $C \Leftrightarrow D$ as an abbreviation for $(C \Rightarrow D) \sqcap (D \Rightarrow C)$. The reduction formula $\phi_{\mathcal{M},w}$ is the conjunction of several formulae. We start with giving conjuncts that set up some basic elements of the tree structure, without synchronizing successor configurations:

- there is always an r -successor, except when we meet the head in a halting configuration:⁶

$$\square (\neg(q_a \sqcup q_r) \sqsubseteq \exists r. \top)$$

Clearly, this enforces the existence of an r -chain without branching, and not a tree (but it also does not disallow branching). For technical reasons, the actual branching can only be enforced much later in the reduction. We still recommend to the reader to think of r as the edges in a tree, as sketched above.

- the A -counter realized by A_0, \dots, A_{k-1} has value 0 at a , and it is incremented along r , modulo 2^k (corresponding to the reset of the counter value to 0 after a complete configuration):

$$\begin{aligned} & \square (a : (C_A = 0)) \\ & \square \left(\top \sqsubseteq \prod_{i < k} \left(\prod_{j < i} A_j \right) \Rightarrow ((A_i \Rightarrow \forall r. \neg A_i) \sqcap (\neg A_i \Rightarrow \forall r. A_i)) \right) \\ & \square \left(\top \sqsubseteq \prod_{i < k} \left(\prod_{j < i} \neg A_j \right) \Rightarrow ((A_i \Rightarrow \forall r. A_i) \sqcap (\neg A_i \Rightarrow \forall r. \neg A_i)) \right) \end{aligned}$$

where $(C_A = 0)$ denotes the concept that is true iff the counter A_0, \dots, A_{k-1} has minimum value, i.e., $\neg A_0 \sqcap \dots \sqcap \neg A_{k-1}$; below, we will use similar concepts $(C_A = i)$ for other fixed (but only polynomially many) values i without further notice;

- I marks the initial configuration, whose first tape cell is represented by the individual a :

$$\begin{aligned} & \square (a : I) \\ & \square (I \sqcap \neg(C_A = 2^k - 1) \sqsubseteq \forall r. I) \end{aligned}$$

- H marks the tape cells that are to the right of the head (recall that the head position is indicated by having a concept from Q at this cell):

$$\square \left((H \sqcup \bigsqcup_{q \in Q} q) \sqcap \neg(C_A = 2^k - 1) \sqsubseteq \forall r. H \right)$$

Now that we have the tree, we can start to enforce conditions that ensure that the tree actually represents an accepting computation of \mathcal{M} on w . Some of these conditions can easily be formalized using the marker concepts I and H , but without referring to the temporal dimension:

⁶In a halting configuration, the tape cells to the right of the head are not represented; in fact, once a halting state is reached, the tape content is irrelevant, and thus there is no need for representing these cells.

—the first configuration (starting at a) is the initial one: \mathcal{M} is in the initial state q_0 , the head is on the left-most tape cell, and the tape content consists of the input $w = \sigma_0 \dots \sigma_{k-1}$ followed by blanks:

$$\begin{aligned} & \Box (a : q_0) \\ & \Box (a : \forall r^i . \sigma_i) \quad \text{for } i < k \\ & \Box (a : \forall r^k . B) \\ & \Box (I \sqcap B \sqcap \neg(C_A = 2^k - 1) \sqsubseteq \forall r . B) \end{aligned}$$

—each tape cell is labelled with exactly one symbol and at most one state:

$$\begin{aligned} & \Box \left(\top \sqsubseteq \bigsqcup_{\sigma \in \Gamma} (\sigma \sqcap \neg \bigsqcap_{\sigma' \in \Gamma \setminus \{\sigma\}} \neg \sigma') \right) \\ & \Box \left(\top \sqsubseteq \bigsqcap_{q, q' \in Q, q \neq q'} \neg (q \sqcap q') \right) \end{aligned}$$

—there is only one head position per configuration:

$$\Box \left(H \sqsubseteq \bigsqcap_{q \in Q} \neg q \right)$$

It remains to synchronize successor configurations, as described above. Note that this includes implementing the transitions. We start with setting up the A' -counter based on the concept names A'_0, \dots, A'_{k-1} , which plays a central role in this part of the reduction:

—for each possible value i of the counter, there is a temporal world in which the concept memberships of a regarding the (flexible) concept names A'_0, \dots, A'_{k-1} encode this value:

$$\begin{aligned} & \Box \left(\bigwedge_{i < k} \left(\bigwedge_{j < i} a : A'_j \rightarrow ((a : A'_i \rightarrow \exists a : \neg A'_i) \wedge (a : \neg A'_i \rightarrow \exists a : A'_i)) \right) \right) \\ & \Box \left(\bigwedge_{i < k} \left(\bigvee_{j < i} a : \neg A'_j \rightarrow ((a : A'_i \rightarrow \exists a : A'_i) \wedge (a : \neg A'_i \rightarrow \exists a : \neg A'_i)) \right) \right) \end{aligned}$$

This is basically the same formula as for the A -counter, but the values of the A' -counter are considered for the fixed initial individual a , and they are incremented along the temporal dimension. It is not necessary to set the value of the A' -counter at the initial time point to zero. It will have some value, and since we count modulo 2^k every value between 0 and $2^k - 1$ will be reached by successively adding 1.

—The value of the A' -counter is preserved along r ; thus, in any fixed world all nodes of the tree agree on the value of the A' -counter (as announced above):

$$\begin{aligned} & \Box (A'_i \sqsubseteq \forall r . A'_i) \\ & \Box (\neg A'_i \sqsubseteq \forall r . \neg A'_i) \end{aligned}$$

Before we can continue describing conjuncts of the reduction formula, we need to introduce some abbreviations. In the following, we use $(C_A = C_{A'})$ to denote the concept $(A_0 \Leftrightarrow A'_0) \sqcap \dots \sqcap (A_{k-1} \Leftrightarrow A'_{k-1})$, which states that the value of

the A -counter coincides with the value of the A' -counter. Accordingly, $(C_A = C_{A'} + 1 \bmod 2^k)$ expresses the condition that the value of the A -counter is equal to the value of the A' -counter plus 1 (modulo 2^k). This can be expressed by a recasting of the incrementation concept given already twice above:

$$(C_A = C_{A'} + 1 \bmod 2^k) := \prod_{i < k} \left(\prod_{j < i} A'_j \Rightarrow ((A'_i \Rightarrow \neg A_i) \sqcap (\neg A'_i \Rightarrow A_i)) \sqcap \prod_{j < k} \left(\prod_{j < i} \neg A'_j \Rightarrow ((A'_i \Rightarrow A_i) \sqcap (\neg A'_i \Rightarrow \neg A_i)) \right)$$

The concept $(C_A = C_{A'} + 2 \bmod 2^k)$, which expresses the condition that the value of the A -counter is equal to the value of the A' -counter plus 2 (modulo 2^k), can be defined similarly, using an auxiliary set A''_0, \dots, A''_{k-1} of flexible concept names.

We are now ready to describe the final parts of the reduction. All remaining conditions require us to propagate information (states, tape contents) between some tape cell and the same cell at the successor configurations, which is done as follows. Assume we are in the i -th tape cell and want to ensure, say, that the symbol σ stored in that cell is also stored in the i -th cell of the successor configuration. We then switch to the world where the A' -counter has value i , i.e., the values of the A -counter and of the A' -counter have the same value at the current node in that world. There, we make the flexible concept name σ' true to memorize σ . We then follow the role r for 2^k steps (still in that world), as this brings us to the i -th tape cell of the successor configurations. We ensure that we do not ‘forget’ the memorized information σ' on the way, i.e., we make σ' true at each node along the traveled path. To know when we have reached our destination, we reuse the A -counter: it suffices to look for the first node along the path where the A -counter and the A' -counter have the same value again (namely i). At the destination node, we can then enforce that the rigid concept name σ holds (no need to switch worlds again since σ is rigid), which is what we wanted to achieve.

More formally, this is implemented using the following formulae:

—symbols not under the head do not change; note that the ‘switching of worlds’ is not done explicitly; instead, it is implicit in the first GCI:

$$\begin{aligned} \square \left(\sigma \sqcap \prod_{q \in Q} \neg q \sqcap (C_A = C_{A'}) \sqsubseteq \forall r. \sigma' \right) & \quad \text{for all } \sigma \in \Gamma \\ \square (\sigma' \sqcap \neg(C_A = C_{A'})) \sqsubseteq \forall r. \sigma' & \quad \text{for all } \sigma \in \Gamma \\ \square (\sigma' \sqcap (C_A = C_{A'})) \sqsubseteq \sigma & \quad \text{for all } \sigma \in \Gamma \end{aligned}$$

—Transitions are implemented in a similar way. It is here where we finally enforce the branching of the tree. Since there are many possible transitions that an ATM can make in a given configuration, we explicitly store the chosen one using the concept names $T_{p,\nu,M}$. More precisely, this concept name is true at the *successor node* of the node where the head is currently located. For an existential state, one of the possible transitions is chosen; since we cannot enforce the condition that a node in the tree has only one successor, we ensure instead that all successors satisfy the same concept name $T_{p,\nu,M}$. For a universal state, we enforce the condition that every possible transition is realized in some successor. Since the concept names $T_{p,\nu,M}$ are true at the successor nodes and since the head may

move during a transition, we sometimes have to travel $2^k - 1$ or $2^k + 1$ steps instead of 2^k steps to reach the intended destination. This is the reason why we sometimes have to manipulate the counter value when comparing the A -counter with the A' -counter:

$$\begin{aligned}
& \Box \left(q \sqcap \sigma \sqsubseteq \bigsqcup_{(p,\nu,M) \in \Theta(q,\sigma)} \forall r. T_{p,\nu,M} \right) && \text{for all } q \in Q_{\exists}, \sigma \in \Sigma \\
& \Box \left(q \sqcap \sigma \sqsubseteq \bigsqcap_{(p,\nu,M) \in \Theta(q,\sigma)} \exists r. T_{p,\nu,M} \right) && \text{for all } q \in Q_{\forall}, \sigma \in \Sigma \\
& \Box (T_{q,\sigma,M} \sqcap (C_A = C_{A'} + 1 \bmod 2^k) \sqsubseteq \forall r. \sigma') && \text{for all } \sigma \in \Gamma, q \in Q, \\
& && M \in \{L, R\} \\
& \Box (T_{q,\sigma,R} \sqcap (C_A = C_{A'}) \sqsubseteq \forall r. q') && \text{for all } \sigma \in \Gamma, q \in Q \\
& \Box (T_{q,\sigma,L} \sqcap (C_A = C_{A'} + 2 \bmod 2^k) \sqsubseteq \forall r. q') && \text{for all } \sigma \in \Gamma, q \in Q \\
& \Box (q' \sqcap \neg(C_A = C_{A'}) \sqsubseteq \forall r. q') && \text{for all } q \in Q \\
& \Box (q' \sqcap (C_A = C_{A'}) \sqsubseteq q) && \text{for all } q \in Q
\end{aligned}$$

It remains to encode the fact that the input $w = \sigma_0 \dots \sigma_{k-1}$ is accepted. Since any computation of \mathcal{M} is terminating, and halting configurations (i.e., configurations with state q_a or q_r) are the only ones without successor configurations, this can be done as follows:

—We can express the fact that the initial configuration for input w is accepting by disallowing the state q_r to occur:

$$\Box (\top \sqsubseteq \neg q_r)$$

This finishes the definition of $\phi_{\mathcal{M},w}$, which is the conjunction of the formulae introduced above. It is easy to see that the size of $\phi_{\mathcal{M},w}$ is polynomial in k . Moreover, given the intuition provided above, it is a routine task to prove that $\phi_{\mathcal{M},w}$ is satisfiable w.r.t. rigid names iff $w \in L(\mathcal{M})$. \square

4.2 The complexity upper bound

Here, we show that the complexity lower bound provided by the above lemma is tight. To be more precise, we prove a 2-EXPTIME upper bound for \mathcal{ALC} -LTL. Obviously, this also establishes the same upper bound for the restricted case of \mathcal{ALC} -LTL $_{|gGCI}$.

LEMMA 4.3. *Satisfiability in \mathcal{ALC} -LTL w.r.t. rigid names is in 2-EXPTIME.*

Proof. Let ϕ be an \mathcal{ALC} -LTL formula. We build its propositional abstraction $\widehat{\phi}$ by replacing each \mathcal{ALC} -axiom by a propositional variable such that there is a 1–1 relationship between the \mathcal{ALC} -axioms $\alpha_1, \dots, \alpha_n$ occurring in ϕ and the propositional variables p_1, \dots, p_n used for the abstraction. We assume in the following that p_i was used to replace α_i ($i = 1, \dots, n$).

Consider a set $\mathcal{S} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$, i.e., a set of subsets of $\{p_1, \dots, p_n\}$. Such a set induces the following (propositional) LTL formula:

$$\widehat{\phi}_{\mathcal{S}} := \widehat{\phi} \wedge \square \left(\bigvee_{X \in \mathcal{S}} \left(\bigwedge_{p \in X} p \wedge \bigwedge_{p \notin X} \neg p \right) \right)$$

If ϕ is satisfiable in an \mathcal{ALC} -LTL structure $\mathfrak{J} = (\mathcal{I}_i)_{i=0,1,\dots}$, then there is an $\mathcal{S} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ such that $\widehat{\phi}_{\mathcal{S}}$ is satisfiable in a propositional LTL structure. In fact, for each \mathcal{ALC} -interpretation \mathcal{I}_i of \mathfrak{J} , we define the set

$$X_i := \{p_j \mid 1 \leq j \leq n \text{ and } \mathcal{I}_i \text{ satisfies } \alpha_j\},$$

and then take $\mathcal{S} = \{X_i \mid i = 0, 1, \dots\}$. We say that \mathcal{S} is *induced* by the \mathcal{ALC} -LTL structure $\mathfrak{J} = (\mathcal{I}_i)_{i=0,1,\dots}$. The fact that \mathfrak{J} satisfies ϕ implies that its propositional abstraction satisfies $\widehat{\phi}_{\mathcal{S}}$, where the *propositional abstraction* $\widehat{\mathfrak{J}} = (w_i)_{i=0,1,\dots}$ of \mathfrak{J} is defined such that world w_i makes variable p_j true iff \mathcal{I}_i satisfies α_j . However, guessing such a set $\mathcal{S} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ and then testing whether the induced propositional LTL formula $\widehat{\phi}_{\mathcal{S}}$ is satisfiable is not sufficient for checking satisfiability w.r.t. rigid names of the \mathcal{ALC} -LTL formula ϕ . We must also check whether the guessed set \mathcal{S} can indeed be induced by some \mathcal{ALC} -LTL structure that respects the rigid concept and role names.

To this end, assume that a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ is given. For every $i, 1 \leq i \leq k$, and every *flexible* concept name A (*flexible* role name r) occurring in $\alpha_1, \dots, \alpha_n$, we introduce a copy $A^{(i)}$ ($r^{(i)}$). We call $A^{(i)}$ ($r^{(i)}$) the i th copy of A (r). The \mathcal{ALC} -axiom $\alpha_j^{(i)}$ is obtained from α_j by replacing every occurrence of a flexible name by its i th copy. The sets X_i ($1 \leq i \leq k$) induce the following Boolean \mathcal{ALC} -knowledge bases:

$$\mathcal{B}_i := \bigwedge_{p_j \in X_i} \alpha_j^{(i)} \wedge \bigwedge_{p_j \notin X_i} \neg \alpha_j^{(i)}$$

Claim. *The \mathcal{ALC} -LTL formula ϕ is satisfiable w.r.t. rigid names iff there is a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ such that the propositional LTL formula $\widehat{\phi}_{\mathcal{S}}$ is satisfiable and the Boolean \mathcal{ALC} -knowledge base $\mathcal{B} := \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ is consistent.*

For the “only if” direction, recall that we have already seen how an \mathcal{ALC} -LTL structure $\mathfrak{J} = (\mathcal{I}_\iota)_{\iota=0,1,\dots}$ satisfying ϕ can be used to define a set $\mathcal{S} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ such that $\widehat{\phi}_{\mathcal{S}}$ is satisfiable. Let $\mathcal{S} = \{X_1, \dots, X_k\}$. For each $\iota = 0, 1, \dots$ there is an index $i_\iota \in \{1, \dots, k\}$ such that \mathcal{I}_ι induces the set X_{i_ι} , i.e.,

$$X_{i_\iota} = \{p_j \mid 1 \leq j \leq n \text{ and } \mathcal{I}_\iota \text{ satisfies } \alpha_j\},$$

and, conversely, for each $i \in \{1, \dots, k\}$ there is an index $\iota \in \{0, 1, 2, \dots\}$ such that $i = i_\iota$. Let $\iota_1, \dots, \iota_k \in \{0, 1, 2, \dots\}$ be such that $i_{\iota_1} = 1, \dots, i_{\iota_k} = k$. The \mathcal{ALC} -interpretation \mathcal{J}_i is obtained from \mathcal{I}_{ι_i} by interpreting the i th copy of each flexible (concept or role) name in \mathcal{J}_i like the original flexible name in \mathcal{I}_{ι_i} , and interpreting all rigid names in \mathcal{J}_i exactly as in \mathcal{I}_{ι_i} . By our construction of \mathcal{J}_i and our definition of the Boolean \mathcal{ALC} -knowledge base \mathcal{B}_i , we have that \mathcal{J}_i is a model of \mathcal{B}_i . Recall

that the interpretations $\mathcal{I}_{\iota_1}, \dots, \mathcal{I}_{\iota_k}$ (and thus also $\mathcal{J}_1, \dots, \mathcal{J}_k$) all have the same domain. In addition, the interpretations of the rigid names coincide in $\mathcal{I}_{\iota_1}, \dots, \mathcal{I}_{\iota_k}$ (and thus also in $\mathcal{J}_1, \dots, \mathcal{J}_k$) and the flexible symbols have been renamed. Thus, the union \mathcal{J} of $\mathcal{J}_1, \dots, \mathcal{J}_k$ is a well-defined \mathcal{ALC} -interpretation, and it is easy to see that it is a model of $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$.

To show the “if” direction, assume that there is a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ such that $\widehat{\phi}_{\mathcal{S}}$ is satisfiable and $\mathcal{B} := \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ is consistent. Let $\widehat{\mathcal{J}} = (w_\iota)_{\iota=0,1,\dots}$ be a propositional LTL structure satisfying $\widehat{\phi}_{\mathcal{S}}$, and let \mathcal{J} be an \mathcal{ALC} -interpretation satisfying \mathcal{B} . By the definition of $\widehat{\phi}_{\mathcal{S}}$, for every world w_ι there is exactly one index $i_\iota \in \{1, \dots, k\}$ such that w_ι satisfies

$$\bigwedge_{p \in X_{i_\iota}} p \wedge \bigwedge_{p \notin X_{i_\iota}} \neg p.$$

For $i \in \{1, \dots, k\}$, we use the \mathcal{ALC} -interpretation \mathcal{J} satisfying \mathcal{B} to define an \mathcal{ALC} -interpretation \mathcal{J}_i as follows: \mathcal{J}_i interprets the rigid names like \mathcal{J} , and it interprets the flexible names just as \mathcal{J} interprets the i th copies of them. Note that the interpretations \mathcal{J}_i are over the same domain and respect the rigid symbols, i.e., they interpret them identically. We can now define an \mathcal{ALC} -LTL structure respecting rigid symbols and satisfying ϕ as follows: $\mathcal{J} := (\mathcal{I}_\iota)_{\iota=0,1,\dots}$ where $\mathcal{I}_\iota := \mathcal{J}_{i_\iota}$.

This completes the proof of the claim. It remains to show that the claim provides us with a decision procedure for satisfiability in \mathcal{ALC} -LTL w.r.t. rigid names that runs in deterministic double-exponential time.

First, note that there are 2^{2^n} many subsets \mathcal{S} of $\mathcal{P}(\{p_1, \dots, p_n\})$ to be tested, where n is of course linearly bounded by the size of ϕ . For each of these subsets $\mathcal{S} = \{X_1, \dots, X_k\}$, whose cardinality k is bounded by 2^n , we need to check satisfiability of $\widehat{\phi}_{\mathcal{S}}$ and consistency of $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$.

The size of $\widehat{\phi}_{\mathcal{S}}$ is at most exponential in the size of ϕ , and the complexity of the satisfiability problem in propositional LTL is in PSPACE, and thus in particular in EXPTIME. Consequently, satisfiability of $\widehat{\phi}_{\mathcal{S}}$ can be tested in double-exponential time in the size of ϕ .

The Boolean \mathcal{ALC} -knowledge base \mathcal{B} is a conjunction of $k \leq 2^n$ Boolean \mathcal{ALC} -knowledge bases \mathcal{B}_i , where the size of each \mathcal{B}_i is polynomial in the size of ϕ . The consistency problem for Boolean \mathcal{ALC} -knowledge bases is EXPTIME-complete (see, e.g., Theorem 2.27 in [Gabbay et al. 2003] or Lemma 6.4 below). Consequently, consistency of \mathcal{B} can also be tested in double-exponential time in the size of the input formula ϕ .

Overall, we thus have double-exponentially many tests, where each test takes double-exponential time. This provides us with a double-exponential bound for testing satisfiability in \mathcal{ALC} -LTL w.r.t. rigid names based on the above claim. \square

The hardness proof given in Section 4.1 shows that the double-exponential upper bound that we have just shown is indeed optimal. However, to get an intuition for what actually makes the problem so hard, let us analyze in more detail where the algorithm sketched above needs to spend double-exponential time. The algorithm needs

- (1) to consider 2^{2^n} many subsets $\mathcal{S} = \{X_1, \dots, X_k\}$ of $\mathcal{P}(\{p_1, \dots, p_n\})$;
- (2) for each such subset test $\widehat{\phi}_{\mathcal{S}}$ for satisfiability;
- (3) test $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ for consistency.

The main culprit is 3. Intuitively, the presence of rigid roles ensures that the consistency test for the conjunction $\bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ needs to be done for the whole conjunction, and cannot be reduced to separate test for the conjuncts (or some enriched versions of the conjuncts). Thus, the Boolean \mathcal{ALC} -knowledge base to be tested for consistency is exponentially large. Since the consistency problem for Boolean \mathcal{ALC} -knowledge bases is EXPTIME-complete, testing this knowledge base for consistency requires in the worst-case double-exponential time. In contrast, 2. is actually harmless. According to what we have said above, it needs exponential space, but we will see in the next section that this can even be reduced to exponential time. Finally, 1. also does not really require double-exponential time since one could guess an appropriate set \mathcal{S} within NEXPTIME.

5. REASONING WITHOUT RIGID NAMES

In this section, we consider the case where we have no rigid names at all. As mentioned in the introduction, this case is also treated in [Gabbay et al. 2003], where it is shown that an EXPTIME upper bound for the satisfiability problem follows from more general results proved in Chapter 11 of [Gabbay et al. 2003] (see the remark following Theorem 14.15 on page 605 of [Gabbay et al. 2003]). For the sake of completeness, and as a warmup for Section 6, we give a direct proof of this upper bound below. To this end, we show that, in this simple case, the claim shown in the proof of Lemma 4.3 implies that satisfiability can be decided in deterministic exponential time.

Since we now consider satisfiability without rigid names, all role and concept names are flexible. Consequently, the Boolean \mathcal{ALC} -knowledge bases \mathcal{B}_i defined in the proof of Lemma 4.3 do not share concept or role names, and can thus be tested for consistency separately.

LEMMA 5.1. *Let $\mathcal{B}_1, \dots, \mathcal{B}_k$ be Boolean \mathcal{ALC} -knowledge bases over disjoint sets of names. Then $\mathcal{B}_1 \wedge \dots \wedge \mathcal{B}_k$ is consistent iff, for each $i = 1, \dots, k$, \mathcal{B}_i is consistent.*

Proof. Obviously, consistency of $\mathcal{B}_1 \wedge \dots \wedge \mathcal{B}_k$ implies consistency of \mathcal{B}_i for all i , $1 \leq i \leq k$. Conversely, if all the knowledge bases \mathcal{B}_i ($i = 1, \dots, k$) are consistent, then each of them has a model with a countably infinite domain. This means that we can assume without loss of generality that these models have the same domain. In addition, since these models obey the UNA, we can also assume that they interpret the individual names in the same way. Putting together the interpretations of all concept and role names from the separate models yields an interpretation that is a model of all the knowledge bases $\mathcal{B}_1, \dots, \mathcal{B}_k$, and thus a model of $\mathcal{B}_1 \wedge \dots \wedge \mathcal{B}_k$. \square

Looking back at the proof of Lemma 4.3, we see that k is exponential in the size of the input formula ϕ , and that each Boolean \mathcal{ALC} -knowledge bases \mathcal{B}_i has a size that is polynomial in the size of ϕ . Thus, the consistency test for each \mathcal{B}_i takes time exponential in the size of ϕ . Consequently, testing all the knowledge bases $\mathcal{B}_1, \dots, \mathcal{B}_k$ for consistency can be achieved in exponential time.

However, if we simply apply the decision procedure suggested by the claim from the proof of Lemma 4.3, we do not obtain an EXPTIME-procedure. In fact, guessing a subset $\mathcal{S} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ would require non-deterministic exponential time (since we have exponentially many sets to choose from), and testing the propositional LTL formula $\widehat{\phi}_{\mathcal{S}}$ for satisfiability would require exponential space (since the size of $\widehat{\phi}_{\mathcal{S}}$ can be exponential in the size of ϕ , and satisfiability in propositional LTL is PSPACE-complete).

The first problem can easily be avoided. Instead of guessing an appropriate set \mathcal{S} , we compute the maximal one: let $\widehat{\mathcal{S}}$ consist of all sets $X \subseteq \{p_1, \dots, p_n\}$ such that the Boolean \mathcal{ALC} -knowledge base

$$\mathcal{B}_X := \bigwedge_{p_j \in X} \alpha_j \wedge \bigwedge_{p_j \notin X} \neg \alpha_j$$

is consistent. Note that we need not rename flexible names here since the knowledge bases \mathcal{B}_X are considered separately.

LEMMA 5.2. *The \mathcal{ALC} -LTL formula ϕ is satisfiable iff the propositional LTL formula $\widehat{\phi}_{\widehat{\mathcal{S}}}$ is satisfiable.*

Proof. The “if” direction is an immediate consequence of Lemma 5.1, the claim shown in the proof of Lemma 4.3, and the definition of $\widehat{\mathcal{S}}$.

For the “only if” direction, assume that ϕ is satisfiable. By the claim shown in the proof of Lemma 4.3, this implies that there is a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ such that $\widehat{\phi}_{\mathcal{S}}$ is satisfiable and $\bigwedge_{1 \leq i \leq k} \mathcal{B}_{X_i}$ is consistent. Consistency of $\bigwedge_{1 \leq i \leq k} \mathcal{B}_{X_i}$ implies that the knowledge bases \mathcal{B}_{X_i} ($i = 1, \dots, k$) are consistent, and thus we have $\mathcal{S} \subseteq \widehat{\mathcal{S}}$. Consequently, satisfiability of $\widehat{\phi}_{\mathcal{S}}$ implies satisfiability of $\widehat{\phi}_{\widehat{\mathcal{S}}}$. \square

The set $\widehat{\mathcal{S}}$ can be computed in time exponential in the size of ϕ . In fact, there are exponentially many sets $X \subseteq \{p_1, \dots, p_n\}$ to be considered, and testing consistency of \mathcal{B}_X for each of these sets can be done in exponential time.

This leaves us with the problem of testing satisfiability of the propositional LTL formula

$$\widehat{\phi}_{\widehat{\mathcal{S}}} = \widehat{\phi} \wedge \square \left(\bigvee_{X \in \widehat{\mathcal{S}}} \left(\bigwedge_{p \in X} p \wedge \bigwedge_{p \notin X} \neg p \right) \right)$$

in time exponential in the size of ϕ . Since the size of $\widehat{\phi}$ is bounded by the size of ϕ , it is sufficient to give an exponential upper bound in the size of $\widehat{\phi}$. To this end, note that the only effect of the box-formula in $\widehat{\phi}_{\widehat{\mathcal{S}}}$ is to restrict the worlds w in a propositional LTL structure satisfying $\widehat{\phi}$ to being induced by one of the elements of $\widehat{\mathcal{S}}$. Given a world w in a propositional LTL structure, we say that it is induced by a set $X \subseteq \{p_1, \dots, p_n\}$ (written $w = w_X$) iff we have $p_i \in X$ iff w makes p_i true ($i = 1, \dots, n$).

LEMMA 5.3. *The propositional LTL structure $\widehat{\mathcal{J}} = (w_\iota)_{\iota=0,1,\dots}$ satisfies $\widehat{\phi}_{\widehat{\mathcal{S}}}$ iff it satisfies $\widehat{\phi}$ and for every world w_ι of $\widehat{\mathcal{J}}$ there is a set $X \in \widehat{\mathcal{S}}$ such that $w_\iota = w_X$.*

One way of deciding satisfiability of a propositional LTL formula $\widehat{\phi}$ is to construct a Büchi automaton $\mathcal{A}_{\widehat{\phi}}$ that accepts the propositional LTL structures satisfying $\widehat{\phi}$ [Wolper et al. 1983]. To be more precise, let $\Sigma := \mathcal{P}(\{p_1, \dots, p_n\})$. Then the propositional LTL structure $\widehat{\mathcal{I}} = (w_\iota)_{\iota=0,1,\dots}$ can be represented by an infinite word $X_0X_1\dots$ over Σ , where X_ι is such that $w_\iota = w_{X_\iota}$. The Büchi automaton $\mathcal{A}_{\widehat{\phi}}$ is built such that it accepts exactly those infinite words over Σ that represent propositional LTL structures satisfying $\widehat{\phi}$. Consequently, $\widehat{\phi}$ is satisfiable iff the language accepted by $\mathcal{A}_{\widehat{\phi}}$ is non-empty. The size of $\mathcal{A}_{\widehat{\phi}}$ is exponential in the size of $\widehat{\phi}$, and the emptiness test for Büchi automata is polynomial in the size of the automaton.

Given such an automaton $\mathcal{A}_{\widehat{\phi}}$ for $\widehat{\phi}$ we can easily modify it into one accepting exactly the words representing propositional LTL structures satisfying $\widehat{\phi}_{\widehat{\mathcal{S}}}$. In fact, we just need to remove all transitions that use a letter from $\Sigma \setminus \widehat{\mathcal{S}}$. Obviously, this modification can be done in time polynomial in the size of $\mathcal{A}_{\widehat{\phi}}$, and thus in time exponential in the size of $\widehat{\phi}$. The size of the resulting automaton is obviously still only exponential in the size of $\widehat{\phi}$, and thus its emptiness can be tested in time exponential in the size of $\widehat{\phi}$. This yields the desired procedure that can check satisfiability of $\widehat{\phi}_{\widehat{\mathcal{S}}}$ in time exponential in the size of $\widehat{\phi}$. Overall, we have thus proved an EXPTIME upper bound for satisfiability in \mathcal{ALC} -LTL without rigid names.

THEOREM 5.4. *Satisfiability without rigid names is EXPTIME-complete both in \mathcal{ALC} -LTL and in \mathcal{ALC} -LTL $_{gGCI}$.*

Proof. We have just shown that satisfiability in \mathcal{ALC} -LTL without rigid names in EXPTIME. Obviously, this also yields an EXPTIME upper bound for the restricted case of \mathcal{ALC} -LTL $_{gGCI}$.

For the hardness part of the theorem, it is obviously sufficient to show EXPTIME-hardness for \mathcal{ALC} -LTL $_{gGCI}$. This follows from the well-known fact that, in \mathcal{ALC} , satisfiability of a concept C w.r.t. a single GCI $C_1 \sqsubseteq C_2$ is EXPTIME-complete [Schild 1991]. In fact, C is satisfiable w.r.t. $C_1 \sqsubseteq C_2$ iff the \mathcal{ALC} -LTL $_{gGCI}$ formula $\Box(C_1 \sqsubseteq C_2) \wedge a : C$ is satisfiable. \square

6. REASONING WITH RIGID CONCEPTS

In this section, we consider the case where rigid concept names are available, but not rigid role names. First, note that, in contrast to temporal DLs where temporal operators may occur inside concept descriptions, rigid concept names cannot easily be expressed within the logic without rigid concept names. In fact, the GCIs $A \sqsubseteq \Box A$ and $\neg A \sqsubseteq \Box \neg A$ express the condition that A must be interpreted in a rigid way. However, they are not allowed by the syntax of \mathcal{ALC} -LTL since the box is applied directly to a concept, and not to an axiom.

We will show below that, for \mathcal{ALC} -LTL, the presence of rigid concept names indeed increases the complexity of the satisfiability problem, unless GCIs are restricted to being global. First, we treat the case of arbitrary GCIs, and then the special case of global GCIs.

THEOREM 6.1. *Satisfiability in \mathcal{ALC} -LTL w.r.t. rigid concepts is NEXPTIME-complete.*

6.1 The complexity lower bound

We show NEXPTIME-hardness of satisfiability in \mathcal{ALC} -LTL w.r.t. rigid concepts by a reduction from the 2^{n+1} -bounded domino problem [Lewis 1978; Börger et al. 1997].

LEMMA 6.2. *Satisfiability in \mathcal{ALC} -LTL w.r.t. rigid concepts is NEXPTIME-hard.*

Proof. First, we introduce the bounded version of the domino problem used in the reduction. A *domino system* is a triple $\mathfrak{D} = (D, H, V)$, where D is a finite set of *domino types* and $H, V \subseteq D \times D$ are the horizontal and vertical matching conditions. Let \mathfrak{D} be a domino system and $I = d_0, \dots, d_{n-1} \in D^*$ an *initial condition*, i.e. a sequence of domino types of length $n > 0$. A mapping $\tau : \{0, \dots, 2^{n+1} - 1\} \times \{0, \dots, 2^{n+1} - 1\} \rightarrow D$ is a *2^{n+1} -bounded solution of \mathfrak{D} respecting the initial condition I* iff, for all $x, y < 2^{n+1}$, the following holds:

- if $\tau(x, y) = d$ and $\tau(x \oplus_{2^{n+1}} 1, y) = d'$, then $(d, d') \in H$;
- if $\tau(x, y) = d$ and $\tau(x, y \oplus_{2^{n+1}} 1) = d'$, then $(d, d') \in V$;
- $\tau(i, 0) = d_i$ for $i < n$;

where $\oplus_{2^{n+1}}$ denotes addition modulo 2^{n+1} .

It is well-known (see [Lewis 1978] and [Börger et al. 1997], Theorem 6.1.2) that there is a domino system $\mathfrak{D} = (D, H, V)$ such that the following problem is NEXPTIME-hard: given an initial condition $I = d_0, \dots, d_{n-1} \in D^*$, does \mathfrak{D} have a 2^{n+1} -bounded solution respecting I or not?

We show that this problem can be reduced in polynomial time to satisfiability in \mathcal{ALC} -LTL w.r.t. rigid concepts. Interestingly, in our reduction we do not use any role names. All we need are the following concept and individual names:

- a single individual name a ;
- the elements of D as *rigid* concept names, and a primed version of them as *flexible* concept names;
- rigid* concept names X_0, \dots, X_n and Y_0, \dots, Y_n that are used to realize two binary counters modulo 2^{n+1} , where the X -counter describes the horizontal and the Y -counter the vertical position of a domino;
- flexible* concept names $Z_0, \dots, Z_{2^{n+1}}$ that are used to realize a binary counter modulo $2^{2^{n+1}}$, whose function will be explained below;
- an auxiliary *flexible* concept name N .

Intuitively, the first $n+1$ bits of the Z -counter are used to represent 2^{n+1} horizontal components $0 \leq x < 2^{n+1}$, and the second $n+1$ bits of the Z -counter are used to represent 2^{n+1} vertical components $0 \leq y < 2^{n+1}$. By counting with the Z -counter up to $2^{2^{n+1}}$ in the temporal dimension, we ensure that every position $(x, y) \in \{0, \dots, 2^{n+1} - 1\} \times \{0, \dots, 2^{n+1} - 1\}$ is represented *in some world*. The counting is done using the individual name a , i.e., we enforce that, for every possible value of the Z -counter, there is a world where a belongs to the concepts from the corresponding subset of $\{Z_0, \dots, Z_{2^{n+1}}\}$. The rigid concept names X_0, \dots, X_n and Y_0, \dots, Y_n are then used to ensure that, *in every world*, there are individuals belonging to subsets of these concepts such that every position $(x, y) \in \{0, \dots, 2^{n+1} - 1\} \times \{0, \dots, 2^{n+1} -$

1} is realized in this world. Appropriate GCIs are used to ensure that (i) every position represented this way carries exactly one domino type; (ii) the horizontal and vertical matching conditions are respected; and (iii) the initial condition is satisfied.

Recall that we use $\phi \rightarrow \psi$ as an abbreviation for $\neg\phi \vee \psi$, $C \Rightarrow D$ as an abbreviation for $\neg C \sqcup D$, and $C \Leftrightarrow D$ as an abbreviation for $(C \Rightarrow D) \sqcap (D \Rightarrow C)$.

The reduction formula $\phi_{\mathcal{D},I}$ is the conjunction of the following formulae:

—for every possible value of the Z -counter, there is a world where a belongs to the concepts from the corresponding subset of $\{Z_0, \dots, Z_{2n+1}\}$:

$$\begin{aligned} \Box & \left(\bigwedge_{i \leq 2n+1} \left(\bigwedge_{j < i} a : Z_j \rightarrow ((a : Z_i \rightarrow \mathbf{X}(a : \neg Z_i)) \wedge (a : \neg Z_i \rightarrow \mathbf{X}(a : Z_i))) \right) \wedge \right. \\ & \left. \bigwedge_{i \leq 2n+1} \left(\bigvee_{j < i} a : \neg Z_j \rightarrow ((a : Z_i \rightarrow \mathbf{X}(a : Z_i)) \wedge (a : \neg Z_i \rightarrow \mathbf{X}(a : \neg Z_i))) \right) \right) \end{aligned}$$

—the value of the Z -counter is shared by all individuals belonging to the current world: for all $i \leq 2n + 1$

$$\Box ((\top \sqsubseteq Z_i) \vee (\top \sqsubseteq \neg Z_i))$$

—in every world, there is at least one individual for which the combined value of the X - and the Y -counter corresponds to the value of the Z -counter for a (and thus every individual) in this world:

$$\begin{aligned} \Box & \left(\neg(\top \sqsubseteq \neg N) \wedge \right. \\ & \bigwedge_{0 \leq i \leq n} (N \sqcap Z_i \sqsubseteq X_i) \wedge \bigwedge_{n+1 \leq i \leq 2n+1} (N \sqcap Z_i \sqsubseteq Y_{i-(n+1)}) \wedge \\ & \left. \bigwedge_{0 \leq i \leq n} (N \sqcap \neg Z_i \sqsubseteq \neg X_i) \wedge \bigwedge_{n+1 \leq i \leq 2n+1} (N \sqcap \neg Z_i \sqsubseteq \neg Y_{i-(n+1)}) \right) \end{aligned}$$

Since the concept names X_i, Y_i are rigid, this actually ensures that in every world every possible combination of values of the X - and Y -counters is realized by some individual. For a given such combination, the corresponding individual obviously must represent the same value combination in every world. Thus, for every position from $\{0, \dots, 2^{n+1} - 1\} \times \{0, \dots, 2^{n+1} - 1\}$ we have a world representing it with the help of the Z -counter, but we also have an individual representing it globally (i.e., in every world) with the help of the X - and Y -counters.

Having represented all positions from $\{0, \dots, 2^{n+1} - 1\} \times \{0, \dots, 2^{n+1} - 1\}$ in this way, we can now start to enforce an admissible tiling of these positions with domino types (i.e., a solution of the domino problem). As with the positions, we again have two copies of the tiling. One of them uses the primed concept names d' for $d \in D$ to tile the positions represented by the worlds with the help of the Z -counter. The other one uses the unprimed concept names $d \in D$ to tile the positions represented by the individuals with the help of the X - and Y -counters.

—every world gets exactly one domino type, expressed using the primed (and thus flexible) variant of the corresponding concept names:

$$\Box \left(\bigvee_{d \in D} (\top \sqsubseteq d') \wedge \bigwedge_{d, e \in D, d \neq e} (\top \sqsubseteq \neg(d' \sqcap e')) \right)$$

—the domino type of a given world is transferred globally to the individuals representing the same position as the world:

$$\Box \left(\bigcap_{0 \leq i \leq n} (Z_i \Leftrightarrow X_i) \cap \bigcap_{n+1 \leq i \leq 2n+1} (Z_i \Leftrightarrow Y_{i-(n+1)}) \sqsubseteq \bigcap_{d \in D} (d \Leftrightarrow d') \right)$$

Since the concept names d for $d \in D$ are rigid, this type is then associated with the individual in every world. Since every world has exactly one “primed” domino type (which is shared by all its individuals), every individual also has exactly one “unprimed” domino type: the one of the world representing the same position.

The two versions of the tiling can now be used to enforce the horizontal and vertical matching conditions. For example, the fact that the individual representing position $(x, y + 1)$ is present in the world representing position (x, y) can be used to formulate the vertical matching condition.

We use the notation $C_X = C_Z^h$ ($C_Y = C_Z^v$) to say that the value of the X -counter agrees with the value represented by the first $n + 1$ bits of the Z -counter (the value of the Y -counter agrees with the value represented by the second $n + 1$ bits of the Z -counter). Accordingly, $(C_X = C_Z^h + 1 \bmod 2^{n+1})$ says that the value of the X -counter is equal to the value represented by the first $n + 1$ bits of the Z -counter plus 1 (modulo 2^{n+1}). The intended meaning of the notation $(C_Y = C_Z^v + 1 \bmod 2^{n+1})$ should now be obvious. Details on how this can actually be expressed using concept descriptions are given in the proof of Lemma 4.2.

—the horizontal and vertical matching conditions are enforced as follows:

$$\begin{aligned} \Box((C_X = C_Z^h) \cap (C_Y = C_Z^v + 1 \bmod 2^{n+1})) &\sqsubseteq \bigsqcup_{(d,e) \in V} (d' \sqcap e) \\ \Box((C_Y = C_Z^v) \cap (C_X = C_Z^h + 1 \bmod 2^{n+1})) &\sqsubseteq \bigsqcup_{(d,e) \in H} (d' \sqcap e) \end{aligned}$$

For example, the first line looks at an individual that represents position $(x, y + 1)$ in a world that represents position (x, y) , and ensures that the domino type e associated with the individual (expressed by the rigid concept name e) is vertically compatible with the domino type d associated with the world (expressed by the flexible concept name d').

It remains to represent the initial condition $I = d_0, \dots, d_{n-1}$. For this end, we employ the notation $C_Z^v = 0$ and $C_Z^h = i$ for $0 = 1, \dots, n - 1$, with the obvious meaning and the obvious representation by concept descriptions.

—for all $i = 0, \dots, n - 1$:

$$\Box((C_Z^v = 0) \cap (C_Z^h = i) \sqsubseteq d'_i)$$

This finishes the definition of the \mathcal{ALC} -LTL formula $\phi_{\mathfrak{D}, I}$, which is the conjunction of the formulae introduced above. It is easy to see that the size of $\phi_{\mathfrak{D}, I}$ is polynomial in n . Moreover, it is a routing task to prove in detail that $\phi_{\mathfrak{D}, I}$ is satisfiable w.r.t. rigid concepts iff \mathfrak{D} has a 2^{n+1} -bounded solution respecting I . \square

6.2 The complexity upper bound

Here, we show that the complexity lower bound provided by the above lemma is tight.

LEMMA 6.3. *Satisfiability in \mathcal{ALC} -LTL w.r.t. rigid concepts is in NEXPTIME.*

Proof. We want to reuse the claim shown in the proof of Lemma 4.3:

Fact. *The \mathcal{ALC} -LTL formula ϕ is satisfiable w.r.t. rigid names iff there is a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ such that the propositional LTL formula $\hat{\phi}_{\mathcal{S}}$ is satisfiable and the Boolean \mathcal{ALC} -knowledge base $\mathcal{B} := \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ is consistent.*

If we apply this claim in the case where only concept names can be rigid, then we know that the Boolean \mathcal{ALC} -knowledge bases \mathcal{B}_i are built over disjoint sets of role names. The only shared names are the rigid concept names. Assume we have guessed a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$, which can clearly be done within NEXPTIME. The formula $\hat{\phi}_{\mathcal{S}}$ is itself of size exponential in the size of ϕ . However, we can use the same approach as in the proof of Theorem 5.4 to show that its satisfiability can actually be tested in time exponential in the size of ϕ .

Instead of testing the consistency of $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ directly (which would provide us with a double-exponential time bound), we try to reduce this test to k separate consistency tests, each requiring time exponential in the size of ϕ . Before we can do this, we need another guessing step. Assume that A_1, \dots, A_r are all the rigid concept names occurring in ϕ , and that a_1, \dots, a_s are all the individual names occurring in ϕ . We guess a set $\mathcal{T} \subseteq \mathcal{P}(\{A_1, \dots, A_r\})$ and a mapping $\mathfrak{t} : \{a_1, \dots, a_s\} \rightarrow \mathcal{T}$. Again, this guess can clearly be done within NEXPTIME.

Given \mathcal{T} and \mathfrak{t} , we extend the knowledge bases \mathcal{B}_i to knowledge bases $\hat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t})$ as follows. For $Y \subseteq \{A_1, \dots, A_r\}$, let C_Y be the concept description $C_Y := \prod_{A \in Y} A \sqcap \prod_{A \notin Y} \neg A$. We define

$$\hat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t}) := \mathcal{B}_i \wedge \bigwedge_{\mathfrak{t}(a)=Y} a : C_Y \wedge \top \sqsubseteq \bigsqcup_{Y \in \mathcal{T}} C_Y \wedge \bigwedge_{Y \in \mathcal{T}} \neg(\top \sqsubseteq \neg C_Y)$$

Claim. *The Boolean \mathcal{ALC} -knowledge base $\mathcal{B} := \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ is consistent iff there is a set $\mathcal{T} \subseteq \mathcal{P}(\{A_1, \dots, A_r\})$ and a mapping $\mathfrak{t} : \{a_1, \dots, a_s\} \rightarrow \mathcal{T}$ such that the Boolean knowledge bases $\hat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t})$ for $i = 1, \dots, k$ are separately consistent.*

For the “only if” direction, assume that $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ has a model $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$. Let \mathcal{T} consist of those sets $Y \subseteq \{A_1, \dots, A_r\}$ such that there is a $d \in \Delta$ with $d \in (C_Y)^{\mathcal{I}}$, and let \mathfrak{t} be the mapping satisfying $\mathfrak{t}(a) = Y$ iff $a^{\mathcal{I}} \in (C_Y)^{\mathcal{I}}$. It is easy to see that, with this choice of \mathcal{T} and \mathfrak{t} , all the knowledge bases $\hat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t})$ for $i = 1, \dots, k$ have \mathcal{I} as model.

To show the “if” direction, let $\mathcal{T} \subseteq \mathcal{P}(\{A_1, \dots, A_r\})$ be a set and $\mathfrak{t} : \{a_1, \dots, a_s\} \rightarrow \mathcal{T}$ a mapping such that the Boolean knowledge bases $\hat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t})$ for $i = 1, \dots, k$ have models $\mathcal{I}_i = (\Delta_i, \cdot^{\mathcal{I}_i})$. We can assume without loss of generality⁷ that

⁷This is an easy consequence of the fact that Boolean \mathcal{ALC} -knowledge bases always have a finite model and the fact that the countably infinite disjoint union of a model of a Boolean \mathcal{ALC} -knowledge base is again a model of this knowledge base.

- the domains Δ_i are countably infinite, and
- in each model \mathcal{I}_i , the sets $Y \in \mathcal{T}$ are realized by countably infinitely many individuals, i.e., there are countably infinitely many elements $d \in \Delta_i$ such that $d \in (C_Y)^{\mathcal{I}_i}$.

Consequently, the domains Δ_i are partitioned into the countably infinite sets $\Delta_i(Y)$ (for $Y \in \mathcal{T}$), which are defined as follows:

$$\Delta_i(Y) := \{d \in \Delta_i \mid d \in (C_Y)^{\mathcal{I}_i}\}$$

In addition, for each individual name $a \in \{a_1, \dots, a_s\}$ we have

$$a^{\mathcal{I}_i} \in \Delta_i(\mathbf{t}(a))$$

We are now ready to define the model $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ of \mathcal{B} . As the domain of \mathcal{I} we take the domain of \mathcal{I}_1 , i.e., $\Delta := \Delta_1$. Accordingly, we define $\Delta(Y) := \Delta_1(Y)$ for all $Y \in \mathcal{T}$. Because of the properties stated above, there exist bijections $\pi_i : \Delta_i \rightarrow \Delta$ such that

- the restriction of π_i to $\Delta_i(Y)$ is a bijection between $\Delta_i(Y)$ and $\Delta(Y)$;
- π_i respects individual names, i.e., $\pi_i(a^{\mathcal{I}_i}) = a^{\mathcal{I}_1}$ holds for all $a \in \{a_1, \dots, a_s\}$. (Note that we have the unique name assumption for individual names.)

We use these bijections to define the interpretation function $\cdot^{\mathcal{I}}$ of \mathcal{I} as follows:

- If A is a flexible concept name, then \mathcal{B} contains its copies $A^{(i)}$ for $i = 1, \dots, k$. Their interpretation is defined as follows:

$$(A^{(i)})^{\mathcal{I}} := \{\pi_i(d) \mid d \in (A^{(i)})^{\mathcal{I}_i}\}.$$

- All role names r are flexible, and \mathcal{B} contains their copies $r^{(i)}$ for $i = 1, \dots, k$. Their interpretation is defined as follows:

$$(r^{(i)})^{\mathcal{I}} := \{(\pi_i(d), \pi_i(e)) \mid (d, e) \in (r^{(i)})^{\mathcal{I}_i}\}.$$

- If A is a rigid concept names, then we define

$$A^{\mathcal{I}} := A^{\mathcal{I}_1}$$

- If a is an individual name, then we define

$$a^{\mathcal{I}} := a^{\mathcal{I}_1}$$

To prove the claim, it remains to show that \mathcal{I} is a model of all the knowledge bases \mathcal{B}_i ($i = 1, \dots, k$). This is an immediate consequence of the fact that π_i is an isomorphism between \mathcal{I}_i and \mathcal{I} w.r.t. the concept and role names occurring in \mathcal{B}_i . The isomorphism condition is satisfied for flexible concepts and roles by our definition of $\cdot^{\mathcal{I}}$, and for individual names by our assumptions on π_i . Now, let A be a rigid concept name. We must show that $d \in A^{\mathcal{I}_i}$ iff $\pi_i(d) \in A^{\mathcal{I}}$ holds for all $d \in \Delta_i$. Since Δ_i is partitioned into the sets $\Delta_i(Y)$ for $Y \in \mathcal{T}$, we know that there is a $Y \in \mathcal{T}$ such that $d \in \Delta_i(Y)$, i.e., $d \in (C_Y)^{\mathcal{I}_i}$. In addition, we know that $\pi_i(d) \in \Delta(Y)$, i.e., $\pi_i(d) \in (C_Y)^{\mathcal{I}_1} = (C_Y)^{\mathcal{I}}$. This implies that $d \in A^{\mathcal{I}_i}$ iff $A \in Y$ iff $d \in A^{\mathcal{I}}$, which finishes the proof that π_i is an isomorphism between \mathcal{I}_i and \mathcal{I} . Consequently, \mathcal{I} is a model of $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$, which in turn finishes the proof of the claim.

To finish the proof of the lemma, we show that consistency of the knowledge bases

$$\widehat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t}) = \mathcal{B}_i \wedge \bigwedge_{\mathfrak{t}(a)=Y} a : C_Y \wedge \top \sqsubseteq \bigsqcup_{Y \in \mathcal{T}} C_Y \wedge \bigwedge_{Y \in \mathcal{T}} \neg(\top \sqsubseteq \neg C_Y)$$

can be decided in time exponential in the size of the input formula ϕ . Note that this is not trivial. In fact, while the size of $\mathcal{B}_i \wedge \bigwedge_{\mathfrak{t}(a)=Y} a : C_Y$ is polynomial in the size of ϕ , the cardinality of \mathcal{T} , and thus the size of

$$\top \sqsubseteq \bigsqcup_{Y \in \mathcal{T}} C_Y \wedge \bigwedge_{Y \in \mathcal{T}} \neg(\top \sqsubseteq \neg C_Y)$$

can be exponential in the size of ϕ . Decidability of the consistency of $\widehat{\mathcal{B}}_i(\mathcal{T}, \mathfrak{t})$ in time exponential in the size of ϕ is, however, an immediate consequence of the next lemma.

Overall, this completes the proof of the current lemma. In fact, after two NEXPTIME guesses, all we have to do are k (i.e., exponentially many) EXPTIME consistency tests. \square

Let $\widehat{\mathcal{B}}$ be a Boolean \mathcal{ALC} -knowledge base of size n , A_1, \dots, A_r concept names occurring in $\widehat{\mathcal{B}}$, and $\mathcal{T} \subseteq \mathcal{P}(\{A_1, \dots, A_r\})$. Note that this implies that the cardinality of \mathcal{T} is at most exponential in n , and the size of each $Y \in \mathcal{T}$ is linear in n . We say that an interpretation $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$ is a *model of $\widehat{\mathcal{B}}$ w.r.t. \mathcal{T}* if it is a model of $\widehat{\mathcal{B}}$ that additionally satisfies

$$\mathcal{T} = \{Y \subseteq \{A_1, \dots, A_r\} \mid \text{there is } d \in \Delta \text{ such that } d \in (C_Y)^{\mathcal{I}}\}$$

Accordingly, we say that $\widehat{\mathcal{B}}$ is *consistent w.r.t. \mathcal{T}* if $\widehat{\mathcal{B}}$ has a model w.r.t. \mathcal{T} . Obviously, $\widehat{\mathcal{B}}$ is consistent w.r.t. \mathcal{T} iff the knowledge base

$$\widehat{\mathcal{B}} \wedge \top \sqsubseteq \bigsqcup_{Y \in \mathcal{T}} C_Y \wedge \bigwedge_{Y \in \mathcal{T}} \neg(\top \sqsubseteq \neg C_Y)$$

is consistent.

LEMMA 6.4. *Let $\widehat{\mathcal{B}}$ be a Boolean \mathcal{ALC} -knowledge base of size n , A_1, \dots, A_r concept names occurring in $\widehat{\mathcal{B}}$, and $\mathcal{T} \subseteq \mathcal{P}(\{A_1, \dots, A_r\})$. Then, consistency of $\widehat{\mathcal{B}}$ w.r.t. \mathcal{T} can be decided in time exponential in n .*

Proof. The proof is an adaptation of the proof of Theorem 2.27 in [Gabbay et al. 2003], which shows that the consistency problem for Boolean \mathcal{ALC} -knowledge bases is in EXPTIME.

In [Gabbay et al. 2003] it is assumed (without loss of generality) that \mathcal{ALC} -concept descriptions contain only the constructors \sqcap , \neg , and \exists , that all GCIs are of the form $\top \sqsubseteq C$, and that Boolean knowledge bases are built from such GCIs and concept and role assertions using only the connectives \wedge and \neg . In the following, we assume that $\widehat{\mathcal{B}}$ satisfies these restrictions.

Let $ind(\widehat{\mathcal{B}})$ be the set of all individual names occurring in $\widehat{\mathcal{B}}$, and let $con(\widehat{\mathcal{B}})$ and $sub(\widehat{\mathcal{B}})$ respectively denote the closure under negation of the set of all concept descriptions (including subdescriptions) occurring in $\widehat{\mathcal{B}}$ and the set of all subformulae

of $\widehat{\mathcal{B}}$. As usual, we identify $\neg\neg E$ with E . Thus, the cardinalities of the three sets introduced above are polynomial in n .

A *concept type for $\widehat{\mathcal{B}}$* is a set $\mathbf{c} \subseteq \text{con}(\widehat{\mathcal{B}})$ such that

— $C \sqcap D \in \mathbf{c}$ iff $C, D \in \mathbf{c}$, for all $C \sqcap D \in \text{con}(\widehat{\mathcal{B}})$;

— $\neg C \in \mathbf{c}$ iff $C \notin \mathbf{c}$, for all $C \in \text{con}(\widehat{\mathcal{B}})$.

A *formula type for $\widehat{\mathcal{B}}$* is a set $\mathbf{f} \subseteq \text{sub}(\widehat{\mathcal{B}})$ such that

— $\psi \wedge \chi \in \mathbf{f}$ iff $\psi, \chi \in \mathbf{f}$, for all $\psi \wedge \chi \in \text{con}(\widehat{\mathcal{B}})$;

— $\neg\psi \in \mathbf{f}$ iff $\psi \notin \mathbf{f}$, for all $\psi \in \text{con}(\widehat{\mathcal{B}})$.

Obviously, the number of concept and formula types is exponential in n .

A *model candidate for $\widehat{\mathcal{B}}$* is a triple $(\mathcal{S}, \iota, \mathbf{f})$ such that \mathcal{S} is a set of concept types for $\widehat{\mathcal{B}}$, $\iota : \text{ind}(\widehat{\mathcal{B}}) \rightarrow \mathcal{S}$ is a function, and \mathbf{f} is a formula type for $\widehat{\mathcal{B}}$ such that

- (a) $\widehat{\mathcal{B}} \in \mathbf{f}$;
- (b) $a : C \in \mathbf{f}$ implies $C \in \iota(a)$;
- (c) $(a, b) : r \in \mathbf{f}$ implies $\{\neg C \mid \neg\exists r.C \in \iota(a)\} \subseteq \iota(b)$.

The model candidate $(\mathcal{S}, \iota, \mathbf{f})$ for $\widehat{\mathcal{B}}$ is called a *quasimodel for $\widehat{\mathcal{B}}$* if it additionally satisfies

- (d) for every $\mathbf{c} \in \mathcal{S}$ and every $\exists r.C \in \mathbf{c}$, there is $\mathbf{d} \in \mathcal{S}$ such that $\{\neg D \mid \neg\exists r.D \in \mathbf{c}\} \cup \{C\} \subseteq \mathbf{d}$;
- (e) for every $\mathbf{c} \in \mathcal{S}$ and every concept C , if $\neg C \in \mathbf{c}$, then $\top \sqsubseteq C \notin \mathbf{f}$;
- (f) for every concept C , if $\neg(\top \sqsubseteq C) \in \mathbf{f}$, then there is a $\mathbf{c} \in \mathcal{S}$ such that $C \in \mathbf{c}$;
- (g) \mathcal{S} is not empty.

In [Gabbay et al. 2003] it is shown that $\widehat{\mathcal{B}}$ is consistent iff there is a quasimodel for $\widehat{\mathcal{B}}$. In order to characterize consistency of $\widehat{\mathcal{B}}$ w.r.t. \mathcal{T} , we need to add two additional conditions. The quasimodel $(\mathcal{S}, \iota, \mathbf{f})$ for $\widehat{\mathcal{B}}$ *respects \mathcal{T}* if it additionally satisfies

- (h) for every concept type $\mathbf{c} \in \mathcal{S}$, there is a set $Y \in \mathcal{T}$ such that $Y = \mathbf{c} \cap \{A_1, \dots, A_r\}$;
- (k) for every set $Y \in \mathcal{T}$, there is a concept type $\mathbf{c} \in \mathcal{S}$ such that $Y = \mathbf{c} \cap \{A_1, \dots, A_r\}$.

A simple adaptation of the proof in [Gabbay et al. 2003] can be used to show that $\widehat{\mathcal{B}}$ is consistent w.r.t. \mathcal{T} iff there is a quasimodel for $\widehat{\mathcal{B}}$ that respects \mathcal{T} .

Next, we show that the EXPTIME-algorithm for checking the existence of a quasimodel described in [Gabbay et al. 2003] can be adapted to check for the existence of a quasimodel that respects \mathcal{T} . The adapted algorithm works as follows. Given $\widehat{\mathcal{B}}$ and \mathcal{T} , it enumerates all model candidates $(\mathcal{S}, \iota, \mathbf{f})$ for $\widehat{\mathcal{B}}$ where \mathcal{S} is the set of *all* concept types for $\widehat{\mathcal{B}}$. Let $\mathfrak{C}_1, \dots, \mathfrak{C}_N$ be these candidates. As shown in [Gabbay et al. 2003], there are at most exponentially many of these candidates and they can be enumerated in exponential time.

Set $i = 1$ and consider $\mathfrak{C}_i = (\mathcal{S}, \iota, \mathbf{f})$.

Step 1. Go through all concept types in \mathcal{S} . We call a concept type $\mathbf{c} \in \mathcal{S}$ *defective* if one of the following three conditions holds:

- (d) is violated for some concept $\exists r.C \in \mathbf{c}$;
- (e) is violated for some concept C with $\neg C \in \mathbf{c}$;
- (h) is violated.

If we have found a defective concept type \mathbf{c} , and this concept type is not in the range of ι , then we set $\mathcal{S} := \mathcal{S} \setminus \{\mathbf{c}\}$ and continue with Step 1 (i.e., again go through all concept types in \mathcal{S} and check whether one of them is defective). If we have found a defective concept type \mathbf{c} that is in the range of ι , then we stop considering \mathbf{C}_i and go to Step 3. If in some iteration of Step 1 we find that none of the concept types in \mathcal{S} is defective, then we go to Step 2.

Step 2. Check whether the triple $(\mathcal{S}', \iota, \mathbf{f})$ obtained through the application of Step 1 satisfies (f), (g), and (k). If it does, then stop with output “quasimodel respecting \mathcal{T} exists.” Otherwise, go to Step 3.

Step 3. Set $i := i + 1$. If $i \leq N$, then go to Step 1. Otherwise, stop with output “no quasimodel respecting \mathcal{T} .”

It is easy to see that this algorithm yields the output “quasimodel respecting \mathcal{T} exists” iff $\widehat{\mathcal{B}}$ indeed has a quasimodel respecting \mathcal{T} .

It is also not hard to see that the algorithm runs in time exponential in n . The index i goes from 1 to N , where N is at most exponential in n . The cardinality of the set of all concept types is exponential in n , and in each iteration of Step 1 for a fixed index $i \in \{1, \dots, N\}$, one defective concept type is removed (or Step 1 is terminated for this index i). Thus, for a fixed i , the number of iterations of Step 1 is at most exponential in n . Finally, every single iteration of Step 1 needs only exponential time. There are at most exponentially many concept types to be considered, and checking for a violation of (d), (e), or (h) can be done in exponential time. Note in particular that this is also true for (h): one just needs to go through the (exponentially many) elements of \mathcal{T} . Similarly, checking for a violation of (f), (g), or (k) in Step 2 can be done in exponential time. \square

This completes the proof of Theorem 6.1, which states that satisfiability in \mathcal{ALC} -LTL w.r.t. rigid concepts is NEXPTIME-complete.

6.3 Satisfiability in \mathcal{ALC} -LTL $_{gGCI}$ w.r.t. rigid concepts

Here, we show that restricting GCIs to global ones decreases the complexity of the satisfiability problem.

THEOREM 6.5. *The satisfiability problem in \mathcal{ALC} -LTL $_{gGCI}$ w.r.t. rigid concepts is EXPTIME-complete.*

EXPTIME-hardness is an immediate consequence of Theorem 5.4, which states that the problem is already EXPTIME-hard without rigid concepts. Before proving the EXPTIME upper bound, we introduce an additional piece of notation. The conjunction of \mathcal{ALC} -axioms \mathcal{B} is said to be *ϕ -exhaustive* if, for every individual name a and every rigid concept name A occurring in ϕ , either $a : A$ or $a : \neg A$ occurs as a conjunct in \mathcal{B} .

LEMMA 6.6. *Satisfiability in $\mathcal{ALC}\text{-LTL}|_{gGCI}$ w.r.t. rigid concepts is in EXPTIME.*

Proof. Consider an $\mathcal{ALC}\text{-LTL}|_{gGCI}$ formula $\phi = \Box\mathcal{B} \wedge \xi$, where \mathcal{B} is a conjunction of \mathcal{ALC} -axioms and ξ is an $\mathcal{ALC}\text{-LTL}$ formula that does not contain GCIs. We can assume without loss of generality that \mathcal{B} is ϕ -exhaustive. In fact, given an arbitrary Boolean \mathcal{ALC} -knowledge base \mathcal{B} , we can build all the ϕ -exhaustive knowledge bases \mathcal{B}' that are obtained from \mathcal{B} by conjoining to it, for every individual name a and every rigid concept name A occurring in ϕ , either $a : A$ or $a : \neg A$. Obviously, $\phi = \Box\mathcal{B} \wedge \xi$ is satisfiable w.r.t. rigid concepts iff $\Box\mathcal{B}' \wedge \xi$ is satisfiable w.r.t. rigid concepts for one of the extension \mathcal{B}' of \mathcal{B} obtained this way. Since the size of each such an extension is polynomial and there are only exponentially many such extensions, it is sufficient to show that testing satisfiability of $\Box\mathcal{B}' \wedge \xi$ w.r.t. rigid concepts for ϕ -exhaustive knowledge bases \mathcal{B}' is in EXPTIME.

Thus, we assume in the following that \mathcal{B} is ϕ -exhaustive. The proof that satisfiability of $\phi = \Box\mathcal{B} \wedge \xi$ can be tested in EXPTIME combines ideas used in the proofs of Lemma 4.3 and Theorem 5.4. Following the approach used in the proof of Lemma 4.3, we abstract every ABox assertion α_i occurring in ξ by a propositional variable p_i , thus building the propositional LTL-formula $\hat{\xi}$. As with the proof of Theorem 5.4, we compute the set $\hat{\mathcal{S}}$, which consists of those $X \subseteq \{p_1, \dots, p_n\}$ for which the Boolean \mathcal{ALC} -knowledge base

$$\mathcal{B}_X := \mathcal{B} \wedge \bigwedge_{p_j \in X} \alpha_j \wedge \bigwedge_{p_j \notin X} \neg \alpha_j$$

is consistent. This computation can be done in exponential time since it requires exponentially many EXPTIME consistency tests.

Claim. *Let $\phi = \Box\mathcal{B} \wedge \xi$ be such that \mathcal{B} is a ϕ -exhaustive conjunction of \mathcal{ALC} -axioms and ξ is an $\mathcal{ALC}\text{-LTL}$ formula not containing GCIs. Then ϕ is satisfiable w.r.t. rigid concepts iff the propositional LTL formula*

$$\hat{\xi}_{\hat{\mathcal{S}}} := \hat{\xi} \wedge \Box \left(\bigvee_{X \in \hat{\mathcal{S}}} \left(\bigwedge_{p_j \in X} p_j \wedge \bigwedge_{p_j \notin X} \neg p_j \right) \right)$$

is satisfiable.

Note that proving this claim actually completes the proof of our lemma. In fact, as shown in the proof of Theorem 5.4, satisfiability of $\hat{\xi}_{\hat{\mathcal{S}}}$ can be decided in exponential time.

The proof of the “only if” direction of the claim is a simple combination of the arguments used in the proofs of (i) the “only if” direction of Lemma 5.2, and (ii) the “only if” direction of the claim shown in the proof of Lemma 4.3.

To show the “if” direction, assume that $\hat{\xi}_{\hat{\mathcal{S}}}$ is satisfiable. Let $\hat{\mathcal{I}} = (w_\iota)_{\iota=0,1,\dots}$ be a propositional LTL structure satisfying $\hat{\xi}_{\hat{\mathcal{S}}}$. By construction, for every ι there is a set $X_\iota \subseteq \{p_1, \dots, p_n\}$ such that

- $w_\iota = w_{X_\iota}$, i.e., $X_\iota = \{p_j \mid j \in \{1, \dots, n\} \text{ and } w_\iota \text{ makes } p_j \text{ true}\}$.
- the Boolean \mathcal{ALC} -knowledge base $\mathcal{B}_{X_\iota} = \mathcal{B} \wedge \bigwedge_{p_j \in X_\iota} \alpha_j \wedge \bigwedge_{p_j \notin X_\iota} \neg \alpha_j$ is consistent.

Let $\mathcal{I}_\iota = (\Delta^{\mathcal{I}_\iota}, \cdot^{\mathcal{I}_\iota})$ be a model of \mathcal{B}_{X_ι} ($\iota = 0, 1, \dots$). To complete the proof of the claim, we use the models \mathcal{I}_ι to construct models $\mathcal{J}_\iota = (\Delta, \cdot^{\mathcal{J}_\iota})$ of \mathcal{B}_{X_ι} that (i) have a common domain Δ , (ii) interpret the individual names in a rigid way, and (iii) respect rigid concept names. This, together with the fact that $\widehat{\mathcal{I}} = (w_\iota)_{\iota=0,1,\dots}$ satisfies $\widehat{\xi}_{\widehat{\mathcal{S}}}$, then implies that the \mathcal{ALC} -LTL structure $(\mathcal{J}_\iota)_{\iota=0,1,\dots}$ satisfies ϕ and respects rigid concept names.

To simplify this construction, we assume that the models \mathcal{I}_ι have a certain restricted shape. In a DL interpretation $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$, we call an element $x \in \Delta^{\mathcal{I}}$ *named* if there is an individual name a such that $x = a^{\mathcal{I}}$; all other elements of $\Delta^{\mathcal{I}}$ are called *anonymous*. The following is easy to see:

Fact. *Any consistent Boolean \mathcal{ALC} -knowledge base has a model $\mathcal{I} = (\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ that satisfies the following property for every role name r : if $(x, y) \in r^{\mathcal{I}}$ and x is anonymous, then y is also anonymous.*

For example, a standard tableau-based algorithm that test consistency of Boolean \mathcal{ALC} -knowledge bases generates models that satisfy the property stated in the fact.

In the following, we assume that the models \mathcal{I}_ι satisfy this property. We can also assume without loss of generality that the domains $\Delta^{\mathcal{I}_\iota}$ of these models are almost disjoint, except for the named individuals. To be more precise, let $\{a_1, \dots, a_\ell\}$ be the individual names occurring in ϕ , and let $\Theta := \{x_1, \dots, x_\ell\}$ be a set of cardinality ℓ . We assume that $\Delta^{\mathcal{I}_\iota} \cap \Delta^{\mathcal{I}_{\iota'}} = \Theta$ for $\iota \neq \iota'$ and that $a_i^{\mathcal{I}_\iota} = x_i$ for $i = 1, \dots, \ell$.

Let us now build the new models \mathcal{J}_ι based on the models \mathcal{I}_ι satisfying these properties. The constant domain Δ is the union of the domains \mathcal{I}_ι , i.e., Δ is obtained as the following disjoint union:

$$\Delta := \Theta \cup \bigcup_{\iota} (\Delta^{\mathcal{I}_\iota} \setminus \Theta).$$

The interpretation function $\cdot^{\mathcal{J}_\iota}$ of \mathcal{J}_ι is defined as follows:

- We define $a_i^{\mathcal{J}_\iota} := x_i$ for $i = 1, \dots, \ell$.
- If A is a concept name and $x \in \Delta^{\mathcal{J}_\iota}$, we distinguish two cases, depending on whether x is named or anonymous:
 - (1) If x is named, then we define $x \in A^{\mathcal{J}_\iota}$ iff $x \in A^{\mathcal{I}_\iota}$.
 - (2) If x is anonymous, then we define $x \in A^{\mathcal{J}_\iota}$ iff $x \in A^{\mathcal{I}_\kappa}$.
- If r is a role name, $x \in \Delta^{\mathcal{J}_\iota}$ and $y \in \Delta^{\mathcal{J}_\lambda}$, then we distinguish four cases:
 - (1) If x and y are named, then we define $(x, y) \in r^{\mathcal{J}_\iota}$ iff $(x, y) \in r^{\mathcal{I}_\iota}$.
 - (2) If x is named and y is anonymous, then we define $(x, y) \in r^{\mathcal{J}_\iota}$ iff $\iota = \lambda$ and $(x, y) \in r^{\mathcal{I}_\iota}$.
 - (3) If x and y are anonymous, then we define $(x, y) \in r^{\mathcal{J}_\iota}$ iff $\kappa = \lambda$ and $(x, y) \in r^{\mathcal{I}_\kappa}$.
 - (4) If x is anonymous and y is named, then $(x, y) \notin r^{\mathcal{J}_\iota}$.

By construction, the interpretations \mathcal{J}_ι have a common domain and they interpret the individual names in a rigid way. Next, we show that they respect rigid concept names, i.e., we have $A^{\mathcal{J}_{\iota_1}} = A^{\mathcal{J}_{\iota_2}}$ for all rigid concept names A and all $\iota_1, \iota_2 \geq 0$.

Thus, assume that $x \in \Delta^{\mathcal{J}_\kappa}$ is given. We must show that $x \in A^{\mathcal{J}_{\iota_1}}$ iff $x \in A^{\mathcal{J}_{\iota_2}}$. If x is named, then we have $x = a^{\mathcal{I}_{\iota_1}} = a^{\mathcal{I}_{\iota_2}}$ for some individual name $a \in \{a_1, \dots, a_\ell\}$. Since \mathcal{B} was assumed to be ϕ -exhaustive, either $a : A$ or $a : \neg A$ occurs as a conjunct in \mathcal{B} . This, together with the fact that \mathcal{I}_{ι_1} and \mathcal{I}_{ι_2} are models of \mathcal{B} , implies that $x \in A^{\mathcal{I}_{\iota_1}}$ iff $x \in A^{\mathcal{I}_{\iota_2}}$, and thus $x \in A^{\mathcal{J}_{\iota_1}}$ iff $x \in A^{\mathcal{J}_{\iota_2}}$. If x is anonymous, then $x \in A^{\mathcal{I}_{\iota_1}}$ iff $x \in A^{\mathcal{I}_\kappa}$ iff $x \in A^{\mathcal{I}_{\iota_2}}$.

It remains to show that, for every $\iota \geq 0$, the interpretation \mathcal{J}_ι is a model of \mathcal{B}_{X_ι} . The proof of this is based on the following two observations, in which C is assumed to be an arbitrary \mathcal{ALC} -concept description:

- (1) If x is named, then $x \in C^{\mathcal{J}_\iota}$ iff $x \in C^{\mathcal{I}_\iota}$.
- (2) If $x \in \Delta^{\mathcal{I}_\kappa}$ is anonymous, then $x \in C^{\mathcal{J}_\iota}$ iff $x \in C^{\mathcal{I}_\kappa}$.

Before proving these observations by induction on the structure of C , we show that they can indeed be used to prove that \mathcal{J}_ι is a model of \mathcal{B}_{X_ι} .

First, assume that $C \sqsubseteq D$ is a GCI that occurs as a conjunct in \mathcal{B} . Consider an element $x \in C^{\mathcal{J}_\iota}$. We must show that $x \in D^{\mathcal{J}_\iota}$. If x is named, then $x \in C^{\mathcal{J}_\iota}$ implies $x \in C^{\mathcal{I}_\iota}$ by the above Observation 1. Since \mathcal{I}_ι is a model of \mathcal{B} , $x \in C^{\mathcal{I}_\iota}$ implies $x \in D^{\mathcal{I}_\iota}$, and thus $x \in D^{\mathcal{J}_\iota}$, again by Observation 1. Now, assume that x is anonymous, and that it belongs to $\Delta^{\mathcal{I}_\kappa}$ for some $\kappa \geq 0$. Then $x \in C^{\mathcal{J}_\iota}$ implies $x \in C^{\mathcal{I}_\kappa}$ by the above Observation 2. Since \mathcal{I}_κ is a model of \mathcal{B} , $x \in C^{\mathcal{I}_\kappa}$ implies $x \in D^{\mathcal{I}_\kappa}$, and thus $x \in D^{\mathcal{J}_\iota}$, again by Observation 2.

Second, assume that $a : C$ is a concept assertion that occurs as a conjunct in \mathcal{B} or in $\bigwedge_{p_j \in X_\iota} \alpha_j \wedge \bigwedge_{p_j \notin X_\iota} \neg \alpha_j$. (Note that negated concept assertions are also just assertions since $\neg(a : D)$ is equivalent to $a : \neg D$.) We know that \mathcal{I}_ι is a model of \mathcal{B}_{X_ι} , and thus of $a : C$, i.e., $a^{\mathcal{I}_\iota} \in C^{\mathcal{I}_\iota}$. In addition, we have $a^{\mathcal{I}_\iota} = a^{\mathcal{J}_\iota}$, and since $a^{\mathcal{I}_\iota}$ is named, $a^{\mathcal{I}_\iota} \in C^{\mathcal{I}_\iota}$ iff $a^{\mathcal{I}_\iota} \in C^{\mathcal{J}_\iota}$. This shows that \mathcal{J}_ι is a model of the assertion $a : C$.

Third, assume that $(a, b) : r$ is a role assertion that occurs as a conjunct in \mathcal{B} or in $\bigwedge_{p_j \in X_\iota} \alpha_j$. We know that \mathcal{I}_ι is a model of \mathcal{B}_{X_ι} , and thus of $(a, b) : r$, i.e., $(a^{\mathcal{I}_\iota}, b^{\mathcal{I}_\iota}) \in r^{\mathcal{I}_\iota}$. In addition, we have $a^{\mathcal{I}_\iota} = a^{\mathcal{J}_\iota}$ and $b^{\mathcal{I}_\iota} = b^{\mathcal{J}_\iota}$. Since $a^{\mathcal{I}_\iota}, b^{\mathcal{I}_\iota}$ are named, $(a^{\mathcal{I}_\iota}, b^{\mathcal{I}_\iota}) \in r^{\mathcal{I}_\iota}$ implies $(a^{\mathcal{J}_\iota}, b^{\mathcal{J}_\iota}) \in r^{\mathcal{J}_\iota}$, by the definition of $r^{\mathcal{J}_\iota}$. Consequently, \mathcal{J}_ι is a model of $(a, b) : r$.

Finally, assume that $\neg((a, b) : r)$ is a negated role assertion that occurs as a conjunct in $\bigwedge_{p_j \notin X_\iota} \neg \alpha_j$. We know that \mathcal{I}_ι is a model of \mathcal{B}_{X_ι} , and thus of $\neg((a, b) : r)$, i.e., $(a^{\mathcal{I}_\iota}, b^{\mathcal{I}_\iota}) \notin r^{\mathcal{I}_\iota}$. In addition, we have $a^{\mathcal{I}_\iota} = a^{\mathcal{J}_\iota}$ and $b^{\mathcal{I}_\iota} = b^{\mathcal{J}_\iota}$. Since $a^{\mathcal{I}_\iota}, b^{\mathcal{I}_\iota}$ are named, $(a^{\mathcal{I}_\iota}, b^{\mathcal{I}_\iota}) \notin r^{\mathcal{I}_\iota}$ implies $(a^{\mathcal{J}_\iota}, b^{\mathcal{J}_\iota}) \notin r^{\mathcal{J}_\iota}$, by the definition of $r^{\mathcal{J}_\iota}$. Consequently, \mathcal{J}_ι is a model of $\neg((a, b) : r)$. This finishes our proof that \mathcal{J}_ι is a model of \mathcal{B}_{X_ι} .

It remains to show that the above observations are in fact correct. We prove this by induction on the structure of C . The base case (C is a concept name) is trivial by the definition of the interpretations \mathcal{J}_ι . For the induction step, it is sufficient to consider conjunction, negation, and existential restrictions. Both conjunction and negation are trivial. Thus, the only interesting case is where C is of the form $\exists r.D$.

First, assume that x is named. If $x \in (\exists r.D)^{\mathcal{I}_\iota}$, then we know that there is an element $y \in \Delta^{\mathcal{I}_\iota}$ such that $(x, y) \in r^{\mathcal{I}_\iota}$ and $y \in D^{\mathcal{I}_\iota}$. If y is named, then this implies $(x, y) \in r^{\mathcal{J}_\iota}$ and $y \in D^{\mathcal{J}_\iota}$, where the first statement holds by the definition

of $r^{\mathcal{J}_\iota}$ and the second by the induction hypothesis. This shows $x \in (\exists r.D)^{\mathcal{J}_\iota}$. Now, assume that y is anonymous. By our assumption that, except for the named individuals, the domains of the models \mathcal{I}_κ are disjoint, we know that $(x, y) \in r^{\mathcal{I}_\iota}$ implies $y \in \Delta^{\mathcal{I}_\iota}$. As before, we obtain $(x, y) \in r^{\mathcal{J}_\iota}$ and $y \in D^{\mathcal{J}_\iota}$ by the definition of $r^{\mathcal{J}_\iota}$ and the induction hypothesis, respectively. This shows $x \in (\exists r.D)^{\mathcal{J}_\iota}$ also in this case.

Conversely, assume that x is named and $x \in (\exists r.D)^{\mathcal{J}_\iota}$. Then we know that there is an element $y \in \Delta$ such that $(x, y) \in r^{\mathcal{J}_\iota}$ and $y \in D^{\mathcal{J}_\iota}$. If y is named, then this implies $(x, y) \in r^{\mathcal{I}_\iota}$ and $y \in D^{\mathcal{I}_\iota}$, by the definition of $r^{\mathcal{J}_\iota}$ and the induction hypothesis, respectively. If y is anonymous, then the definition of $r^{\mathcal{J}_\iota}$ yields $y \in \Delta^{\mathcal{I}_\iota}$ and $(x, y) \in r^{\mathcal{I}_\iota}$. Since $y \in \Delta^{\mathcal{I}_\iota}$, $y \in D^{\mathcal{J}_\iota}$ implies $y \in D^{\mathcal{I}_\iota}$, by the induction hypothesis. Thus, we have $x \in (\exists r.D)^{\mathcal{I}_\iota}$ in both cases.

Second, assume that $x \in \Delta^{\mathcal{I}_\kappa}$ is anonymous. If $x \in (\exists r.D)^{\mathcal{I}_\kappa}$, then we know that there is an element $y \in \Delta^{\mathcal{I}_\kappa}$ such that $(x, y) \in r^{\mathcal{I}_\kappa}$ and $y \in D^{\mathcal{I}_\kappa}$. By our assumption on the models \mathcal{I}_κ , we know that y is also anonymous. The definition of $r^{\mathcal{J}_\iota}$ thus yields $(x, y) \in r^{\mathcal{J}_\iota}$, and the induction hypothesis yields $y \in \Delta^{\mathcal{J}_\iota}$. Thus, we have $x \in (\exists r.D)^{\mathcal{J}_\iota}$.

Conversely, assume that $x \in \Delta^{\mathcal{I}_\kappa}$ is anonymous, and $x \in (\exists r.D)^{\mathcal{J}_\iota}$. Then we know that there is an element $y \in \Delta$ such that $(x, y) \in r^{\mathcal{J}_\iota}$ and $y \in D^{\mathcal{J}_\iota}$. The definition of $r^{\mathcal{J}_\iota}$ implies that y is also anonymous, and that $y \in \Delta^{\mathcal{I}_\kappa}$ and $(x, y) \in r^{\mathcal{I}_\kappa}$. The induction hypothesis thus yields $y \in D^{\mathcal{I}_\kappa}$. This shows $x \in (\exists r.D)^{\mathcal{I}_\kappa}$. \square

When defining the notion of an \mathcal{ALC} -LTL formula *with global GCIs*, we have restricted these formulae to being of the form $\phi = \Box \mathcal{B} \wedge \xi$ where \mathcal{B} is a conjunction of \mathcal{ALC} -axioms and ξ is an \mathcal{ALC} -LTL formula that does not contain GCIs. Allowing also for negated \mathcal{ALC} -axioms as conjuncts in \mathcal{B} would only add syntactic sugar, but would not increase the expressiveness of $\mathcal{ALC}\text{-LTL}|_{gGCI}$. In fact, a negated assertion $\neg\alpha$ occurring as a conjunct in \mathcal{B} can be moved as a conjunct $\Box\neg\alpha$ to ξ , and a negated GCI $\neg(C \sqsubseteq D)$ can be replaced by an assertion $a : C \sqcap \neg D$ in \mathcal{B} , where a is a new individual name.

One might think that it is even possible to relax the condition such that \mathcal{B} is an arbitrary Boolean \mathcal{ALC} -knowledge base. This is, however, not the case since it would increase the complexity of the satisfiability problem to NEXPTIME.

DEFINITION 6.7. *We say that ϕ is an \mathcal{ALC} -LTL formula with global Boolean knowledge base iff it is of the form $\phi = \Box \mathcal{B} \wedge \xi$ where \mathcal{B} is a Boolean \mathcal{ALC} -knowledge base and ξ is an \mathcal{ALC} -LTL formula that does not contain GCIs.*

A careful analysis of the proof of Lemma 6.2 shows that the \mathcal{ALC} -LTL formula constructed in the reduction is actually an \mathcal{ALC} -LTL formula with global Boolean knowledge base.

COROLLARY 6.8. *Satisfiability of \mathcal{ALC} -LTL w.r.t. rigid concepts and with global Boolean knowledge bases is NEXPTIME-complete.*

7. RESTRICTING THE TEMPORAL COMPONENT

In this section, we consider the fragment $\mathcal{ALC}\text{-LTL}|_\diamond$ of \mathcal{ALC} -LTL, in which \diamond is the only temporal operator. Our aim is to prove that satisfiability in $\mathcal{ALC}\text{-LTL}|_\diamond$

w.r.t. rigid names is in EXPTIME. The main reason for this is that we can restrict the attention to \mathcal{ALC} -LTL structures respecting rigid concept and role names that consist of only polynomially many distinct \mathcal{ALC} -interpretations. Before we can formulate this fact more formally in the next lemma,⁸ we need to introduce some more notations. The *weight* of the \mathcal{ALC} -LTL structure $\mathfrak{I} = (\mathcal{I}_i)_{i=0,1,\dots}$ is defined to be the cardinality of the set $\{\mathcal{I}_i \mid i = 0, 1, \dots\}$.⁹ The *set of subformulae* $\text{sub}(\phi)$ of the \mathcal{ALC} -LTL $_{\diamond}$ formula ϕ is defined in the obvious way, i.e., $\text{sub}(\phi) = \{\phi\}$ if ϕ is an \mathcal{ALC} -axiom, $\text{sub}(\diamond\phi) = \{\diamond\phi\} \cup \text{sub}(\phi)$, $\text{sub}(\phi \wedge \psi) = \{\phi \wedge \psi\} \cup \text{sub}(\phi) \cup \text{sub}(\psi)$, etc. The size of the \mathcal{ALC} -LTL $_{\diamond}$ formula ϕ is denoted by $|\phi|$.

LEMMA 7.1. *If the \mathcal{ALC} -LTL $_{\diamond}$ formula ϕ is satisfiable w.r.t. rigid names, then there is an \mathcal{ALC} -LTL structure \mathfrak{K} respecting rigid concept and role names of weight at most $|\phi| + 2$ such that $\mathfrak{K}, 0 \models \phi$.*

Proof. Let \mathfrak{I} be an \mathcal{ALC} -LTL structure respecting rigid concept and role names such that $\mathfrak{I}, 0 \models \phi$. For each $\psi \in \text{sub}(\phi)$ we use $P(\psi)$ to denote the set of time points at which \mathfrak{I} makes ψ true, i.e.,

$$P(\psi) := \{i \mid i \geq 0 \text{ and } \mathfrak{I}, i \models \psi\}.$$

We claim that there is an $\ell > 0$ such that, for all $i \geq \ell$, we have that $\mathfrak{I}, i \models \psi$ implies that $P(\psi)$ is infinite. In fact, since $\text{sub}(\phi)$ is finite, the set

$$P_{\text{fin}} := \bigcup_{\substack{\psi \in \text{sub}(\phi) \\ P(\psi) \text{ finite}}} P(\psi)$$

is also finite. Thus, we can choose ℓ to be the least positive integer not belonging to P_{fin} .

For every $\psi \in \text{sub}(\phi)$, we choose a time point $p(\psi) \geq 0$ as follows:

—If $P(\psi)$ is finite, then $p(\psi)$ is the maximal element of $P(\psi)$, i.e., $p(\psi)$ is the maximal j with $\mathfrak{I}, j \models \psi$. Note that, in this case, we have $p(\psi) < \ell$ and $\mathfrak{I}, p(\psi) \models \psi$.

—If $P(\psi)$ is infinite, then $p(\psi)$ is an arbitrary time point $j \geq \ell$ such that $\mathfrak{I}, j \models \psi$.

Note that, in this case, we have $p(\psi) \geq \ell$ and $\mathfrak{I}, p(\psi) \models \psi$.

Let $\text{ran}(p)$ denote the range of the function $p : \text{sub}(\phi) \rightarrow \{0, 1, 2, \dots\}$, and let $k_0, \dots, k_{m-1}, k_m, \dots, k_{n-1}$ be an enumeration of $\text{ran}(p) \cup \{0, \ell\}$ such that

— $k_0 < \dots < k_{m-1} < k_m < \dots < k_{n-1}$ and

— $m \in \{1, \dots, n-1\}$ is chosen such that $k_i < \ell$ iff $i < m$.

We define the \mathcal{ALC} -LTL structure \mathfrak{K} as the sequence of \mathcal{ALC} -interpretations

$$\mathcal{I}_{k_0} \cdots \mathcal{I}_{k_{m-1}} (\mathcal{I}_{k_m} \cdots \mathcal{I}_{k_{n-1}})^\omega,$$

⁸Note that this lemma is a generalization to \mathcal{ALC} -LTL $_{\diamond}$ of a very similar lemma for LTL $_{\diamond}$, the restriction of propositional LTL to its diamond fragment (see, e.g., Lemma 6.40 in [Blackburn et al. 2001]).

⁹Recall that all the \mathcal{ALC} -interpretations within one \mathcal{ALC} -LTL structure have the same domain. For this reason, we can use exact equality of interpretations rather than equality up to isomorphism when defining the weight of an \mathcal{ALC} -LTL structure.

i.e., $\mathfrak{R} = (\mathcal{J}_i)_{i=0,1,\dots}$ with $\mathcal{J}_i = \mathcal{I}_{f(i)}$, where

— $f(i) = k_i$ for $i < m$;

— $f(i) = k_{m+((i-m) \bmod (n-m))}$ for $i \geq m$.

Clearly, the fact that \mathfrak{J} respects rigid concept names and role names implies that \mathfrak{R} does the same. In addition, since the number of subformulae of ϕ is bounded by the size $|\phi|$ of ϕ , we have $n \leq |\phi| + 2$, and the weight of \mathfrak{R} is bounded by n . Thus, it remains to show that $\mathfrak{R}, 0 \models \phi$. This is an immediate consequence of the following claim and the fact that $f(0) = k_0 = 0$ and $\phi \in \text{sub}(\phi)$.

Claim. *For all $\psi \in \text{sub}(\phi)$ and all $i \geq 0$, we have $\mathfrak{R}, i \models \psi$ iff $\mathfrak{J}, f(i) \models \psi$.*

The proof of the claim is by induction on the structure of ψ . We concentrate on the only non-trivial case, i.e., the case where ψ is of the form $\diamond\chi$.

To show the “if” direction, assume that $\mathfrak{J}, f(i) \models \diamond\chi$. First assume that $P(\chi)$ is finite. Let $s \geq 0$ be maximal such that $\mathfrak{J}, s \models \chi$. Since $\mathfrak{J}, f(i) \models \diamond\chi$, we have $s \geq f(i)$. By the definition of the function p , we also have $p(\chi) = s$. In addition, by the definition of the function f , and since $s \geq f(i)$, there is a $j \geq i$ with $f(j) = s$. The induction hypothesis yields $\mathfrak{R}, j \models \chi$, and thus, by the semantics of the diamond operator, $\mathfrak{R}, i \models \diamond\chi$.

Now assume that $P(\chi)$ is infinite. By the definition of the function p , we have $p(\chi) \geq \ell$. In addition, by the definition of the function f , there thus is a $j > i$ with $f(j) = p(\chi)$. Since $\mathfrak{J}, p(\chi) \models \chi$, the induction hypothesis yields $\mathfrak{R}, j \models \chi$, and thus $\mathfrak{R}, i \models \diamond\chi$. This completes the proof of the “if” direction.

To show the “only if” direction, assume that $\mathfrak{R}, i \models \diamond\chi$. First, assume that $P(\chi)$ is finite. We know that there is a $j \geq i$ with $\mathfrak{R}, j \models \chi$. The induction hypothesis yields $\mathfrak{J}, f(j) \models \chi$. Since $P(\chi)$ is finite, we have $f(j) < \ell$, and the definition of f together with $j \geq i$ yields $f(j) \geq f(i)$. By the semantics of the diamond operator, this implies $\mathfrak{J}, f(i) \models \diamond\chi$.

If $P(\chi)$ is infinite, then we have $\mathfrak{J}, s \models \diamond\chi$ for every $s \geq 0$, and thus $\mathfrak{J}, f(i) \models \diamond\chi$ is trivially satisfied. This completes the proof of the claim, and thus of the lemma. \square

Given this lemma, we can now prove that satisfiability of $\mathcal{ALC}\text{-LTL}|_{\diamond}$ formulae w.r.t. rigid names can be decided within deterministic exponential time.

LEMMA 7.2. *Satisfiability in $\mathcal{ALC}\text{-LTL}|_{\diamond}$ w.r.t. rigid names is in EXPTIME.*

Proof. The proof of this lemma is very similar to the one of Lemma 4.3. We use the same notation as in that proof. The first step is to establish the following claim, where the set of propositional $\text{LTL}|_{\diamond}$ formulae is defined in the obvious way, and where we use n to denote the number of \mathcal{ALC} -axioms occurring in ϕ .

Claim. *The $\mathcal{ALC}\text{-LTL}|_{\diamond}$ formula ϕ is satisfiable w.r.t. rigid names iff there is a set $\mathcal{S} = \{X_1, \dots, X_k\} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ of cardinality $k \leq |\phi| + 2$ such that the propositional $\text{LTL}|_{\diamond}$ formula $\hat{\phi}_{\mathcal{S}}$ is satisfiable and the Boolean \mathcal{ALC} -knowledge base $\mathcal{B} := \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ is consistent.*

The “if” direction of this claim is an immediate consequence of the corresponding claim shown in the proof of Lemma 4.3. For the “only if” direction, we can use the proof of the “only if” direction of the corresponding claim shown in the proof of Lemma 4.3. The only difference is that we start with an \mathcal{ALC} -LTL structure \mathfrak{J} respecting rigid concept and role names of *weight at most* $|\phi|+2$ such that $\mathfrak{R}, 0 \models \phi$. The existence of such a structure is guaranteed by Lemma 7.1. It is easy to see that then the set $\mathcal{S} = \{X_1, \dots, X_k\}$ defined in the proof of Lemma 4.3 indeed is of cardinality $k \leq |\phi| + 2$.

This completes the proof of the claim. It remains to show that the claim provides us with a decision procedure for satisfiability in \mathcal{ALC} -LTL $_{\diamond}$ w.r.t. rigid names that runs in deterministic exponential time. Let $m := |\phi|$. There are $\leq (2^n)^{m+2} \leq 2^{m(m+2)}$ subsets $\mathcal{S} \subseteq \mathcal{P}(\{p_1, \dots, p_n\})$ of cardinality $\leq m+2$ to be considered. The size of each such subset $\mathcal{S} = \{X_1, \dots, X_k\}$ of cardinality $k \leq m+2$ is polynomial in m , and thus, the size of both $\hat{\phi}_{\mathcal{S}}$ and $\mathcal{B} = \bigwedge_{1 \leq i \leq k} \mathcal{B}_i$ is polynomial in m . Since satisfiability in propositional LTL is in PSPACE and the consistency problems for Boolean \mathcal{ALC} -knowledge bases is in EXPTIME, this completes the proof of the lemma. \square

EXPTIME-hardness of satisfiability in \mathcal{ALC} -LTL $_{\diamond}$ can be shown as in the proof of Theorem 5.4 by a reduction of the well-known EXPTIME-hard problem of satisfiability of an \mathcal{ALC} -concept C w.r.t. a single GCI $C_1 \sqsubseteq C_2$. In fact, C is satisfiable w.r.t. $C_1 \sqsubseteq C_2$ iff the \mathcal{ALC} -LTL $_{\diamond}$ formula $a : C \wedge \neg \diamond \neg (C_1 \sqsubseteq C_2)$ is satisfiable w.r.t. rigid names.

THEOREM 7.3. *Satisfiability in \mathcal{ALC} -LTL $_{\diamond}$ w.r.t. rigid names is an EXPTIME-complete problem.*

Obviously, the above reduction does not depend on the availability of rigid names, and the EXPTIME decision procedure described in the proof of Lemma 7.2 also works in case there are only rigid concepts or no rigid names at all.

COROLLARY 7.4. *Satisfiability in \mathcal{ALC} -LTL $_{\diamond}$ w.r.t. rigid concepts (without rigid names) is an EXPTIME-complete problem.*

8. CONCLUSION

We have analyzed the complexity of combinations of \mathcal{ALC} with LTL where temporal operators can only be applied to TBox axioms and ABox assertions, but not to concepts. One main result is that, under this syntactic restriction, reasoning remains decidable (in elementary time) even in the presence of rigid roles. The 2-EXPTIME complexity can be reduced to NEXPTIME by disallowing rigid roles, or alternatively to EXPTIME by disallowing the use of temporal (and Boolean!) operators on TBox axioms, i.e., switching to global TBoxes. Further reductions in complexity are possible by disallowing rigid concept and restricting the temporal component to the operators \diamond and \square . We once more refer to Table II for full details.

An interesting direction for future research is to consider weakenings of the temporalized DLs studied in this paper that are obtained by replacing \mathcal{ALC} with more lightweight description logics such as \mathcal{EL} and DL-Lite. For the \mathcal{EL} case, we suspect

that the lower bounds presented in this paper can be modified to go through, which would imply that the complexity of reasoning does not decrease.¹⁰ On the other hand, it seems reasonable to expect that a replacement of \mathcal{ALC} with DL-Lite does result in a significant decrease of complexity.

Another topic of interest is to replace constant domains with expanding domains, decreasing domains, or varying domains. In traditional temporalized DLs, this choice usually has no considerable impact on the complexity of reasoning. Due to the absence of temporal operators on concepts, though, it is not clear that this is also the case for the temporalized DLs studied in this paper. In contrast, it is quite conceivable that adopting, say, varying domains leads to a decrease of complexity for some of the DLs defined here.

REFERENCES

- ARTALE, A. AND FRANCONI, E. 1998. A temporal description logic for reasoning about actions and plans. *J. of Artificial Intelligence Research* 9, 463–506.
- ARTALE, A. AND FRANCONI, E. 2000. A survey of temporal extensions of description logics. *Ann. of Mathematics and Artificial Intelligence* 30, 171–210.
- ARTALE, A. AND FRANCONI, E. 2001. Temporal description logics. In *Handbook of Time and Temporal Reasoning in Artificial Intelligence*, D. Gabbay, M. Fisher, and L. Vila, Eds. The MIT Press.
- ARTALE, A., FRANCONI, E., WOLTER, F., AND ZAKHARYASCHEV, M. 2002. A temporal description logic for reasoning over conceptual schemas and queries. In *Proc. of the 8th Eur. Workshop on Logics in Artificial Intelligence (JELIA'02)*. Lecture Notes in Computer Science, vol. 2424. Springer-Verlag, 98–110.
- ARTALE, A., KONTCHAKOV, R., LUTZ, C., WOLTER, F., AND ZAKHARYASCHEV, M. 2007. Temporalising tractable description logics. In *Proceedings of the Fourteenth International Symposium on Temporal Representation and Reasoning*. IEEE Computer Society Press.
- ARTALE, A., KONTCHAKOV, R., RYZHIKOV, V., AND ZAKHARYASCHEV, M. 2009. Extending DL-Lite sometime in the future. In *Proceedings of the 2009 International Workshop on Description Logics (DL2009)*, B. C. Grau, I. Horrocks, B. Motik, and U. Sattler, Eds. CEUR Electronic Workshop Proceedings, <http://ceur-ws.org/Vol-477/>.
- ARTALE, A., LUTZ, C., AND TOMAN, D. 2007. A description logic of change. In *Proc. of the 20th Int. Joint Conf. on Artificial Intelligence (IJCAI 2007)*. 218–223.
- BAADER, F., BAUER, A., AND LIPPMANN, M. 2009. Runtime verification using a temporal description logic. In *Proc. of the 7th Int. Symp. on Frontiers of Combining Systems (FroCoS 2009)*, S. Ghilardi and R. Sebastiani, Eds. Lecture Notes in Computer Science, vol. 5749. Springer-Verlag, 149–164.
- BAADER, F., CALVANESE, D., MCGUINNESS, D., NARDI, D., AND PATEL-SCHNEIDER, P. F., Eds. 2003. *The Description Logic Handbook: Theory, Implementation, and Applications*. Cambridge University Press.
- BAADER, F., GHILARDI, S., AND LUTZ, C. 2008. LTL over description logic axioms. In *Proc. of the 11th Int. Conf. on Principles of Knowledge Representation and Reasoning (KR 2008)*, G. Brewka and J. Lang, Eds. Morgan Kaufmann, Los Altos, 684–694.
- BLACKBURN, P., DE RIJKE, M., AND VENEMA, Y. 2001. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science, vol. 53. Cambridge University Press.
- BÖRGER, E., GRÄDEL, E., AND GUREVICH, Y. 1997. *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer-Verlag.
- CHANDRA, A. K., KOZEN, D. C., AND STOCKMEYER, L. J. 1981. Alternation. *J. of the ACM* 28, 1, 114–133.

¹⁰This refers to the complexity of implication/subsumption rather than satisfiability, as \mathcal{EL} cannot express contradictions.

- FINGER, M. AND GABBAY, D. 1992. Adding a temporal dimension to a logic system. *J. of Logic, Language and Information* 2, 203–233.
- GABBAY, D., KURUSZ, A., WOLTER, F., AND ZAKHARYASCHEV, M. 2003. *Many-dimensional Modal Logics: Theory and Applications*. Elsevier.
- HORROCKS, I., PATEL-SCHNEIDER, P. F., AND VAN HARMELEN, F. 2003. From SHIQ and RDF to OWL: The making of a web ontology language. *Journal of Web Semantics* 1, 1, 7–26.
- LEWIS, H. R. 1978. Complexity of solvable cases of the decision problem for the predicate calculus. In *19th Annual Symposium on Foundations of Computer Science (FOCS)*. IEEE, 35–47.
- LUTZ, C. 2001. Interval-based temporal reasoning with general TBoxes. In *Proc. of the 17th Int. Joint Conf. on Artificial Intelligence (IJCAI 2001)*. 89–94.
- LUTZ, C., WOLTER, F., AND ZAKHARYASCHEV, M. 2008. Temporal description logics: A survey. In *Proceedings of the Fifteenth International Symposium on Temporal Representation and Reasoning*. IEEE Computer Society Press.
- PNUELI, A. 1977. The temporal logic of programs. In *Proc. of the 18th Annual Symp. on the Foundations of Computer Science (FOCS'77)*. 46–57.
- SCHILD, K. 1991. A correspondence theory for terminological logics: Preliminary report. In *Proc. of the 12th Int. Joint Conf. on Artificial Intelligence (IJCAI'91)*. 466–471.
- SCHILD, K. 1993. Combining terminological logics with tense logic. In *Proc. of the 6th Portuguese Conf. on Artificial Intelligence (EPIA'93)*. Lecture Notes in Computer Science, vol. 727. Springer-Verlag, 105–120.
- SCHMIDT-SCHAUSS, M. AND SMOLKA, G. 1991. Attributive concept descriptions with complements. *Artificial Intelligence* 48, 1, 1–26.
- SCHMIEDEL, A. 1990. A temporal terminological logic. In *Proc. of the 8th Nat. Conf. on Artificial Intelligence (AAAI'90)*. 640–645.
- SCHULZ, S., MARKÓ, K., AND SUNTISRIVARAPORN, B. 2006. Complex occurments in clinical terminologies and their representation in a formal language. In *Proceedings of the First European Conference on SNOMED CT (SMCS'06)*. Copenhagen, Denmark.
- TOBIES, S. 1999. A NEXPTIME-complete description logic strictly contained in C^2 . In *Proc. of the Annual Conf. of the Eur. Assoc. for Computer Science Logic (CSL'99)*, J. Flum and M. Rodríguez-Artalejo, Eds. Lecture Notes in Computer Science, vol. 1683. Springer-Verlag, 292–306.
- WOLPER, P., VARDI, M. Y., AND SISTLA, A. P. 1983. Reasoning about infinite computation paths. In *Proc. of the 24th Annual Symp. on the Foundations of Computer Science (FOCS'83)*. IEEE Computer Society Press, 185–194.
- WOLTER, F. AND ZAKHARYASCHEV, M. 1999. Temporalizing description logics. In *Proc. of the 2nd Int. Workshop on Frontiers of Combining Systems (FroCoS'98)*, M. de Rijke and D. Gabbay, Eds. Wiley, Amsterdam.