

Keys, Nominals, and Concrete Domains

Carsten Lutz

*Theoretical Computer Science, TU Dresden,
D-01062 Dresden, Germany*

LUTZ@TCS.INF.TU-DRESDEN.DE

Carlos Areces

*INRIA Lorraine, Nancy,
54602 Villers les Nancy Cedex, France*

ARECES@LORIA.FR

Ian Horrocks

*Department of Computer Science,
University of Manchester
Oxford Road, Manchester M13 9PL, UK*

HORROCKS@CS.MAN.AC.UK

Ulrike Sattler

*Department of Computer Science,
University of Manchester
Oxford Road, Manchester M13 9PL, UK*

SATTLER@CS.MAN.AC.UK

Abstract

Many description logics (DLs) combine knowledge representation on an abstract, logical level with an interface to “concrete” domains like numbers and strings with built-in predicates such as $<$, $+$, and *prefix-of*. These hybrid DLs have turned out to be useful in several application areas, such as providing the basis for expressive ontology languages. We propose to further extend such DLs with *key constraints* that allow the expression of statements like “US citizens are uniquely identified by their social security number”. Based on this idea, we introduce a number of natural description logics and perform a detailed analysis of their decidability and computational complexity. It turns out that naive extensions with key constraints easily lead to undecidability, whereas more careful extensions yield NEXPTIME-complete DLs for a variety of useful concrete domains.

1. Motivation

Description logics (DLs) are a family of formalisms that allow the representation of and reasoning about conceptual knowledge in a structured and semantically well-understood manner (Baader, McGuinness, Nardi, & Patel-Schneider, 2003). The central entities for representing such knowledge are *concepts*, which are constructed from atomic *concept names* (unary predicates) and *role names* (binary relations) by means of the concept and role constructors offered by a particular DL. For example, in the basic propositionally closed description logic \mathcal{ALC} (Schmidt-Schauß & Smolka, 1991), we can describe a company that has part-time employees but only full-time managers using the concept

$$\text{Company} \sqcap \exists \text{employee. Parttime} \sqcap \forall \text{employee.} (\neg \text{Manager} \sqcup \neg \text{Parttime}).$$

In this example, all uppercase words denote concept names while the lowercase *employee* denotes a role name.

Rather than being viewed only as conceptual entities in a knowledge base, concepts can, more generally, be understood as the central notion in various kinds of class-centered formalisms. In the last decade, this observation has given rise to various new and exciting applications of description logics such as reasoning about database conceptual models expressed in entity-relationship diagrams or object-oriented schemas, and reasoning about ontologies for use in the semantic web, see (Calvanese, Lenzerini, & Nardi, 1998; Calvanese, De Giacomo, & Lenzerini, 1998) and (Baader, Horrocks, & Sattler, 2002a; Horrocks, 2002; Horrocks, Patel-Schneider, & van Harmelen, 2002), respectively. These new applications have, in turn, stimulated research in description logics since the expressive power of existing DLs was insufficient for the new tasks. One important extension of “classical” description logics is with expressive means that allow the integration of numbers and other datatypes: assume, e.g., that we want to continue our example from above by equipping companies with a founding year and employees with a hiring year. Then, we may want to describe companies that were founded before 1970 and state that the hiring year of employees is not prior to the founding year of the employing company. To do this, we obviously need a way to talk about natural numbers (such as 1970) and comparisons between natural numbers.

Nowadays, the standard approach to integrate datatypes into description logics is to extend DLs with *concrete domains* as first proposed by Baader and Hanschke (1991a), see also the survey (Lutz, 2003). More precisely, a concrete domain \mathcal{D} consists of a set (such as the natural numbers) and predicates which are associated with a *fixed* extension over this set¹ (such as the unary $=_0$, the binary $<$, and the ternary $+$). The integration of concrete domains into, say, the description logic \mathcal{ALC} is achieved by adding

1. *abstract features*, i.e. functional roles;
2. *concrete features*, i.e. (partial) functions associating values from the concrete domain (e.g., natural numbers) with logical objects;
3. a concrete domain-based concept constructor.

The DL obtained by extending \mathcal{ALC} in this way is called $\mathcal{ALC}(\mathcal{D})$, where \mathcal{D} denotes a concrete domain that can be viewed as a parameter to the logic. For example, using a suitable concrete domain \mathcal{D} we can describe the constraints formulated above: the concept

$$\text{Employee} \sqcap \exists \text{employer}. (\exists \text{foundingyear}. <_{1970}) \sqcap \exists \text{hiringyear}, \text{employer foundingyear}. \geq$$

describes the set of employees who are employed by a company founded before 1970 and who have a hiring year not prior to the company’s founding year. Here, the term inside parenthesis and the third conjunct are instances of the concrete domain concept constructor (not to be confused with the existential value restriction as in $\exists \text{employee}. \text{Parttime}$), *employer* is an abstract feature, and *foundingyear* and *hiringyear* are concrete features.

Concrete domains are rather important in the applications of DLs mentioned above:

- The standard way of using description logics for reasoning about conceptual database models is to translate a given model into a DL representation and then use a description logic reasoning procedure for detecting inconsistencies and inferring consequences

1. This fixed extension is why these predicates are often called “built-in”.

of the information provided explicitly in the model such as additional, implicit containments between entities/classes (Calvanese et al., 1998). Since most databases store “concrete” data like numbers and strings, constraints concerning such data are usually part of the conceptual model and should thus also be captured by the description logic used for reasoning. Indeed, the above example concepts can be viewed as the DL encoding of constraints from a database about companies and their employees. As discussed in (Lutz, 2002c), description logics with concrete domains are well-suited for conceptual modeling applications involving concrete datatypes.

- In the construction of ontologies for the semantic web, so-called *concrete datatypes* play a prominent role (Horrocks et al., 2002). Say, for example, that we want to construct an ontology which can be used for describing car dealers’ web pages and web services. In such an ontology, concrete datatypes such as prices, manufacturing years, and names of car models will doubtlessly be very important. To formulate this ontology using a DL, we need a way to represent these concrete datatypes. Consequently, almost all DLs that have been proposed as an ontology language for the semantic web are equipped with some form of concrete domain (Fensel, van Harmelen, Horrocks, McGuinness, & Patel-Schneider, 2001; Horrocks et al., 2002; Dean, Connolly, van Harmelen, Hendler, Horrocks, McGuinness, Patel-Schneider, & Stein, 2002).

In this paper, we propose to further extend the expressive power of description logics with concrete domains in a way that is useful both for knowledge representation and the two applications sketched above. Let us describe the basic idea — which is to use concrete features for defining “key constraints”— using three examples:

1. Suppose that, in a knowledge representation application, we represent nationalities by concept names such as `US` and `German` and, for `US` citizens, we store the social security number using a concrete feature `ssn`. Then it would be natural to state that `US` citizens are uniquely identified by their social security number, i.e. any two distinct instances of

$$\text{Human} \sqcap \exists \text{nationality.US}$$

must have different values for the `ssn` feature. In our extension of DLs with concrete domains, this can be expressed by using the *key definition*

$$(\text{ssn keyfor Human} \sqcap \exists \text{nationality.US}).$$

2. Returning to our database about companies and employees, it could be useful to equip every employee with (i) a concrete feature `branch` storing the branch-ID in which she is working and (ii) a concrete feature `id` storing her personnel-ID. It would then be natural to enforce that even though personnel-IDs are not unique, the branch-ID together with the personnel-ID uniquely identifies employees. We can do this by using the *composite key definition*

$$(\text{branch, id keyfor Employee}).$$

3. In the car dealers’ ontology, we may assume that cars as well as manufacturers are equipped with identification numbers and that every car is uniquely identified by the combination of its own identification number and its manufacturers one. To express this, we could employ a composite key definition referring to *sequences* of features:

$$(\text{id, manufacturer id keyfor Car}).$$

More formally, we propose to extend DLs to provide for concrete domains with *key boxes*, which are sets of key definitions of the form

$$(u_1, \dots, u_n \text{ keyfor } C),$$

where each u_i is a sequence $f_1 \dots f_n g$ of abstract features f_j followed by a single concrete feature g , and C is a concept. As the above examples illustrate, the idea of key constraints is very natural. Since, moreover, keys play an important role in databases and, as mentioned above, reasoning about database conceptual models is an important, challenging application of description logics, several approaches to extend description logics with keys have already been investigated (Borgida & Weddell, 1997; Calvanese, De Giacomo, & Lenzerini, 2000; Khizder, Toman, & Weddell, 2001). What distinguishes our approach from existing ones, however, is the idea to use concrete domains for constructing key constraints, rather than defining keys on an abstract, logical level.

The goal of this paper is *to provide a comprehensive analysis of the effects on decidability and computational complexity of adding key boxes to description logics with concrete domains*. To this end, we extend the two description logics $\mathcal{ALC}(\mathcal{D})$ and $\mathcal{SHOQ}(\mathcal{D})$ with key boxes, in this way obtaining $\mathcal{ALCK}(\mathcal{D})$ and $\mathcal{SHOQK}(\mathcal{D})$, respectively. While the basic DL with concrete domains $\mathcal{ALC}(\mathcal{D})$ has already been discussed above, $\mathcal{SHOQ}(\mathcal{D})$ was proposed as an ontology language in (Horrocks & Sattler, 2001). It provides a wealth of expressive possibilities such as general concept inclusion axioms (GCIs), transitive roles, role hierarchies, nominals, and qualifying number restrictions. Moreover, it offers a restricted variant of the concrete domain constructor that disallows the use of *sequences* of features in order to avoid undecidability of reasoning.

The main outcome of our investigations is that key constraints can have a dramatic impact on the decidability and complexity of reasoning: for example, whereas satisfiability of $\mathcal{ALC}(\mathcal{D})$ -concepts is known to be PSPACE-complete (Lutz, 2002b), we will show that satisfiability of $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. key boxes is, in general, undecidable. Decidability can be regained if we restrict the concepts used in key boxes to Boolean combinations of concept names (Boolean key boxes). Interestingly, satisfiability of $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. Boolean key boxes is still NEXPTIME-complete even for very simple concrete domains. In the case of $\mathcal{SHOQ}(\mathcal{D})$ and $\mathcal{SHOQK}(\mathcal{D})$, the leap in complexity is somewhat less dramatic since $\mathcal{SHOQ}(\mathcal{D})$ -concept satisfiability is already EXPTIME-complete: again, the addition of key boxes results in NEXPTIME-complete reasoning problems.

It is interesting to note that there exists a close connection between key definitions and so-called *nominals*, i.e. concept names that can have *at most one* instance, such as Pope. Nominals are a standard means of expressivity in description logics and sometimes appear in disguise as the “one-of” operator (Borgida & Patel-Schneider, 1994; Horrocks et al., 2002). It is not hard to see that key boxes can “simulate” nominals: if, for example,

we use a concrete domain based on the natural numbers and providing unary predicates $=_n$ for equality with $n \in \mathbb{N}$, then the key definition (g keyfor \top), where \top stands for logical truth, obviously makes the concept $\exists g.=_n$ behave like a nominal, for each $n \in \mathbb{N}$. For this reason, we also consider $\mathcal{ALCO}(\mathcal{D})$, the extension of $\mathcal{ALC}(\mathcal{D})$ with nominals, and $\mathcal{ALCOK}(\mathcal{D})$, the extension of $\mathcal{ALCK}(\mathcal{D})$ with nominals.² Our main result concerning nominals is that, although in general being of lower expressive power than key boxes, they already make reasoning NEXPTIME-hard if combined with concrete domains: there exist concrete domains \mathcal{D} such that $\mathcal{ALCO}(\mathcal{D})$ -concept satisfiability is NEXPTIME-complete. We should like to stress that all NEXPTIME-hardness results obtained in this paper are in accordance with the observation made in (Lutz, 2004), namely that the PSPACE-upper bound for reasoning with $\mathcal{ALC}(\mathcal{D})$ is not robust w.r.t. extensions of the logic: there exist several “seemingly harmless” extensions of $\mathcal{ALC}(\mathcal{D})$ (for example with acyclic TBoxes and with inverse roles) which make the complexity of reasoning leap from PSPACE-completeness to NEXPTIME-completeness for many natural concrete domains.

The remainder of this paper is organized as follows: in Section 2, we formally introduce concrete domains, key boxes, and the DL $\mathcal{ALCOK}(\mathcal{D})$ together with its fragments $\mathcal{ALCK}(\mathcal{D})$ and $\mathcal{ALCO}(\mathcal{D})$. Moreover, we define Boolean key boxes and some other important properties of key boxes: *path-free* key boxes prohibit the use of sequences of features in key definitions; in *unary* key boxes, each key definition involves exactly one sequence of features; and *composite* key boxes are simply non-unary ones.

Section 3 is devoted to establishing lower bounds for extensions of $\mathcal{ALC}(\mathcal{D})$ with key boxes or nominals. In Section 3.1, we use a reduction of the Post Correspondence Problem to prove that $\mathcal{ALCK}(\mathcal{D})$ -concept satisfiability w.r.t. (non-Boolean) key boxes is undecidable if the concrete domain \mathcal{D} includes the natural numbers, a unary predicate for equality with zero, binary equality and inequality, and a binary incrementation predicate. We then shift our attention towards Boolean key boxes since, in Section 4, we show that this restriction restores decidability. In Section 3.2, we introduce a NEXPTIME-complete variant of the domino problem and three concrete domains that are useful for the reduction of this problem to concept satisfiability in DLs with Boolean key boxes or nominals. In Section 3.3, we use these concrete domains to prove that $\mathcal{ALCK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean and path-free key boxes is NEXPTIME-hard if \mathcal{D} includes two unary predicates denoting disjoint singleton sets. We then strengthen this result to unary key boxes, but, to compensate for the weaker key box formalism, we use more expressive concrete domains. For example, it suffices that the concrete domain \mathcal{D} includes the natural numbers, a unary predicate $=_n$ for each $n \in \mathbb{N}$, and ternary addition. In Section 3.4, we prove that there exists concrete domains \mathcal{D} such that $\mathcal{ALCO}(\mathcal{D})$ -concept satisfiability without reference to key boxes is already NEXPTIME-hard. For this result, we use rather powerful concrete domains, e.g., one with predicates for multiplication and exponentiation. However, we are able to show that there exist concrete domains \mathcal{D} that are computationally very simple (PTIME) if considered in isolation, but lead to NEXPTIME-hardness of $\mathcal{ALCO}(\mathcal{D})$.

The purpose of Section 4 is to develop reasoning procedures for description logics with key boxes and to prove upper complexity bounds matching the NEXPTIME lower bounds established in the previous section. We start in Section 4.1 with a tableau algorithm that

2. Note that the logic $\mathcal{SHOQ}(\mathcal{D})$ already provides for nominals.

decides $\mathcal{ALCCOK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean key boxes, provided that the concrete domain \mathcal{D} is *key-admissible*. Intuitively, a concrete domain \mathcal{D} is key-admissible if there exists an algorithm that takes a finite conjunction c of predicates from \mathcal{D} over some set of variables, decides whether this conjunction is satisfiable and, if c is satisfiable, chooses a solution of c and returns the information on which variables take the same values in it. We call such an algorithm a \mathcal{D} -tester. We have chosen a tableau algorithm since this type of reasoning procedure has the potential to be implemented in efficient reasoners and has been shown to behave well in practice (Horrocks, Sattler, & Tobies, 2000; Haarslev & Möller, 2001). Our algorithm implies the following upper complexity bound: if \mathcal{D} is a key-admissible concrete domain for which a non-deterministic polynomial time \mathcal{D} -tester exists, then $\mathcal{ALCCOK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean key boxes is in NEXPTIME.

In Section 4.2, we devise a tableau algorithm for $\mathcal{SHOQK}(\mathcal{D})$ -concept satisfiability w.r.t. path-free key boxes which might involve *non-Boolean* concepts. The restriction to Boolean concepts in key boxes was necessary for $\mathcal{ALCCOK}(\mathcal{D})$ in order to avoid undecidability. For $\mathcal{SHOQK}(\mathcal{D})$, this restriction is not necessary since $\mathcal{SHOQK}(\mathcal{D})$'s concrete domain constructor is weaker than the one provided by $\mathcal{ALCCOK}(\mathcal{D})$: it does not admit the use of *sequences* of features as arguments. As a by-product of the correctness proof of the algorithm, we obtain a bounded model property for $\mathcal{SHOQK}(\mathcal{D})$, which implies that $\mathcal{SHOQK}(\mathcal{D})$ -concept satisfiability w.r.t. path-free key boxes is in NEXPTIME if \mathcal{D} is a key-admissible concrete domain for which a non-deterministic polynomial time \mathcal{D} -tester exists.

In Section 5, we summarize the results obtained and give an outlook to possible future research.

2. Description Logics with Concrete Domains

In the following, we introduce the description logic $\mathcal{ALCCOK}(\mathcal{D})$. Let us start with defining concrete domains:

Definition 2.1 (Concrete Domain). A *concrete domain* \mathcal{D} is a pair $(\Delta_{\mathcal{D}}, \Phi_{\mathcal{D}})$, where $\Delta_{\mathcal{D}}$ is a set and $\Phi_{\mathcal{D}}$ a set of predicate names. Each predicate name $P \in \Phi_{\mathcal{D}}$ is associated with an arity n and an n -ary predicate $P^{\mathcal{D}} \subseteq \Delta_{\mathcal{D}}^n$.

Based on concrete domains, we can now define $\mathcal{ALCCOK}(\mathcal{D})$ -concepts and key boxes.

Definition 2.2 ($\mathcal{ALCCOK}(\mathcal{D})$ Syntax). Let N_C , N_O , N_R , N_{cF} be pairwise disjoint and countably infinite sets of *concept names*, *nominals*, *role names*, and *concrete features*. Furthermore, we assume that N_R contains a countably infinite subset N_{aF} of *abstract features*. A *path* u is a composition $f_1 \cdots f_n g$ of n abstract features f_1, \dots, f_n ($n \geq 0$) and a concrete feature g . Let \mathcal{D} be a concrete domain. The set of $\mathcal{ALCCOK}(\mathcal{D})$ -concepts is the smallest set such that

- every concept name and every nominal is a concept, and
- if C and D are concepts, R is a role name, g is a concrete feature, u_1, \dots, u_n are paths, and $P \in \Phi_{\mathcal{D}}$ is a predicate of arity n , then the following expressions are also concepts:

$$\neg C, C \sqcap D, C \sqcup D, \exists R.C, \forall R.C, \exists u_1, \dots, u_n.P, \text{ and } g \uparrow.$$

A *key definition* is an expression

$$(u_1, \dots, u_k \text{ keyfor } C),$$

where u_1, \dots, u_k ($k \geq 1$) are paths and C is a concept. A finite set of key definitions is called *key box*.

As usual, we use \top as abbreviation for an arbitrary propositional tautology, \perp as abbreviation for $\neg\top$, $C \rightarrow D$ as abbreviation for $\neg C \sqcup D$, and $C \leftrightarrow D$ as abbreviation for $(C \rightarrow D) \sqcap (D \rightarrow C)$. Throughout this paper, we will also consider several fragments of the description logic $\mathcal{ALCCOK}(\mathcal{D})$. The DL $\mathcal{ALCCO}(\mathcal{D})$ is obtained from $\mathcal{ALCCOK}(\mathcal{D})$ by admitting only empty key boxes. In particular, the set of $\mathcal{ALCCO}(\mathcal{D})$ -concepts is just the set of $\mathcal{ALCCOK}(\mathcal{D})$ -concepts. Furthermore, by disallowing the use of nominals, we obtain the fragment $\mathcal{ALC}(\mathcal{D})$ of $\mathcal{ALCCO}(\mathcal{D})$ and $\mathcal{ALCK}(\mathcal{D})$ of $\mathcal{ALCCOK}(\mathcal{D})$.

The description logic $\mathcal{ALCCOK}(\mathcal{D})$ is equipped with a Tarski-style set-theoretic semantics. Along with the semantics, we introduce the two standard inference problems: concept satisfiability and concept subsumption.

Definition 2.3 ($\mathcal{ALCCOK}(\mathcal{D})$ Semantics). An *interpretation* \mathcal{I} is a pair $(\Delta_{\mathcal{I}}, \cdot^{\mathcal{I}})$, where $\Delta_{\mathcal{I}}$ is a non-empty set, called the *domain*, and $\cdot^{\mathcal{I}}$ is the *interpretation function*. The interpretation function maps

- each concept name C to a subset $C^{\mathcal{I}}$ of $\Delta_{\mathcal{I}}$,
- each nominal N to a singleton subset $N^{\mathcal{I}}$ of $\Delta_{\mathcal{I}}$,
- each role name R to a subset $R^{\mathcal{I}}$ of $\Delta_{\mathcal{I}} \times \Delta_{\mathcal{I}}$,
- each abstract feature f to a partial function $f^{\mathcal{I}}$ from $\Delta_{\mathcal{I}}$ to $\Delta_{\mathcal{I}}$, and
- each concrete feature g to a partial function $g^{\mathcal{I}}$ from $\Delta_{\mathcal{I}}$ to $\Delta_{\mathcal{D}}$.

If $u = f_1 \cdots f_n g$ is a path, then $u^{\mathcal{I}}(d)$ is defined as $g^{\mathcal{I}}(f_n^{\mathcal{I}} \cdots (f_1^{\mathcal{I}}(d)) \cdots)$. The interpretation function is extended to arbitrary concepts as follows:

$$\begin{aligned} (\neg C)^{\mathcal{I}} &:= \Delta_{\mathcal{I}} \setminus C^{\mathcal{I}} \\ (C \sqcap D)^{\mathcal{I}} &:= C^{\mathcal{I}} \cap D^{\mathcal{I}} \\ (C \sqcup D)^{\mathcal{I}} &:= C^{\mathcal{I}} \cup D^{\mathcal{I}} \\ (\exists R.C)^{\mathcal{I}} &:= \{d \in \Delta_{\mathcal{I}} \mid \text{There is } e \in \Delta_{\mathcal{I}} \text{ with } (d, e) \in R^{\mathcal{I}} \text{ and } e \in C^{\mathcal{I}}\} \\ (\forall R.C)^{\mathcal{I}} &:= \{d \in \Delta_{\mathcal{I}} \mid \text{For all } e \in \Delta_{\mathcal{I}}, \text{ if } (d, e) \in R^{\mathcal{I}}, \text{ then } e \in C^{\mathcal{I}}\} \\ (\exists u_1, \dots, u_n.P)^{\mathcal{I}} &:= \{d \in \Delta_{\mathcal{I}} \mid \exists x_1, \dots, x_n \in \Delta_{\mathcal{D}} : u_i^{\mathcal{I}}(d) = x_i \text{ and } (x_1, \dots, x_n) \in P^{\mathcal{D}}\} \\ (g \uparrow)^{\mathcal{I}} &:= \{d \in \Delta_{\mathcal{I}} \mid g^{\mathcal{I}}(d) \text{ undefined}\}. \end{aligned}$$

Let \mathcal{I} be an interpretation. Then \mathcal{I} is a *model* of a concept C iff $C^{\mathcal{I}} \neq \emptyset$. Moreover, \mathcal{I} *satisfies* a key definition $(u_1, \dots, u_n \text{ keyfor } C)$ if, for any $a, b \in C^{\mathcal{I}}$,

$$\text{if, for all } 1 \leq i \leq n, u_i^{\mathcal{I}}(a) = u_i^{\mathcal{I}}(b), \text{ then } a = b.$$

\mathcal{I} is a *model* of a key box \mathcal{K} iff \mathcal{I} satisfies all key definitions in \mathcal{K} . A concept C is *satisfiable w.r.t. a key box \mathcal{K}* iff C and \mathcal{K} have a common model. C is *subsumed by a concept D w.r.t. a key box \mathcal{K}* (written $C \sqsubseteq_{\mathcal{K}} D$) iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all models \mathcal{I} of \mathcal{K} .

It is well-known that, in description logics providing for all Boolean operators, subsumption can be reduced to (un)satisfiability and vice versa: $C \sqsubseteq_{\mathcal{K}} D$ iff $C \sqcap \neg D$ is unsatisfiable w.r.t. \mathcal{K} and C is satisfiable w.r.t. \mathcal{K} iff $C \not\sqsubseteq_{\mathcal{K}} \perp$. This allows us to concentrate on concept satisfiability when devising complexity bounds for reasoning with description logics: lower and upper complexity bounds for concept satisfiability imply corresponding bounds for concept subsumption—only for the complementary complexity class.

If decision procedures for description logics with concrete domains are to be devised without committing to a *particular* concrete domain, then a well-defined interface between the decision procedure and a concrete domain reasoner is needed. Usually, this interface is based on the assumption that the concrete domain is *admissible* (Baader & Hanschke, 1991a; Lutz, 2002a, 2003):

Definition 2.4 (\mathcal{D} -conjunction, Admissibility). Let \mathcal{D} be a concrete domain and \mathbf{V} a set of variables. A \mathcal{D} -conjunction is a (finite) predicate conjunction of the form

$$c = \bigwedge_{i < k} (x_0^{(i)}, \dots, x_{n_i}^{(i)}) : P_i,$$

where P_i is an n_i -ary predicate for $i < k$ and the $x_j^{(i)}$ are variables from \mathbf{V} . A \mathcal{D} -conjunction c is *satisfiable* iff there exists a function δ mapping the variables in c to elements of $\Delta_{\mathcal{D}}$ such that $(\delta(x_0^{(i)}), \dots, \delta(x_{n_i}^{(i)})) \in P_i^{\mathcal{D}}$ for each $i < k$. Such a function is called a *solution* for c . We say that the concrete domain \mathcal{D} is *admissible* iff

1. $\Phi_{\mathcal{D}}$ contains a name $\top_{\mathcal{D}}$ for $\Delta_{\mathcal{D}}$;
2. $\Phi_{\mathcal{D}}$ is closed under negation, i.e., for each n -ary predicate $P \in \Phi_{\mathcal{D}}$, there is a predicate $\overline{P} \in \Phi_{\mathcal{D}}$ of arity n such that $\overline{P}^{\mathcal{D}} = \Delta_{\mathcal{D}}^n \setminus P^{\mathcal{D}}$;
3. satisfiability of \mathcal{D} -conjunctions is decidable.

We refer to the satisfiability of \mathcal{D} -conjunctions as *\mathcal{D} -satisfiability*.

As we shall see, it sometimes makes a considerable difference w.r.t. complexity and decidability to restrict key boxes in various ways, for example to admit only the concept \top on the right-hand side of key definitions or to disallow paths of length greater than one. Therefore, we introduce some useful notions.

Definition 2.5 (Boolean, Path-free, Simple). A key box \mathcal{K} is called

- *Boolean* if all concepts appearing in (key definitions in) \mathcal{K} are Boolean combinations of concept names;
- *path-free* if, for all key definitions $(g_1, \dots, g_n \text{ keyfor } C)$ in \mathcal{K} , $g_1, \dots, g_n \in \mathbf{N}_{\text{cF}}$;
- *simple* if it is both path-free and Boolean;

- *unary* if all key definitions in \mathcal{K} are *unary key definitions*, i.e. of the form (u keyfor C).

A concept C is called *path-free* if, in all its subconcepts of the form $\exists u_1, \dots, u_n.P$, u_1, \dots, u_n are concrete features.

To emphasize that a key box must *not* necessarily be Boolean or path-free, we sometimes call such a key box *general*. Similarly, to emphasize that a key box is not necessarily a unary key box, we sometimes call such a key box *composite key box*.

3. Lower Bounds

In this section, we prove lower complexity bounds for description logics with concrete domains and key boxes and/or nominals. In Section 3.1, we start with showing that the satisfiability of $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. (general) key boxes is undecidable for many interesting concrete domains. This discouraging result is relativized by the fact that, in Section 4.1, we shall prove that the restriction to Boolean key boxes restores decidability. It is thus interesting to look for lower complexity bounds that apply under this restriction. In preparation for this, we introduce in Section 3.2 a NEXPTIME-complete variant of the domino problem and three concrete domains that are well-suited for reductions of this problem.

In Section 3.3, we then prove that satisfiability of path-free $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. simple key boxes is NEXPTIME-hard for a large class of concrete domains \mathcal{D} and that, for many concrete domains, this does hold even if we restrict key boxes to unary ones. Finally, we consider the description logic $\mathcal{ALCO}(\mathcal{D})$ in Section 3.4 and identify several concrete domains such that $\mathcal{ALCO}(\mathcal{D})$ -concept satisfiability (without key boxes!) is NEXPTIME-hard. As we will explain, key boxes and nominals are closely related: key boxes can express nominals, but are in general more powerful.

3.1 Undecidability of $\mathcal{ALCK}(\mathcal{D})$ with General Key Boxes

We prove that satisfiability of $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. key boxes is undecidable for a large class of concrete domains if we allow complex $\mathcal{ALCK}(\mathcal{D})$ -concepts to occur in key definitions. The proof is by a reduction of the well-known undecidable Post Correspondence Problem (Post, 1946; Hopcroft & Ullman, 1979).

Definition 3.1 (PCP). An *instance* P of the *Post Correspondence Problem* (PCP) is given by a finite, non-empty list $(\ell_1, r_1), \dots, (\ell_k, r_k)$ of pairs of words over some alphabet Σ . A sequence of integers i_1, \dots, i_m , with $m \geq 1$, is called a *solution* for P if $\ell_{i_1} \cdots \ell_{i_m} = r_{i_1} \cdots r_{i_m}$. The PCP is to decide whether a given instance P has a solution.

For reducing the PCP to the satisfiability of our DLs, we need an appropriate concrete domain. It is obviously natural to use a concrete domain based on words and concatenation. We will later see that the results obtained for this concrete domain carry over to other concrete domains based on numbers and arithmetics. The following concrete domain was introduced in (Lutz, 2004). Its definition presupposes a fixed alphabet Σ that is at least binary.

Definition 3.2 (Concrete domain \mathcal{W}). The concrete domain \mathcal{W} is defined by setting $\Delta_{\mathcal{W}} := \Sigma^*$ and defining $\Phi_{\mathcal{W}}$ as the smallest set containing the following predicates:

$$\begin{aligned}
 \text{Step} &:= \prod_{1 \leq i \leq k} \exists f_i. (\neg A \sqcap \exists g. =_\epsilon \sqcap \exists \ell, r. \neq) \\
 &\quad \sqcap \prod_{1 \leq i \leq k} (\exists \ell, f_i \ell. \text{conc}_{\ell_i} \sqcap \exists r, f_i r. \text{conc}_{r_i}) \\
 C_P &:= \exists \ell. =_\epsilon \sqcap \exists r. =_\epsilon \\
 &\quad \sqcap \exists R. (A \sqcap \exists g. =_\epsilon \sqcap \neg \text{Step}) \\
 &\quad \sqcap \text{Step} \\
 \mathcal{K}_P &:= \{g \text{ keyfor } \neg \text{Step}\}
 \end{aligned}$$

Figure 1: The $\mathcal{ALCK}(\mathcal{W})$ reduction concept C_P and key box \mathcal{K}_P .

- unary predicates `word` and `nword` with $\text{word}^{\mathcal{W}} = \Delta_{\mathcal{W}}$ and $\text{nword}^{\mathcal{W}} = \emptyset$,
- unary predicates $=_\epsilon$ and \neq_ϵ with $=_\epsilon^{\mathcal{W}} = \{\epsilon\}$ and $\neq_\epsilon^{\mathcal{W}} = \Sigma^+$,
- a binary equality predicate $=$ and a binary inequality predicate \neq with the obvious interpretation, and
- for each $w \in \Sigma^+$, two binary predicates conc_w and nconc_w with

$$\text{conc}_w^{\mathcal{W}} = \{(u, v) \mid v = uw\} \text{ and } \text{nconc}_w^{\mathcal{W}} = \{(u, v) \mid v \neq uw\}.$$


It is readily checked that \mathcal{W} satisfies properties 1 and 2 of admissibility (see Definition 2.4). Moreover, \mathcal{W} -satisfiability is decidable:

Theorem 3.1 (Lutz, 2004). *\mathcal{W} -satisfiability is in PTIME.*

Thus, \mathcal{W} is admissible and even of low complexity. This is important since our aim is to demonstrate that the undecidability of $\mathcal{ALCK}(\mathcal{W})$ -concept satisfiability is due to the presence of keys, and not due to the high complexity of \mathcal{W} -satisfiability.

We can now discuss the reduction of the PCP. A given instance $(\ell_1, r_1), \dots, (\ell_k, r_k)$ is translated into an $\mathcal{ALCK}(\mathcal{D})$ -concept C_P and key box \mathcal{K}_P as defined in Figure 1 such that P has a solution iff C_P is unsatisfiable w.r.t. \mathcal{K}_P . The idea behind the reduction is that a common model of C_P and \mathcal{K}_P encodes all *potential solutions* for P (i.e., sequences i_1, \dots, i_n of integers i_j between 1 and k) and to make sure that none of them is indeed a solution. In Figure 1, f_1, \dots, f_k denote abstract features while g, ℓ , and r denote concrete features. The definition of the concept `Step` just serves as an abbreviation and should not be confused with so-called TBoxes (see Section 4.2 for the definition of TBoxes). Models of C_P and \mathcal{K}_P , such as the one displayed in Figure 2, have the form of an infinite k -ary tree whose root is connected to an “extra node” x via the role R . Intuitively, each node of the tree represents one partial solution i_1, \dots, i_n , its ℓ -successor represents the corresponding left concatenation $\ell_{i_1} \cdots \ell_{i_n}$, and its r -successor the corresponding right concatenation $r_{i_1} \cdots r_{i_n}$.

To enforce the existence of the infinite tree, we employ the key box \mathcal{K}_P : consider for example the root node’s f_1 -successor in Figure 2—let us call this node y . Due to Line 3


 Figure 2: An example model of C_P and \mathcal{K}_P .

of C_P and Line 1 of **Step**, we have $y \in (\exists g. =_{\epsilon})^{\mathcal{I}}$. Due to Line 2 of C_P , we also have $x \in (\exists g. =_{\epsilon})^{\mathcal{I}}$ and $x \in (\neg \text{Step})^{\mathcal{I}}$, where x is the “extra node” mentioned above. In view of the key box \mathcal{K}_P , this implies that either (i) $x = y$ or (ii) $y \in \text{Step}^{\mathcal{I}}$. It is easy to see that (i) is impossible since Line 2 of C_P and Line 1 of **Step** imply that $x \in A^{\mathcal{I}}$ and $y \in (\neg A)^{\mathcal{I}}$. Hence $y \in \text{Step}^{\mathcal{I}}$ and, by Line 2 of **Step**, y has the appropriate f_i -successors for $1 \leq i \leq n$. In the same way, the construction of the tree can be continued ad infinitum. The second line in the definition of **Step** enforces that $\ell^{\mathcal{I}}(z) = \ell_{i_1} \cdots \ell_{i_n}$ and $r^{\mathcal{I}}(z) = r_{i_1} \cdots r_{i_n}$ for z an $f_{i_1} \cdots f_{i_n}$ -successor of the root node. Finally, the concept $\exists \ell, r. \neq$ in Line 1 of **Step** implies that $\ell^{\mathcal{I}}(z) \neq r^{\mathcal{I}}(z)$ holds at all nodes z of the tree (except for the root), which implies that no potential solution is a solution.

Since the size of C_P and \mathcal{K}_P is clearly polynomial in k and the key box \mathcal{K}_P is a unary key box, we obtain the following proposition.

Proposition 3.2. *The satisfiability of $\mathcal{ALCK}(\mathcal{W})$ -concepts w.r.t. (non-Boolean) path-free unary key boxes is undecidable.*

To emphasize that this undecidability result was obtained using a very simple concrete domain, let us combine Theorem 3.1 with Proposition 3.2.

Theorem 3.3. *There exists a concrete domain \mathcal{D} such that \mathcal{D} -satisfiability is in PTIME and satisfiability of $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. (non-Boolean) path-free unary key boxes is undecidable.*

On first sight, the concrete domain \mathcal{W} might look artificial and one may question the relevance of lower bounds that have been obtained using \mathcal{W} . However, it is straightforward to encode words as natural numbers and to define concatenation of words as rather simple operations on the natural numbers (Baader & Hanschke, 1992): a word $w \neq \epsilon$ over the alphabet Σ of cardinality $\#\Sigma$ can be interpreted as a number written at base $\#\Sigma + 1$ in which the symbol that is the “0 digit” does not occur. Hence, we can use the corresponding natural number (e.g. at base 10) to represent a word w , and the number 0 to represent

the empty word. The concatenation of two words v and w can then be expressed as $vw = v \cdot (\#\Sigma + 1)^{|w|} + w$, where $|w|$ denotes the length of the word w . Moreover, exponentiation can be expressed as multiple multiplications, multiplication as multiple additions, and addition as multiple incrementation: this is shown in Section 5.6 of (Lutz, 2004) for the case of $\mathcal{ALC}(\mathcal{D})$ extended with TBoxes (c.f. Section 4.2) and can easily be adapted to $\mathcal{ALC}(\mathcal{D})$ with non-Boolean key boxes. This observation gives rise to the following theorem:

Theorem 3.4. *Let \mathcal{D} be a concrete domain such that $\cdot \subseteq \Delta_{\mathcal{D}}$, $\Phi_{\mathcal{D}}$ contains a unary predicate $=_0$ with $(=_0)^{\mathcal{D}} = \{0\}$, binary equality and inequality predicates, and a binary predicate incr with $\text{incr}^{\mathcal{D}} \cap \{(n, x) \mid n \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} = \{(k, k + 1) \mid k \in \cdot\}$. Then satisfiability of $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. (non-Boolean) path-free unary key boxes is undecidable.*

3.2 Domino Problems and Concrete Domains

In this section, we introduce a NEXPTIME-complete variant of the well-known, undecidable domino problem (Berger, 1966; Knuth, 1968), and then define three concrete domains \mathcal{D}_1 , \mathcal{D}_2 , and \mathcal{D}_3 . These concrete domains will be used in Sections 3.3 and 3.4 to establish lower bounds for reasoning with $\mathcal{ALCK}(\mathcal{D})$ and Boolean key boxes, and for reasoning with $\mathcal{ALCO}(\mathcal{D})$.

In general, a domino problem is given by a finite set of *tile types*. Intuitively, all tile types are of the same size, each type having a square shape and colored edges. An unlimited number of tiles of each type is available. In the NEXPTIME-hard variant of the domino problem that we use, the task is to tile a $2^{n+1} \times 2^{n+1}$ -torus (i.e., a $2^{n+1} \times 2^{n+1}$ -rectangle whose borders are “glued” together) where neighboring edges have the same color.

Definition 3.3 (Domino System). A *domino system* \mathfrak{D} is a triple (T, H, V) , where $T \subsetneq \cdot$ is a finite set of *tile types* and $H, V \subseteq T \times T$ represent the horizontal and vertical matching conditions. Let \mathfrak{D} be a domino system and $a = a_0, \dots, a_{n-1}$ an *initial condition*, i.e. an n -tuple of tiles. A mapping $\tau : \{0, \dots, 2^{n+1}\} \times \{0, \dots, 2^{n+1}\} \rightarrow T$ is a *solution* for \mathfrak{D} and a iff, for all $x, y < 2^{n+1}$, the following holds:

- if $\tau(x, y) = t$ and $\tau(x \oplus_{2^{n+1}} 1, y) = t'$, then $(t, t') \in H$
- if $\tau(x, y) = t$ and $\tau(x, y \oplus_{2^{n+1}} 1) = t'$, then $(t, t') \in V$
- $\tau(i, 0) = a_i$ for $i < n$.

where \oplus_i denotes addition modulo i .

It follows from results in (Börger, Grädel, & Gurevich, 1997) that the above variant of the domino problem is NEXPTIME-complete.

We define the concrete domain \mathcal{D}_1 to be used in the reduction of the NEXPTIME-complete domino problem to $\mathcal{ALCK}(\mathcal{D}_1)$ -concept satisfiability w.r.t. Boolean key boxes.

Definition 3.4 (Concrete Domain \mathcal{D}_1). The concrete domain \mathcal{D}_1 is defined by setting $\Delta_{\mathcal{D}_1} := \{0, 1\}$ and $\Phi_{\mathcal{D}_1}$ to the (smallest) set containing the following predicates:

- unary predicates $\top_{\mathcal{D}_1}$ with $(\top_{\mathcal{D}_1})^{\mathcal{D}_1} = \Delta_{\mathcal{D}_1}$ and $\perp_{\mathcal{D}_1}$ with $(\perp_{\mathcal{D}_1})^{\mathcal{D}_1} = \emptyset$;
- unary predicates $=_0$ and $=_1$ with $(=_i)^{\mathcal{D}_1} = \{i\}$, $i \in \{0, 1\}$.

The second concrete domain D_2 will be used for a reduction of the NEXPTIME-complete domino problem to $\mathcal{ALCK}(D_2)$ -concept satisfiability w.r.t. Boolean *unary* key boxes. For this reduction we need to “store” vectors of bits in single concrete domain elements.

Definition 3.5 (Concrete Domain D_2). For every $n \in \mathbb{N}$, a function $v : \{0, \dots, n-1\} \rightarrow \{0, 1\}$ is called a *bit vector* of dimension n . We use BV_n to denote the set of all bit vectors of dimension n . The concrete domain D_2 is defined by setting $\Delta_{D_2} := \bigcup_{i>0} BV_i$ and Φ_{D_2} to the (smallest) set containing the following predicates:

- unary predicates \top_{D_2} with $(\top_{D_2})^{D_2} = \Delta_{D_2}$ and \perp_{D_2} with $(\perp_{D_2})^{D_2} = \emptyset$;
- for every $k, i \in \mathbb{N}$ with $i < k$, unary predicates $\text{bit}0_k^i$ and $\text{bit}1_k^i$ with

$$(\text{bit}n_k^i)^{D_2} = \{v \in \Delta_{D_2} \mid v \in BV_k \text{ and } v(i) = n\}.$$

Moreover, unary predicates $\overline{\text{bit}0_k^i}$ and $\overline{\text{bit}1_k^i}$ with $(\overline{\text{bit}n_k^i})^{D_2} = \Delta_{D_2} \setminus (\text{bit}n_k^i)^{D_2}$.

The last concrete domain D_3 is used in the reduction of the NEXPTIME-complete domino problem to $\mathcal{ALCO}(D_3)$ -concept satisfiability. In this reduction, the concrete domain D_3 contains two kinds of elements: firstly, there are elements of Δ_{D_3} that can represent the *whole* $2^{n+1} \times 2^{n+1}$ -torus, so-called *domino arrays*. Secondly, there are elements of Δ_{D_3} that represent positions in the torus. For technical reasons to be discussed later, these elements are vectors of natural numbers rather than bit vectors, and in the following we shall just call them *vectors*. A domino array is then a function mapping each pair of vectors (of certain length) to a natural number which represents a tile type.

Definition 3.6 (Concrete Domain D_3). For every $k \in \mathbb{N}$, a function $v : \{0, \dots, k-1\} \rightarrow \mathbb{N}$ is called a *vector* of dimension k . We use VE_k to denote the set of all vectors of dimension k . For every $k \in \mathbb{N}$, a function $k : VE_k \times VE_k \rightarrow \mathbb{N}$ is called a *domino array* of dimension k . We use DA_k to denote the set of all domino arrays of dimension k . The concrete domain D_3 is defined by setting $\Delta_{D_3} := \bigcup_{i>0} VE_i \cup \bigcup_{i>0} DA_i$ and Φ_{D_3} to the (smallest) set containing the following predicates:

- unary predicates \top_{D_3} with $(\top_{D_3})^{D_3} = \Delta_{D_3}$ and \perp_{D_3} with $(\perp_{D_3})^{D_3} = \emptyset$;
- for every $k, i \in \mathbb{N}$ with $i < k$, unary predicates $\text{pos}0_k^i$ and $\text{pos}1_k^i$ with

$$(\text{pos}n_k^i)^{D_3} = \{v \in \Delta_{D_3} \mid v \in VE_k \text{ and } v(i) = n\}$$

and unary predicates $\overline{\text{pos}0_k^i}$ and $\overline{\text{pos}1_k^i}$ with $(\overline{\text{pos}n_k^i})^{D_3} = \Delta_{D_3} \setminus (\text{pos}n_k^i)^{D_3}$;

- for every $k, i \in \mathbb{N}$, a predicate tile_k^i of arity 3 with

$$(\text{tile}_k^i)^{D_3} = \{(v_x, v_y, d) \mid v_x, v_y \in VE_k, d \in DA_k, \text{ and } d(v_x, v_y) = i\}$$

and a predicate $\overline{\text{tile}_k^i}$ of arity 3 with $(\overline{\text{tile}_k^i})^{D_3} = (\Delta_{D_3})^3 \setminus (\text{tile}_k^i)^{D_3}$.

The reason for using vectors of natural numbers rather than bit vectors in the definition of D_3 is that we want D_3 -satisfiability to be of low complexity, preferably in PTIME: consider the D_3 -conjunction

$$\text{pos}0_2^0(x) \wedge \text{pos}0_2^0(y) \wedge \text{pos}0_2^0(z) \wedge \\ \text{tile}_2^7(x, v, d) \wedge \text{tile}_2^8(y, v, d) \wedge \text{tile}_2^9(z, v, d).$$

If we use bit vectors rather than vectors of natural numbers, then the upper line enforces that at least two out of the three variables x , y , and z must take the same value. Since the value of v is fixed, the lower line makes the conjunction unsatisfiable: it tries to assign the three different values 7,8,9 to two different positions in the domino array. It seems unlikely that this kind of inconsistency can be detected in polynomial time. This problem is circumvented by using vectors of natural numbers in the definition of D_3 (but enforcing them to be bit vectors in the reduction): in this case, the above conjunction is clearly satisfiable.

Proposition 3.5. *For each $i \in \{1, 2, 3\}$, the concrete domain D_i is admissible and satisfiability of D_i -conjunctions is in PTIME.*

For D_1 , this is trivial. For D_2 , a proof can be found in Appendix A. And for D_3 , a proof can be found in (Lutz, Areces, Horrocks, & Sattler, 2002).

3.3 NEXPTIME-hardness of $\mathcal{ALCK}(\mathcal{D})$ with Boolean Key Boxes

In this section, we prove two NEXPTIME-lower bounds for $\mathcal{ALCK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean key boxes by reducing the NEXPTIME-complete domino problem introduced in the previous section. The first reduction uses the very simple concrete domain D_1 , but depends on composite key definitions. The second reduction uses the slightly more complex concrete domain D_2 , but only needs unary key definitions. As we will see, the two reductions yield different, incomparable results.

We first reduce the NEXPTIME-complete domino problem to $\mathcal{ALCK}(D_1)$ -concept satisfiability w.r.t. Boolean composite key boxes. Each domino system $\mathfrak{D} = (T, H, V)$ with initial condition $a = a_0, \dots, a_{n-1}$ is translated into an $\mathcal{ALCK}(D_1)$ -concept $C_{\mathfrak{D},a}$ as displayed in Figure 3. Names such as *TreeX* and *TreeY* are used as abbreviations, only. We use $\forall R^i.C$ as an abbreviation for the n -fold nesting $\forall R. \dots \forall R.C$. The names $x\text{pos}_i$ and $y\text{pos}_i$ used in the figure denote concrete features. In the definition of the *Init* concept, for each $n \in \dots$, $\text{bit}_i(n)$ is supposed to denote the i 'th bit of the binary representation of n . We claim that $C_{\mathfrak{D},a}$ is satisfiable w.r.t. the key box

$$\{(x\text{pos}_0, \dots, x\text{pos}_n, y\text{pos}_0, \dots, y\text{pos}_n \text{ keyfor } \top)\}$$

iff there exists a solution for \mathfrak{D} and a . To substantiate this claim, let us walk through the reduction and explain the various parts of the concept $C_{\mathfrak{D},a}$. The first step towards understanding the structure of models of $C_{\mathfrak{D},a}$ (which is the key to understanding the reduction itself) is to note that the purpose of the first line of $C_{\mathfrak{D},a}$ is to enforce a tree structure of depth $2(n+1)$, whose leaves correspond to positions in the $2^{n+1} \times 2^{n+1}$ -torus. More precisely, the *TreeX* concept guarantees that, in every model \mathcal{I} of $C_{\mathfrak{D},a}$, there exists a binary tree of depth $n+1$. Moreover, the *DistX_k* concepts (there exists one for each

$$\begin{aligned}
 \text{TreeX} &:= \exists R.X_0 \sqcap \exists R.\neg X_0 \sqcap \prod_{i=1..n} \forall R^i.(\text{DistX}_{i-1} \sqcap \exists R.X_i \sqcap \exists R.\neg X_i) \\
 \text{TreeY} &:= \text{DistX}_n \sqcap \exists R.Y_0 \sqcap \exists R.\neg Y_0 \sqcap \\
 &\quad \prod_{i=1..n} \forall R^i.(\text{DistY}_{i-1} \sqcap \text{DistX}_n \sqcap \exists R.Y_i \sqcap \exists R.\neg Y_i) \\
 \text{DistX}_k &:= \prod_{i=0..k} ((X_i \rightarrow \forall R.X_i) \sqcap (\neg X_i \rightarrow \forall R.\neg X_i)) \\
 \text{DistY}_k &:= \prod_{i=0..k} ((Y_i \rightarrow \forall R.Y_i) \sqcap (\neg Y_i \rightarrow \forall R.\neg Y_i)) \\
 \text{TransXPos} &:= \prod_{i=0..n} (X_i \rightarrow \exists \text{xpos}_i. =_1) \sqcap (\neg X_i \rightarrow \exists \text{xpos}_i. =_0) \\
 \text{TransYPos} &:= \prod_{i=0..n} (Y_i \rightarrow \exists \text{ypos}_i. =_1) \sqcap (\neg Y_i \rightarrow \exists \text{ypos}_i. =_0) \\
 \text{Succs} &:= \exists R_x.(\text{TransXPos} \sqcap \text{TransYPos}) \sqcap \exists R_y.(\text{TransXPos} \sqcap \text{TransYPos}) \\
 \text{XSuccOk} &:= \prod_{i=0..n} ((Y_i \rightarrow \forall R_x.Y_i) \sqcap (\neg Y_i \rightarrow \forall R_x.\neg Y_i)) \\
 &\quad \prod_{k=0..n} \left(\prod_{j=0..k} X_j \right) \rightarrow ((X_k \rightarrow \forall R_x.\neg X_k) \sqcap (\neg X_k \rightarrow \forall R_x.X_k)) \\
 &\quad \prod_{k=0..n} \left(\bigsqcup_{j=0..k} \neg X_j \right) \rightarrow ((X_k \rightarrow \forall R_x.X_k) \sqcap (\neg X_k \rightarrow \forall R_x.\neg X_k)) \\
 \text{YSuccOk} &:= \prod_{i=0..n} ((X_i \rightarrow \forall R_y.X_i) \sqcap (\neg X_i \rightarrow \forall R_y.\neg X_i)) \\
 &\quad \prod_{k=0..n} \left(\prod_{j=0..k} Y_j \right) \rightarrow ((Y_k \rightarrow \forall R_y.\neg Y_k) \sqcap (\neg Y_k \rightarrow \forall R_y.Y_k)) \\
 &\quad \prod_{k=0..n} \left(\bigsqcup_{j=0..k} \neg Y_j \right) \rightarrow ((Y_k \rightarrow \forall R_y.Y_k) \sqcap (\neg Y_k \rightarrow \forall R_y.\neg Y_k)) \\
 \text{Label} &:= \bigsqcup_{i \in T} D_i \sqcap \prod_{i,j \in T, i \neq j} \neg(D_i \sqcap D_j) \\
 \text{CheckMatch} &:= \bigsqcup_{(i,j) \in H} (D_i \sqcap \forall R_x.D_j) \sqcap \bigsqcup_{(i,j) \in V} (D_i \sqcap \forall R_y.D_j) \\
 \text{Init} &:= \prod_{i=0..n-1} \left(\left(\prod_{j=0..n, \text{bit}_j(i)=0} \neg X_j \sqcap \prod_{j=0..n, \text{bit}_j(i)=1} X_j \sqcap \prod_{j=0..n} \neg Y_j \right) \rightarrow D_{a_i} \right) \\
 C_{\mathcal{D},a} &:= \text{TreeX} \sqcap \forall R^{n+1}.\text{TreeY} \\
 &\quad \sqcap \forall R^{2(n+1)}.(\text{TransXPos} \sqcap \text{TransYPos} \sqcap \text{Succs} \sqcap \text{XSuccOk} \sqcap \text{YSuccOk}) \\
 &\quad \sqcap \forall R^{2(n+1)}.(\text{Label} \sqcap \text{CheckMatch} \sqcap \text{Init})
 \end{aligned}$$

 Figure 3: The $\mathcal{ALCK}(\mathcal{D}_1)$ reduction concept $C_{\mathcal{D},a}$.

$k \in \{0, \dots, n\}$ ensure that the leaves of this tree are binarily numbered (from 0 to $2^{n+1} - 1$) by the concept names X_0, \dots, X_n . More precisely, for a domain object $d \in \Delta_{\mathcal{I}}$, set

$$\text{xpsn}(d) = \sum_{i=0}^n \alpha_i(d) * 2^i \quad \text{where} \quad \alpha_i(d) = \begin{cases} 1 & \text{if } d \in X_i^{\mathcal{I}} \\ 0 & \text{otherwise.} \end{cases}$$

The **TreeX** and **DistX** concepts ensure that there exist nodes $d_0, \dots, d_{2^{n+1}-1}$ at level $n+1$ of the tree such that $\text{xpsn}(d_i) = i$. Intuitively, this numbering represents the horizontal positions in the $2^{n+1} \times 2^{n+1}$ -torus. The vertical positions are coded in a similar way by the Y_0, \dots, Y_n concept names. More specifically, the concepts **TreeY**, **DistX**, and **DistY** ensure that every d_i ($i \leq 2^{n+1} - 1$) is the root of another tree, in which (i) every node has the same “ X_0, \dots, X_n -configuration” as its root node, and (ii) the leaves are numbered binarily using the concept names Y_0, \dots, Y_n (note that the **TreeY** concept appears in $C_{\mathcal{D},a}$ inside a $\forall R^{n+1}$ value restriction). Define

$$\text{ypsn}(d) = \sum_{i=0}^n \beta_i(d) * 2^i \quad \text{where} \quad \beta_i(d) = \begin{cases} 1 & \text{if } d \in Y_i^{\mathcal{I}} \\ 0 & \text{otherwise.} \end{cases}$$

In the set of leaf nodes of all the trees enforced by the **TreeY** concept, there exists, for each $i, j < 2^{n+1}$, an³ object $e_{i,j} \in \Delta_{\mathcal{I}}$ such that $\text{xpsn}(e_{i,j}) = i$ and $\text{ypsn}(e_{i,j}) = j$, i.e., each $e_{i,j}$ represents the position (i, j) in the $2^{n+1} \times 2^{n+1}$ -torus.

The next step is to translate the individual bits of the numbering of the $e_{i,j}$ -objects, which are up to now represented by concept names, into concrete domain values. This is done by the **TransXPos** and **TransYPos** concepts which ensure that, for all $\ell \leq n$, we have $\text{xpos}_{\ell}^{\mathcal{I}}(e_{i,j}) = 0$ if $e_{i,j} \in \neg X_{\ell}$, $\text{xpos}_{\ell}^{\mathcal{I}}(e_{i,j}) = 1$ if $e_{i,j} \in X_{\ell}$, and similar for ypos_{ℓ} and Y_{ℓ} . Since \mathcal{I} is a model for the key box

$$\{(\text{xpos}_0, \dots, \text{xpos}_n, \text{ypos}_0, \dots, \text{ypos}_n \text{ keyfor } \top)\},$$

grid positions are *uniquely* represented by domain elements from $(\text{TransXPos} \sqcap \text{TransYPos})^{\mathcal{I}}$, i.e., if $d, e \in (\text{TransXPos} \sqcap \text{TransYPos})^{\mathcal{I}}$ such that $\text{xpsn}(d) = \text{xpsn}(e)$ and $\text{ypsn}(d) = \text{ypsn}(e)$, then $d = e$. This fact is used in the concepts **Succs**, **XSuccOk**, and **YSuccOk** to enforce that, for the two roles R_x and R_y and each $i, j \leq n$, the following holds:

$$\begin{aligned} R_x^{\mathcal{I}} \cap (\{e_{i,j}\} \times \Delta_{\mathcal{I}}) &= \{(e_{i,j}, e_{(i \oplus_{2^{n+1}} 1), j})\} \\ R_y^{\mathcal{I}} \cap (\{e_{i,j}\} \times \Delta_{\mathcal{I}}) &= \{(e_{i,j}, e_{i, (j \oplus_{2^{n+1}} 1)})\}. \end{aligned} \quad (*)$$

The **Succs** concept ensures that, for each $e_{i,j}$, there exists an R_x -successor and an R_y -successor, and that both are in $(\text{TransXPos} \sqcap \text{TransYPos})^{\mathcal{I}}$. Let d be an R_x -successor of $e_{i,j}$. Then the **XSuccOk** concept ensures that $\text{xpsn}(d) = i \oplus_{2^{n+1}} 1$ and $\text{ypsn}(d) = j$. Before we explain how it does this, let us note that, since all $e_{i,j}$ are in $(\text{TransXPos} \sqcap \text{TransYPos})^{\mathcal{I}}$ and the grid positions are uniquely represented by elements of $(\text{TransXPos} \sqcap \text{TransYPos})^{\mathcal{I}}$, this implies $d = e_{(i \oplus_{2^{n+1}} 1), j}$ which shows that the upper line of $(*)$ does indeed hold.

Let us now consider the **XSuccOk** concept in some more detail. It is essentially the DL-formulation of the well-known propositional formula

$$\bigwedge_{k=0}^n \left(\bigwedge_{j=0}^{k-1} x_j = 1 \right) \rightarrow (x_k = 1 \leftrightarrow x'_k = 0) \quad \wedge \quad \bigwedge_{k=0}^n \left(\bigvee_{j=0}^{k-1} x_j = 0 \right) \rightarrow (x_k = x'_k)$$

which encodes incrementation modulo 2^{n+1} , i.e., if t is the number (binarily) encoded by the propositional variables x_0, \dots, x_n and t' is the number encoded by the propositional

3. It does not matter if there is more than one such object.

variables x'_0, \dots, x'_n , then we have $t' = t + 1$ modulo 2^{n+1} , see (Börger et al., 1997). Taking into account the $\forall R_x$ quantifiers in XSuccOk , it is readily checked that this concept has just the desired effect: to ensure that, for every R_x -successor d of $e_{i,j}$, we have $\text{xpsn}(d) = \text{xpsn}(e_{(i \oplus_{2^{n+1}} 1), j}) = i \oplus_{2^{n+1}} 1$. The explanation of YSuccOk and how it enforces the lower line of $(*)$ is analogous to the XSuccOk case.

It remains to ensure that every grid position is labeled with precisely one tile and that the initial condition as well as the horizontal and vertical matching conditions are satisfied. The tiles are represented by concept names D_i (where i is from the set of tiles T) and the described tasks are accomplished in the standard way by the concepts Label , Init , and CheckMatch .

It is worth noting that the reduction concept is path-free and the key box is simple, i.e., path-free and Boolean. Path-freeness of concepts is often used to tame the complexity of description logics with concrete domains (although it largely sacrifices their expressive power) (Lutz, 2003; Baader, Lutz, Sturm, & Wolter, 2002b; Haarslev, Möller, & Wessel, 2001; Horrocks & Sattler, 2001). For example, if $\mathcal{ALC}(\mathcal{D})$ is augmented with general TBoxes, then reasoning with arbitrary concepts is undecidable while reasoning with path-free concepts is EXPTIME -complete if \mathcal{D} is admissible and \mathcal{D} -satisfiability is in EXPTIME (Lutz, 2002a). This “taming approach” does not work in the presence of key boxes since, as we have just seen, satisfiability of $\mathcal{ALC}(\mathcal{D})$ -concepts w.r.t. key boxes is (under some natural assumptions) NEXPTIME -hard, even if both concept and key box are path-free.

Since the size of $C_{\mathcal{D},a}$ and of the used key box is clearly polynomial in n , we obtain the following proposition.

Proposition 3.6. *The satisfiability of path-free $\mathcal{ALCK}(\mathcal{D}_1)$ -concepts w.r.t. simple key boxes is NEXPTIME -hard.*

In (Lutz, 2002b), it is proved that (non path-free) $\mathcal{ALC}(\mathcal{D})$ -concept satisfiability is PSPACE -complete if \mathcal{D} -satisfiability is in PSPACE . In particular, it thus follows from Proposition 3.5 that $\mathcal{ALC}(\mathcal{D}_1)$ -concept satisfiability is PSPACE -complete. Thus, there is a rather dramatic increase of complexity if key boxes are added to $\mathcal{ALC}(\mathcal{D}_1)$. To stress that this increase is due to the key boxes themselves and not to the complexity of \mathcal{D}_1 -satisfiability, we reformulate Proposition 3.6:

Theorem 3.7. *There exists a concrete domain \mathcal{D} such that \mathcal{D} -satisfiability is in PTIME and satisfiability of path-free $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. simple key boxes is NEXPTIME -hard.*

Although, due to its very low expressivity, the concrete domain \mathcal{D}_1 itself is not very natural for knowledge representation, it is a fragment of many concrete domains that have been proposed in the literature (Baader & Hanschke, 1992; Haarslev & Möller, 2001; Lutz, 2003, 2002b). Indeed, the presented reduction strategy can be adapted to quite many “standard” concrete domains. Let us formulate a (very weak) condition that a concrete domain must satisfy in order for the presented reduction strategy to be applicable.

Theorem 3.8. *Let \mathcal{D} be a concrete domain. If there exist $a, b \in \Delta_{\mathcal{D}}$ with $a \neq b$ and $P_1, P_2 \in \Phi_{\mathcal{D}}$ such that $P_1^{\mathcal{D}} = \{a\}$ and $P_2^{\mathcal{D}} = \{b\}$, then the satisfiability of path-free $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. simple key boxes is NEXPTIME -hard.*

$$\begin{aligned}
\text{Succs2} &:= \exists R_x. \text{TransPos} \sqcap \exists R_y. \text{TransPos} \\
\text{TransPos} &:= \prod_{i=0..n} ((X_i \rightarrow \exists \text{bv}. \text{bit}1_{2(n+1)}^i) \sqcap \neg X_i \rightarrow \exists \text{bv}. \text{bit}0_{2(n+1)}^i) \sqcap \\
&\quad \prod_{i=0..n} ((Y_i \rightarrow \exists \text{bv}. \text{bit}1_{2(n+1)}^{n+i+1}) \sqcap \neg Y_i \rightarrow \exists \text{bv}. \text{bit}0_{2(n+1)}^{n+i+1}) \\
C_{\mathcal{D},a} &:= \text{TreeX} \sqcap \forall R^{n+1}. \text{TreeY} \\
&\quad \sqcap \forall R^{2(n+1)}. (\text{TransPos} \sqcap \text{Succs2} \sqcap \text{XSuccOk} \sqcap \text{YSuccOk}) \\
&\quad \sqcap \forall R^{2(n+1)}. (\text{Label} \sqcap \text{CheckMatch} \sqcap \text{Init})
\end{aligned}$$

Figure 4: The $\mathcal{ALCK}(\mathcal{D}_2)$ reduction concept $C_{\mathcal{D},a}$.

We now present the second NEXPTIME -hardness result for $\mathcal{ALCK}(\mathcal{D})$ -concept satisfiability. This time, we reduce the NEXPTIME -complete domino problem to the satisfiability of path-free $\mathcal{ALCK}(\mathcal{D}_2)$ -concepts w.r.t. simple *unary* key boxes. The reduction is very similar to the previous one and we only discuss the differences.

In the first reduction, we represented the individual bits of grid positions by individual concrete features xpos_i and ypos_i and used a composite key box to ensure that each point in the torus is represented by at most one element. In the second reduction, we use a single concrete feature bv and represent an entire position (i, j) in the torus using a bit vector from the concrete domain \mathcal{D}_2 . This allows us to enforce the above mentioned uniqueness of representations using a *unary* key box.

The modified reduction concept $C_{\mathcal{D},a}$ can be found in Figure 4, where the concepts TreeX , TreeY , DistX_k , DistY_k , XSuccOk , YSuccOk , Label , CheckMatch , and Init are defined as in Figure 3. The translation of the position in the torus encoded by $X_0, \dots, X_n, Y_0, \dots, Y_n$ into a bit vector is done by the TransPos concept in a straightforward manner. Given what was said about the first reduction, it is not hard to see that $C_{\mathcal{D},a}$ is satisfiable w.r.t. the key box $\{(\text{bv keyfor } \top)\}$ iff there exists a solution for \mathcal{D} and a . We thus obtain the following proposition.

Proposition 3.9. *The satisfiability of path-free $\mathcal{ALCK}(\mathcal{D}_2)$ -concepts w.r.t. simple unary key boxes is NEXPTIME -hard.*

Again, we relate the NEXPTIME lower bound to the complexity of \mathcal{D}_2 -complexity, which is determined in Proposition 3.5.

Theorem 3.10. *There exists a concrete domain \mathcal{D} such that \mathcal{D} -satisfiability is in PTIME and the satisfiability of path-free $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. simple unary key boxes is NEXPTIME -hard.*

Since the elements of $\Delta_{\mathcal{D}_2}$ are bit vectors, the concrete domain \mathcal{D}_2 cannot be considered a natural choice for many application areas. But, in the reduction, \mathcal{D}_2 can be replaced by several natural concrete domains.

The central observation is that we use bit vectors only to injectively translate sequences of bits into values of the concrete domain, i.e., we translate sequences of $2(n+1)$ bits (represented by the concept names X_0, \dots, X_n and Y_0, \dots, Y_n) into elements of $\Delta_{\mathcal{D}_2}$ such

that, for *distinct* sequences, the results of the translation are also distinct. Due to this restricted use of bit vectors, there are several ways to replace them by natural numbers. For example, we can replace TransPos with the following concept $\text{TransPos}'$ which ensures that, for each $d \in \text{TransPos}'^{\mathcal{I}}$, $s_{2n+1}^{\mathcal{I}}(d) = \text{xpsn}(d) + 2^{n+1} \cdot \text{ypsn}(d)$:

$$\begin{aligned} \text{TransPos}' := & \exists \text{zero}.\text{=}_0 \sqcap \prod_{i=1..2n+1} (\exists t_i.\text{=}_{2^i}) \sqcap (\neg X_0 \rightarrow \exists s_0.\text{=}_0) \sqcap (X_0 \rightarrow \exists s_0.\text{=}_1) \sqcap \\ & \prod_{i=1..n} \left((\neg X_i \rightarrow \exists (s_{i-1}, \text{zero}, s_i).\text{+}) \sqcap (X_i \rightarrow \exists (s_{i-1}, t_i, s_i).\text{+}) \right) \sqcap \\ & \prod_{i=n+1..2n+1} \left((\neg Y_i \rightarrow \exists (s_{i-1}, \text{zero}, s_i).\text{+}) \sqcap (Y_i \rightarrow \exists (s_{i-1}, t_i, s_i).\text{+}) \right) \end{aligned}$$

where zero , s_i , and t_i are concrete features, =_k (with $k \in \mathbb{N}$) denotes a unary predicate with the obvious extension, and + denotes a ternary addition predicate.

It is easy to check that, whenever two objects $d, e \in \text{TransPos}'^{\mathcal{I}}$ do *not* agree on the interpretation of the $X_0, \dots, X_n, Y_0, \dots, Y_n$, then $s_{2n+1}^{\mathcal{I}}(d) \neq s_{2n+1}^{\mathcal{I}}(e)$, and thus the key box $\{(s_{2n+1} \text{ keyfor } \top)\}$ can be used for the reduction. The size of $\text{TransPos}'$ is obviously polynomial in n if the numbers k appearing in =_k predicates are coded in binary. We thus obtain the following theorem:

Theorem 3.11. *Let \mathcal{D} be a concrete domain such that*

1. $\mathbb{N} \subseteq \Delta_{\mathcal{D}}$,
2. $\Phi_{\mathcal{D}}$ contains, for each $k \in \mathbb{N}$, a predicate =_k with $(\text{=}_k)^{\mathcal{D}} = \{k\}$ where the size of (the representation of) =_k is logarithmic in k , and
3. $\Phi_{\mathcal{D}}$ contains a predicate + with $(\text{+})^{\mathcal{D}} \cap \{(k_1, k_2, x) \mid k_1, k_2 \in \mathbb{N} \text{ and } x \in \Delta_{\mathcal{D}}\} = \{(k_1, k_2, k_1 + k_2) \mid k_1, k_2 \in \mathbb{N}\}$.

Then the satisfiability of path-free $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. simple unary key boxes is NEXPTIME-hard.

For example, this theorem yields NEXPTIME-lower bounds for $\mathcal{ALCK}(\mathcal{D})$ instantiated with the concrete domains proposed in (Baader & Hanschke, 1992; Haarslev & Möller, 2001; Lutz, 2003, 2002b). An alternative to the addition predicate is to use multiplication to injectively translate sequences of bits into natural numbers. More precisely, let p_1, \dots, p_{2n+1} be the first $n + 1$ prime numbers and define another version of TransPos as follows:

$$\begin{aligned} \text{TransPos}'' := & \exists \text{one}.\text{=}_1 \sqcap \prod_{i=1..2n+1} (\exists t_i.\text{=}_{p_i}) \sqcap (\neg X_0 \rightarrow \exists s_0.\text{=}_0) \sqcap (X_0 \rightarrow \exists s_0.\text{=}_1) \sqcap \\ & \prod_{i=1..n} \left((\neg X_i \rightarrow \exists (s_{i-1}, \text{one}, s_i).\text{*}) \sqcap (X_i \rightarrow \exists (s_{i-1}, t_i, s_i).\text{*}) \right) \sqcap \\ & \prod_{i=n+1..2n+1} \left((\neg Y_i \rightarrow \exists (s_{i-1}, \text{one}, s_i).\text{*}) \sqcap (Y_i \rightarrow \exists (s_{i-1}, t_i, s_i).\text{*}) \right) \end{aligned}$$

where * is a ternary multiplication predicate.

Since the factorization of natural numbers into prime numbers is unique, we can again use the key box $\{(s_{2n+1} \text{ keyfor } \top)\}$ for the reduction. Moreover, it is well-known that the k 'th prime is polynomial in k (Graham, Knuth, & Patashnik, 1990), and thus the size of the concept $\text{TransPos}''$ is polynomial in n even if the numbers k in =_k predicates are coded unarily. We thus obtain another theorem concerning quite natural concrete domains:

Theorem 3.12. *Let \mathcal{D} be a concrete domain such that*

1. $\cdot \subseteq \Delta_{\mathcal{D}}$,
2. $\Phi_{\mathcal{D}}$ contains, for each $k \in \cdot$, a predicate $=_k$ with $(=_k)^{\mathcal{D}} = \{k\}$, and
3. $\Phi_{\mathcal{D}}$ contains a predicate $*$ with $(*)^{\mathcal{D}} \cap \{(k_1, k_2, x) \mid k_1, k_2 \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} = \{(k_1, k_2, k_1 \cdot k_2) \mid k_1, k_2 \in \cdot\}$.

Then the satisfiability of path-free $\mathcal{ALCK}(\mathcal{D})$ -concepts w.r.t. simple unary key boxes is NEXPTIME-hard.

3.4 NEXPTIME-hardness of $\mathcal{ALCO}(\mathcal{D})$

As we already pointed out in Section 1, the relationship between key boxes and nominals is rather close: the latter can be “simulated” by the former if the concrete domain provides predicates that can be used to uniquely describe elements of $\Delta_{\mathcal{D}}$. For example, in $\mathcal{ALCK}(\mathcal{D}_1)$ the concept $\exists g.=_0$ behaves as a nominal if we use the key definition (g keyfor \top). We can even define n nominals using n single concrete features in unary-key definitions. In the logics $\mathcal{ALCK}(\mathcal{D}_2)$ and $\mathcal{ALCK}(\mathcal{D}_3)$, a single concrete feature and unary key definitions are sufficient to simulate an arbitrary number of nominals: for example, in $\mathcal{ALCK}(\mathcal{D}_2)$ the concept $C = \exists g.\text{bit}0_2^0 \sqcap \exists g.\text{bit}1_2^1$ uniquely describes the bit vector $(0, 1) \in \text{BV}_2 \subseteq \Delta_{\mathcal{D}_2}$, i.e., $a \in C^{\mathcal{I}}$ implies $g^{\mathcal{I}}(a) = (0, 1)$. Obviously, any other bit vector (of any length!) can be described in a similar way.

This illustrates that, for most non-trivial concrete domains \mathcal{D} , the logic $\mathcal{ALCK}(\mathcal{D})$ is (at least) as expressive as $\mathcal{ALCO}(\mathcal{D})$. Although the converse does not hold, the expressive power of $\mathcal{ALCO}(\mathcal{D})$ is still sufficient to prove NEXPTIME-hardness of concept satisfiability, provided that a suitable concrete domain \mathcal{D} is used. In this section, we reduce the NEXPTIME-complete domino-problem to $\mathcal{ALCO}(\mathcal{D}_3)$ -concept satisfiability.

Again, let $\mathfrak{D} = (T, H, V)$ be a domino system and $a = a_0, \dots, a_{n-1}$ an initial condition. Then the reduction concept $C_{\mathfrak{D}, a}$ is defined in Figure 5, where $\text{bv}x$, $\text{bv}y$, $\text{bv}x_s$, $\text{bv}y_s$, and darr denote concrete features, N denotes a nominal, and the concepts $\text{Tree}X$, $\text{Tree}Y$, $\text{Dist}X_k$, and $\text{Dist}Y_k$ are defined as in Figure 3. As in the previous reductions, we now give a detailed explanation of the reduction strategy to show that $C_{\mathfrak{D}, a}$ is satisfiable iff there exists a solution for \mathfrak{D} and a . Formal details can then easily be worked out by the interested reader.

Let \mathcal{I} be a model for $C_{\mathfrak{D}, a}$. To explain the structure of \mathcal{I} , it is convenient to start with the first line of $C_{\mathfrak{D}, a}$. As in the previous reductions, the $\text{Tree}X$ and $\text{Tree}Y$ concepts are used to ensure that \mathcal{I} contains a tree-shaped substructure of depth $n + 1$ whose leaf nodes are the roots of additional trees of depth $n + 1$ such that the set of the leaf nodes of the latter trees correspond to the positions in the $2^{n+1} \times 2^{n+1}$ -torus, i.e., for each position, there is a leaf node representing it. The torus positions are binarily encoded by the concept names X_0, \dots, X_n and Y_0, \dots, Y_n and we use $e_{i,j}$ to refer to the leaf with $\text{xpsn}(e_{i,j}) = i$ and $\text{yposn}(e_{i,j}) = j$ (see Section 3.3).

As in the previous reductions, the numbers coded by X_0, \dots, X_n and Y_0, \dots, Y_n are translated into concrete domain values, which is done by the $\text{Trans}X\text{Pos}$ and $\text{Trans}Y\text{Pos}$ concepts. Note that, in contrast to the $\mathcal{ALCK}(\mathcal{D}_2)$ -reduction, the x-position and the y-position are not stored in the same bit vector, but rather in the two distinct ones $\text{bv}x$


$$\begin{aligned}
 \text{Nominal} &:= \exists f. N \\
 \text{XSucc} &:= \prod_{k=0..n} \left(\prod_{j=0..k} X_j \right) \rightarrow (X_k \leftrightarrow \neg X'_k) \sqcap \prod_{k=0..n} \left(\bigsqcup_{j=0..k} \neg X_j \right) \rightarrow (X_k \leftrightarrow X'_k) \\
 \text{YSucc} &:= \prod_{k=0..n} \left(\prod_{j=0..k} Y_j \right) \rightarrow (Y_k \leftrightarrow \neg Y'_k) \sqcap \prod_{k=0..n} \left(\bigsqcup_{j=0..k} \neg Y_j \right) \rightarrow (Y_k \leftrightarrow Y'_k) \\
 \text{TransXPos} &:= \prod_{i=0..n} \left((X_i \rightarrow \exists \text{bv}x. \text{pos}1_{n+1}^i) \sqcap (\neg X_i \rightarrow \exists \text{bv}x. \text{pos}0_{n+1}^i) \right) \\
 \text{TransYPos} &:= \prod_{i=0..n} \left((Y_i \rightarrow \exists \text{bv}y. \text{pos}1_{n+1}^i) \sqcap (\neg Y_i \rightarrow \exists \text{bv}y. \text{pos}0_{n+1}^i) \right) \\
 \text{TransXSucc} &:= \prod_{i=0..n} \left((X'_i \rightarrow \exists \text{bv}x_s. \text{pos}1_{n+1}^i) \sqcap (\neg X'_i \rightarrow \exists \text{bv}x_s. \text{pos}0_{n+1}^i) \right) \\
 \text{TransYSucc} &:= \prod_{i=0..n} \left((Y'_i \rightarrow \exists \text{bv}y_s. \text{pos}1_{n+1}^i) \sqcap (\neg Y'_i \rightarrow \exists \text{bv}y_s. \text{pos}0_{n+1}^i) \right) \\
 \text{CheckHMatch} &:= \bigsqcup_{i,j \in H} (\exists (\text{bv}x, \text{bv}y, f \circ \text{darr}). \text{tile}_{n+1}^i \sqcap \exists (\text{bv}x_s, \text{bv}y, f \circ \text{darr}). \text{tile}_{n+1}^j) \\
 \text{CheckVMatch} &:= \bigsqcup_{i,j \in V} (\exists (\text{bv}x, \text{bv}y, f \circ \text{darr}). \text{tile}_{n+1}^i \sqcap \exists (\text{bv}x, \text{bv}y_s, f \circ \text{darr}). \text{tile}_{n+1}^j) \\
 \text{Init2} &:= \prod_{i=0..n-1} \left(\left(\prod_{j=0..n, \text{bit}_j(i)=0} \neg X_j \sqcap \prod_{j=0..n, \text{bit}_j(i)=1} X_j \sqcap \prod_{j=0..n} \neg Y_j \right) \right. \\
 &\quad \left. \rightarrow \exists (\text{bv}x, \text{bv}y, f \circ \text{darr}). \text{tile}_{n+1}^{a_i} \right) \\
 C_{\mathcal{D},a} &:= \text{TreeX} \sqcap \forall R^{n+1}. \text{TreeY} \sqcap \forall R^{2(n+1)}. \text{Nominal} \sqcap \\
 &\quad \forall R^{2(n+1)}. (\text{TransXPos} \sqcap \text{TransYPos} \sqcap \\
 &\quad \text{XSucc} \sqcap \text{YSucc} \sqcap \text{TransXSucc} \sqcap \text{TransYSucc} \sqcap \\
 &\quad \text{Init2} \sqcap \text{CheckHMatch} \sqcap \text{CheckVMatch})
 \end{aligned}$$

 Figure 5: The $\mathcal{ALCO}(\mathcal{D}_3)$ reduction concept $C_{\mathcal{D},a}$.

and $\text{bv}y$. Also in contrast to the previous reduction, the actual tiling of the torus is *not* represented by the leaf nodes $e_{i,j}$, but rather by a domino array: the last conjunct in the first line of $C_{\mathcal{D},a}$ ensures that every leaf $e_{i,j}$ is connected via the abstract feature f to the (unique) element $w \in N^{\mathcal{I}}$.

The domain element w is associated with a domino array via the concrete feature darr (as we shall see later, this is guaranteed by the CheckHMatch and CheckVMatch concepts). This domino array represents the tiling of the $2^{n+1} \times 2^{n+1}$ -torus. Summing up, the structure of \mathcal{I} is roughly as shown in Figure 6.

Since the tiling is stored in a domino array, we need to explain the purpose of the leaf nodes $e_{i,j}$: these nodes are used to enforce the initial condition and the horizontal and vertical matching condition. Let us discuss the horizontal matching condition (the vertical matching condition is enforced analogously): the XSucc concept is the DL reformulation

Figure 6: The structure of models of $C_{\mathcal{D},a}$.

of the propositional logic formula for incrementation modulo 2^{n+1} and ensures that, for each $e_{i,j}$, the concept names X'_0, \dots, X'_n encode the number $i \oplus_{2^{n+1}} 1$, i.e., the horizontal position of $e_{i,j}$'s horizontal neighbor. In addition to the storage of the horizontal and vertical position of $e_{i,j}$ in $\text{bv}_x(e_{i,j})$ and $\text{bv}_y(e_{i,j})$, we also store the horizontal position $i \oplus_{2^{n+1}} 1$ of $e_{i,j}$'s horizontal successor in $\text{bv}_{xs}(e_{i,j})$. Finally, `CheckHMatch` verifies that the tiles of the positions (i, j) and $(i \oplus_{2^{n+1}} 1, j)$, which are both stored in the domino array, are compatible with the horizontal matching condition.

Note that `CheckHMatch` also ensures that the domain element w (with $\{w\} = N^{\mathcal{I}}$) has a domino array attached via the concrete feature `darr` and that, for each position (i, j) , the (unique!) tile stored in the domino array is from the set T . The initial condition is ensured via the `Init2` concept in a similar way. We (again) use $\text{bit}_j(i)$ to denote the j 'th bit of the binary encoding of the natural number i .

Using the above considerations, the correctness of the reduction is readily checked. Moreover, the size of $C_{\mathcal{D},a}$ is at most polynomial in n . Note that $C_{\mathcal{D},a}$ is *not* path-free: paths of length two appear in the concepts `CheckHMatch`, `CheckVMatch`, and `Init2`. Summing up, the reduction described yields the following result:

Proposition 3.13. *The satisfiability of $\mathcal{ALCO}(\mathcal{D}_3)$ -concepts is NEXPTIME-hard.*

Again, we relate the NEXPTIME lower bound to the complexity of \mathcal{D}_3 -complexity, which is determined in Proposition 3.5.

Theorem 3.14. *There exists a concrete domain \mathcal{D} such that \mathcal{D} -satisfiability is in PTIME and the satisfiability of $\mathcal{ALCO}(\mathcal{D})$ -concepts is NEXPTIME-hard.*

Note that the reduction uses only a *single nominal* N . This is a dramatic increase of complexity since it was shown in (Lutz, 2002b) that satisfiability of $\mathcal{ALC}(\mathcal{D})$ -concepts (i.e.,

without nominals and key boxes) is PSPACE-complete provided that \mathcal{D} is admissible and \mathcal{D} -satisfiability is in PSPACE.

As in previous sections, we note that D_3 can be replaced by more natural concrete domains in the NEXPTIME-hardness proof presented. The idea is to represent the whole domino array by a single natural number and then to use arithmetic operations to access the individual positions: a natural number k can be viewed as a domino array by partitioning its binary representation into $2^{n+1} \cdot 2^{n+1} = 2^{2(n+1)}$ “sections” of length $\lceil \log(\#T) \rceil$, where $\#T$ denotes the cardinality of the set of tile types T . Each such section describes the tile of a single position in the torus. The sections can be accessed by using integer division and remainder operations: if k is the natural number representing the torus, then the tile of position i is computed by

$$(k \operatorname{div} 2^{i \cdot \lceil \log(\#T) \rceil}) \bmod 2^{\lceil \log(\#T) \rceil} + 1.$$

Thus, we introduce ternary predicates div for integer division and mod for computing the remainder of a division, and a binary predicate 2^\times expressing exponentiation with basis 2. Then we modify the reduction as follows: we replace $\operatorname{TransXPos}$ and $\operatorname{TransYPos}$ by the $\operatorname{TransPos}'$ concept from Section 3.3 to translate the two numbers encoded by X_1, \dots, X_n and Y_1, \dots, Y_n into a single natural number that is stored in the concrete feature s_{2n+1} . We then devise a new concept $\operatorname{Tile}[i]$ (for each $i \in T$) enforcing that the position identified by the feature s_{2n+1} is labeled with tile i :

$$\begin{aligned} \operatorname{Tile}[i] := & \exists r. =_{\lceil \log(\#T) \rceil} \sqcap \exists s_{2n+1}, r, r'. * \sqcap \exists r', r''. 2^\times \sqcap \exists \text{one}. =_1 \sqcap \exists r, \text{one}, t. + \sqcap \exists t, t'. 2^\times \\ & \sqcap \exists f_{\text{torus}}, r'', u. \operatorname{div} \sqcap \exists u, t'', \operatorname{tile}. \operatorname{mod} \sqcap \exists \operatorname{tile}. =_i. \end{aligned}$$

Here, $r, r', r'', t, t', u, \text{one}, \text{torus}$, and tile are concrete features. The torus feature is the counterpart of the darr feature in the original reduction, i.e., it stores the natural number that represents the tiling array. We can use the $\operatorname{Tile}[i]$ concept in the obvious way inside the $\operatorname{CheckHMatch}$, $\operatorname{CheckVMatch}$, and $\operatorname{Init2}$ concepts. The size of the resulting reduction concept is polynomial in n if the numbers k appearing in $=_k$ predicates are coded in binary. We thus obtain the following theorem:

Theorem 3.15. *Let \mathcal{D} be a concrete domain such that*

1. $\cdot \subseteq \Delta_{\mathcal{D}}$,
2. $\Phi_{\mathcal{D}}$ contains the predicates a predicate $=_k$ (for each $k \in \cdot$), 2^\times , $+$, $*$, div , mod with the following extensions

$$\begin{aligned} (=k)^{\mathcal{D}} &= \{k\} \\ (2^\times)^{\mathcal{D}} \cap \{(k, x) \mid k \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} &= \{(k, 2^k) \mid k \in \cdot\} \\ (+)^{\mathcal{D}} \cap \{(k_1, k_2, x) \mid k_1, k_2 \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} &= \{(k_1, k_2, k_1 + k_2) \mid k_1, k_2 \in \cdot\} \\ (*)^{\mathcal{D}} \cap \{(k_1, k_2, x) \mid k_1, k_2 \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} &= \{(k_1, k_2, k_1 \cdot k_2) \mid k_1, k_2 \in \cdot\} \\ (\operatorname{div})^{\mathcal{D}} \cap \{(k_1, k_2, x) \mid k_1, k_2 \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} &= \{(k_1, k_2, k_1 \operatorname{div} k_2) \mid k_1, k_2 \in \cdot\} \\ (\operatorname{mod})^{\mathcal{D}} \cap \{(k_1, k_2, x) \mid k_1, k_2 \in \cdot \text{ and } x \in \Delta_{\mathcal{D}}\} &= \{(k_1, k_2, k_1 \bmod k_2) \mid k_1, k_2 \in \cdot\} \end{aligned}$$

Then the satisfiability of $\mathcal{ALCO}(\mathcal{D})$ -concepts is NEXPTIME-hard.

4. Reasoning Procedures

This section is devoted to developing reasoning procedures for DLs with concrete domains, nominals, and keys. We start with devising a tableau algorithm that decides the satisfiability of $\mathcal{ALCOK}(\mathcal{D})$ -concepts w.r.t. Boolean key boxes. This algorithm yields a NEXPTIME upper complexity bound matching the lower bounds established in Section 3.3.

Then we consider the rather powerful description logic $\mathcal{SHOQK}(\mathcal{D})$. This DL, which is an extension of $\mathcal{SHOQ}(\mathcal{D})$ as introduced in (Horrocks & Sattler, 2001; Pan & Horrocks, 2002), provides a wealth of expressive means such as transitive roles, role hierarchies, nominals, and qualifying number restrictions. Moreover, $\mathcal{SHOQK}(\mathcal{D})$ is equipped with a restricted variant of the concrete domain constructor and with key boxes. We develop a tableau algorithm for deciding the satisfiability of $\mathcal{SHOQK}(\mathcal{D})$ -concepts w.r.t. path-free key boxes. Due to the restrictedness of $\mathcal{SHOQK}(\mathcal{D})$'s concrete domain constructor, we can even admit general rather than only Boolean key boxes. Again, the algorithm yields a tight NEXPTIME upper complexity bound.

4.1 A Tableau Algorithm for $\mathcal{ALCOK}(\mathcal{D})$ with Boolean Key Boxes

Tableau algorithms decide the satisfiability of the input concept (in our case w.r.t. the input key box) by attempting to construct a model for it. More precisely, a tableau algorithm starts with an initial data structure induced by the input concept and then repeatedly applies so-called completion rules to it. This rule application can be thought of as attempting to construct a model for the input concept. Finally, either the algorithm will find an obvious contradiction or it will encounter a situation that is contradiction-free and in which no more completion rules are applicable. In the former case, the input concept is unsatisfiable, while it is satisfiable in the latter.

When devising a tableau algorithm for a description logic with concrete domains but without committing to a *particular* concrete domain, it is commonly assumed that the concrete domain is admissible, which in particular implies decidability of the satisfiability of \mathcal{D} -conjunctions. In the presence of keys, however, this is not enough: if a \mathcal{D} -conjunction is satisfiable, we also want to know which of its variables take the same values in an arbitrary but fixed solution. As an example, consider the concrete domain $\mathbf{N} = (\mathbb{N}, \{<_n \mid n \in \mathbb{N}\})$ and the \mathbf{N} -conjunction

$$c = <_2(v_1) \wedge <_2(v_2) \wedge <_2(v_3).$$

Obviously, one solution δ for c satisfies $\delta(v_1) = \delta(v_2)$, another satisfies $\delta(v_1) = \delta(v_3)$, and so on. Our tableau algorithm uses such identity information passed from the concrete domain reasoner since, in the presence of key boxes, it can have an impact on the structure of the constructed model. For example, this information reveals the unsatisfiability of

$$\exists R.A \sqcap \exists R.(\neg A \sqcap B) \sqcap \exists R.(\neg A \sqcap \neg B) \sqcap \forall R.\exists g.<_2 \text{ w.r.t. } (g \text{ key for } \top).$$

To formalize this requirement, we strengthen the notion of admissibility into *key-admissibility*. Since the tableau algorithm developed in this section is non-deterministic, we formulate key-admissibility in a non-deterministic way.

Definition 4.1 (Key-admissible). A concrete domain \mathcal{D} is *key-admissible* iff it satisfies the following properties:

1. $\Phi_{\mathcal{D}}$ contains a name $\top_{\mathcal{D}}$ for $\Delta_{\mathcal{D}}$;
2. $\Phi_{\mathcal{D}}$ is closed under negation;
3. there exists an algorithm that takes as input a \mathcal{D} -conjunction c , returns **clash** if c is unsatisfiable, and otherwise non-deterministically outputs an equivalence relation \sim on the set of variables V used in c such that there exists a solution δ for c with the following property: for all $v, v' \in V$

$$\delta(v) = \delta(v') \text{ iff } v \sim v'.$$

An algorithm showing the behaviour described in item 3 above is called a \mathcal{D} -*tester*, and the equivalence relations \sim are called *concrete equivalences*. We say that *extended \mathcal{D} -satisfiability is in NP* if there exists a \mathcal{D} -tester running in polynomial time.

Please note that key-admissibility is less esoteric than it might seem: any concrete domain that is admissible and provides for an equality predicate is also key-admissible. Due to admissibility, the presence of an equality predicate implies that an inequality predicate is also available. We can thus construct a \mathcal{D} -tester from an algorithm for \mathcal{D} -satisfiability: when presented with a predicate conjunction c , we simply “guess” an equivalence relation \sim on the set of variables used in c . Then we decide the (non-extended) satisfiability of the conjunction $c \wedge \bigwedge_{v \sim v'} = (v, v') \wedge \bigwedge_{v \not\sim v'} \neq (v, v')$, return **clash** if it is unsatisfiable and \sim otherwise. The rather weak condition that an equality predicate should be present is satisfied by almost all concrete domains proposed in the literature, see e.g. (Lutz, 2003; Baader & Hanschke, 1991b; Kamp & Wache, 1996; Haarslev, Lutz, & Möller, 1998; Baader & Sattler, 1998).

Throughout this chapter, we assume that any concrete domain is equipped with an equality predicate. This assumption is w.l.o.g. since any \mathcal{D} -conjunction using equality can be translated into an equivalent one without equality by identifying variables according to the stated equalities. This assumption must not be confused with what was discussed in the previous paragraph: even if the concrete domain \mathcal{D} is admissible and its set of predicates is thus closed under negation, this assumption does *not* imply the presence of an inequality predicate.

We need some more prerequisites before we can start the presentation of the tableau algorithm: a concept is in *negation normal form (NNF)* if negation occurs only in front of concept names and nominals. It is easily seen that, if the concrete domain \mathcal{D} is admissible, then every $\mathcal{ALCOK}(\mathcal{D})$ -concept can be converted into an equivalent one in NNF by exhaustively applying the rewrite rules displayed in Figure 7. We use $\dot{\neg}C$ to denote the result of converting $\neg C$ to NNF. A key box is in NNF if all concepts occurring in key definitions are in NNF. In what follows, we generally assume input concepts and key boxes to be in NNF. Let C be an $\mathcal{ALCOK}(\mathcal{D})$ -concept and \mathcal{K} a key box. We use $\text{sub}(C)$ to denote the set of subconcepts of C (including C itself) and $\text{con}(\mathcal{K})$ to denote the set of concepts appearing on the right-hand side of key definitions in \mathcal{K} . For a set of concepts Γ , $\text{sub}(\Gamma)$ denotes the set $\bigcup_{C \in \Gamma} \text{sub}(C)$. Moreover, we write $\text{cl}(C, \mathcal{K})$ as abbreviation for the set

$$\text{sub}(C) \cup \text{sub}(\text{con}(\mathcal{K})) \cup \{\dot{\neg}D \mid D \in \text{sub}(\text{con}(\mathcal{K}))\}.$$

| |
|--|
| $\begin{aligned} \neg(C \sqcap D) &\rightsquigarrow \neg C \sqcup \neg D & \neg(C \sqcup D) &\rightsquigarrow \neg C \sqcap \neg D & \neg\neg C &\rightsquigarrow C \\ \neg(\exists R.C) &\rightsquigarrow \forall R.\neg C & \neg(\forall R.C) &\rightsquigarrow \exists R.\neg C \\ \\ \neg(\exists u_1, \dots, u_n.P) &\rightsquigarrow \exists u_1, \dots, u_n.\overline{P} \sqcup u_1 \uparrow \sqcup \dots \sqcup u_n \uparrow \\ \neg(g \uparrow) &\rightsquigarrow \exists g.\top_{\mathcal{D}} \end{aligned}$ |
|--|

Figure 7: The NNF rewrite rules.

We now start the presentation of the tableau algorithm by introducing the underlying data structure.

Definition 4.2 (Completion System). Let O_a and O_c be disjoint and countably infinite sets of *abstract* and *concrete nodes*. A *completion tree* for an $\mathcal{ALCOK}(\mathcal{D})$ -concept C and a key box \mathcal{K} is a finite, labeled tree $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$ with nodes $V_a \cup V_c$ such that $V_a \subseteq O_a$, $V_c \subseteq O_c$, and all nodes from V_c are leaves. The tree is labeled as follows:

- each node $a \in V_a$ is labeled with a subset $\mathcal{L}(a)$ of $\text{cl}(C, \mathcal{K})$;
- each edge $(a, b) \in E$ with $a, b \in V_a$ is labeled with a role name $\mathcal{L}(a, b)$ occurring in C or \mathcal{K} ;
- each edge $(a, x) \in E$ with $a \in V_a$ and $x \in V_c$ is labeled with a concrete feature $\mathcal{L}(a, x)$ occurring in C or \mathcal{K} .

For $a \in V_a$, we use $\text{lev}_{\mathbf{T}}(a)$ to denote the depth on which a occurs in \mathbf{T} (starting with the root node on depth 0). A *completion system* for an $\mathcal{ALCOK}(\mathcal{D})$ -concept C and a key box \mathcal{K} is a tuple $(\mathbf{T}, \mathcal{P}, \prec, \sim)$, where

- $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$ is a completion tree for C and \mathcal{K} ,
- \mathcal{P} is a function mapping each $P \in \Phi_{\mathcal{D}}$ of arity n in C to a subset of V_c^n ,
- \prec is a linear ordering of V_a such that $\text{lev}_{\mathbf{T}}(a) \leq \text{lev}_{\mathbf{T}}(b)$ implies $a \prec b$, and
- \sim is an equivalence relation on V_c .

Let $(V_a, V_c, E, \mathcal{L})$ be a completion tree. A node $b \in V_a$ is an *R-successor* of a node $a \in V_a$ if $(a, b) \in E$ and $\mathcal{L}(a, b) = R$, while a node $x \in V_c$ is a *g-successor* of a if $(a, x) \in E$ and $\mathcal{L}(a, x) = g$. For a path u , the notion *u-successor* is defined in the obvious way.

Intuitively, the relation \sim records equalities between concrete nodes found during the (non-deterministic) model construction process. The recording is necessary since equalities between concrete nodes can induce equalities between abstract nodes which, in turn, can imply more equalities between concrete nodes. This can be seen in the following example: assume the completion tree contains, for $i \in \{1, 2\}$, an abstract node a_i with a concrete g -successor x_i and a concrete g' -successor y_i . Now assume that the key box contains (g key for \top), and that the \mathcal{D} -tester returns $x_1 \sim x_2$. As a consequence, a_1 and a_2 represent the same element and thus functionality of g' implies that also y_1 and y_2 represent the same (concrete) element. To deal with such effects, we define an equivalence relation \approx_a on abstract nodes and a second equivalence relation \approx_c on concrete nodes.

Definition 4.3 (\approx_a and \approx_c Relations). Let $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ be a completion system for a concept C and a key box \mathcal{K} with $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$, and let \approx be an equivalence relation on V_a . For each $R \in \mathbf{N}_R$, a node $b \in V_a$ is an R/\approx -neighbor of a node $a \in V_a$ if there exists a node $c \in V_a$ such that $a \approx c$ and b is an R -successor of c . Similarly, for each $g \in \mathbf{N}_{cF}$, a node $x \in V_c$ is a g/\approx -neighbor of a if there exists a node $c \in V_a$ such that $a \approx c$ and x is a g -successor of c . For paths u , the notion u/\approx -neighbor is defined in the obvious way.

We define a sequence of equivalence relations $\approx_a^0 \subseteq \approx_a^1 \subseteq \dots$ on V_a as follows:

$$\begin{aligned} \approx_a^0 &= \{(a, a) \in V_a^2\} \cup \\ &\quad \{(a, b) \in V_a^2 \mid \text{there is an } N \in \mathbf{N}_O \text{ such that } N \in \mathcal{L}(a) \cap \mathcal{L}(b)\} \\ \approx_a^{i+1} &= \approx_a^i \cup \\ &\quad \{(a, b) \in V_a^2 \mid \text{there is a } c \in V_a \text{ and an } f \in \mathbf{N}_{aF} \text{ such that} \\ &\quad \quad \quad a \text{ and } b \text{ are } f/\approx_a^i\text{-neighbors of } c\} \cup \\ &\quad \{(a, b) \in V_a^2 \mid \text{there is a } (u_1, \dots, u_n \text{ keyfor } C) \in \mathcal{K}, \\ &\quad \quad \quad u_i/\approx_a^i\text{-neighbors } x_i \text{ of } a \text{ for } 1 \leq i \leq n, \text{ and} \\ &\quad \quad \quad u_i/\approx_a^i\text{-neighbors } y_i \text{ of } b \text{ for } 1 \leq i \leq n \\ &\quad \quad \quad \text{such that } C \in \mathcal{L}(a) \cap \mathcal{L}(b) \text{ and } x_i \sim y_i \text{ for } 1 \leq i \leq n\} \end{aligned}$$

Finally, set $\approx_a = \bigcup_{i \geq 0} \approx_a^i$. Then define

$$\approx_c = \sim \cup \{(x, y) \in V_c^2 \mid \text{there is an } a \in V_a \text{ and a } g \in \mathbf{N}_{cF} \text{ such that} \\ x \text{ and } y \text{ are } g/\approx_a\text{-neighbors of } a\}.$$

This definition reflects the above mentioned tight coupling between the concrete and abstract equalities: if the \mathcal{D} -tester finds (or guesses) that two concrete nodes are equal, the tableau algorithm may use this to deduce (via the computation of \approx_a and \approx_c) even more equalities between concrete nodes.

Let \mathcal{D} be a key-admissible concrete domain. To decide the satisfiability of an $\mathcal{ALCOK}(\mathcal{D})$ -concept C_0 w.r.t. a Boolean key box \mathcal{K} (both in NNF), the tableau algorithm is started with the *initial completion tree*

$$\mathbf{T}_{C_0} = (\{a_0\}, \emptyset, \emptyset, \{a_0 \mapsto \{C_0\}\})$$

in the *initial completion system*

$$S_{C_0} = (\mathbf{T}_{C_0}, \mathcal{P}_\emptyset, \emptyset, \emptyset),$$

where \mathcal{P}_\emptyset maps each $P \in \Phi_{\mathcal{D}}$ occurring in C_0 to \emptyset . We now introduce an operation that is used by the completion rules to add new nodes to completion trees.

Definition 4.4 (“+” Operation). An abstract or concrete node is called *fresh* in a completion tree \mathbf{T} if it does not appear in \mathbf{T} . Let $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ be a completion system with $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$. We use the following notions:

- Let $a \in V_a$, $b \in O_a$ fresh in \mathbf{T} , and $R \in \mathbf{N}_R$. We write $S + aRb$ to denote the completion system S' that can be obtained from S by adding b to V_a and (a, b) to E and setting $\mathcal{L}(a, b) = R$ and $\mathcal{L}(b) = \emptyset$. Moreover, b is inserted into \prec such that $b \prec c$ implies $\text{lev}_{\mathbf{T}}(b) \leq \text{lev}_{\mathbf{T}}(c)$.

| | |
|---------------|---|
| R \sqcap | if $C_1 \sqcap C_2 \in \mathcal{L}(a)$ and $\{C_1, C_2\} \not\subseteq \mathcal{L}(a)$ then $\mathcal{L}(a) := \mathcal{L}(a) \cup \{C_1, C_2\}$ |
| R \sqcup | if $C_1 \sqcup C_2 \in \mathcal{L}(a)$ and $\{C_1, C_2\} \cap \mathcal{L}(a) = \emptyset$ then $\mathcal{L}(a) := \mathcal{L}(a) \cup \{C\}$ for some $C \in \{C_1, C_2\}$ |
| R \exists | if $\exists R.C \in \mathcal{L}(a)$ and there is no R/\approx_a -neighbor b of a such that $C \in \mathcal{L}(b)$, then set $S := S + aRb$ for a fresh $b \in \mathcal{O}_a$ and $\mathcal{L}(b) := \{C\}$ |
| R \forall | if $\forall R.C \in \mathcal{L}(a)$, b is an R/\approx_a -neighbor of a , and $C \notin \mathcal{L}(b)$ then set $\mathcal{L}(b) := \mathcal{L}(b) \cup \{C\}$ |
| R $\exists c$ | if $\exists u_1, \dots, u_n.P \in \mathcal{L}(a)$ and there exist no $x_1, \dots, x_n \in V_c$ such that x_i is u_i/\approx_a -neighbor of a for $1 \leq i \leq n$ and $(x_1, \dots, x_n) \in \mathcal{P}(P)$ then set $S := (S + au_1x_1 + \dots + au_nx_n)$ with $x_1, \dots, x_n \in \mathcal{O}_c$ fresh and $\mathcal{P}(P) := \mathcal{P}(P) \cup \{(x_1, \dots, x_n)\}$ |
| Rch | if $(u_1, \dots, u_n \text{ keyfor } C) \in \mathcal{K}$ and there exist $x_1, \dots, x_n \in V_c$ such that x_i is u_i/\approx_a -neighbor of a for $1 \leq i \leq n$ and $\{C, \dot{C}\} \cap \mathcal{L}(a) = \emptyset$ then set $\mathcal{L}(a) := \mathcal{L}(a) \cup \{D\}$ for some $D \in \{C, \dot{C}\}$ |
| Rp | if $\mathcal{L}(b) \not\subseteq \mathcal{L}(a)$ and $a \in V_a$ is minimal w.r.t. \prec such that $a \approx_a b$ then set $\mathcal{L}(a) := \mathcal{L}(a) \cup \mathcal{L}(b)$ |

 Figure 8: Completion rules for $\mathcal{ALCCOK}(\mathcal{D})$.

- Let $a \in V_a$, $x \in \mathcal{O}_c$ fresh in \mathbf{T} and $g \in \mathcal{N}_{\text{CF}}$. We write $S + agx$ to denote the completion system S' that can be obtained from S by adding x to V_c and (a, x) to E and setting $\mathcal{L}(a, x) = g$.

When nesting the $+$ operation, we omit brackets writing, e.g., $S + aR_1b + bR_2c$ for $(S + aR_1b) + bR_2c$. Let $u = f_1 \dots f_n g$ be a path. With $S + aux$, where $a \in V_a$ and $x \in \mathcal{O}_c$ is fresh in \mathbf{T} , we denote the completion system S' that can be obtained from S by taking distinct nodes $b_1, \dots, b_n \in \mathcal{O}_a$ which are fresh in \mathbf{T} and setting

$$S' := S + af_1b_1 + \dots + b_{n-1}f_nb_n + b_n g x.$$

Strictly speaking, the $S + aRb$ operation is non-deterministic since we did not specify how precisely the node b is inserted into \prec . However, since this is *don't care* non-determinism, we will view the “ $+$ ” operation as being deterministic.

The completion rules can be found in Figure 8. Note that the R \sqcup and Rch rules are non-deterministic, i.e., they have more than one possible outcome (this is true *don't know* non-determinism). Some further remarks on the completion rules are in order: the upper five rules are well-known from existing tableau algorithms for $\mathcal{ALCC}(\mathcal{D})$ -concept satisfiability, see for example (Lutz, 2002a). Only the use of R/\approx_a -neighbors and u/\approx_a -neighbors in the rules R \forall , R \exists , and R $\exists c$ deserves a comment. Take for example R \forall : intuitively, if we have $a \approx_a b$ for two abstract nodes a and b of the completion tree, then a and b describe the same domain element of the constructed model (and similarly for the \approx_c relation on concrete

nodes). Thus if $a \approx_a b$ and c is an R -successor of a , then c should also be an R -successor of b . However, since we want the completion tree to be a tree, we do not make the latter successorship explicit. To compensate for this, the $R\forall$ rule talks about R/\approx_a -neighbors rather than about R -successors.

The lower two rules are necessary for dealing with key boxes. The Rch rule is a so-called “choose rule” (Hollunder & Baader, 1991; Horrocks et al., 2000): intuitively, it guesses whether or not an abstract node a satisfies C if there exists a key definition $(u_1, \dots, u_n \text{ keyfor } C) \in \mathcal{K}$ such that there are neighbors of a for all the paths u_i . This is necessary since both possibilities may have ramifications: if a satisfies C , then it must be taken into account in the construction of the relation \approx_a ; if a does not satisfy C , then we must deal with the consequences of it satisfying $\neg C$ (e.g. in case that C is \top).

The Rp rule is dealing with equalities between abstract nodes as recorded by the \approx_a relation: since $a \approx_a b$ means that a and b describe the same node in the constructed model, their node labels should be identical. It suffices, however, to choose one representative for each equivalence class of \approx_a and make sure that this representative’s node label contains the labels of all its \approx_a -equivalent nodes. As representative, we use the node that is minimal w.r.t. the ordering \prec , which has been introduced for solely this reason. The Rp rule does the appropriate copying of node labels.

Let us now formalize what it means for a completion system to contain a contradiction.

Definition 4.5 (Clash). Let $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ be a completion system for a concept C and a key box \mathcal{K} with $\mathbf{T} = (V_a, V_c, \prec, \sim)$. We say that the completion system S is *concrete domain satisfiable* iff the conjunction

$$\zeta_S = \bigwedge_{P \text{ used in } C} \bigwedge_{(x_1, \dots, x_n) \in \mathcal{P}(P)} P(x_1, \dots, x_n) \wedge \bigwedge_{x \approx_c y} \neg(x, y)$$

is satisfiable. S is said to contain a *clash* iff

1. there is an $a \in V_a$ and an $A \in \mathbf{N}_C$ such that $\{A, \neg A\} \subseteq \mathcal{L}(a)$,
2. there are $a \in V_a$ and $x \in V_c$ such that $g \uparrow \in \mathcal{L}(a)$ and x is g/\approx_a -neighbor of a ,
3. S is not concrete domain satisfiable.

If S does not contain a clash, S is called *clash-free*. S is called *complete* iff no completion rule is applicable to S .

The tableau algorithm is described in Figure 9 in pseudo-code notation. In this figure, `test` calls a \mathcal{D} -tester as specified in Definition 4.1. Let us spend a few words on the while loop. There obviously exist close relationships between the relations \sim and \approx_c and the predicate conjunction ζ_S :

- $\sim \subseteq \approx_c$ (note that both \approx_a and \approx_c depend on \sim and are thus recomputed in each step of the while loop);
- by definition of ζ_S and \mathcal{D} -tester, the result of `test(ζ_S)` yields a relation containing \approx_c (and thus also \sim).

```

define procedure sat( $S$ )
  do
    if  $S$  contains a clash then
      return unsatisfiable
     $\sim := \text{test}(\zeta_S)$ 
    compute  $\approx_a$ 
    compute  $\approx_c$ 
  while  $\sim \neq \approx_c$ 
  if  $S$  contains a clash then
    return unsatisfiable
  if  $S$  is complete then
    return satisfiable
   $S' :=$  the application of a completion rule to  $S$ 
  return sat( $S'$ )

```

Figure 9: The $\mathcal{ALCOK}(\mathcal{D})$ tableau algorithm.

Using these facts, one may check that, in each step of the while loop, new tuples are added to the \sim relation, but none are deleted (see the proof of Lemma B.2 in the appendix). The while loop is due to the facts that (i) \approx_a is defined using \sim , (ii), \approx_c is defined using \approx_a , and (iii) new concrete equalities in \approx_c may then imply even more concrete and/or abstract equalities, and so on.

A similar concrete-abstract interplay takes place in the course of several recursion steps: equalities between concrete nodes provided by the \mathcal{D} -tester may make new rules applicable (for example Rp and $\text{R}\exists\text{c}$) which changes \mathcal{P} and thus also ζ_S . This may subsequently lead to the detection of more equalities between concrete nodes by the \mathcal{D} -tester, and so forth. These considerations show that, in the presence of keys, there exists a close interplay between the concrete domain reasoner and the tableau algorithm, which is not needed if keys are not present: without keys, it suffices to apply the concrete domain satisfiability check only once after the completion rules have been exhaustively applied (Baader & Hanschke, 1991a).

The detailed proof of termination, soundness, and completeness together with a complexity analysis of the tableau algorithm defined in this section is given in Appendix B.

Theorem 4.1. *Let \mathcal{D} be a key-admissible concrete domain. If extended \mathcal{D} -satisfiability is in NP, then $\mathcal{ALCOK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean key boxes is in NEXPTIME.*

We should note that, in the way it is presented here, the algorithm leaves quite some room for optimizations. One possible optimization concerns the “re-use” of f -successors (for abstract features f): for example, when applying the $\text{R}\exists$ rule to a concept $\exists f.C \in \mathcal{L}(a)$, where a already has an f -successor b , we could simply add C to $\mathcal{L}(b)$ instead of adding a new f -successor c and recording that $b \approx_a c$.

Another candidate for optimizations is the test function. Recall that this function takes a predicate conjunction c with set of variables V and non-deterministically returns a concrete equivalence, i.e., a relation \sim such that there exists a solution δ for c with $v_i \sim v_j$ iff $\delta(v_i) = \delta(v_j)$ (see Definition 4.1). It is not hard to devise an $\mathcal{ALC}(\mathcal{D})$ -concept that enforces

completion systems to have exponentially many concrete nodes by slightly adapting well-known \mathcal{ALC} -concepts that enforce models of exponential size (Halpern & Moses, 1992). Hence, the size of input conjunctions c to **test** can be exponential in the size of the input concept. Now note that, even for trivial \mathcal{D} -conjunctions

$$c = \top_{\mathcal{D}}(v_1) \wedge \cdots \wedge \top_{\mathcal{D}}(v_k)$$

we have an exponential number of distinct concrete equivalences \sim . Thus, the number of possible outcomes of a call to the **test** function may be *double exponential* in the size of the input concept. Considering the above example, a natural approach to attack this problem is to require **test** to return only *minimal* concrete equivalences: intuitively, an equivalence is minimal if only those variables are equivalent whose equality is enforced by the conjunction.

More precisely, \sim is called *minimal* if there exists no concrete equivalence \sim' such that (i) $x \sim' y$ implies $x \sim y$ and (ii) there are x, y with $x \sim y$ and $x \not\sim' y$. We conjecture that restricting **test** in this way does not destroy soundness and completeness of the tableau algorithm. However, although this definitely is a worthwhile optimization, it does not help to overcome the existence of double exponentially many outcomes of **test** in the worst case—at least not for all concrete domains \mathcal{D} : consider the concrete domain \mathbf{N} from Page 24 and conjunctions of the form

$$c_i = <_i(v_1) \wedge \cdots \wedge <_i(v_{2i}).$$

It is readily checked that, for each $i \geq 1$, the number of minimal concrete equivalences for c_i is exponential in i . Moreover, it is not hard to devise a concept C_i of size logarithmic in i that enforces completion systems S such that $\zeta_S = c_i$. Hence, there are still double exponentially many outcomes of the **test** function.

In the example just discussed, the exponential branching of **test** is clearly due to the discreteness of the natural numbers. Indeed, if we use a dense structure for defining concrete domains, it seems that the restriction to minimal concrete equivalences can have the desired effect, namely that the number of **test**'s possible outcomes becomes polynomial in the size of its input and thus exponential in the size of the input concept. For example, consider the concrete domain \mathbf{Q} , which is defined as follows:

- $\Delta_{\mathbf{Q}}$ is the set \mathbb{Q} of rational numbers;
- $\Phi_{\mathbf{Q}}$ contains unary predicates $\top_{\mathbf{Q}}$ and its negation $\perp_{\mathbf{Q}}$, unary predicates $=_q$ and \neq_q for each $q \in \mathbb{Q}$, binary comparison predicates $\{<, \leq, =, \neq, \geq, >\}$, a ternary addition predicate $+$, and its negation $\bar{+}$ (all with the obvious semantics).

It is readily checked that \mathbf{Q} is key-admissible (note that it provides a binary equality predicate) and thus falls into our framework. We conjecture that there exists only one minimal concrete equivalence for every \mathbf{Q} -predicate conjunction c : intuitively, it seems possible to (inductively) determine a relation \sim on the set of variables V used in c such that (i) $x \sim y$ implies that $\delta(x) = \delta(y)$ for *every* solution δ for c and (ii) there exists a solution δ for c such that $v \not\sim v'$ implies $\delta(v) \neq \delta(v')$. Clearly, \sim is a minimal concrete equivalence. Moreover, due to (i) it is the only one.

4.2 A Tableau Algorithm for $\mathcal{SHOQK}(\mathcal{D})$

Although $\mathcal{ALCOK}(\mathcal{D})$ is a quite powerful DL, it lacks several expressive means that can be found in most state-of-the-art description logic systems such as FaCT and RACER (Horrocks, 1998; Horrocks et al., 2000; Haarslev & Möller, 2001). In this section, we consider the very expressive description logic $\mathcal{SHOQK}(\mathcal{D})$ which provides for concrete domains, key boxes, and nominals, but also for many other means of expressivity such as transitive roles, role hierarchies, qualifying number restrictions, and general TBoxes. Modulo some details, $\mathcal{SHOQK}(\mathcal{D})$ can be viewed as the extension of the DL $\mathcal{SHOQ}(\mathcal{D})$ with key boxes. $\mathcal{SHOQK}(\mathcal{D})$ was proposed in (Horrocks & Sattler, 2001) and extended in (Pan & Horrocks, 2002) as a tool for ontology reasoning in the context of the semantic web (Berners-Lee, Hendler, & Lassila, 2001; Baader et al., 2002a).

One very important feature of $\mathcal{SHOQK}(\mathcal{D})$ are so-called TBoxes, i.e. concept equations of the form $C \doteq D$ that are used as a “background theory” in reasoning. Since it is well-known that combining general TBoxes and the concrete domain constructor easily leads to undecidability (Baader & Hanschke, 1992; Lutz, 2004), $\mathcal{SHOQK}(\mathcal{D})$ only offers a path-free variant of the concrete domain constructor—i.e. only concrete features are admitted inside this constructor rather than paths of arbitrary length. (Haarslev et al., 2001; Horrocks & Sattler, 2001) show that this restriction regains decidability. Path-freeness of the concrete domain constructor obviously renders abstract features unnecessary, and thus this syntactic type is not available in $\mathcal{SHOQK}(\mathcal{D})$.

4.2.1 THE DESCRIPTION LOGIC $\mathcal{SHOQK}(\mathcal{D})$

Let us now define $\mathcal{SHOQK}(\mathcal{D})$ in a formal way, starting with the syntax.

Definition 4.6 ($\mathcal{SHOQK}(\mathcal{D})$ Syntax). A *role axiom* is either a *role inclusion*, which is of the form $R \sqsubseteq S$ with $R, S \in \mathbf{N}_R$, or a *transitivity axiom* $\text{Trans}(R)$ where $R \in \mathbf{N}_R$. A *role box* \mathcal{R} is a finite set of role axioms. Let $\underline{*}$ be the reflexive-transitive closure of the role inclusions in \mathcal{R} . A role name R is called *simple* if $S \underline{*} R$ implies $\text{Trans}(S) \notin \mathcal{R}$ for all role names S . Let \mathcal{D} be a concrete domain. The set of $\mathcal{SHOQK}(\mathcal{D})$ -concepts is the smallest set such that

- every concept name and every nominal is a concept, and
- if C and D are concepts, R is a role name, S a simple role name, n and k are natural numbers, g_1, \dots, g_n are concrete features, and $P \in \Phi_{\mathcal{D}}$ is a predicate of arity n , then the following expressions are also concepts:

$$\neg C, C \sqcap D, C \sqcup D, \exists R.C, \forall R.C, (\geq k S C), (\leq k S C), \exists g_1, \dots, g_n.P, \text{ and } g_1 \uparrow.$$

A *concept equation* is an expression $C \doteq D$ with C and D concepts. A *TBox* is a finite set of concept equations.

For $\mathcal{SHOQK}(\mathcal{D})$, we consider key boxes that differ in two aspects from the ones we considered for $\mathcal{ALCOK}(\mathcal{D})$: in the following, we assume key boxes to be path-free, but we admit complex concepts to occur in key definitions. Note that abstract features and paths do not occur in the syntax of $\mathcal{SHOQK}(\mathcal{D})$ —as will become clear after the semantics has been defined, the former can be “simulated” by the more general number restrictions ($\leq n R C$).

As usual in description logics of the $\mathcal{SHIQ}/\mathcal{SHOQ}$ family, we require role names in number restrictions to be simple since admitting arbitrary roles yields undecidability of reasoning (Horrocks et al., 2000; Horrocks & Sattler, 2001). If the role box \mathcal{R} is clear from the context, we will usually write $\text{Trans}(R)$ instead of $\text{Trans}(R) \in \mathcal{R}$. We now introduce the semantics of $\mathcal{SHOQK}(\mathcal{D})$ and the relevant reasoning problems.

Definition 4.7 ($\mathcal{SHOQK}(\mathcal{D})$ Semantics). Interpretations $\mathcal{I} = (\Delta_{\mathcal{I}}, \cdot^{\mathcal{I}})$ are defined as in Definition 2.3, where the function $\cdot^{\mathcal{I}}$ is extended to the novel $\mathcal{SHOQK}(\mathcal{D})$ -concepts as follows:

$$\begin{aligned} (\leq k R C)^{\mathcal{I}} &:= \{d \in \Delta_{\mathcal{I}} \mid \#\{e \mid (d, e) \in R^{\mathcal{I}}\} \leq k\} \text{ and} \\ (\geq k R C)^{\mathcal{I}} &:= \{d \in \Delta_{\mathcal{I}} \mid \#\{e \mid (d, e) \in R^{\mathcal{I}}\} \geq k\}. \end{aligned}$$

Let \mathcal{I} be an interpretation. Then \mathcal{I} *satisfies* a concept equation $C \doteq D$ if $C^{\mathcal{I}} = D^{\mathcal{I}}$. \mathcal{I} is a *model* of a TBox \mathcal{T} if \mathcal{I} satisfies all concept equations in \mathcal{T} . Similarly, \mathcal{I} satisfies a role inclusion $R \sqsubseteq S$ if $R^{\mathcal{I}} \subseteq S^{\mathcal{I}}$ and a transitivity axiom $\text{Trans}(R)$ if $R^{\mathcal{I}}$ is a transitive relation. \mathcal{I} is a *model* of a role box \mathcal{R} if \mathcal{I} satisfies all role inclusions and transitivity axioms in \mathcal{R} .

Let \mathcal{T} be a TBox, \mathcal{R} a role box, and \mathcal{K} a key box. A concept C is *satisfiable w.r.t.* \mathcal{T} , \mathcal{R} , and \mathcal{K} iff C , \mathcal{T} , \mathcal{R} , and \mathcal{K} have a common model. C is *subsumed by* a concept D w.r.t. \mathcal{T} , \mathcal{R} , and \mathcal{K} (written $C \sqsubseteq_{\mathcal{T}, \mathcal{R}, \mathcal{K}} D$) iff $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$ for all common models \mathcal{I} of \mathcal{T} , \mathcal{R} , and \mathcal{K} .

Note that, due to the requirement that role names used inside number restrictions should be simple, existential and universal value restrictions are no syntactic sugar: in contrast to number restrictions, they can be used on *all* roles.

It is well-known that, in many expressive description logics, reasoning *with* TBoxes can be reduced to reasoning *without* them (Schild, 1991; Horrocks & Sattler, 2001): for $\mathcal{SHOQK}(\mathcal{D})$, a concept C is satisfiable w.r.t. \mathcal{T} , \mathcal{R} , and \mathcal{K} iff the concept

$$\exists R.C \sqcap \forall R. \left(\prod_{D \doteq E \in \mathcal{T}} D \leftrightarrow E \right) \sqcap \prod_{\substack{\text{nominal } N \text{ used} \\ \text{in } C, \mathcal{T}, \text{ or } \mathcal{K}}} (\exists R.N)$$

is satisfiable w.r.t. to \mathcal{R}' , \mathcal{K} , and the empty TBox, where R is a fresh role not appearing in C , \mathcal{R} , and \mathcal{T} , and

$$\mathcal{R}' := \mathcal{R} \cup \{\text{Trans}(R)\} \cup \bigcup_{\substack{\text{role name } S \text{ used} \\ \text{in } C, \mathcal{T}, \mathcal{R}, \text{ or } \mathcal{K}}} \{S \sqsubseteq R\}.$$

Since subsumption can be reduced to satisfiability as described in Section 2, in the following we will only consider satisfiability of concepts w.r.t. role boxes and key boxes, but without TBoxes. We will also generally assume role boxes \mathcal{R} to be *acyclic*, i.e. to satisfy the following condition: for each role name R , there are no role names R_1, \dots, R_k such that $R = R_1 = R_k$ and $R_i \sqsubseteq R_{i+1} \in \mathcal{R}$ for $1 \leq i < k$. It is not hard to see that this is no restriction since cycles can be eliminated: if R_1, \dots, R_k is a cycle in \mathcal{R} , then we have $R_1^{\mathcal{I}} = \dots = R_k^{\mathcal{I}}$ for all interpretations \mathcal{I} . Thus we can simply remove the cycle from \mathcal{R} and replace every occurrence of R_2, \dots, R_k in C , \mathcal{R} , and \mathcal{K} with R_1 . Moreover, we have to possibly add $\text{Trans}(R_1)$ if, before the cycle elimination, we had $\text{Trans}(R_i)$ for some i with $1 \leq i \leq n$.

| |
|--|
| $\neg(\geq n R C) \rightsquigarrow (\leq (n-1) R C) \text{ if } n \geq 1$ $\neg(\geq 0 R C) \rightsquigarrow \perp$ $\neg(\leq n R C) \rightsquigarrow (\geq (n+1) R C)$ |
|--|

 Figure 10: The $\mathcal{SHOQK}(\mathcal{D})$ NNF rewrite rules.

Before we turn our attention to the construction of a tableau algorithm for $\mathcal{SHOQK}(\mathcal{D})$, let us comment on a few minor differences between $\mathcal{SHOQK}(\mathcal{D})$ as introduced here and the original version of $\mathcal{SHOQ}(\mathcal{D})$ as described in (Horrocks & Sattler, 2001). The main difference is that our logic, like the extensions investigated in (Haarslev et al., 2001; Pan & Horrocks, 2002), allows n -ary predicates while Horrocks and Sattler restrict themselves to unary predicates. Moreover, $\mathcal{SHOQ}(\mathcal{D})$ as introduced in (Horrocks & Sattler, 2001) uses *concrete roles* rather than concrete features, the difference being that concrete roles are not necessarily functional. Due to this non-functionality, the original $\mathcal{SHOQ}(\mathcal{D})$ admits two variants $\exists T.P$ and $\forall T.P$ of the concrete domain constructor (where T is a concrete role and P a unary predicate). In $\mathcal{SHOQK}(\mathcal{D})$, we can simulate the universal variant by writing $\exists g.P \sqcup g\uparrow$ since concrete features g are interpreted as partial functions and, in contrast to Horrocks and Sattler, we have the undefinedness constructor $g\uparrow$ available. Except for the n -ary predicates which provide important additional expressivity, we view these deviations as minor ones since they are easily seen to not affect decidability and complexity of reasoning.

4.2.2 A TABLEAU ALGORITHM FOR $\mathcal{SHOQK}(\mathcal{D})$

The basic intuitions of the $\mathcal{SHOQK}(\mathcal{D})$ tableau algorithm are similar to the $\mathcal{ALCOK}(\mathcal{D})$ algorithm, with one exception: to deal with the various expressive means of $\mathcal{SHOQK}(\mathcal{D})$, it is convenient to introduce a certain abstraction of models, so-called *tableaux*. The main difference between tableaux and models is that, in tableaux, roles declared to be transitive must not necessarily be described by transitive relations. We show that there exists a tableau for a given concept, key box, and TBox if and only if they have a common model. The aim of the $\mathcal{SHOQK}(\mathcal{D})$ algorithm is then to construct a tableau for its input rather than trying to construct a model. To do this, the algorithm employs completion forests as its underlying data structure.

We first introduce tableaux. Let us start with discussing some preliminaries. As for $\mathcal{ALCOK}(\mathcal{D})$, we assume all concepts and key boxes to be in NNF, i.e. negation occurs only in front of concept names and nominals. We again use $\dot{\neg}C$ to denote the NNF of $\neg C$. The additional NNF rewrite rules for $\mathcal{SHOQK}(\mathcal{D})$ can be found in Figure 10 and complete those given for $\mathcal{ALCOK}(\mathcal{D})$ in Figure 7.

For a concept D , role box \mathcal{R} , and key box \mathcal{K} , we define

$$\begin{aligned} \text{cl}(D, \mathcal{K}) &:= \text{sub}(D) \cup \text{sub}(\text{con}(\mathcal{K})) \cup \{\dot{\neg}C \mid C \in \text{sub}(D) \cup \text{sub}(\text{con}(\mathcal{K}))\} \\ \text{cl}(D, \mathcal{R}, \mathcal{K}) &:= \text{cl}(D, \mathcal{K}) \cup \{\forall R.C \mid R \underline{\equiv} S \text{ and } \forall S.C \in \text{cl}(D, \mathcal{K})\}, \end{aligned}$$

Obviously, the cardinality of $\text{cl}(D, \mathcal{R}, \mathcal{K})$ is linear in the size of D , \mathcal{R} , and \mathcal{K} . In what follows, we write $\mathbf{N}_R^{D, \mathcal{R}, \mathcal{K}}$ to denote the set of role names occurring in D , \mathcal{R} , or \mathcal{K} , and $\mathbf{N}_{\text{cF}}^{D, \mathcal{K}}$ to denote the sets of concrete features occurring in D or \mathcal{K} . We are now ready to define tableaux.

Definition 4.8 (Tableau). Let D be a $\mathcal{SHOQK}(D)$ -concept in NNF, \mathcal{R} a role box, and \mathcal{K} a path-free key box in NNF. A *tableau* T for D w.r.t. \mathcal{R} and \mathcal{K} is a tuple $(\mathbf{S}_a, \mathbf{S}_c, \mathcal{L}, E, e, \mathcal{P})$ such that

- $\mathbf{S}_a, \mathbf{S}_c$ are sets of *abstract* and *concrete individuals*,
- $\mathcal{L} : \mathbf{S}_a \rightarrow 2^{\text{cl}(D, \mathcal{R}, \mathcal{K})}$ maps each abstract individual to a subset of $\text{cl}(D, \mathcal{R}, \mathcal{K})$,
- $E : \mathbf{S}_a \times \mathbf{S}_a \rightarrow 2^{\mathbb{N}_R^{D, \mathcal{R}, \mathcal{K}}}$ maps pairs of abstract individuals to sets of roles,
- $e : \mathbf{S}_a \times \mathbb{N}_{\text{cF}}^{D, \mathcal{K}} \rightarrow \mathbf{S}_c$ maps pairs of abstract individuals and concrete features to concrete individuals,
- \mathcal{P} maps each n -ary concrete predicate in $\text{cl}(D, \mathcal{R}, \mathcal{K})$ to a set of n -tuples over \mathbf{S}_c ,
- there is an abstract individual $s_0 \in \mathbf{S}_a$ such that $D \in \mathcal{L}(s_0)$, and

for all $s, t \in \mathbf{S}_a$, $C, C_1, C_2 \in \text{cl}(D, \mathcal{R}, \mathcal{K})$, $R, S \in \mathbb{N}_R^{D, \mathcal{R}, \mathcal{K}}$, and for

$$S^T(s, C) := \{t \in \mathbf{S}_a \mid S \in E(s, t) \text{ and } C \in \mathcal{L}(t)\},$$

it holds that:

- (T1) if $C \in \mathcal{L}(s)$, then $\neg C \notin \mathcal{L}(s)$,
- (T2) if $C_1 \sqcap C_2 \in \mathcal{L}(s)$, then $C_1 \in \mathcal{L}(s)$ and $C_2 \in \mathcal{L}(s)$,
- (T3) if $C_1 \sqcup C_2 \in \mathcal{L}(s)$, then $C_1 \in \mathcal{L}(s)$ or $C_2 \in \mathcal{L}(s)$,
- (T4) if $R \in E(s, t)$ and $R \sqsubseteq S$, then $S \in E(s, t)$,
- (T5) if $\forall R. C \in \mathcal{L}(s)$ and $R \in E(s, t)$, then $C \in \mathcal{L}(t)$,
- (T6) if $\exists R. C \in \mathcal{L}(s)$, then there is some $t \in \mathbf{S}_a$ such that $R \in E(s, t)$ and $C \in \mathcal{L}(t)$,
- (T7) if $\forall S. C \in \mathcal{L}(s)$ and $R \in E(s, t)$ for some $R \sqsubseteq S$ with $\text{Trans}(R)$, then $\forall R. C \in \mathcal{L}(t)$,
- (T8) if $(\geq n S C) \in \mathcal{L}(s)$, then $\#S^T(s, C) \geq n$,
- (T9) if $(\leq n S C) \in \mathcal{L}(s)$, then $\#S^T(s, C) \leq n$,
- (T10) if either $(\leq n S C) \in \mathcal{L}(s)$ and $S \in E(s, t)$ or $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$ and $e(t, g_i)$ is defined for all $1 \leq i \leq n$, then $\{C, \dot{\neg} C\} \cap \mathcal{L}(t) \neq \emptyset$,
- (T11) if $N \in \mathcal{L}(s) \cap \mathcal{L}(t)$, then $s = t$,
- (T12) if $\exists g_1, \dots, g_n. P \in \mathcal{L}(s)$, then there are $x_1, \dots, x_n \in \mathbf{S}_c$ with $e(s, g_i) = x_i$ and $(x_1, \dots, x_n) \in \mathcal{P}(P)$,
- (T13) \bigwedge_P used in $D, \mathcal{K} \bigwedge_{(x_1, \dots, x_n) \in \mathcal{P}(P)} P(x_1, \dots, x_n) \wedge \bigwedge_{x \neq y} x \neq y$ is satisfiable,
- (T14) if $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$, $C \in \mathcal{L}(s) \cap \mathcal{L}(t)$, and $e(s, g_i) = e(t, g_i)$ for all $1 \leq i \leq n$, then $s = t$,

(T15) if $g \uparrow \in \mathcal{L}(s)$, then $e(s, g)$ is undefined.

Note that the predicate conjunction in **(T13)** uses a binary inequality predicate. In general, we do not require the concrete domain \mathcal{D} to be equipped with such a predicate and thus this predicate conjunction is not necessarily a \mathcal{D} -conjunction. However, it is nevertheless “safe” to use **(T13)** in the given form since tableaux are only used in proofs and we do not need a concrete domain reasoner that is capable of deciding the satisfiability of this conjunction. The following lemma, whose proof is provided in Appendix C, shows that our definition of tableaux provides an adequate abstraction of models.

Lemma 4.2. *Let D be a $\mathcal{SHOQK}(\mathcal{D})$ -concept in NNF, \mathcal{R} a role box, and \mathcal{K} a key box in NNF. Then D is satisfiable w.r.t. \mathcal{R} and \mathcal{K} iff D has a tableau w.r.t. \mathcal{R} and \mathcal{K} .*

Given Lemma 4.2, in order to decide satisfiability of $\mathcal{SHOQK}(\mathcal{D})$ -concepts w.r.t. role and key boxes, we may use a (tableau) algorithm that tries to construct a tableau for the input. In the following, we will describe such an algorithm in detail.

As in the previous section, the algorithm works on completion systems. However, in the case of $\mathcal{SHOQK}(\mathcal{D})$ the core component of completion systems is a completion forest rather than a completion tree. The reason for this is that some completion rules remove nodes from the completion system and in this way can disconnect one tree into two subtrees.

Definition 4.9 (Completion System). Let D be a $\mathcal{SHOQK}(\mathcal{D})$ -concept in NNF, \mathcal{R} a role box, and \mathcal{K} a path-free key box in NNF. For each concept $(\geq n R C) \in \text{cl}(D, \mathcal{R}, \mathcal{K})$ and $1 \leq i \leq n$, we reserve a concept name A_i^{nRC} not appearing in $\text{cl}(D, \mathcal{R}, \mathcal{K})$ and define an *extended closure*

$$\text{cl}^+(D, \mathcal{R}, \mathcal{K}) := \text{cl}(D, \mathcal{R}, \mathcal{K}) \cup \{A_1^{nRC}, \dots, A_n^{nRC} \mid (\geq n R C) \in \text{cl}(D, \mathcal{R}, \mathcal{K})\},$$

Let O_a and O_c be disjoint and countably infinite sets of *abstract* and *concrete nodes*. A *completion forest* for D , \mathcal{R} , and \mathcal{K} is a finite forest $\mathbf{F} = (V_a, V_c, E, \mathcal{L})$ with nodes $V_a \cup V_c$ such that $V_a \subseteq O_a$, $V_c \subseteq O_c$, and all nodes from V_c are leaves. The forest is labelled as follows:

- each node $a \in V_a$ is labelled with a subset $\mathcal{L}(a)$ of $\text{cl}^+(D, \mathcal{R}, \mathcal{K})$,
- each edge $(a, b) \in E$ with $a, b \in V_a$ is labeled with a non-empty set of role names $\mathcal{L}(a, b)$ occurring in D , \mathcal{R} , or \mathcal{K} , and
- each edge $(a, x) \in E$ with $a \in V_a$ and $x \in V_c$ is labeled with a concrete feature $\mathcal{L}(a, x)$ occurring in D , \mathcal{R} , or \mathcal{K} .

A *completion system* for D , \mathcal{R} , and \mathcal{K} is a tuple $S = (\mathbf{F}, \mathcal{P}, \sim_c, \prec)$ such that

- $\mathbf{F} = (V_a, V_c, E, \mathcal{L})$ is a completion forest for D , \mathcal{R} , and \mathcal{K} ,
- \mathcal{P} maps each n -ary concrete predicate in $\text{cl}(D, \mathcal{R}, \mathcal{K})$ to a set of n -tuples in V_c ,
- \sim_c is an equivalence relation on V_c , and
- \prec is a linear ordering on V_a .

A node $t \in V_a$ is called an R -successor of a node $s \in V_a$ if, for some R' with $R' \sqsubseteq^* R$, we have $R' \in \mathcal{L}(s, t)$. A node $x \in V_c$ is called a g -successor of a node $s \in V_a$ if $\mathcal{L}(s, x) = g$. Finally, we write $s \neq t$ if s and t are R -successors of the same node and there is some $A_i^{nRC} \in \mathcal{L}(s)$ and $A_j^{nRC} \in \mathcal{L}(t)$ with $i \neq j$.

Some remarks are in order here. Firstly, in contrast to the $\mathcal{ALCOK}(\mathcal{D})$ case, the relation \prec is no longer required to respect the level of a node. This is due to the fact that (a) we have to enforce termination artificially anyway, and this property of \prec is not used to prove termination, and (b) the level of a node might change anyway since a node might become a root node because some completion rules will remove edges.

Secondly, the relation \sim_c will be returned by a \mathcal{D} -tester, and is used to compute a relation \approx_a which is then used by the tableau algorithm. However, we do not need to compute the relation \approx_c from \approx_a since, in contrast to the $\mathcal{ALCOK}(\mathcal{D})$ case, all concepts and key boxes are assumed to be path-free.

Thirdly, the new concept names A_i^{nRC} will be used to ensure that successors of a node x generated for some $(\geq n R C) \in \mathcal{L}(x)$ will not be merged later due to a concept $(\leq n' R C') \in \mathcal{L}(x)$: each generated successor is labelled with a different concept A_i^{nRC} ; since merging two nodes means unifying their node labels, it then suffices to disallow the occurrence of distinct concepts A_i^{nRC} in the same node label through a suitable definition of “clash”.

Since $\mathcal{SHOQK}(\mathcal{D})$ provides for transitive roles, we need some cycle-detection mechanism in order to guarantee termination of our algorithm: roughly spoken, if we encounter a node which is “similar” to an already existing one, then this node does not need to be further explored. Speaking in terms of (Horrocks et al., 2000; Baader & Sattler, 2000), we employ a mechanism called *subset blocking*.

Definition 4.10 (Blocked). Let \preceq be the reflexive closure of \prec . A node $t \in V_a$ is *blocked* by a node $s \in V_a$ if $\mathcal{L}(t) \subseteq \mathcal{L}(s)$, and $s \preceq s'$, for all s' with $\mathcal{L}(t) \subseteq \mathcal{L}(s')$.

Note that, unlike to what is done, e.g., in (Horrocks et al., 2000), the blocking node is not necessarily an ancestor of the blocked node, but can be anywhere in the forest. It may even be that blocked nodes have unblocked successors. This modification is used later to obtain a NEXPTIME upper bound.

To decide the satisfiability of an $\mathcal{ALCOK}(\mathcal{D})$ -concept D w.r.t. a role box \mathcal{R} and a path-free key box \mathcal{K} (where D and \mathcal{K} are in NNF), the tableau algorithm is started with the *initial completion system*

$$\begin{aligned} S_D &= (\mathbf{F}_D, \mathcal{P}_\emptyset, \emptyset, \emptyset), \text{ where} \\ \mathbf{F}_D &= (\{s_0\}, \emptyset, \emptyset, \{s_0 \mapsto \{D\}\}) \text{ and} \\ \mathcal{P}_\emptyset &\quad \text{maps each } P \in \Phi_D \text{ occurring in } D \text{ and } \mathcal{K} \text{ to } \emptyset. \end{aligned}$$

Then the algorithm repeatedly applies completion rules. Before the actual rules are given, we introduce some new notions: firstly, we define the equivalence relation \approx_a over V_a as follows: $s \approx_a t$ if one of the following conditions is satisfied:

- $N \in \mathcal{L}(s) \cap \mathcal{L}(t)$ for some nominal N or

- $(g_1 \dots, g_n \text{ keyfor } C) \in \mathcal{K}$, $C \in \mathcal{L}(s) \cap \mathcal{L}(t)$, and there are x_i, y_i such that $g_i \in E(s, x_i) \cap E(t, y_i)$ and $x_i \sim_c y_i$ for $1 \leq i \leq n$.

Intuitively, two abstract nodes related via the \approx_a relation describe the same individual in a tableau.

Secondly, we use the following abbreviations in the formulation of the rules (written in italic):

- To *remove an abstract node s and all its incoming and outgoing edges*, remove s from V_a and each (s, t) and (t, s) from E for all $t \in V_a \cup V_c$.
- *Adding a g -successor* of an abstract node s means doing nothing if there exists a g -successor $x \in V_c$ of s and, otherwise, adding $E(s, x) = g$ for some $x \in O_c$ that does not yet occur in the completion forest.
- To *update the relation \sim_c* , a \mathcal{D} -tester is asked to decide the satisfiability of the \mathcal{D} -conjunction

$$\zeta_S := \bigwedge_{\substack{P \text{ used in } \mathcal{D}, \mathcal{K} \\ (x_1, \dots, x_n) \in \mathcal{P}(P)}} P(x_1, \dots, x_n) \wedge \bigwedge_{x \sim_c y} x = y$$

and returns, in case that this conjunction is satisfiable, an “updated” concrete equivalence \sim_c as defined in Definition 4.1.

Concerning the predicate conjunction used in updates, recall that we can w.l.o.g. assume the concrete domain to contain an equality predicate as discussed after Definition 4.1.

The completion rules are given in Figure 11. We generally assume that new nodes x are introduced into the completion forest such that $y \prec x$ for all already existing nodes y . Before further describing the tableau algorithm, we comment on the completion rules. The rules $R\sqcup$, $R\leq$, $R\exists c$, and Rch are non-deterministic, i.e., their application has more than one possible outcome. For the $R\exists c$ rule, this is due to the update operation performed on \sim_c using the \mathcal{D} -tester: as discussed at the end of Section 4.1, computing a concrete equivalence for a given \mathcal{D} -conjunction may result in a high degree of non-determinism. Please note that, in contrast to $\mathcal{ALCOK}(\mathcal{D})$, we now only need to call the \mathcal{D} -tester in this rule—and not after each rule application.

Next, the $R\approx_a$ rule takes care of abstract nodes related via \approx_a . Since all the nodes from a \approx_a -equivalent class denote the same individual, we choose only one representative whose node label contains the labels of all other nodes in the class. This representative simply is the \prec -minimal node of the equivalence class and the $R\approx_a$ rule performs the appropriate copying of node labels.

The $R\leq$ rule is the only rule to remove nodes: it removes a surplus R -successor t of a node s with $(\leq n R C) \in \mathcal{L}(s)$. Since the subtree below t is not removed, t 's successors are new, additional root nodes. This behavior is the reason why we work on a completion forest.

As in the $\mathcal{ALCOK}(\mathcal{D})$ case, the tableau algorithm stops applying rules if it finds an obvious contradiction, a “clash”, or if no more completion rules are applicable.

Definition 4.11 (Clash). Let $S = (\mathbf{F}, \mathcal{P}, \sim_c, \prec)$ be a completion system for \mathcal{D} , \mathcal{R} , and \mathcal{K} and $\mathbf{F} = (V_a, V_c, E, \mathcal{L})$. Then S is said to contain a *clash* if one of the following conditions applies:

| | |
|---------------|---|
| R \sqcap | if $C_1 \sqcap C_2 \in \mathcal{L}(s)$, s is not blocked, and $\{C_1, C_2\} \not\subseteq \mathcal{L}(s)$, then $\mathcal{L}(s) := \mathcal{L}(s) \cup \{C_1, C_2\}$ |
| R \sqcup | if $C_1 \sqcup C_2 \in \mathcal{L}(s)$, s is not blocked, and $\{C_1, C_2\} \cap \mathcal{L}(s) = \emptyset$, then $\mathcal{L}(s) := \mathcal{L}(s) \cup \{C\}$ for some $C \in \{C_1, C_2\}$ |
| R \exists | if $\exists R.C \in \mathcal{L}(s)$, s is not blocked, and s has no R -successor t with $C \in \mathcal{L}(t)$ then create a new node t such that $t' \prec t$ for all $t' \in V_a$ and set $E(s, t) := \{R\}$ and $\mathcal{L}(t) := \{C\}$ |
| R \geq | if $(\geq n S C) \in \mathcal{L}(s)$, s is not blocked, and there are no n S -successors t_1, \dots, t_n of s with $C \in \mathcal{L}(t_i)$ and $t_i \neq t_j$ for $1 \leq i < j \leq n$, then create n new nodes t_1, \dots, t_n s.t. $t' \prec t_i$ for $1 \leq i \leq n$ and all $t' \in V_a$, and set $E(s, t_i) := \{S\}$ and $\mathcal{L}(t_i) := \{C, A_i^{nSC}\}$ for $1 \leq i \leq n$ |
| R \leq | if $(\leq n S C) \in \mathcal{L}(s)$, s is not blocked, s has $n + 1$ S -successors t_0, \dots, t_n with $C \in \mathcal{L}(t_i)$ for $0 \leq i \leq n$, then choose i, j such that $t_i \prec t_j$, set $\mathcal{L}(t_i) := \mathcal{L}(t_i) \cup \mathcal{L}(t_j)$, $\mathcal{L}(s, t_i) := \mathcal{L}(s, t_i) \cup \mathcal{L}(s, t_j)$, and remove t_j and all its incoming and outgoing edges |
| R $\exists c$ | if $\exists g_1, \dots, g_n.P \in \mathcal{L}(s)$, s is not blocked, and there are no g_i -successors x_i with $(x_1, \dots, x_n) \in \mathcal{P}(P)$ then add a g_i -successor of s for each $1 \leq i \leq n$, for y_i the g_i -successor of s , add (y_1, \dots, y_n) to $\mathcal{P}(P)$, and update \sim_c |
| R \forall | if $\forall R.C \in \mathcal{L}(s)$, s is not blocked, and there is an R -successor t of s with $C \notin \mathcal{L}(t)$, then $\mathcal{L}(t) := \mathcal{L}(t) \cup \{C\}$ |
| R \forall_+ | if $\forall S.C \in \mathcal{L}(s)$, s is not blocked, there is some R with $\text{Trans}(R)$ and $R \boxtimes S$, and an R -successor t of s with $\forall R.C \notin \mathcal{L}(t)$, then $\mathcal{L}(t) := \mathcal{L}(t) \cup \{\forall R.C\}$ |
| Rch | if s is an S -successor of s' and $(\leq n S C) \in \mathcal{L}(s')$ or s has g_i -successors x_i for all $1 \leq i \leq n$ and $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$ and s is not blocked and $\{C, \dot{\vdash} C\} \cap \mathcal{L}(s) = \emptyset$, then $\mathcal{L}(s) := \mathcal{L}(s) \cup \{E\}$ for some $E \in \{C, \dot{\vdash} C\}$ |
| R \approx_a | if $s \approx_a t$, $\mathcal{L}(t) \not\subseteq \mathcal{L}(s)$, $s \prec t$, and s is not blocked, then set $\mathcal{L}(s) := \mathcal{L}(s) \cup \mathcal{L}(t)$ |

 Figure 11: The completion rules for $\mathcal{SHOQK}(\mathcal{D})$.

- (C1) for some concept name $A \in \mathbf{N}_C$ and some node $s \in V_a$, $\{A, \neg A\} \subseteq \mathcal{L}(s)$;
- (C2) the \mathcal{D} -conjunction ζ_S defined above is not satisfiable;
- (C3) $s \neq s$ for some $s \in V_a$;
- (C4) for some $s \in V_a$ and $g \in \mathbf{N}_{cF}$, we have $g\uparrow \in \mathcal{L}(s)$ and s has a g -successor.

A completion system not containing a clash is called *clash-free*. The completion system is *complete* if none of the completion rules is applicable.

Due to the simplicity of the algorithm, we refrain from describing it in pseudo-code notation: the algorithm starts with the initial completion system and then repeatedly applies the completion rules, checking for clashes after each rule application. If a clash is detected, it returns *unsatisfiable*. If a complete and clash-free completion system is found, then the algorithm returns *satisfiable*. Note that, since some of the completion rules are non-deterministic, the algorithm is also non-deterministic.

Details of the proof of termination, soundness, and completeness are given in Appendix C. Unfortunately, we have to leave the complexity of the algorithm as an open problem: it is not hard to prove that it runs in double exponential time, but it is not clear whether exponential time also suffices. However, we can still use the algorithm to obtain a tight complexity bound for $\mathcal{SHOQK}(\mathcal{D})$: the following corollary is an easy by-product of the correctness proofs, for a proof see again Appendix C.

Corollary 4.3. *If a $\mathcal{SHOQK}(\mathcal{D})$ -concept D is satisfiable w.r.t. a role box \mathcal{R} and a path-free key box \mathcal{K} , then D is satisfiable w.r.t. \mathcal{R} and \mathcal{K} in a model of size at most $|\Delta_{\mathcal{I}}| \leq 2^m$ for $m = \#\text{cl}^+(D, \mathcal{R}, \mathcal{K})$.*

Thus the following is an alternative algorithm for deciding satisfiability of a $\mathcal{SHOQK}(\mathcal{D})$ -concept D w.r.t. a role box \mathcal{R} and a path-free key box \mathcal{K} : first, guess an interpretation \mathcal{I} with cardinality of $\Delta_{\mathcal{I}}$ bounded by 2^m , using placeholder variables from \mathbf{O}_c instead of concrete values in the interpretation of concrete features. Let V_c be the set of variables from \mathbf{O}_c occurring in \mathcal{I} . Additionally guess an “interpretation” \mathcal{P} for the concrete domain predicates: just as in completion forests, \mathcal{P} maps each n -ary concrete predicate used in D or \mathcal{K} to an n -ary relation on V_c . Then perform standard (polynomial-time) model checking to ensure that \mathcal{I} is a model of D . In doing this, treat concepts of the form $\exists g_1, \dots, g_n.P$ using the interpretation of predicates \mathcal{P} . It is easily checked in polynomial time that \mathcal{I} is also a model of \mathcal{R} and \mathcal{K} —for the latter, assume that all placeholder variables stand for different values. Finally, use the concrete domain \mathcal{D} -tester to check whether the conjunction

$$\bigwedge_{\substack{P \text{ used in } D, \mathcal{K} \\ (x_1, \dots, x_n) \in \mathcal{P}(P)}} P(x_1, \dots, x_n)$$

is satisfiable. Answer “yes” if it is and “no” otherwise. Since this algorithm can clearly be implemented in NEXPTIME provided that there is a \mathcal{D} -tester running in non-deterministic polynomial time, we obtain the following:

Theorem 4.4. *Let \mathcal{D} be a key-admissible concrete domain such that extended \mathcal{D} -satisfiability is in NP, then $\mathcal{SHOQK}(\mathcal{D})$ -concept satisfiability w.r.t. TBoxes, role boxes, and path-free key boxes is in NEXPTIME.*

5. Conclusion

In this paper, we have identified key constraints as an interesting extension of description logics with concrete domains. Starting from this observation, we introduced a number of natural description logics and provided a comprehensive analysis of the decidability and complexity of reasoning. The main observation of our investigations is that key boxes can have dramatic consequences on the complexity of reasoning: for example, the PSPACE-complete DL $\mathcal{ALC}(\mathcal{D})$ becomes NEXPTIME-complete if extended with path-free, unary, Boolean key boxes and undecidable if extended with path-free, unary, non-Boolean key boxes. We show that the restriction to Boolean key boxes (in the $\mathcal{ALCOK}(\mathcal{D})$ case) and to path-free key boxes (in the $\mathcal{SHOQK}(\mathcal{D})$ case) yields a decidable and NEXPTIME-complete reasoning problems.

We selected $\mathcal{ALC}(\mathcal{D})$ and $\mathcal{SHOQ}(\mathcal{D})$ as the basis for our analysis since, in our opinion, these are the most fundamental description logics with concrete domains. Going one step further, it would be interesting to combine key boxes with other extensions of concrete domains, such as the ones presented in (Lutz, 2003) and (Lutz, 2004). To name only one possibility, the extension of both $\mathcal{ALCOK}(\mathcal{D})$ and $\mathcal{SHOQ}(\mathcal{D})$ with inverse roles seems to be a natural idea. Note that inverse roles interact with several of the available means of expressivity: while \mathcal{ALC} with inverse roles is PSPACE complete (Horrocks, Sattler, & Tobies, 1999), \mathcal{ALCO} with inverse roles is EXPTIME-complete (Areces, Blackburn, & Marx, 1999) and $\mathcal{ALC}(\mathcal{D})$ with inverse roles even NEXPTIME-complete (Lutz, 2004).

Other options for future research are more closely related to the material presented in this paper. For example, is $\mathcal{SHOQK}(\mathcal{D})$ -concept satisfiability still decidable if we drop the requirement of key boxes to be path-free? Moreover, we had to leave the exact time requirements of our tableau algorithm as an open problem. If this algorithm runs in (non-deterministic) exponential time, it directly yields Theorem 4.4 rather than via a bounded model property.

References

- Areces, C., Blackburn, P., & Marx, M. (1999). A road-map on complexity for hybrid logics. In Flum, J., & Rodríguez-Artalejo, M. (Eds.), *Computer Science Logic*, No. 1683 in Lecture Notes in Computer Science, pp. 307–321. Springer-Verlag.
- Baader, F., Horrocks, I., & Sattler, U. (2002a). Description logics for the semantic web. *KI – Künstliche Intelligenz*, 3. To appear.
- Baader, F., Lutz, C., Sturm, H., & Wolter, F. (2002b). Fusions of description logics and abstract description systems. *Journal of Artificial Intelligence Research (JAIR)*, 16, 1–58.
- Baader, F., & Sattler, U. (1998). Description logics with concrete domains and aggregation. In Prade, H. (Ed.), *Proceedings of the Thirteenth European Conference on Artificial Intelligence (ECAI'98)*, pp. 336–340. John Wiley & Sons.
- Baader, F., & Hanschke, P. (1991a). A scheme for integrating concrete domains into concept languages. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 452–457, Sydney, Australia.

- Baader, F., & Hanschke, P. (1991b). A scheme for integrating concrete domains into concept languages. DFKI research report RR-91-10, German Research Center for Artificial Intelligence (DFKI).
- Baader, F., & Hanschke, P. (1992). Extensions of concept languages for a mechanical engineering application. In *Proceedings of the 16th German AI-Conference (GWAI-92)*, Vol. 671 of *Lecture Notes in Computer Science*, pp. 132–143. Springer-Verlag.
- Baader, F., McGuinness, D. L., Nardi, D., & Patel-Schneider, P. (2003). *The Description Logic Handbook: Theory, implementation and applications*. Cambridge University Press.
- Baader, F., & Sattler, U. (2000). Tableau algorithms for description logics. In Dyckhoff, R. (Ed.), *Proceedings of the International Conference on Automated Reasoning with Tableaux and Related Methods (Tableaux 2000)*, Vol. 1847 of *Lecture Notes in Artificial Intelligence*, pp. 1–18. Springer-Verlag.
- Berger, R. (1966). The undecidability of the domino problem. *Memoirs of the American Mathematical Society*, 66.
- Berners-Lee, T., Hendler, J., & Lassila, O. (2001). The semantic web. *Scientific American*, 284(5), 34–43.
- Börger, E., Grädel, E., & Gurevich, Y. (1997). *The Classical Decision Problem*. Perspectives in Mathematical Logic. Springer-Verlag.
- Borgida, A., & Patel-Schneider, P. F. (1994). A semantics and complete algorithm for subsumption in the classic description logic. *Journal of Artificial Intelligence Research*, 1, 277–308.
- Borgida, A., & Weddell, G. E. (1997). Adding uniqueness constraints to description logics (preliminary report). In Bry, F., Ramakrishnan, R., & Ramamohanarao, K. (Eds.), *Proceedings of the 5th International Conference on Deductive and Object-Oriented Databases (DOOD97)*, Vol. 1341 of *LNCS*, pp. 85–102. Springer.
- Calvanese, D., De Giacomo, G., & Lenzerini, M. (1998). On the decidability of query containment under constraints. In *Proceedings of the 17th ACM SIGACT SIGMOD SIGART Symposium on Principles of Database Systems (PODS'98)*, pp. 149–158.
- Calvanese, D., De Giacomo, G., & Lenzerini, M. (2000). Keys for free in description logics. In Baader, F., & Sattler, U. (Eds.), *Proceedings of the 2000 International Workshop in Description Logics (DL2000)*, No. 33 in CEUR-WS (<http://ceur-ws.org/>), pp. 79–88.
- Calvanese, D., Lenzerini, M., & Nardi, D. (1998). Description logics for conceptual data modeling. In Chomicki, J., & Saake, G. (Eds.), *Logics for Databases and Information Systems*, pp. 229–263. Kluwer Academic Publisher.
- Dean, M., Connolly, D., van Harmelen, F., Hendler, J., Horrocks, I., McGuinness, D. L., Patel-Schneider, P. F., & Stein, L. A. (2002). Web ontology language (OWL) reference version 1.0. W3C Working Draft.
- Fensel, D., van Harmelen, F., Horrocks, I., McGuinness, D., & Patel-Schneider, P. F. (2001). OIL: An ontology infrastructure for the semantic web. *IEEE Intelligent Systems*, 16(2), 38–45.

- Graham, R. L., Knuth, D. E., & Patashnik, O. (1990). *Concrete Mathematics*. Addison Wesley Publ. Co., Reading, Massachusetts.
- Haarslev, V., Lutz, C., & Möller, R. (1998). Foundations of spatioterminological reasoning with description logics. In Cohn, A., Schubert, L., & S.C.Shapiro (Eds.), *Proceedings of the Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98)*, pp. 112–124. Morgan Kaufman.
- Haarslev, V., & Möller, R. (2001). RACER system description. In Goré, R., Leitsch, A., & Nipkow, T. (Eds.), *Proceedings of the First International Joint Conference on Automated Reasoning (IJCAR'01)*, No. 2083 in Lecture Notes in Artificial Intelligence, pp. 701–705. Springer-Verlag.
- Haarslev, V., Möller, R., & Wessel, M. (2001). The description logic \mathcal{ALCNH}_{R+} extended with concrete domains: A practically motivated approach. In Goré, R., Leitsch, A., & Nipkow, T. (Eds.), *Proceedings of the First International Joint Conference on Automated Reasoning IJCAR'01*, No. 2083 in Lecture Notes in Artificial Intelligence, pp. 29–44. Springer-Verlag.
- Halpern, J. Y., & Moses, Y. (1992). A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(3), 319–380.
- Hollunder, B., & Baader, F. (1991). Qualifying number restrictions in concept languages. In *Proceedings of the Second International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pp. 335–346, Boston, MA, USA.
- Hopcroft, J. E., & Ullman, J. D. (1979). *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley.
- Horrocks, I., Sattler, U., & Tobies, S. (2000). Practical reasoning for very expressive description logics. *Logic Journal of the IGPL*, 8(3), 239–264.
- Horrocks, I. (1998). Using an expressive description logic: Fact or fiction?. In *Proceedings of the Sixth International Conference on the Principles of Knowledge Representation and Reasoning (KR98)*, pp. 636–647.
- Horrocks, I. (2002). Reasoning with expressive description logics: Theory and practice. In Voronkov, A. (Ed.), *Proceedings of the 18th International Conference on Automated Deduction (CADE 2002)*, No. 2392 in Lecture Notes in Artificial Intelligence, pp. 1–15. Springer.
- Horrocks, I., Patel-Schneider, P. F., & van Harmelen, F. (2002). Reviewing the design of DAML+OIL: An ontology language for the semantic web. In *Proceedings of the 18th National Conference on Artificial Intelligence (AAAI 2002)*.
- Horrocks, I., & Sattler, U. (2001). Ontology reasoning in the $\mathcal{SHOQ}(D)$ description logic. In Nebel, B. (Ed.), *Proceedings of the Seventeenth International Joint Conference on Artificial Intelligence (IJCAI'01)*, pp. 199–204. Morgan-Kaufmann.
- Horrocks, I., Sattler, U., & Tobies, S. (1999). Practical reasoning for expressive description logics. In Ganzinger, H., McAllester, D., & Voronkov, A. (Eds.), *Proceedings of the 6th International Conference on Logic for Programming and Automated Reasoning (LPAR'99)*, No. 1705 in Lecture Notes in Artificial Intelligence, pp. 161–180. Springer-Verlag.

- Kamp, G., & Wache, H. (1996). CTL - a description logic with expressive concrete domains. Tech. rep. LKI-M-96/01, Laboratory for Artificial Intelligence (LKI), University of Hamburg, Germany.
- Khizder, V. L., Toman, D., & Weddell, G. E. (2001). On decidability and complexity of description logics with uniqueness constraints. In den Bussche, J. V., & Vianu, V. (Eds.), *Proceedings of the 8th International Conference on Database Theory (ICDT2001)*, Vol. 1973 of *LNCS*, pp. 54–67. Springer.
- Knuth, D. (1968). *The Art of Computer Programming*, Vol. 1. Addison-Wesley.
- Lutz, C. (2003). Description logics with concrete domains—a survey. In *Advances in Modal Logics Volume 4*. World Scientific Publishing Co. Pte. Ltd. To appear.
- Lutz, C. (2002a). *The Complexity of Reasoning with Concrete Domains*. Ph.D. thesis, LuFG Theoretical Computer Science, RWTH Aachen, Germany.
- Lutz, C. (2002b). PSPACE reasoning with the description logic $\mathcal{ALCF}(\mathcal{D})$. *Logic Journal of the IGPL*, 10(5), 535–568.
- Lutz, C. (2002c). Reasoning about entity relationship diagrams with complex attribute dependencies. In Horrocks, I., & Tessaris, S. (Eds.), *Proceedings of the International Workshop in Description Logics 2002 (DL2002)*, No. 53 in CEUR-WS (<http://ceur-ws.org/>), pp. 185–194.
- Lutz, C. (2004). NExpTime-complete description logics with concrete domains. *ACM Transactions on Computational Logic*, To appear.
- Lutz, C., Areces, C., Horrocks, I., & Sattler, U. (2002). Keys, nominals, and concrete domains. LTCS-report 02-04, Technical University Dresden. see <http://lat.inf.tu-dresden.de/research/reports.html>.
- Pan, J. Z., & Horrocks, I. (2002). Reasoning in the $\mathcal{SHOQ}(\mathcal{D}_n)$ description logic. In Horrocks, I., & Tessaris, S. (Eds.), *Proceedings of the International Workshop in Description Logics 2002 (DL2002)*, No. 53 in CEUR-WS (<http://ceur-ws.org/>), pp. 53–62.
- Post, E. M. (1946). A variant of a recursively unsolvable problem. *Bulletin of the American Mathematical Society*, 52, 264–268.
- Schild, K. D. (1991). A correspondence theory for terminological logics: Preliminary report. In Mylopoulos, J., & Reiter, R. (Eds.), *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence (IJCAI-91)*, pp. 466–471. Morgan Kaufmann.
- Schmidt-Schauß, M., & Smolka, G. (1991). Attributive concept descriptions with complements. *Artificial Intelligence*, 48(1), 1–26.

Appendix A. Proofs of Section 3.2

We prove that D_2 -satisfiability can be decided in PTIME.

Proposition A.1. *D_2 -satisfiability is in PTIME.*

Proof. Let c be a D_2 -conjunction. We show that c is satisfiable iff none of the following conditions applies:

1. c contains a conjunct $\perp_{D_2}(x)$;
2. c contains conjuncts $\text{bit}0_k^i(x)$ and $\text{bit}1_k^i(x)$;
3. c contains conjuncts $\text{bit}n_k^i(x)$ and $\text{bit}m_\ell^j(x)$ with $k \neq \ell$;
4. c contains conjuncts $\text{bit}n_k^i(x)$ and $\overline{\text{bit}n_k^i(x)}$;
5. c contains conjuncts $\text{bit}n_k^i(x)$, $\overline{\text{bit}0_k^j(x)}$, and $\overline{\text{bit}1_k^j(x)}$.

It is easily seen that c is unsatisfiable if one of the conditions applies. Assume now that Conditions 1 to 5 do *not* apply to c and let X be the set of variables used in c . For each $x \in X$, set $t(x) = k$ if $\text{bit}n_k^i(x) \in c$ for some $n, i \in \dots$.⁴ If $\text{bit}n_k^i(x) \notin c$ for all $n, i, k \in \dots$, then set $t(x) = r$ for some r not appearing as an index \cdot_r to a predicate in c . The mapping t is well-defined since c is finite, Condition 3 does not apply, and the only predicates available are $\text{bit}n_k^i(\cdot)$, $\perp_{D_2}(\cdot)$, and $\top_{D_2}(\cdot)$. We define a solution δ for c as follows: for each $x \in X$, set $\delta(x)$ to the bit vector $v \in \text{BV}_{t(x)}$ in which the i 'th bit is 1 if $\text{bit}1_{t(x)}^i(x) \in c$ or $\overline{\text{bit}0_{t(x)}^i(x)} \in c$, and 0 otherwise. It remains to prove that δ is indeed a solution for c :

- Let $\text{bit}0_k^i(x) \in c$. Then $t(x) = k$ and thus $\delta(x) \in \text{BV}_k$. Since Condition 2 does not apply, we have $\text{bit}1_k^i(x) \notin c$. Moreover, non-applicability of Condition 4 implies $\overline{\text{bit}0_k^i(x)} \notin c$. By definition of δ , the i 'th bit of $\delta(x)$ is thus 0.
- Let $\text{bit}1_k^i(x) \in c$. Then $t(x) = k$ and $\delta(x) \in \text{BV}_k$. By definition of δ , the i 'th bit of $\delta(x)$ is 1.
- Let $\overline{\text{bit}0_k^i(x)} \in c$. If $t(x) \neq k$, then $\delta(x) \notin \text{BV}_k$. Thus $\delta(x) \in (\overline{\text{bit}0_k^i})^{D_2}$ and we are done. If $t(x) = k$, then the i 'th bit of $\delta(x)$ is 1 by definition of δ and thus again $\delta(x) \in (\overline{\text{bit}0_k^i})^{D_2}$.
- Let $\overline{\text{bit}1_k^i(x)} \in c$. If $t(x) \neq k$, then $\delta(x) \notin \text{BV}_k$ and we are done. If $t(x) = k$, then $\text{bit}n_k^j(x) \in c$ for some $n, j \in \dots$. Since Condition 5 does not apply, we thus have $\overline{\text{bit}0_k^i(x)} \notin c$. Moreover, non-applicability of Condition 4 yields $\text{bit}1_k^i(x) \notin c$. Thus, by definition of δ , the i 'th bit of $\delta(x)$ is 0.

It is obvious that the listed properties can be checked in polynomial time. □

4. We use “ $P(x) \in c$ ” as an abbreviation for “ $P(x)$ is a conjunct in c ”.

Appendix B. Proofs of Section 4.1

We prove termination, soundness, and completeness of the $\mathcal{ALCOK}(\mathcal{D})$ tableau algorithm presented in Section 4.1, starting with termination. We start with establishing a few notions and technical lemmas.

Let C be a concept and \mathcal{K} a key box. We use $|C|$ to denote the length of C , i.e. the number of symbols used to write it down, and $|\mathcal{K}|$ to denote $\sum_{(u_1, \dots, u_k \text{ key for } C) \in \mathcal{K}} |C|$. For a path $u = f_1 \cdots f_k g$, we use $|u|$ to denote $k + 1$. The *role depth* of concepts is defined inductively as follows:

$$\begin{aligned} \text{rd}(A) &= \text{rd}(N) = \text{rd}(g\uparrow) = 0 \\ \text{rd}(\exists u_1, \dots, u_n.P) &= \max\{|u_i| \mid 1 \leq i \leq n\} - 1 \\ \text{rd}(C \sqcap D) &= \text{rd}(C \sqcup D) = \max\{\text{rd}(C), \text{rd}(D)\} \\ \text{rd}(\exists R.C) &= \text{rd}(\forall R.C) = \text{rd}(C) + 1 \end{aligned}$$

The following series of lemmas will eventually allow us to prove termination.

Lemma B.1. *There is a constant k such that, if the tableau algorithm is started on input C_0, \mathcal{K} and $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$ is a completion tree constructed during the run of the algorithm, then $\#V_a \leq 2^{|C_0|^k}$ and $\#V_c \leq 2^{|C_0|^k}$.*

Proof. Using induction on the number of rule applications and a case distinction according to the applied rule, it is straightforward to show that

$$C \in \mathcal{L}(a) \text{ implies } \text{rd}(C) \leq |C_0| - \text{lev}_{\mathbf{T}}(a) \quad (*)$$

for all constructed completion trees \mathbf{T} . We omit the details but note that, (1) for treating the Rch rule, one needs to employ the fact that \mathcal{K} is Boolean and thus only adds concepts of role depth 0 to node labels, and (2) for treating the Rp rule, we use that $a \prec b$ implies $\text{lev}_{\mathbf{T}}(a) \leq \text{lev}_{\mathbf{T}}(b)$.

This implies an upper bound on the depth of constructed completion trees: first, only the $R\exists$ and $R\exists c$ rules generate new nodes, and an application of either rule to a node $a \in V_a$ implies $\mathcal{L}(a) \neq \emptyset$ and thus $\text{lev}_{\mathbf{T}}(a) \leq |C_0|$ by (*). Second, each new (abstract or concrete) node b generated by an application of these rules to a node $a \in V_a$ clearly satisfies $\text{lev}_{\mathbf{T}}(b) \leq \text{lev}_{\mathbf{T}}(a) + \max(1, \text{mpl}(C_0))$, where $\text{mpl}(C_0)$ denotes the maximum length of paths in C_0 (note that concepts in \mathcal{K} may not contain any paths since it is Boolean). Since $\text{mpl}(C_0) \leq |C_0|$, the above observations imply that the depth of constructed completion trees is bounded by $2 \cdot |C_0|$.

Now for the out-degree. If a node a is generated, then this is due to the application of a rule $R\exists$ or $R\exists c$ and, initially, a has at most one successor. Let us analyze the number of successors generated by later applications of the rules $R\exists$ and $R\exists c$ to a : these rules can be applied at most once for each concept of the form $\exists R.C$ or $\exists u_1, \dots, u_n.P$ in $\mathcal{L}(a)$. By definition of $\text{cl}(C_0, \mathcal{K})$ and since \mathcal{K} is Boolean, the number of such concepts per node label is bounded by $\#\text{sub}(C_0) \leq |C_0|$. Moreover, each rule application creates at most $|C_0|$ successors. Hence, the out-degree of constructed completion trees is bounded by $|C_0|^2 + 1$. \square

Lemma B.2. *There is a constant k such that, if the tableau algorithm is started with C_0, \mathcal{K} , then, in every recursion step, the while loop terminates after at most $2^{|C_0|^k}$ steps.*

Proof. Fix an argument $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ with $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$ passed to the `sat` function, let \sim_1, \sim_2, \dots be the sequence of concrete equivalences computed in the while loop, and let $\approx_c^1, \approx_c^2, \dots$ be the corresponding \approx_c relations. Since `test`(ζ_S) calls a \mathcal{D} -tester, each of these calls indeed terminates.

We show that

$$\sim_1 \subsetneq \sim_2 \subsetneq \dots, \quad (*)$$

which implies Lemma B.2: by Lemma B.1, there exists a constant k such that $\#V_c \leq 2^{|C_0|^k}$. Hence, we have $\#\sim_i \leq 2^{2 \cdot |C_0|^k}$ which, together with (*), implies that the number of steps performed by the while loop is also bounded by $2^{2 \cdot |C_0|^k}$.

Now for the proof of (*). If the while loop reaches the i -th step, then we had $\sim_{i-1} \neq \approx_c^{i-1}$ after step $i - 1$. Since $\sim_{i-1} \subseteq \approx_c^{i-1}$ by definition, this implies $\sim_{i-1} \subsetneq \approx_c^{i-1}$. By definition of ζ_S , it is easy to see that $\approx_c^{i-1} \subseteq \sim_i$ for $i \geq 0$. Hence $\sim_{i-1} \subsetneq \sim_i$. \square

Lemma B.3. *There is a constant k such that, if the tableau algorithm is started with C_0, \mathcal{K} , then the number of recursion calls is bounded by $2^{(|C_0|+|\mathcal{K}|)^k}$.*

Proof. It obviously suffices to establish an appropriate upper bound on the number of rule applications. The $R\sqcap$, $R\sqcup$, $R\exists$, and $R\exists c$ rules can be applied at most once for each concept in a node label. By Lemma B.1, the number of nodes is at most exponential in $|C_0| + |\mathcal{K}|$. Since neither nodes nor concepts in node labels are ever deleted, the fact that node labels are subsets of $\text{cl}(C_0, \mathcal{K})$ thus implies that the number of applications of these rules is at most exponential in $|C_0| + |\mathcal{K}|$. The same holds for the rules $R\forall$ and Rp , which can be applied at most once for every concept $C \in \text{cl}(C_0, \mathcal{K})$ and every pair of (abstract) nodes. Finally, the number of Rch applications is at most exponential in $|C_0| + |\mathcal{K}|$ since this rule can be applied at most once for every abstract node and every key definition in \mathcal{K} . \square

Termination is now an obvious consequence of Lemmas B.2 and B.3.

Corollary B.4 (Termination). *The tableau algorithm terminates on any input.*

Let us now prove soundness of the algorithm.

Lemma B.5 (Soundness). *If the tableau algorithm returns `satisfiable`, then the input concept C_0 is satisfiable w.r.t. the input key box \mathcal{K} .*

Proof. If the tableau algorithm returns `satisfiable`, then there exists a complete and clash-free completion system $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ for C_0 . Let $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$. By definition of the tableau algorithm, there is a completion system $S' = (\mathbf{T}, \mathcal{P}, \prec, \sim')$ such that a call to `test`($\zeta_{S'}$) returned \sim . Moreover, we have $\sim = \approx_c$ in S . Thus, there exists a solution δ for $\zeta_{S'}$ such that

$$\delta(x) = \delta(y) \text{ iff } x \approx_c y. \quad (\dagger)$$

Clearly, δ is also a solution for ζ_S : since the second component \mathcal{P} is the same in S and S' , δ is a solution for the first part $\bigwedge_{P \text{ used in } C} \bigwedge_{(x_1, \dots, x_n) \in \mathcal{P}(P)} P(x_1, \dots, x_n)$ of ζ_S . Moreover, for

each conjunct $=(x, y)$ from the second part of ζ_S , we have $x \approx_c y$ by definition of ζ_S and thus $\delta(x) = \delta(y)$ by (\dagger) .

We now use S and δ to construct an interpretation \mathcal{I} by setting

$$\begin{aligned} \Delta_{\mathcal{I}} &= \{a \in V_a \mid \text{there is no } b \in V_a \text{ such that } a \approx_a b \text{ and } b \prec a\} \cup \{w\} \\ A^{\mathcal{I}} &= \{a \in \Delta_{\mathcal{I}} \mid A \in \mathcal{L}(a)\} \\ N^{\mathcal{I}} &= \begin{cases} \{a \in \Delta_{\mathcal{I}} \mid N \in \mathcal{L}(a)\} & \text{if there is an } a \in \Delta_{\mathcal{I}} \text{ such that } N \in \mathcal{L}(a) \\ \{w\} & \text{otherwise} \end{cases} \\ R^{\mathcal{I}} &= \{(a, b) \in \Delta_{\mathcal{I}} \times \Delta_{\mathcal{I}} \mid \text{there are } a', b' \in V_a \text{ such that } a \approx_a a', b \approx_a b', \text{ and} \\ &\quad b' \text{ is an } R\text{-successor of } a'\} \\ g^{\mathcal{I}} &= \{(a, \delta(x)) \in \Delta_{\mathcal{I}} \times \Delta_{\mathcal{D}} \mid x \text{ is a } g/\approx_a\text{-neighbor of } a\} \end{aligned}$$

for all $A \in \mathbf{N}_{\mathbf{C}}$, $N \in \mathbf{N}_{\mathbf{O}}$, $R \in \mathbf{N}_{\mathbf{R}}$, and $g \in \mathbf{N}_{\mathbf{CF}}$. We first show that \mathcal{I} is well-defined:

- $N^{\mathcal{I}}$ is a singleton for each $N \in \mathbf{N}_{\mathbf{O}}$. For assume that there exist $a, b \in \Delta_{\mathcal{I}}$ such that $a \neq b$ and $N \in \mathcal{L}(a) \cap \mathcal{L}(b)$. By definition of \approx_a (Definition 4.3), $N \in \mathcal{L}(a) \cap \mathcal{L}(b)$ implies $a \approx_a b$. This, together with $a, b \in \Delta_{\mathcal{I}}$, yields $a \prec b$ and $b \prec a$, contradicting \prec being a linear ordering.
- $f^{\mathcal{I}}$ is functional for each $f \in \mathbf{N}_{\mathbf{aF}}$. For assume that there exist $a, b, c \in \Delta_{\mathcal{I}}$ such that $\{(a, b), (a, c)\} \subseteq f^{\mathcal{I}}$ and $b \neq c$. Then there exist $a_1, a_2, b', c' \in V_a$ such that $a \approx_a a_1 \approx_a a_2$, $b \approx_a b'$, $c \approx_a c'$, b' is an f -successor of a_1 , and c' is an f -successor of a_2 . By definition of \approx_a , we thus have $b' \approx_a c'$ implying $b \approx_a c$. Since $b, c \in \Delta_{\mathcal{I}}$, this yields $b \prec c$ and $c \prec b$, a contradiction.
- $g^{\mathcal{I}}$ is functional for each $g \in \mathbf{N}_{\mathbf{CF}}$. For assume that there exist an $a \in \Delta_{\mathcal{I}}$ and $x, y \in V_c$ such that $\{(a, \delta(x)), (a, \delta(y))\} \subseteq g^{\mathcal{I}}$ and $\delta(x) \neq \delta(y)$. Then x and y are both g/\approx_a -neighbors of a . By definition of \approx_c , we thus have $x \approx_c y$ implying $\delta(x) = \delta(y)$ by (\dagger) , a contradiction.

Now we show the following claim. In the proof, we use the notion of $f_1 \cdots f_k/\approx_a$ -neighbors (with f_1, \dots, f_k abstract features), which is defined analogous to u/\approx_a -neighbors for paths u .

Claim 1: For all $a \in \Delta_{\mathcal{I}}$ and all paths u , we have $u^{\mathcal{I}}(a) = \alpha$ iff there is a u_i/\approx_a -neighbor x of a with $\delta(x) = \alpha$.

Proof: Let $u = f_1 \cdots f_k g$. Using induction on i it is easily proved that, for all i and all $b \in \Delta_{\mathcal{I}}$, we have $f_i^{\mathcal{I}}(\cdots (f_1^{\mathcal{I}}(a)) \cdots) = b$ iff there is a $f_1 \cdots f_i/\approx_a$ -neighbor b' of a with $b \approx_a b'$. Thus we have in particular that $f_k^{\mathcal{I}}(\cdots (f_1^{\mathcal{I}}(a)) \cdots) = b$ iff there is a $f_1 \cdots f_k/\approx_a$ -neighbor b' of a with $b \approx_a b'$. To prove the claim, it hence remains to use the definition of $g^{\mathcal{I}}$ together with (\dagger) .

The following claim is central for showing that \mathcal{I} is a model for C_0 and \mathcal{K} .

Claim 2: For all $a \in \Delta_{\mathcal{I}}$ and $C \in \text{cl}(C_0, \mathcal{K})$, if $C \in \mathcal{L}(a)$, then $a \in C^{\mathcal{I}}$.

Since C_0 is in the label of the root node, Claim 2 clearly implies that \mathcal{I} is a model for C_0 . Moreover, we can use it to prove that \mathcal{I} satisfies all key definitions (u_1, \dots, u_n keyfor C) in \mathcal{K} : fix $a, b \in C^{\mathcal{I}}$ such that $u_i^{\mathcal{I}}(a) = u_i^{\mathcal{I}}(b)$ for $1 \leq i \leq n$. Non-applicability of Rch yields

$\{C, \dot{\vdash} C\} \cap \mathcal{L}(a) \neq \emptyset$. If $\dot{\vdash} C \in \mathcal{L}(a)$, then Claim 2 implies $a \in (\dot{\vdash} C)^{\mathcal{I}}$ in contradiction to $a \in C^{\mathcal{I}}$. Thus we obtain $C \in \mathcal{L}(a)$. In an analogous way, we can argue that $C \in \mathcal{L}(b)$. Since $u_i^{\mathcal{I}}(a)$ and $u_i^{\mathcal{I}}(b)$ are defined for $1 \leq i \leq n$, Claim 1 yields that a has an u_i/\approx_a -neighbor x_i with $\delta(x_i) = u_i^{\mathcal{I}}(a)$ and b an u_i/\approx_a -neighbor y_i with $\delta(y_i) = u_i^{\mathcal{I}}(b)$ for $1 \leq i \leq n$. Thus the fact that $u_i^{\mathcal{I}}(a) = u_i^{\mathcal{I}}(b)$ yields $\delta(x_i) = \delta(y_i)$ for $1 \leq i \leq n$. By (\dagger) we obtain $x_i \approx_c y_i$ and thus $x_i \sim y_i$ for $1 \leq i \leq n$. By definition of \approx_a , we thus get $a \approx_a b$. Since $a, b \in \Delta_{\mathcal{I}}$, we obtain $a \not\sim b$ and $b \not\sim a$ by definition of $\Delta_{\mathcal{I}}$ and thus $a = b$.

It remains to prove Claim 2, which we do using structural induction:

- C is a concept name or a nominal. Easy by construction of \mathcal{I} .
- $C = \neg D$. Since $C \in \text{cl}(C_0, \mathcal{K})$, C is in NNF and D is a concept name. Since S is clash-free, $C \in \mathcal{L}(a)$ implies $D \notin \mathcal{L}(a)$. Thus, $a \notin D^{\mathcal{I}}$ by construction of \mathcal{I} , which yields $a \in (\neg D)^{\mathcal{I}}$.
- $C = \exists u_1, \dots, u_n. P$. Since the $R\exists c$ rule is not applicable, there exist $x_1, \dots, x_n \in V_c$ such that x_i is a u_i/\approx_a -neighbor of a for $1 \leq i \leq n$ and $(x_1, \dots, x_n) \in \mathcal{P}(P)$. Claim 1 yields $u_i^{\mathcal{I}}(a) = \delta(x_i)$ for $1 \leq i \leq n$. Since $(x_1, \dots, x_n) \in \mathcal{P}(P)$ and δ is a solution for ζ_S , we have $(\delta(x_1), \dots, \delta(x_n)) \in P^{\mathcal{D}}$ and thus $a \in C^{\mathcal{I}}$.
- $C = g\uparrow$. Since S is clash-free, there exists no $x \in V_c$ such that x is g/\approx_a -neighbor of a . Thus by Claim 1 there is no α such that $(a, \alpha) \in g^{\mathcal{I}}$.
- $C = D \sqcap E$ or $C = D \sqcup E$. Straightforward using completeness and the induction hypothesis.
- $C = \exists R.D$. Since the $R\exists$ rule is not applicable, a has an R/\approx_a -neighbor b such that $D \in \mathcal{L}(b)$. Let b' be minimal w.r.t. \prec such that $b \approx_a b'$. By definition of \mathcal{I} , we have $(a, b') \in R^{\mathcal{I}}$. Non-applicability of the $R\text{p}$ rule yields $D \in \mathcal{L}(b')$. By induction, we get $b' \in D^{\mathcal{I}}$ and thus $a \in C^{\mathcal{I}}$.
- $C = \forall R.D$. Let $(a, b) \in R^{\mathcal{I}}$. By definition of \mathcal{I} , this implies that there exist $a', b' \in V_a$ such that a is minimal w.r.t. \prec and $a \approx_a a'$, b is minimal w.r.t. \prec and $b \approx_a b'$, and b' is an R -successor of a' . Since b' is clearly an R/\approx_a -neighbor of a , non-applicability of $R\forall$ yields $D \in \mathcal{L}(b')$, which implies $D \in \mathcal{L}(b)$ due to the non-applicability of $R\text{p}$. By induction, we get $b \in D^{\mathcal{I}}$. Since this holds independently of the choice of b , we obtain $a \in (\forall R.D)^{\mathcal{I}}$.

□

Lemma B.6 (Completeness). *If the input concept C_0 is satisfiable w.r.t. the input key box \mathcal{K} , then the tableau algorithm returns satisfiable.*

Proof. Let \mathcal{I} be a model of C_0 and \mathcal{K} . We use \mathcal{I} to “guide” the (non-deterministic parts of) the algorithm such that it constructs a complete and clash-free completion system. A completion system $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ with $\mathbf{T} = (V_a, V_c, E, \mathcal{L})$ is called \mathcal{I} -compatible if there exist mappings $\pi : V_a \rightarrow \Delta_{\mathcal{I}}$ and $\tau : V_c \rightarrow \Delta_{\mathcal{D}}$ such that

$$(Ca) \quad C \in \mathcal{L}(a) \Rightarrow \pi(a) \in C^{\mathcal{I}}$$

- (Cb) b is an R -successor of $a \Rightarrow (\pi(a), \pi(b)) \in R^{\mathcal{I}}$
- (Cc) x is a g -successor of $a \Rightarrow g^{\mathcal{I}}(\pi(a)) = \tau(x)$
- (Cd) $(x_1, \dots, x_n) \in \mathcal{P}(P) \Rightarrow (\tau(x_1), \dots, \tau(x_n)) \in P^{\mathcal{D}}$
- (Ce) $x \sim y \Rightarrow \tau(x) = \tau(y)$

We first establish the following claim:

Claim 1: If a completion system S is \mathcal{I} -compatible, then (i) $a \approx_a b$ implies $\pi(a) = \pi(b)$ and (ii) $x \approx_c y$ implies $\tau(x) = \tau(y)$.

Proof: We show by induction on i that $a \approx_a^i b$ implies $\pi(a) = \pi(b)$ (see Definition 4.3), which yields (i).

- Start. If $a \approx_a^0 b$, then there exists a nominal N such that $N \in \mathcal{L}(a) \cap \mathcal{L}(b)$. By (Ca) we obtain $\pi(a) \in N^{\mathcal{I}}$ and $\pi(b) \in N^{\mathcal{I}}$, which yields $\pi(a) = \pi(b)$ by definition of the semantics.
- Step. For $a \approx_a^i b$, we distinguish three cases:
 1. If $a \approx_a^{i-1} b$, then $\pi(a) = \pi(b)$ by induction.
 2. There is a $c \in V_a$ and an $f \in \mathbf{N}_{aF}$ such that both a and b are f/\approx_a^{i-1} -neighbors of c . Hence, there exist $c_1, c_2 \in V_a$ such that $c \approx_a^{i-1} c_1 \approx_a^{i-1} c_2$, a is an f -successor of c_1 , and b is an f -successor of c_2 . By induction, we have $\pi(c) = \pi(c_1) = \pi(c_2)$. Thus (Cb) yields $\{(\pi(c), \pi(a)), (\pi(c), \pi(b))\} \subseteq f^{\mathcal{I}}$, which implies $\pi(a) = \pi(b)$ by definition of the semantics.
 3. There exist $(u_1, \dots, u_n \text{ keyfor } C) \in \mathcal{K}$, u_i/\approx_a^{i-1} -neighbors x_i of a and u_i/\approx_a^{i-1} -neighbors y_i of b for $1 \leq i \leq n$ such that $C \in \mathcal{L}(a) \cap \mathcal{L}(b)$ and $x_i \sim y_i$ for $1 \leq i \leq n$. (Ca) yields $a, b \in C^{\mathcal{I}}$. Using induction, (Cb), and (Cc), it is straightforward to show that $u_i^{\mathcal{I}}(\pi(a)) = \tau(x_i)$ and $u_i^{\mathcal{I}}(\pi(b)) = \tau(y_i)$ for $1 \leq i \leq n$. By (Ce), this implies $u_i^{\mathcal{I}}(\pi(a)) = u_i^{\mathcal{I}}(\pi(b))$ for $1 \leq i \leq k$. Since \mathcal{I} is a model of the key box \mathcal{K} , this yields $\pi(a) = \pi(b)$ by definition of the semantics.

Now for Part (ii) of Claim 1. If $x \approx_c y$, then either $x \sim y$ or there is an $a \in V_a$ and a $g \in \mathbf{N}_{cF}$ such that both x and y are g/\approx_a -neighbors of a . In the former case, (Ce) yields $\tau(x) = \tau(y)$. In the latter case, Part (i) of the claim and (Cc) yields $\{(\pi(a), \tau(x)), (\pi(a), \tau(y))\} \subseteq g^{\mathcal{I}}$ which implies $\tau(x) = \tau(y)$. This finishes the proof of Claim 1.

We can now show that the completion rules can be applied such that \mathcal{I} -compatibility is preserved.

Claim 2: If a completion system S is \mathcal{I} -compatible and a rule R is applicable to S , then R can be applied such that an \mathcal{I} -compatible completion system S' is obtained.

Proof: Let S be an \mathcal{I} -compatible completion system, let π and τ be functions satisfying (Ca) to (Ce), and let R be a completion rule applicable to S . We make a case distinction according to the type of R .

- R \sqcap The rule is applied to a concept $C_1 \sqcap C_2 \in \mathcal{L}(a)$. By (Ca), $C_1 \sqcap C_2 \in \mathcal{L}(a)$ implies $\pi(a) \in (C_1 \sqcap C_2)^{\mathcal{I}}$ and hence $\pi(a) \in C_1^{\mathcal{I}}$ and $\pi(a) \in C_2^{\mathcal{I}}$. Since the rule adds C_1 and C_2 to $\mathcal{L}(a)$, it yields a completion system that is \mathcal{I} -compatible via the same π and τ .
- R \sqcup The rule is applied to $C_1 \sqcup C_2 \in \mathcal{L}(a)$. $C_1 \sqcup C_2 \in \mathcal{L}(a)$ implies $\pi(a) \in C_1^{\mathcal{I}}$ or $\pi(a) \in C_2^{\mathcal{I}}$. Since the rule adds either C_1 or C_2 to $\mathcal{L}(a)$, it can be applied such that it yields a completion system that is \mathcal{I} -compatible via the same π and τ .
- R \exists The rule is applied to a concept $\exists R.C \in \mathcal{L}(a)$, generates a new R -successor b of a and sets $\mathcal{L}(b) = \{C\}$. By (Ca), we have $\pi(a) \in (\exists R.C)^{\mathcal{I}}$ and, hence, there exists a $d \in \Delta_{\mathcal{I}}$ such that $(\pi(a), d) \in R^{\mathcal{I}}$ and $d \in C^{\mathcal{I}}$. Set $\pi' := \pi \cup \{b \mapsto d\}$. It is readily checked that the resulting completion system is \mathcal{I} -compatible via π' and τ .
- R \forall The rule is applied to a concept $\forall R.C \in \mathcal{L}(a)$ and adds C to $\mathcal{L}(b)$ of an existing R/\approx_a -neighbor b of a . Hence, there exists an a' such that $a \approx_a a'$ and b is an R -successor of a' . By Part (i) of Claim 1, we have $\pi(a) = \pi(a')$. Thus, by (Ca) we have $\pi(a') \in (\forall R.C)^{\mathcal{I}}$ while (Cb) yields $(\pi(a'), \pi(b)) \in R^{\mathcal{I}}$. By definition of the semantics, $\pi(b) \in C^{\mathcal{I}}$ and thus the resulting completion system is \mathcal{I} -compatible via π and τ .
- R $\exists c$ The rule is applied to a concept $\exists u_1, \dots, u_n.P \in \mathcal{L}(a)$ with $u_i = f_1^{(i)} \dots f_{k_i}^{(i)} g_i$ for $1 \leq i \leq n$. The rule application generates new abstract nodes $b_j^{(i)}$ and x_j for $1 \leq i \leq n$ and $1 \leq j \leq k_i$ such that
- $b_1^{(i)}$ is an $f_1^{(i)}$ -successor of a for $1 \leq i \leq n$,
 - $b_j^{(i)}$ is an $f_j^{(i)}$ -successor of $b_{j-1}^{(i)}$ for $1 \leq i \leq n$ and $2 < j \leq k_i$,
 - x_i is g_i -successor of $b_{k_i}^{(i)}$ for $1 \leq i \leq n$, and
 - $(x_1, \dots, x_n) \in \mathcal{P}(P)$.
- By (Ca), we have $\pi(a) \in (\exists u_1, \dots, u_n.P)^{\mathcal{I}}$. Hence, there exist $d_j^{(i)} \in \Delta_{\mathcal{I}}$ for $1 \leq i \leq n$ and $1 \leq j \leq k_i$ and $\alpha_1, \dots, \alpha_n \in \Delta_{\mathcal{D}}$ such that
- $(\pi(a), d_1^{(i)}) \in (f_1^{(i)})^{\mathcal{I}}$ for $1 \leq i \leq n$,
 - $(d_{j-1}^{(i)}, d_j^{(i)}) \in (f_j^{(i)})^{\mathcal{I}}$ for $1 \leq i \leq n$ and $2 < j \leq k_i$,
 - $g_i^{\mathcal{I}}(d_{k_i}^{(i)}) = \alpha_i$ for $1 \leq i \leq n$, and
 - $(\alpha_1, \dots, \alpha_n) \in P^{\mathcal{D}}$.
- Set $\pi' := \pi \cup \bigcup_{1 \leq i \leq n \text{ and } 1 \leq j \leq k_i} \{b_j^{(i)} \mapsto d_j^{(i)}\}$ and $\tau' := \tau \cup \bigcup_{1 \leq i \leq n} \{x_i \mapsto \alpha_i\}$. The resulting completion system is \mathcal{I} -compatible via π' and τ' .
- Rch The rule is applied to an abstract node a and a key definition $(u_1, \dots, u_n \text{ keyfor } C) \in \mathcal{K}$ and non-deterministically adds either C or $\dot{\neg} C$. By definition of the semantics, $\pi(a) \in C^{\mathcal{I}}$ or $\pi(a) \in (\dot{\neg} C)^{\mathcal{I}}$. Hence, Rch can be applied such that the resulting completion system is \mathcal{I} -compatible via π and τ .

Rp The rule is applied to a concept $C \in \mathcal{L}(a)$ and adds C to the label $\mathcal{L}(b)$ of a node b with $a \approx_a b$. By **(Ca)**, we have $\pi(a) \in C^{\mathcal{I}}$. Since Claim 1 yields $\pi(a) = \pi(b)$, it follows that the resulting completion system is \mathcal{I} -compatible via π and τ .

Finally, we show that \mathcal{I} -compatibility implies clash-freeness.

Claim 3: Every \mathcal{I} -compatible completion system is clash-free.

Proof: Let $S = (\mathbf{T}, \mathcal{P}, \prec, \sim)$ be an \mathcal{I} -compatible completion system. Consider the three kinds of a clash:

- Due to **(Ca)**, a clash of the form $\{A, \neg A\} \in \mathcal{L}(a)$ clearly contradicts the semantics.
- Assume that there are $a \in V_a$ and $x \in V_c$ such that $g\uparrow \in \mathcal{L}(a)$ and x is g/\approx_a -neighbor of a . Then there exists a $b \in V_a$ such that $a \approx_a b$ and x is g -successor of b . By Claim 1, $a \approx_a b$ implies $\pi(a) = \pi(b)$. Thus, $g\uparrow \in \mathcal{L}(a)$ and **(Ca)** give $\pi(b) \in (g\uparrow)^{\mathcal{I}}$. We obtain a contradiction since **(Cc)** yields $(\pi(b), \tau(x)) \in g^{\mathcal{I}}$.
- Properties **(Cd)** and **(Ce)** and Part (ii) of Claim 1 imply that τ is a solution for ζ_S . Thus, S is concrete domain satisfiable.

We can now describe the “guidance” of the tableau algorithm by the model \mathcal{I} in detail: we ensure that, at all times, the considered completion systems are \mathcal{I} -compatible. This does obviously hold for the initial completion system

$$S_{C_0} = (\mathbf{T}_{C_0}, \mathcal{P}_\emptyset, \emptyset, \emptyset) \text{ with } \mathbf{T}_{C_0} = (\{a_0\}, \emptyset, \emptyset, \{a_0 \mapsto \{C\}\}).$$

We guide the non-deterministic test function such that, when given a predicate conjunction ζ_S with set of variables $V_c \subseteq \mathbf{O}_c$ as input, it returns the relation \sim defined by setting $x \sim y$ iff $\tau(x) = \tau(y)$ for all $x, y \in V$. The relation \sim is a concrete equivalence since τ is a solution for ζ_S (see above). With this guidance, **(Ce)** is obviously satisfied after each call to **test**, and the other properties are not affected by such a call. According to Claim 2, we can apply the completion rules such that \mathcal{I} -compatibility is preserved. By Corollary B.4, the algorithm terminates always, hence also when guided in this way. Since, by Claim 3, we will not find a clash, the algorithm returns **satisfiable**. \square

The tableau algorithm yields decidability and a tight upper complexity bound for $\mathcal{ALCOK}(\mathcal{D})$ -concept satisfiability w.r.t. key boxes.

Theorem B.7 (Theorem 4.1 of Section 4.1). *Let \mathcal{D} be a key-admissible concrete domain. If extended \mathcal{D} -satisfiability is in NP, then $\mathcal{ALCOK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean key boxes is in NEXPTIME.*

Proof. Corollary B.4 and Lemmas B.5 and B.6 yield decidability of $\mathcal{ALCOK}(\mathcal{D})$ -concept satisfiability w.r.t. Boolean key boxes. For complexity, Lemma B.3 provides an exponential bound on the number of recursion calls. Hence, it remains to show that each single recursion step needs at most exponential time. By Lemma B.2, the while loop terminates after at most exponentially many steps. In each such step, we compute the relations \approx_a and \approx_c , which are used in the construction of the predicate conjunction ζ_S and for checking termination of the while loop. Since, by Lemma B.1, there exists an exponential bound on the number

of abstract and concrete nodes in the completion system S , this can obviously be done in exponential time. Moreover, Lemma B.1 implies that the size of ζ_S is at most exponential. This together with the fact that extended \mathcal{D} -satisfiability is in NP implies that the call to `test` needs at most exponential time. All remaining tasks (checking for clashes, completeness, and rule applicability) can clearly also be performed in exponential time. \square

Appendix C. Proofs of Section 4.2

We first provide the proof of Lemma 4.2 which shows that the notion of tableaux introduced in Section 4.2 is an adequate abstraction of models.

Lemma C.1 (Lemma 4.2 in Section 4.2). *Let D be a $\text{SHOQK}(\mathcal{D})$ -concept in NNF, \mathcal{R} a role box, and \mathcal{K} a path-free key box in NNF. Then D is satisfiable w.r.t. \mathcal{R} and \mathcal{K} iff D has a tableau w.r.t. \mathcal{R} and \mathcal{K} .*

Proof. For the “only-if” direction, we construct a tableau T from a common model \mathcal{I} of D , \mathcal{R} , and \mathcal{K} as follows:

$$\begin{aligned} \mathbf{S}_a &:= \Delta_{\mathcal{I}} \\ \mathbf{S}_c &:= \{x \in \Delta_{\mathcal{D}} \mid g^{\mathcal{I}}(s) = x \text{ for some } s \in \mathbf{S}_a\} \\ \mathcal{L}(s) &:= \{C \in \text{cl}(D, \mathcal{R}, \mathcal{K}) \mid s \in C^{\mathcal{I}}\} \\ E(s, t) &:= \{S \in \mathbf{N}_{\mathcal{R}}^{D, \mathcal{R}, \mathcal{K}} \mid (s, t) \in S^{\mathcal{I}}\} \\ e(s, g) &:= g^{\mathcal{I}}(s) \text{ if } g^{\mathcal{I}}(s) \text{ is defined} \\ \mathcal{P}(P) &:= \{(x_1, \dots, x_n) \in \mathbf{S}_c^n \mid (x_1, \dots, x_n) \in P^{\mathcal{D}}\}. \end{aligned}$$

It can be easily verified that T is a tableau for D w.r.t. \mathcal{R} and \mathcal{K} : the proof that T satisfies **(T1)** – **(T9)** is identical to the corresponding cases in (Horrocks et al., 2000; Horrocks & Sattler, 2001); **(T10)** holds by definition of \mathcal{L} ; **(T11)** by definition of \mathcal{L} and the fact that nominals are interpreted as singleton sets; **(T12)** by definition of \mathcal{L} , e , and \mathcal{P} together with the semantics of concepts $\exists g_1, \dots, g_n.P$; **(T13)** since the identity function on \mathbf{S}_c is clearly a solution for the listed predicate conjunction; **(T14)** by definition of \mathcal{L} and e together with the semantics of key constraints; and finally **(T15)** by definition of \mathcal{L} and e together with the semantics of concepts $g\uparrow$.

For the “if” direction, let $T = (\mathbf{S}_a, \mathbf{S}_c, \mathcal{L}, E, e, \mathcal{P})$ be a tableau for D w.r.t. \mathcal{R} and \mathcal{K} and let δ be a solution for the predicate conjunction in **(T13)**. We construct a model \mathcal{I} for D as follows:

$$\begin{aligned} \Delta_{\mathcal{I}} &:= \mathbf{S}_a \\ A^{\mathcal{I}} &:= \{s \in \Delta_{\mathcal{I}} \mid A \in \mathcal{L}(s)\} \quad \text{for concept names } A \\ N^{\mathcal{I}} &:= \{s \in \Delta_{\mathcal{I}} \mid N \in \mathcal{L}(s)\} \quad \text{for nominals } N \\ R^{\mathcal{I}} &:= \begin{cases} \bigcup_{\substack{S \in \mathbf{N}_{\mathcal{R}} \\ S \neq R}} S^{\mathcal{I}} \cup \{(s, t) \mid R \in E(s, t)\} & \text{for } R \in \mathbf{N}_{\mathcal{R}} \setminus \mathbf{N}_{\text{cF}} \text{ with } \textit{not Trans}(R) \\ \{(s, t) \mid R \in E(s, t)\}^+ & \text{for } R \in \mathbf{N}_{\mathcal{R}} \setminus \mathbf{N}_{\text{cF}} \text{ with } \textit{Trans}(R) \end{cases} \\ g^{\mathcal{I}}(s) &:= \begin{cases} \delta(x) & \text{if } e(s, g) = x \\ \text{undefined} & \text{if } e(s, g) \text{ is undefined} \end{cases} \quad \text{for } g \in \mathbf{N}_{\text{cF}}. \end{aligned}$$

Due to **(T11)**, the interpretation of nominals is a singleton. Moreover, the interpretation of roles is well-defined since role boxes are acyclic. The following claim is central for proving that \mathcal{I} is indeed a model for D , \mathcal{R} , and \mathcal{K} :

Claim: For each $C \in \text{cl}(D, \mathcal{R}, \mathcal{K})$, $C \in \mathcal{L}(s)$ implies $s \in C^{\mathcal{I}}$.

Proof: We proceed by induction on the *norm* $\|C\|$ of C , which is defined as follows:

$$\begin{aligned} \|A\| &:= \|\neg A\| && := 0 \text{ for } A \text{ concept name} \\ \|g\uparrow\| &:= \|\exists u_1, \dots, u_n. P\| && := 0 \\ \|C_1 \sqcap C_2\| &:= \|C_1 \sqcup C_2\| && := 1 + \|C_1\| + \|C_2\| \\ \|(\geq n R C)\| &:= \|(\leq n R C)\| && := 1 + \|C\| \end{aligned}$$

For concept names A and nominals N , the claim follows by definition of $A^{\mathcal{I}}$ and $N^{\mathcal{I}}$. For the negation of concept names A and nominals N (note that C is in NNF), the claim follows by definition of $A^{\mathcal{I}}$ and $N^{\mathcal{I}}$ together with **(T1)**. Concepts C of the form $C_1 \sqcap C_2$ and $C_1 \sqcup C_2$ can be treated using **(T2)** and **(T3)** together with the induction hypothesis. For existential, universal, and number restrictions, the proof is analogous to the one for \mathcal{SHIQ} in (Horrocks et al., 2000). For concepts of the form $C = \exists g_1, \dots, g_n. P \in \mathcal{L}(s)$, $s \in C^{\mathcal{I}}$ is an immediate consequence of **(T12)**, the definition of $g_i^{\mathcal{I}}$, and the fact that $(x_1, \dots, x_n) \in \mathcal{P}(P)$ implies $(\delta(x_1), \dots, \delta(x_n)) \in P^{\mathcal{D}}$ by **(T13)**. Finally, for concepts $C = g\uparrow$, $s \in C^{\mathcal{I}}$ is an immediate consequence of the definition of $g^{\mathcal{I}}$ together with **(T15)**. This finishes the proof of the claim.

By definition of tableaux, there exists an $s_0 \in \mathbf{S}_a$ such that $C \in \mathcal{L}(s_0)$. By the claim, $s_0 \in C^{\mathcal{I}}$ and thus \mathcal{I} is a model of C .

Next, we show that \mathcal{I} is a model of \mathcal{R} . By definition of $R^{\mathcal{I}}$, it is obvious that $\text{Trans}(R) \in \mathcal{R}$ implies that $R^{\mathcal{I}}$ is a transitive relation. Now let $S \sqsubseteq R \in \mathcal{R}$. If $\text{Trans}(R) \notin \mathcal{R}$, then we have $S^{\mathcal{I}} \subseteq R^{\mathcal{I}}$ by definition of $R^{\mathcal{I}}$. Now let $\text{Trans}(R) \in \mathcal{R}$ and $(s, t) \in S^{\mathcal{I}}$. If $S \in E(s, t)$, then **(T4)** implies $R \in E(s, t)$, and thus $(s, t) \in R^{\mathcal{I}}$. Otherwise, there is an $S' \sqsubseteq^* S$ with $\text{Trans}(S') \in \mathcal{R}$ and $(s, t) \in \{(u, v) \mid S' \in E(u, v)\}^+$. Now **(T4)** together with $S' \sqsubseteq^* R$ implies that $\{(u, v) \mid S' \in E(u, v)\} \subseteq \{(u, v) \mid R \in E(u, v)\}$, and thus $\text{Trans}(R) \in \mathcal{R}$ implies that $(s, t) \in R^{\mathcal{I}}$.

It remains to show that \mathcal{I} is a model of \mathcal{K} . To this end, let $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$ and $s, t \in C^{\mathcal{I}}$ such that $g_i^{\mathcal{I}}(s) = g_i^{\mathcal{I}}(t)$ for $1 \leq i \leq n$. Since the predicate conjunction in **(T13)** contains explicit inequalities for all distinct concrete individuals, this implies that $e(s, g_i) = e(t, g_i)$ for $1 \leq i \leq n$. **(T10)** implies $\{C, \dot{\neg} C\} \cap \mathcal{L}(s) \neq \emptyset$ and $\{C, \dot{\neg} C\} \cap \mathcal{L}(t) \neq \emptyset$. If $\dot{\neg} C \in \mathcal{L}(s)$, then the claim yields $s \in (\dot{\neg} C)^{\mathcal{I}}$ contradicting $s \in C^{\mathcal{I}}$. Thus we obtain $C \in \mathcal{L}(s)$ and, in a similar way, $C \in \mathcal{L}(t)$. Finally, **(T14)** implies that $s = t$, and thus \mathcal{I} satisfies \mathcal{K} . \square

We now proceed to prove termination, soundness, and completeness of the tableau algorithm presented in Section 4.2, starting with termination. In the following, we use $|D, \mathcal{R}, \mathcal{K}|$ to denote $|\text{cl}^+(D, \mathcal{R}, \mathcal{K})|$. Recall that this number is polynomial in the size of $D, \mathcal{R}, \mathcal{K}$.

Lemma C.2 (Termination). *Let \mathcal{D} be a key-admissible concrete domain. When started with a $\mathcal{SHOQK}(\mathcal{D})$ concept D in NNF, a role box \mathcal{R} , and a path-free key box \mathcal{K} in NNF, the tableau algorithm terminates.*

Proof. Assume that there are D , \mathcal{R} , and \mathcal{K} such that the tableau algorithm does not terminate. Since \mathcal{D} is key-admissible, this means that there is an infinite sequence S_0, S_1, \dots of completion systems such that (a) S_0 is the initial completion system S_D and (b) S_{i+1} is the result of applying a completion rule to S_i .

This is only possible if the $R\exists$ or the $R\geq$ rules are applied infinitely often: it is easily seen that the rules $R\sqcap$, $R\sqcup$, $R\leq$, $R\exists c$, $R\forall$, $R\forall_+$, Rch , and $R\approx_a$ can only be applied finitely often to completion systems whose set of abstract nodes V_a does not increase since they either add concepts into node labels (whose size is bounded), or they add concrete nodes (whose number is bounded linearly by the number of abstract nodes), or they remove abstract nodes from the forest. Hence there is a sub-sequence S_{i_1}, S_{i_2}, \dots of S_0, S_1, \dots such that S_{i_j} is the result of applying the $R\exists$ or the $R\geq$ rule to $S_{i_{j-1}}$. Let s_{i_ℓ} be the abstract node to which the $R\exists$ or the $R\geq$ rule was applied in $S_{i_{\ell-1}}$. Since $s \prec t$ implies that t was *not* generated before s , the linear ordering \prec is well-founded. Thus, we find an infinite sub-sequence S_{j_1}, S_{j_2}, \dots of S_{i_1}, S_{i_2}, \dots such that either $s_{j_\ell} = s_{j_{\ell+1}}$ for each $\ell \geq 1$ or $s_{j_\ell} \prec s_{j_{\ell+1}}$ for each $\ell \geq 1$. The former, however, is not possible since the $R\exists$ and the $R\geq$ rules are applied at most once per node and concept in $\text{cl}(D, \mathcal{R}, \mathcal{K})$: even if a node is removed, the label copying performed by the $R\leq$ rule together with clashes of type **(C3)** ensures that the $R\geq$ rule is not re-applied to the same concept and node. Thus only the second option remains: there is a subsequence S_{j_1}, S_{j_2}, \dots of S_{i_1}, S_{i_2}, \dots such that $s_{j_\ell} \prec s_{j_{\ell+1}}$ for each $\ell \geq 1$. Let \mathcal{L}_j be the labeling function in S_j . Since each abstract node is labeled with a subset \mathcal{L}_j of $\text{cl}^+(D, \mathcal{R}, \mathcal{K})$, there are nodes $s_{j_k} \prec s_{j_\ell}$ with $k \leq \ell$ and $\mathcal{L}_{j_k}(s_{j_k}) = \mathcal{L}_{j_\ell}(s_{j_\ell})$. Now node labels can only increase and, if a node t is removed, its label is conjoined to the label of a node \hat{t} with $\hat{t} \prec t$. Thus there is a node t in the completion system S_{j_ℓ} with $t \prec s_{j_\ell}$ and $\mathcal{L}_{j_\ell}(s_{j_\ell}) \subseteq \mathcal{L}_{j_\ell}(t)$. By definition, s_{j_ℓ} is thus blocked in S_{j_ℓ} , contradicting the assumption that the $R\exists$ or the $R\geq$ rule is applied to s_{j_ℓ} in S_{j_ℓ} . \square

Lemma C.3 (Soundness). *If the expansion rules can be applied to a $\text{SHOQK}(\mathcal{D})$ concept D in NNF, a role box \mathcal{R} , and a path-free key box \mathcal{K} such that they yield a complete and clash-free completion forest, then D has a tableau w.r.t. \mathcal{R} and \mathcal{K} .*

Proof. Let $S = ((V_a, V_c, E, \mathcal{L}), \mathcal{P}, \sim_c, \prec)$ be a complete and clash-free completion system. We can find a solution δ for ζ_S such that $\delta(x) = \delta(y)$ iff $x \sim_c y$: only the $R\exists c$ rule updates the predicate conjunction ζ_S , and after each rule application the \sim_c relation is updated using the concrete equivalence that a \mathcal{D} -tester returns for ζ_S (note that ζ_S is satisfiable due to clash-freeness). According to Definition 4.1, we can thus find a solution δ as required. From S and δ , we define a finite tableau $T = (\mathbf{S}_a, \mathbf{S}_c, \hat{E}, \hat{\mathcal{L}}, \hat{\mathcal{P}})$ as follows:

$$\begin{aligned}
 \mathbf{S}_a &:= \{s \in V_a \mid s \text{ occurs in } S \text{ and is not blocked}\} \\
 \mathbf{S}_c &:= \{\delta(x) \mid (s, x) \in E(g) \text{ for some } s \in \mathbf{S}_a \text{ and some } g\} \\
 \hat{\mathcal{L}}(s) &:= \mathcal{L}(s) \cap \text{cl}(D, \mathcal{R}, \mathcal{K}) \text{ (the intersection is due to the auxiliary concepts } A_i^{nRC}), \\
 \hat{E}(s, t) &:= \{R \mid t \text{ is an } R\text{-successor of } s \text{ or } t \text{ blocks an } R\text{-successor } t' \text{ of } s\} \\
 e(s, g) &:= \begin{cases} \delta(x) & \text{if } x \text{ is a } g\text{-successor of } s \\ \text{undefined} & \text{if } x \text{ has no } g\text{-successor} \end{cases} \\
 \hat{\mathcal{P}} &:= \text{the restriction of } \mathcal{P} \text{ to } \mathbf{S}_c.
 \end{aligned}$$

Note that the function e is well-defined due to the definition of *adding g -successors*. It remains to show that T satisfies **(T1)**–**(T14)**, which is basically a consequence of S being clash-free and complete.

- **(T1)** is satisfied since S does not contain a clash **(C1)**.
- **(T2)** is satisfied since the $R\sqcap$ rule cannot be applied, and thus $C_1 \sqcap C_2 \in \hat{\mathcal{L}}(s)$ implies $C_1, C_2 \in \hat{\mathcal{L}}(s)$.
- **(T3)** is satisfied since the $R\sqcup$ rule cannot be applied, and thus $C_1 \sqcup C_2 \in \hat{\mathcal{L}}(s)$ implies $\{C_1, C_2\} \cap \hat{\mathcal{L}}(s) \neq \emptyset$.
- For **(T4)**, consider $s, t \in \mathbf{S}_a$ with $R \in \hat{E}(s, t)$ and $R \boxtimes R'$. Then $R \in \hat{E}(s, t)$ implies that t is or blocks an R -successor of s . By definition of “successor”, t is or blocks an R' -successor of s , and thus $R' \in \hat{E}(s, t)$.
- For **(T5)**, let $\forall R.C \in \hat{\mathcal{L}}(s)$ and $R \in \hat{E}(s, t)$. If t is an R -successor of s , then s not being blocked implies $C \in \mathcal{L}(t)$ since the $R\forall$ rule cannot be applied. If t blocks an R -successor t' of s , then s not being blocked and the fact that the $R\forall$ rule cannot be applied yields $C \in \mathcal{L}(t')$, and the blocking condition implies $C \in \mathcal{L}(t)$.

In both cases, we thus have $C \in \hat{\mathcal{L}}(t)$.

- **(T6)** and **(T7)** are satisfied for the same reasons as **(T5)** with $R\forall$ replaced with $R\exists$ and $R\forall_+$.
- For **(T8)**, consider s with $(\geq n R C) \in \hat{\mathcal{L}}(s)$. Hence $(\geq n R C) \in \mathcal{L}(s)$ and completeness of S implies the existence of R -successors t_1, \dots, t_n of s with $C \in \mathcal{L}(t_i)$ and $t_i \neq t_j$ for $i \neq j$. The latter implies, for each $i \neq j$, the existence of integers k, ℓ such that $k \neq \ell$, $A_k^{nRC} \in \mathcal{L}(t_i)$, and $A_\ell^{nRC} \in \mathcal{L}(t_j)$. For **(T8)** to be satisfied, it remains to verify that

- no t_i can block a t_j : if this was the case, the blocking condition would imply that $\{A_k^{nRC}, A_\ell^{nRC}\} \subseteq \mathcal{L}(t_i)$.
- no t can block both t_i and t_j with $i \neq j$: similarly, this would imply that $\{A_k^{nRC}, A_\ell^{nRC}\} \subseteq \mathcal{L}(t)$.

In each case, we would have a clash **(C3)**, in contradiction to S being clash-free.

- For **(T9)**, consider s with $(\leq n R C) \in \hat{\mathcal{L}}(s)$. Hence $(\leq n R C) \in \mathcal{L}(s)$ and, since the $R\leq$ rule cannot be applied, there are at most n R -successors t_i of s . Since each t_i is either not blocked or blocked by exactly one other node (due to \prec being a linear ordering), there are at most n $u_i \in \mathbf{S}_a$ with $R \in \hat{E}(s, u_i)$ and $C \in \hat{\mathcal{L}}(u_i)$.
- For **(T10)**, let $(\leq n R C) \in \hat{\mathcal{L}}(s)$ and $R \in \hat{E}(s, t)$. Hence $(\leq n R C) \in \mathcal{L}(s)$ and t either is an R -successor of s or blocks an R -successor of s . In the first case, non-applicability of the Rch rule implies that $\{C, \dot{\div} C\} \cap \mathcal{L}(t) \neq \emptyset$. In the second case, $\{C, \dot{\div} C\} \cap \mathcal{L}(t') \neq \emptyset$ for t' the R -successor of s blocked by t , and thus the blocking condition yields $\{C, \dot{\div} C\} \cap \mathcal{L}(t) \neq \emptyset$. In both cases, this implies $\{C, \dot{\div} C\} \cap \hat{\mathcal{L}}(t) \neq \emptyset$.

Next, consider $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$ and s such that $e(s, g_i)$ is defined for each i . Hence s has a g_i -successor for each i , and thus s not being blocked and the non-applicability of the Rch rule imply that $\{C, \dot{\vdash} C\} \cap \hat{\mathcal{L}}(t) \neq \emptyset$.

- For **(T11)**, consider $N \in \hat{\mathcal{L}}(s) \cap \hat{\mathcal{L}}(t)$. By definition, $N \in \mathcal{L}(s) \cap \mathcal{L}(t)$ and thus $s \approx_a t$. Moreover, totality of \prec implies that we can assume without loss of generality that $s \prec t$ or $s = t$. Thus non-applicability of the $R\approx_a$ rule implies that $\mathcal{L}(t) \subseteq \mathcal{L}(s)$, and thus t not being blocked implies $s = t$.
- **(T12)** is satisfied since the rule $R\exists c$ cannot be applied.
- For **(T13)**, clash-freeness implies the satisfiability of

$$\bigwedge_{P \text{ used in } D, \mathcal{K}} \bigwedge_{(x_1, \dots, x_n) \in \mathcal{P}(P)} P(x_1, \dots, x_n). \quad (1)$$

By choice of δ , $\delta(x) = \delta(y)$ iff $x \sim_c y$, and thus **(T13)** is satisfied.

- For **(T14)**, let $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$, $C \in \hat{\mathcal{L}}(s) \cap \hat{\mathcal{L}}(t)$, and $e(s, g_i) = e(t, g_i)$, for all $1 \leq i \leq n$. Thus $C \in \mathcal{L}(s) \cap \mathcal{L}(t)$ and, by choice of e and δ , we have $x_i \sim_c y_i$ for $g_i \in E(s, x_i) \cap E(t, y_i)$. Hence $s \approx_a t$. Without loss of generality, we assume that $s \prec t$ or $s = t$. Thus non-applicability of the $R\approx_a$ rule implies that $\mathcal{L}(t) \subseteq \mathcal{L}(s)$, and thus t not being blocked implies $s = t$.
- **(T15)** is satisfied by definition of T and since S does not contain a clash **(C4)**.

□

Lemma C.4 (Completeness). *If a SHOQK(\mathcal{D})-concept D in NNF has a tableau w.r.t. a role box \mathcal{R} and a path-free key box \mathcal{K} , then the expansion rules can be applied to D , \mathcal{R} , and \mathcal{K} such that they yield a complete and clash-free completion forest.*

Proof. Given a tableau $T = (\mathbf{S}_a, \mathbf{S}_c, \hat{\mathcal{L}}, \hat{E}, e, \hat{\mathcal{P}})$ for D w.r.t. \mathcal{R} and \mathcal{K} , we can “guide” the non-deterministic rules $R\sqcup$, Rch, and $R\approx_a$ in such a way that each rule application preserves clash-freeness. This together with termination from Lemma C.2 finishes the proof.

Along with rule application, we perform a stepwise construction of a total mapping π that takes abstract nodes of the completion forest to elements of \mathbf{S}_a and concrete nodes of the completion forest to elements of \mathbf{S}_c .

- $\mathcal{L}(s) \cap \text{cl}(D, \mathcal{R}, \mathcal{K}) \subseteq \hat{\mathcal{L}}(\pi(s))$ for each $s \in V_a$,
- if t is an R -successor of s , then $R \in \hat{E}(\pi(s), \pi(t))$,
- if x is a g -successor of s , then $e(\pi(s), g) = \pi(x)$,
- $x \sim_c y$ iff $\pi(x) = \pi(y)$, and
- if $s \neq t$, then $\pi(s) \neq \pi(t)$.

A mapping satisfying these four conditions is called *correct* in the following. Note that a completion system for which there exists a correct mapping does not contain a clash: due to **(T1)** and the first property, we do not encounter a clash **(C1)**. A clash **(C3)** cannot occur due to the last property. The first and the third property together with **(T15)** ensure that a clash **(C4)** does not occur. Finally, Clashes **(C2)** cannot occur for the following reason: by construction of \mathcal{P} and since edges labelled with abstract features are never removed, for each tuple $(x_1, \dots, x_n) \in \mathcal{P}(P)$, we find an abstract node s and paths u_1, \dots, u_n such that $\exists u_1, \dots, u_n. P \in \mathcal{L}(s)$ and x_i is a u_i -successor of s for $1 \leq i \leq n$. Thus, the first, second, and third property together with **(T12)** and **(T13)** ensure that the conjunction

$$\bigwedge_{\substack{P \text{ used in } D, \mathcal{K} \\ (x_1, \dots, x_n) \in \mathcal{P}(P)}} P(\pi(x_1), \dots, \pi(x_n)) \quad (*)$$

has a solution δ with $\delta(\pi(x)) \neq \delta(\pi(y))$ iff $\pi(x) \neq \pi(y)$. By the fourth property, setting $\delta'(x) := \delta(\pi(x))$ for all $x \in V_c$ thus yields a solution δ' for ζ_S .

The total mapping π is inductively defined as follows: let δ be a solution for the equation in **(T13)**. Choose a node \hat{s}_0 with $D \in \hat{\mathcal{L}}(\hat{s}_0)$, and set $\pi(s_0) := \hat{s}_0$ for s_0 the (only) node of the initial completion forest. Obviously, π is correct. We will now show that each completion rule can be applied in such a way that π either is still correct or that π can be extended to a correct mapping.

- An application of the rule $R\sqcap$ preserves correctness of π due to **(T2)**.
- Due to **(T3)**, the rule $R\sqcup$ can be applied such that correctness is preserved.
- If the rule $R\exists$ adds a new node t for $\exists R.C \in \mathcal{L}(s)$, then correctness implies $\exists R.C \in \hat{\mathcal{L}}(\pi(s))$, and thus **(T6)** implies the existence of some $\hat{t} \in \mathbf{S}_a$ with $R \in E(\pi(s), \hat{t})$ and $C \in \hat{\mathcal{L}}(\hat{t})$. Thus extending π with $\pi(t) := \hat{t}$ obviously yields a correct mapping.
- If the rule $R\geq$ adds n nodes t_i for $(\geq n R C) \in \mathcal{L}(s)$, then correctness implies $(\geq n R C) \in \hat{\mathcal{L}}(\pi(s))$, and thus **(T8)** implies the existence of $\hat{t}_1, \dots, \hat{t}_n \in \hat{\mathbf{S}}_a$ with $\hat{t}_i \neq \hat{t}_j$ for $i \neq j$, $R \in E(\pi(s), \hat{t}_i)$, and $C \in \hat{\mathcal{L}}(\hat{t}_i)$. Thus extending π with $\pi(t_i) := \hat{t}_i$ obviously yields a correct mapping.
- Assume that the $R\leq$ rule is applicable to a node s with $(\leq n R C) \in \mathcal{L}(s)$ and more than n R -successors t_i with $C \in \mathcal{L}(t_i)$. Then correctness implies that $(\leq n R C) \in \hat{\mathcal{L}}(\pi(s))$, $R \in \hat{E}(\pi(s), \pi(t_i))$, and $C \in \hat{\mathcal{L}}(t_i)$. Thus, by **(T9)**, there are $i \neq j$ with $\pi(t_i) = \pi(t_j)$. Again, correctness implies that *not* $t_i \neq t_j$ and, without loss of generality, we can assume that $t_i \prec t_j$. Hence applying the rule and thereby merging $\mathcal{L}(t_j)$ into $\mathcal{L}(t_i)$ preserves correctness.
- For the rule $R\exists c$, π can be extended in a similar way as for $R\exists$: if a new g_i -successor x_i of s is added, then extending π with $\pi(x_i) := e(\pi(s), g_i)$ yields a correct π . Moreover, **(T13)** ensures that \sim_c can be updated in such a way that the fourth condition is preserved.
- For the $R\forall$ rule, π does not need to be extended, and **(T5)**, **(T4)**, and the definition of R -successors imply that correctness is preserved.

- The $R\forall_+$ rule is similar, with the only difference that **(T7)** takes the place of **(T5)**.
- Due to **(T10)**, the rule Rch can be applied without violating correctness.
- For $R\approx_a$, we consider two reasons for $R\approx_a$ to be applicable:
 - $N \in \mathcal{L}(s) \cap \mathcal{L}(t)$. Then correctness of π and **(T11)** imply that $\pi(s) = \pi(t)$.
 - $(g_1, \dots, g_n \text{ keyfor } C) \in \mathcal{K}$, $C \in \mathcal{L}(s) \cap \mathcal{L}(t)$, and $g_i \in E(s, x_i) \cap E(t, y_i)$ and $x_i \sim_c y_i$ for $1 \leq i \leq n$. Then correctness implies that $\hat{e}(\pi(s), g_i) = e(\pi(t), g_i)$, and thus **(T14)** together with the first property of correctness imply that $\pi(s) = \pi(t)$.

In both cases, applying $R\approx_a$ to s and t preserves correctness. □

As an immediate consequence of Lemmas 4.2, C.2, C.3, and C.4, the tableau algorithm always terminates and answers “ D is satisfiable w.r.t. \mathcal{R} and \mathcal{K} ” if and only if the input concept D is satisfiable w.r.t. the input role box \mathcal{R} and the input key box \mathcal{K} . Since concept satisfiability w.r.t. TBoxes can be reduced to concept satisfiability without TBoxes, we obtain the following result:

Proposition C.5. *Let \mathcal{D} be a key-admissible concrete domain. The tableau algorithm decides satisfiability of $\mathcal{SHOQK}(\mathcal{D})$ concepts w.r.t. TBoxes, role boxes, and path-free key boxes.*

It is not hard to verify that the proof of Lemma C.4 together with Lemmas 4.2 and C.2 yield a *bounded model property* for $\mathcal{SHOQK}(\mathcal{D})$, where the bound is exponential.

Corollary C.6. *If a $\mathcal{SHOQK}(\mathcal{D})$ -concept D is satisfiable w.r.t. a role box \mathcal{R} and a path-free key box \mathcal{K} , then D is satisfiable w.r.t. \mathcal{R} and \mathcal{K} in a model of size at most $|\Delta_{\mathcal{I}}| \leq 2^m$ for $m = \# \text{cl}^+(D, \mathcal{R}, \mathcal{K})$.*

Proof. If a $\mathcal{SHOQK}(\mathcal{D})$ -concept D is satisfiable w.r.t. a role box \mathcal{R} and a path-free key box \mathcal{K} , Lemma C.4 implies that the tableau algorithm constructs a complete and clash-free completion forest for D , \mathcal{R} , and \mathcal{K} . By the definition of blocking, the number of abstract nodes in a completion forest that are not blocked is bounded by 2^m , where m is defined as in the corollary: if $s \neq t \in V_a$ are abstract nodes in a completion forest and $\mathcal{L}(s) = \mathcal{L}(t)$, then either s blocks t , t blocks s , or they are both blocked by another node u . Moreover, it is easily seen that the number of concrete successors per abstract node is bounded by the number of concrete features in C , \mathcal{R} , and \mathcal{K} . Now, in the proof of Lemma C.4, the abstract nodes in the tableau constructed from a complete and clash-free completion forest coincide with the nodes that are not blocked in the completion forest. Finally, in the proof of Lemma 4.2 the interpretation domain of a model constructed from a tableau coincides with the abstract nodes in the tableau. Summing up, a $\mathcal{SHOQK}(\mathcal{D})$ -concept that is satisfiable w.r.t. \mathcal{R} and \mathcal{K} has a model of size $|\Delta_{\mathcal{I}}| \leq 2^m$. □