TECHNISCHE UNIVERSITÄT DRESDEN

Dresden University of Technology
Institute for Theoretical Computer Science
Chair for Automata Theory

LTCS–Report

The Complexity of Finite Model Reasoning in Description Logics

Carsten Lutz, Ulrike Sattler, Lidia Tendera

LTCS-Report 02-05

Lehrstuhl für Automatentheorie
Institut für Theoretische Informatik
TU Dresden
http://lat.inf.tu-dresden.de

Hans-Grundig-Str. 25
01062 Dresden
Germany

The Complexity of Finite Model Reasoning in Description Logics

Carsten Lutz, Ulrike Sattler, and Lidia Tendera
Institute for Theoretical Computer Science
TU Dresden, Germany
lastname@tcs.inf.tu-dresden.de

Abstract

We analyze the complexity of finite model reasoning in the description logic $\mathcal{ALCQI}$, i.e. $\mathcal{ALC}$ augmented with qualifying number restrictions, inverse roles, and general TBoxes. It turns out that all relevant reasoning tasks such as concept satisfiability and ABox consistency are ExpTime-complete, regardless of whether the numbers in number restrictions are coded unarily or binarily. Thus, finite model reasoning with $\mathcal{ALCQI}$ is not harder than standard reasoning with $\mathcal{ALCQI}$.

1 Motivation

Description logics (DLs) are a family of logical formalisms that originated in the field of knowledge representation and are nowadays used in a wide range of applications [1]. Similar to many modal logics (to which DLs are closely related), most description logics enjoy the finite model property (FMP). This is, for example, the case for the basic propositionally closed DL $\mathcal{ALC}$ that is well-known to be a notational variant of the multi-modal logic K [13]: satisfiability of $\mathcal{ALC}$-concepts (the DL equivalent of a formula) w.r.t. finite models coincides with the satisfiability of $\mathcal{ALC}$-concepts w.r.t. arbitrary models [13]. However, there also exist description logics that do not enjoy FMP. A rather important example for such a DL is $\mathcal{ALCQI}$, which is obtained from $\mathcal{ALC}$ by adding qualifying number restrictions (corresponding to graded modalities in modal logic), the inverse role constructor (inverse modalities), and general TBoxes (roughly corresponding to the universal modality).

The fact that $\mathcal{ALCQI}$ lacks FMP becomes particularly important if we consider this logic's most prominent application, which is reasoning about conceptual database models: if such a model is described by one of the standard formalisms—namely ER diagrams for relational databases and UML diagrams for object-oriented databases—then it can be translated into a DL TBox, i.e. a set of concept equations; afterwards, a description logic reasoner such as FaCT and RACER can be used to detect inconsistencies and to infer implicit IS-A relationships between entities/classes [3]. This useful and original application has already led to the implementation of tools that provide a GUI for specifying conceptual models, automatize the translation into description logics, and display the information returned by the DL reasoner [8]. When doing reasoning about databases, one is clearly interested in reasoning w.r.t. *finite* models since models describe databases, and these are finite objects. However, all available DL reasoning systems are performing reasoning w.r.t. arbitrary (as opposed to finite) models. Since it is well-known that there exist ER and UML diagrams which are satisfiable only in infinite models [15], this means that some inconsistencies and IS-A relationships will not be detected if existing DL reasoners are used for reasoning about conceptual models.

The main reason for existing DL reasoners to perform only reasoning w.r.t. arbitrary models is that finite model reasoning in description logics such as $\mathcal{ALCQI}$ is not yet well-understood. The only known algorithm is presented by Calvanese in [4], where he proves that reasoning in $\mathcal{ALCQI}$ is decidable in 2-ExpTime. *The purpose of this paper is to improve the understanding of finite model reasoning in description logics by establishing tight* ExpTime *complexity bounds for finite model reasoning in the DL* $\mathcal{ALCQI}$*.* More precisely, in this paper we present the following results:

In Section 3, we develop an algorithm that is capable of deciding the finite satisfiability of $\mathcal{ALCQI}$-concepts w.r.t. TBoxes. Similar to Calvanese's approach, the core idea behind our algorithm is to translate a given satisfiability problem into a set of linear equations that can then be solved by linear programming methods. The main difference to Calvanese's approach is that our equation systems talk about certain components of models, so-called *mosaics*, which allows us to keep the size of equation systems exponential in the size of the input. In this way, we improve the best-known 2-ExpTime upper bound to a tight ExpTime one.

Since the approach presented in Section 3 presupposes unary coding of the numbers occurring inside qualifying number restrictions, in Section 4 we consider finite model reasoning in $\mathcal{ALCQI}$ with numbers coded in binary. We give a polynomial reduction of $\mathcal{ALCQI}$-concept satisfiability w.r.t. TBoxes to the satisfiability of $\mathcal{ALCFI}$-concept satisfiability w.r.t. TBoxes, where $\mathcal{ALCFI}$ is obtained from $\mathcal{ALCQI}$ by allowing only the number 1 to be used in number restrictions. Since finite model reasoning in $\mathcal{ALCFI}$ is in ExpTime by the results from Section 3 (the coding of numbers is not an issue here), we obtain a tight ExpTime bound for finite model reasoning in $\mathcal{ALCQI}$ with numbers coded in binary.

Finally, in Section 5 we consider the finite satisfiabiliy of ABoxes w.r.t. TBoxes. Intuitively, ABoxes describe a particular state of affairs, a "snapshot" of the world. By a reduction to (finite) concept satisfiability, we are able to show that this reasoning task is also ExpTime-complete, independently of the way in which numbers are coded.

2 Preliminaries

We introduce syntax and semantics of $\mathcal{ALCQI}$.

Definition 1 ($\mathcal{ALCQI}$ Syntax) *Let* $\mathbf{R}$ *and* $\mathbf{C}$ *be disjoint and countably infinite sets of* role *and* concept names*. A* role *is either a role name* $R \in \mathbf{R}$ *or the inverse* R^- *of a role name* $R \in \mathbf{R}$*. The set of* $\mathcal{ALCQI}$*-concepts is the smallest set satisfying the following properties:*

- *each concept name* $A \in \mathbf{C}$ *is an* $\mathcal{ALCQI}$*-concept;*
- *if* C *and* D *are* $\mathcal{ALCQI}$*-concepts,* R *is a role, and* n *a natural number, then* $\neg C$*,* $C \sqcap D$*,* $C \sqcup D$*,* $(\leqslant n\ R\ C)$*, and* $(\geqslant n\ R\ C)$ *are also* $\mathcal{ALCQI}$*-concepts.*

A concept equation *is of the form* $C \doteq D$ *for* C, D *two* $\mathcal{ALCQI}$*-concepts. A* TBox *is a finite set of concept equations.* ♢

As usual, we use the standard abbreviations $\rightarrow$ and $\leftrightarrow$ as well as $\exists R.C$ for $(\geqslant 1\ R\ C)$, $\forall R.C$ for $(\leqslant 0\ R\ \neg C)$, $\top$ to denote an arbitrary propositional tautology, and $\bot$ as abbreviation for $\neg\top$. To avoid roles like $(R^-)^-$, we define a function Inv on roles such that $\mathsf{Inv}(R) = R^-$ if R is a role name, and $\mathsf{Inv}(R) = S$ if $R = S^-$. The fragment $\mathcal{ALCFI}$ of $\mathcal{ALCQI}$ is obtained by admitting only atmost restrictions $(\leqslant n\ R\ C)$ with $n \in \{0, 1\}$ and only atleast restrictions $(\geqslant n\ R\ C)$ with $n \in \{1, 2\}$.

Definition 2 ($\mathcal{ALCQI}$ Semantics) *An interpretation* $\mathcal{I}$ *is a pair* $(\Delta^{\mathcal{I}}, \cdot^{\mathcal{I}})$ *where* $\Delta^{\mathcal{I}}$ *is a non-empty set and* $\cdot^{\mathcal{I}}$ *is a mapping which associates*

- *with each concept name* A *a set* $A^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}}$ *and*
- *with each role name* R*, a binary relation* $R^{\mathcal{I}} \subseteq \Delta^{\mathcal{I}} \times \Delta^{\mathcal{I}}$*.*

The interpretation of inverse roles and complex concepts is then defined as follows:

$$
\begin{array}{rcl}
(R^-)^{\mathcal{I}} & = & \{\langle e, d\rangle \mid \langle d, e\rangle \in R^{\mathcal{I}}\} \\
(\neg C)^{\mathcal{I}} & = & \Delta^{\mathcal{I}} \setminus C^{\mathcal{I}} \\
(C \sqcap D)^{\mathcal{I}} & = & C^{\mathcal{I}} \cap D^{\mathcal{I}} \\
(C \sqcup D)^{\mathcal{I}} & = & C^{\mathcal{I}} \cup D^{\mathcal{I}} \\
(\leqslant n\ R\ C)^{\mathcal{I}} & = & \{d \mid \#\{e \in C^{\mathcal{I}} \mid \langle d, e\rangle \in R^{\mathcal{I}}\} \leq n\} \\
(\geqslant n\ R\ C)^{\mathcal{I}} & = & \{d \mid \#\{e \in C^{\mathcal{I}} \mid \langle d, e\rangle \in R^{\mathcal{I}}\} \geq n\}
\end{array}
$$

An interpretation $\mathcal{I}$ satisfies *a concept equation* $C \doteq D$ *if* $C^{\mathcal{I}} = D^{\mathcal{I}}$*, and* $\mathcal{I}$ *is called a* model *of a TBox* $\mathcal{T}$ *if* $\mathcal{I}$ *satisfies all concept equations in* $\mathcal{T}$*.*

A concept C *is* satisfiable w.r.t. a TBox $\mathcal{T}$ *if there is a model* $\mathcal{I}$ *of* $\mathcal{T}$ *with* $C^{\mathcal{I}} \neq \emptyset$*. A concept* C *is* finitely satisfiable w.r.t. a TBox $\mathcal{T}$ *if there is a model* $\mathcal{I}$ *of* $\mathcal{T}$ *with* $C^{\mathcal{I}} \neq \emptyset$ *and* $\Delta^{\mathcal{I}}$ *finite.* ♢

Let us consider a witness for the fact that $\mathcal{ALCQI}$ lacks FMP: the concept $\neg A \sqcap \exists R.A$ is satisfiable w.r.t. the TBox $\{A \doteq \exists R.A \sqcap (\leqslant 1\ R^- \top)\}$, but each of its models contains an infinite R-chain.

There exists another important reasoning problem on concepts and TBoxes: *subsumption*. However, since subsumption can be reduced to (un)satisfiability and vice versa, we just note that all complexity bounds obtained in this paper also apply to subsumption.

In what follows, we will only consider TBoxes of the rather simple form $\{\top \doteq C\}$. This can be done w.l.o.g. since an interpretation $\mathcal{I}$ is a model of a TBox $\mathcal{T} = \{C_i \doteq D_i \mid 1 \leq i \leq n\}$ iff it is a model of $\{\top \doteq \bigsqcap_{1 \leq i \leq n}(C_i \leftrightarrow D_i)\}$.

3 Unary Coding of Numbers

In this section, we present a decision procedure for finite satisfiability of $\mathcal{ALCQI}$-concepts w.r.t. TBoxes that runs in deterministic exponential time, provided that numbers in number restrictions are coded unarily. In Section 4, we will generalize this upper bound to binary coding of numbers.

As observed by Calvanese in [4], combinatorics is an important issue when deciding finite satisfiability of $\mathcal{ALCQI}$-concepts. To illustrate this, consider the TBox

$$\mathcal{T} := \{A \doteq (\geqslant 2\ R\ B), \quad B \doteq (\leqslant 1\ R^-\ A)\}.$$

It should be clear that, in any model of $\mathcal{T}$, there are at least twice as many objects satisfying $B \sqcap (\leqslant 1\ R^-\ A)$ as there are objects satisfying $A \sqcap (\geqslant 2\ R\ B)$. This simple example suggests that (i) *types* (i.e., sets of concepts satisfied by a particular object in a particular model) such as $\{A, (\geqslant 2\ R\ B)\}$ are a natural notion for dealing with finite satisfiability, and (ii) the combinatorics introduced by finite domains can be addressed with inequalities like $2 \cdot x_T \leq x_{T'}$, where the variable x_T describes the number of instance of a type T (e.g. $\{A, (\geqslant 2\ R\ B)\}$), while $x_{T'}$ describes the number of instances of another type T' (e.g. $\{B, (\leqslant 1\ R^-\ A)\}$). These combinatorial constraints are not an issue if infinite domains are admitted: in this case, we can always find a model where all types that have instances at all have the same number of instances, namely countably infinitely many.

Considering the above two points, a first idea to devise a decision procedure for finite satisfiability of $\mathcal{ALCQI}$-concepts w.r.t. TBoxes is to translate an input concept and TBox into a system of inequalities with one variable for each type, and then to use existing algorithms to check whether the equation system has a non-negative integer solution. For example, the satisfiability problem of the concept A w.r.t. the TBox $\mathcal{T}$ above can be translated into the two inequalities

$$\sum_{\{T \mid (\geqslant 2\ R\ B) \in T\}} 2 \cdot x_T \leq \sum_{\{T \mid (\leqslant 1\ \mathsf{Inv}(R)\ A) \in T\}} x_T \quad \text{and} \quad \sum_{\{T \mid A \in T\}} x_T > 0$$

where the sums range over all types induced by the input concept A and TBox $\mathcal{T}$. It is not hard to see that any non-negative integer solution to this equation system can be used to construct a finite model for A and $\mathcal{T}$ and vice versa.

$$
\begin{array}{rcllrcl}
\neg(C \sqcap D) & \rightsquigarrow & \neg C \sqcup \neg D & & \neg(C \sqcup D) & \rightsquigarrow & \neg C \sqcap \neg D \\
\neg\neg C & \rightsquigarrow & C & & \neg(\leqslant n\ R\ C) & \rightsquigarrow & (\geqslant n+1\ R\ C) \\
\neg(\geqslant n\ R\ C)) & \rightsquigarrow & (\leqslant n-1\ R\ C) & \text{if } n > 0 & & & \\
\neg(\geqslant n\ R\ C)) & \rightsquigarrow & \bot & \text{if } n = 0 & & &
\end{array}
$$

Figure 1: The NNF rewrite rules.

Unfortunately, there is a problem with this approach: assume that the input concept and TBox induce types T_1 to T_5 as follows: $(\geqslant 1\ R\ C) \in T_1$, $(\geqslant 1\ R\ D) \in T_2$, $(\leqslant 1\ \mathsf{Inv}(R)\ \top) \in T_3 \cap T_4 \cap T_5$, $C \in T_3 \cap T_4$, and $D \in T_4 \cap T_5$. The translation described above yields the inequalities

$$x_{T_1} \leq x_{T_3} + x_{T_4} \text{ and } x_{T_2} \leq x_{T_4} + x_{T_5},$$

which have $x_{T_1} = x_{T_2} = x_{T_4} = 1$ and $x_{T_3} = x_{T_5} = 0$ as an integer solution. Trying to construct a model with a_1, a_2, and a_4 instances of T_1, T_2, and T_4, respectively, we have to use a_4 as a witness of a_1 being an instance of $(\geqslant 1\ R\ C)$ and a_2 being an instance of $(\geqslant 1\ R\ D)$. However, this violates the $(\leqslant 1\ \mathsf{Inv}(R)\ \top)$ concept in T_4.

This example illustrates that "counting types" does not suffice: conflicts may arise if a type containing an atmost restriction (T_4) can be used as a witness for atleast restrictions in more than one type (T_1 and T_2). In such a situation, it is thus necessary to (additionally) fix the types that are actually used as witnesses for atleast restrictions. We achieve this by defining systems of inequalities based on small chunks of models called *mosaics*, rather than based directly on types. Intuitively, a mosaic describes the type of an object and fixes the type of certain "important" witnesses.

Before defining mosaics, we introduce some preliminaries. In the remainder of this paper, we assume concepts (also those appearing inside TBoxes) to be in negation normal form (NNF), i.e., negation is only allowed in front of concept names. Every $\mathcal{ALCQI}$-concept can be transformed into an equivalent one in NNF by exhaustively applying the rewrite rules displayed in Figure 1. We use $\dot{\neg}C$ to denote the NNF of $\neg C$. For a concept C_0 and a TBox $\mathcal{T} = \{\top \doteq C_\mathcal{T}\}$, $\mathsf{cl}(C_0, \mathcal{T})$ is the smallest set containing all sub-concepts of C_0 and $C_\mathcal{T}$ that is closed under $\dot{\neg}$. It can easily be shown that the cardinality of $\mathsf{cl}(C_0, \mathcal{T})$ is linear in the size of C_0 and $\mathcal{T}$. We use $\mathsf{rol}(C_0, \mathcal{T})$ to denote the set of role names R and their inverses R^- occurring in C_0 or $\mathcal{T}$.

Definition 3 (Type) *A* type T *for* $C_0, \mathcal{T} = \{\top \doteq C_\mathcal{T}\}$ *is a set* $T \subseteq \mathsf{cl}(C_0, \mathcal{T})$ *such that, for each* $D, E \in \mathsf{cl}(C_0, \mathcal{T})$*, we have*

1. $D \in T$ *iff* $\dot{\neg}D \notin T$,
2. *if* $D \sqcap E \in \mathsf{cl}(C_0, \mathcal{T})$*, then* $D \sqcap E \in T$ *iff* $D \in T$ *and* $E \in T$,

3. *if* $D \sqcup E \in \mathsf{cl}(C_0, \mathcal{T})$*, then* $D \sqcup E \in T$ *iff* $D \in T$ *or* $E \in T$*, and*

4. $C_{\mathcal{T}} \in T$.

We use $\mathsf{type}(C_0, \mathcal{T})$ *to denote the set of all types over* $C_0, \mathcal{T}$*. Let* T *be a type and* $\bowtie \in \{\leqslant, \geqslant\}$*. Then we use the following abbreviations:*

$$\begin{aligned} \mathsf{max}^{\bowtie}(T) &:= \max\{n \mid (\bowtie n\ R\ C) \in T\} \\ \mathsf{sum}^{\bowtie}(T) &:= \sum_{(\bowtie n\ R\ C) \in T} n. \end{aligned}$$

◊

We are now ready to define the core notion of our approach: mosaics.

Definition 4 (Mosaic) *For two types* T_1, T_2 *and a role* R*, we write* $\lim_R(T_1, T_2)$ *(*T_2 *is a* limited ressource *for* T_1 *w.r.t.* R*) if* $C \in T_1$ *and* $(\leqslant n\ \mathsf{Inv}(R)\ C) \in T_2$ *for some* $C \in \mathsf{cl}(C_0, \mathcal{T})$ *and* $n \in \mathbb{N}$.

A mosaic *for* $C_0, \mathcal{T}$ *is a triple* $M = (T_M, L_M, E_M)$ *where*

- $T_M \in \mathsf{type}(C_0, \mathcal{T})$,
- L_M *is a function from* $\mathsf{rol}(C_0, \mathcal{T}) \times \mathsf{type}(C_0, \mathcal{T})$ *to* $\mathbb{N}$*, and*
- E_M *is a function from* $\mathsf{rol}(C_0, \mathcal{T}) \times \mathsf{type}(C_0, \mathcal{T})$ *to* $\mathbb{N}$

such that the following conditions are satisfied:

1. *if* $L_M(R, T) > 0$*, then* $\lim_R(T_M, T)$ *and not* $\lim_{\mathsf{Inv}(R)}(T, T_M)$*,*
2. *if* $E_M(R, T) > 0$*, then* $\lim_{\mathsf{Inv}(R)}(T, T_M)$*,*
3. *if* $(\leqslant n\ R\ C) \in T_M$*, then* $n \geq \sum_{\{T \mid C \in T\}} E_M(R, T)$*,*
4. $\#\{(R, T) \mid L_M(R, T) > 0\} \leq \mathsf{sum}^{\geqslant}(T_M)$ *and* $\max(\mathsf{ran}(L_M)) \leq \mathsf{max}^{\geqslant}(T_M)$.

◊

Let us spend a few words on the intuition behind mosaics. Consider a mosaic M and one of its "instances" d in some interpretation. While T_M is simply the type of d, L_M and E_M are used to describe certain "neighbors" of d, i.e. objects e reachable from d via some role. For simplicity, fix a role R. There exist three possibilities for the relationship between T_M and T, the type of e:

1. Not $\lim_R(T_M, T)$ and not $\lim_{\mathsf{Inv}(R)}(T, T_M)$; Then d may have an arbitrary number of R-neighbors of type T and every instance of T may have an arbitrary number of $\mathsf{Inv}(R)$-neighbors of type T_M. Intuitively, R-neighbors of type T are "uncritical" and not recorded in the mosaic.

2. $\lim_R(T_M, T)$ and not $\lim_{\mathsf{Inv}(R)}(T, T_M)$. Then d may have an arbitrary number of R-neighbors of type T, but every instance of T may only have a limited number of $\mathsf{Inv}(R)$-neighbors of type T_M. Thus, R-neighbors of type T are a limited ressource and we record in L_M the *minimal* number of R-neighbors of type T that d needs ("L" for "lower bound").

3. $\lim_{\mathsf{Inv}(R)}(T, T_M)$. Then d may only have a limited number of R-neighbors of type T. To prevent the violation of atmost restrictions in T_M, we record the *exact* number of d's R-neighbors of type T in E_M.

(M1) and (M2) ensure that L_M and E_M record information for the "correct" types as described above; (M3) ensures that atmost restrictions are not violated—by definition, this concerns only neighbors with E_M-types; finally, (M4) puts upper bounds on L_M to ensure that there exist only exponentially many mosaics (see below). Atleast restrictions are not mentioned in the definition of mosaics and will be treated by the systems of inequalities to be defined later.

Now for the number of mosaics. The cardinality of $\mathsf{type}(C_0, \mathcal{T})$ is exponential in the size of C_0 and $\mathcal{T}$. Next, (M2) and (M3) imply $\#\{(R,T) \mid E_M(R,T) > 0\} \leq \mathsf{sum}^{\leqslant}(T_M)$ and $\max(\mathsf{ran}(E_M)) \leq \mathsf{max}^{\leqslant}(T_M)$. This, together with (M4) and the fact that $\mathsf{max}^{\bowtie}(T)$ and $\mathsf{sum}^{\bowtie}(T)$ are linear in the size of C_0 and $\mathcal{T}$ for $\bowtie \in \{\leqslant, \geqslant\}$ (since we assume numbers to be coded in unary), clearly implies that the number of mosaics is bounded exponentially in the size of C_0 and $\mathcal{T}$.

We are now ready to define, for an input concept C_0 and TBox $\mathcal{T}$, a corresponding system of inequalities.

Definition 5 (Equation System) *For C_0 an $\mathcal{ALCQI}$-concept and $\mathcal{T}$ a TBox, we introduce a variable x_M for each mosaic M over $C_0, \mathcal{T}$ and define the equation system $\mathcal{E}_{C_0,\mathcal{T}}$ by taking (i) the equation*

$$\sum_{\{M \mid C_0 \in T_M\}} x_M \geq 1, \tag{E1}$$

(ii) for each pair of types $T,\ T' \in \mathsf{type}(C_0, \mathcal{T})$ and role R such that $\lim_R(T, T')$ and not $\lim_{\mathsf{Inv}(R)}(T', T)$ the equation

$$\sum_{\{M \mid T_M = T\}} L_M(R, T') \cdot x_M \leq \sum_{\{M \mid T_M = T'\}} E_M(\mathsf{Inv}(R), T) \cdot x_M, \tag{E2}$$

and (iii) for each pair of types $T,\ T' \in \mathsf{type}(C_0, \mathcal{T})$ and role R such that $\lim_R(T, T')$ and $\lim_{\mathsf{Inv}(R)}(T', T)$ the equation

$$\sum_{\{M \mid T_M = T\}} E_M(R, T') \cdot x_M = \sum_{\{M \mid T_M = T'\}} E_M(\mathsf{Inv}(R), T) \cdot x_M. \tag{E3}$$

A solution of $\mathcal{E}_{C_0,\mathcal{T}}$ is admissible *if it is a non-negative integer solution and satisfies the following conditions:*

(i) for each pair of types $T,\ T' \in \mathsf{type}(C_0, \mathcal{T})$ and role R such that $\lim_R(T, T')$

and not $\lim_{\mathsf{Inv}(R)}(T', T)$,

$$\text{if} \quad \sum_{\{M \mid T_M = T'\}} E_M(\mathsf{Inv}(R), T) \cdot x_M > 0, \quad \text{then} \quad \sum_{\{M \mid T_M = T\}} x_M > 0. \qquad (A1)$$

(ii) for each mosaic M and each role R, if $x_M > 0$, $(\geqslant n\ R\ C) \in T_M$, and

$$m = \sum_{\{T \mid C \in T\}} L_M(R, T) + \sum_{\{T \mid C \in T\}} E_M(R, T) < n,$$

then (A2)

$$\sum_{\substack{\{M' \mid \ C \in T_{M'},\ \text{not } \lim_R(T_M, T_{M'}), \\ \text{and not } \lim_{\mathsf{Inv}(R)}(T_{M'}, T_M)\}}} x_{M'} > 0,$$

◊

While inequality (E1) guarantees the existence of an instance of C_0, inequalities (E2) and (E3) enforce the lower and exact bounds on the number of neighbors as described by L_M and E_M. A special case is treated by condition (A1): in inequality (E2), it may happen that the left-hand side is zero while the right-hand side is non-zero. In this case, there is an instance of a mosaic M' with $T_{M'} = T'$ and $E_M(\mathsf{Inv}(R), T) > 0$ (counted on the right-hand side), but there is no instance of a mosaic M with $T_M = T$ (counted on the left-hand side)—thus we cannot find any neighbors as required by $E_M(\mathsf{Inv}(R), T)$. To cure this defect, condition (A1) ensures that, if the right-hand side of (E2) is non-zero, then there is at least one instance of a mosaic M with $T_M = T$.[1] Finally, (A2) takes care of atleast restrictions in types T_M: if the number of R-neighbors enforced by L_M and E_M is not enough for some $(\geqslant n\ R\ C) \in T_M$, then we make sure that there is at least one instance of a mosaic M' such that $C \in T_{M'}$ and, for instances of M (M'), the number of R-neighbors ($\mathsf{Inv}(R)$-neighbors) that are instances of M' (M) is not limited.[1]

Lemma 6 *If C_0 is finitely satisfiable w.r.t. $\mathcal{T}$, then the equation system $\mathcal{E}_{C_0,\mathcal{T}}$ has an admissible solution.*

Proof. Let $\mathcal{I}$ be a finite model of C_0 w.r.t. $\mathcal{T}$. From $\mathcal{I}$, we can construct an admissible solution for $\mathcal{E}_{C_0,\mathcal{T}}$. First, let us introduce some notions: for $e \in \Delta^{\mathcal{I}}$, we define the type $t(e)$ that e is instance of as

$$t(e) := \{D \in \mathsf{cl}(C_0, \mathcal{T}) \mid e \in D^{\mathcal{I}}\}.$$

Obviously, $t(e) \in \mathsf{type}(C_0, \mathcal{T})$. For $T \in \mathsf{type}(C_0, \mathcal{T})$, define

$$T^{\mathcal{I}} := \{e \in \Delta^{\mathcal{I}} \mid t(e) = T\}.$$

[1] To see why a single instance suffices, consult the proof sketch of Lemma 6.

Now fix a choice function $\mathsf{ch}(\cdot,\cdot)$ which associates, with each $e \in \Delta^{\mathcal{I}}$ and role R, some $\mathsf{ch}(e,R) \subseteq \Delta^{\mathcal{I}}$ such that

1. $\langle e, e' \rangle \in R^{\mathcal{I}}$, for all $e' \in \mathsf{ch}(e,R)$,
2. $\#\mathsf{ch}(e,R) \leq \mathsf{sum}^{\geqslant}(t(e))$,
3. $\#(\mathsf{ch}(e,R) \cap T^{\mathcal{I}}) \leq \mathsf{max}^{\geqslant}(t(e))$ for each $T \in \mathsf{type}(C_0, \mathcal{T})$, and
4. if $(\geqslant nR.C) \in t(e)$, then $\#(\mathsf{ch}(e,R) \cap C^{\mathcal{I}}) \geq n$.

Using the semantics and the definitions of $\mathsf{max}^{\geqslant}$ and $\mathsf{sum}^{\geqslant}$, it is readily checked that such a choice function exists. For each mosaic M, we then define a set

$$\begin{array}{rl} M^{\mathcal{I}} := \{e \in \Delta^{\mathcal{I}} \mid & t(e) = T_M \text{ and, for all roles } R \text{ and types } T', \\ & \text{if } \lim_R(T_M, T'), \text{ and not } \lim_{\mathsf{Inv}(R)}(T', T_M) \\ & \quad \text{then } L_M(R, T') = \#(\mathsf{ch}(e,R) \cap T'^{\mathcal{I}}), \\ & \text{if } \lim_{\mathsf{Inv}(R)}(T', T_M) \\ & \quad \text{then } E_M(R, T') = \#\{e' \in T'^{\mathcal{I}} \mid \langle e, e' \rangle \in R^{\mathcal{I}}\}\} \end{array}$$

Finally, we set $\hat{x}_M := \# M^{\mathcal{I}}$.

In what follows, we show that $\{\hat{x}_M \mid M \text{ a mosaic}\}$ is an admissible solution. We first establish the following claim:

Claim: For each $e \in \Delta^{\mathcal{I}}$, there exists exactly one mosaic M such that $e \in M^{\mathcal{I}}$.

Using the definition of mosaics, of the sets $M^{\mathcal{I}}$, and of the choice function ch, it is straightforward to show that, for each $e \in \Delta^{\mathcal{I}}$, there exists *at least* one mosaic M such that $T_M = t(e)$ and $e \in M^{\mathcal{I}}$. Now suppose that there exists an $e \in \Delta^{\mathcal{I}}$ and mosaics M_1 and M_2 with $e \in M_1^{\mathcal{I}} \cap M_2^{\mathcal{I}}$. By definition of the sets $M^{\mathcal{I}}$, this implies (i) $T_{M_1} = T_{M_2}$, (ii) $L_{M_1}(R, T') = L_{M_2}(R, T')$ for all roles R and types T' with $\lim_R(T_{M_1}, T')$ and not $\lim_{\mathsf{Inv}(R)}(T', T_{M_1})$, and (iii) $E_{M_1}(R, T') = E_{M_2}(R, T')$ for all roles R and types T' with $\lim_{\mathsf{Inv}(R)}(T', T_{M_1})$. Using Properties 1 and 2 of mosaics, it is now easy to show that $M_1 = M_2$. Thus, for each $e \in \Delta^{\mathcal{I}}$, there exists *exactly* one mosaic M such that $e \in M^{\mathcal{I}}$, for which obviously $T_M = t(e)$.

We now argue that $\{\hat{x}_M \mid M \text{ a mosaic}\}$ satisfies Equations $(E1)$ to $(E3)$ and admissibility conditions $(A1)$ and $(A2)$. Equation $(E1)$ is satisfied since $\mathcal{I}$ is a model for C_0: there is some $e_0 \in C_0^{\mathcal{I}}$ implying $C_0 \in t(e_0)$ and, due to the claim, we have $\hat{x}_M \geq 1$ for some mosaic M with $C_0 \in T_M$.

For $(E2)$, let T, T' be types, R a role with $\lim_R(T, T')$ and *not* $\lim_{\mathsf{Inv}(R)}(T', T)$, and fix some $e_M \in M^{\mathcal{I}}$ for each $M^{\mathcal{I}} \neq \emptyset$. We claim that the following (in)equalities hold which clearly implies (E2).

$$
\begin{aligned}
\sum_{\{M \mid T_M = T\}} L_M(R,T') \cdot \hat{x}_M = & \sum_{\{M \mid T_M = T \wedge M^{\mathcal{I}} \neq \emptyset\}} L_M(R,T') \cdot \hat{x}_M = \\
& \sum_{\{M \mid T_M = T \wedge M^{\mathcal{I}} \neq \emptyset\}} \#\mathsf{ch}(e_M, R) \cap T'^{\mathcal{I}} \cdot \hat{x}_M \leq \\
& \sum_{\{M \mid T_M = T' \wedge M^{\mathcal{I}} \neq \emptyset\}} \#\{e \in T^{\mathcal{I}} \mid \langle e_M, e\rangle \in \mathsf{Inv}(R)^{\mathcal{I}}\} \cdot \hat{x}_M = \\
& \sum_{\{M \mid T_M = T'\}} E_M(\mathsf{Inv}(R), T) \cdot \hat{x}_M
\end{aligned}
$$

The first equality is obvious. The second holds since, for each mosaic M and $e \in M^{\mathcal{I}}$, we have $\#(\mathsf{ch}(e,R) \cap T'^{\mathcal{I}}) = L_M(R,T')$ by definition of $M^{\mathcal{I}}$ (in particular this holds if $e = e_M$). The first inequality holds due to

- the claim,
- Property 1 in the definition of ch, which implies that $\mathsf{ch}(e_M, R) \cap T'^{\mathcal{I}}$ is a lower bound on $\#\{e' \in T'^{\mathcal{I}} \mid \langle e, e'\rangle \in R^{\mathcal{I}}\}$ for each $e \in M^{\mathcal{I}}$,
- a simple graph-theoretic reason: the number of R edges from $T^{\mathcal{I}}$ into $T'^{\mathcal{I}}$ is the same as the number of $\mathsf{Inv}(R)$ edges from $T'^{\mathcal{I}}$ into $T^{\mathcal{I}}$, and
- the fact that each $e \in M^{\mathcal{I}}$ with $T_M = T'$ has the same number of incoming R-edges by definition of $M^{\mathcal{I}}$.

Finally, the last equation is clearly valid by definition of the sets $M^{\mathcal{I}}$.

Equation $(E3)$ is satisfied with a similar yet simpler argument: let T, T' be types, R a role with $\lim_R(T,T')$ and $\lim_{\mathsf{Inv}(R)}(T',T)$, and fix some $e_M \in M^{\mathcal{I}}$ for each $M^{\mathcal{I}} \neq \emptyset$. Then we have

$$
\begin{aligned}
\sum_{\{M \mid T_M = T\}} E_M(R,T') \cdot \hat{x}_M = & \sum_{\{M \mid T_M = T \wedge M^{\mathcal{I}} \neq \emptyset\}} \#\{e' \in T'^{\mathcal{I}} \mid \langle e_M, e'\rangle \in R^{\mathcal{I}}\} \cdot \hat{x}_M = \\
& \sum_{\{M \mid T_M = T' \wedge M^{\mathcal{I}} \neq \emptyset\}} \#\{e \in T^{\mathcal{I}} \mid \langle e_M, e\rangle \in \mathsf{Inv}(R)^{\mathcal{I}}\} \cdot \hat{x}_M = \\
& \sum_{\{M \mid T_M = T'\}} E_M(\mathsf{Inv}(R), T) \cdot \hat{x}_M
\end{aligned}
$$

using similar arguments as for the (E2) case.

Now for the admissibility of our solution. Obviously it is a non-negative integer solution. For $(A1)$, consider types T, T' and a role R with $\lim_R(T,T')$, *not* $\lim_{\mathsf{Inv}(R)}(T',T)$, and

$$
\sum_{\{M \mid T_M = T'\}} E_M(\mathsf{Inv}(R), T) \cdot \hat{x}_M > 0.
$$

Hence there is, by definition of $M^{\mathcal{I}}$, some $\langle e', e\rangle \in \mathsf{Inv}(R)^{\mathcal{I}}$ with $e' \in {T'}^{\mathcal{I}}$ and $e \in T^{\mathcal{I}}$. The claim yields

$$\sum_{\{M \mid T_M = T\}} \hat{x}_M > 0,$$

and thus $(A1)$ is satisfied.

Finally, for $(A2)$, let M be a mosaic with $\hat{x}_M > 0$, $(\geqslant nR.C) \in T_M$, and

$$m = \sum_{\{T \mid C \in T\}} L_M(R,T) + \sum_{\{T \mid C \in T\}} E_M(R,T) < n.$$

Hence there is some $e_M \in T_M^{\mathcal{I}}$ and $e_1, \dots, e_n$ with $e_i \neq e_j$ for all $i \neq j$ and, for all $1 \leq i \leq n$, $\langle e_M, e_i\rangle \in R^{\mathcal{I}}$ and $e_i \in C^{\mathcal{I}}$. By definition of $M^{\mathcal{I}}$ and Property 4 of ch, $m < n$ implies that there is some e_ℓ such that *not* $\lim_{\mathsf{Inv}(R)}(t(e_M), t(e_\ell))$ and *not* $\lim_R(t(e_\ell), t(e_M))$. Since $C \in t(e_\ell)$, the claim yields

$$\sum_{\substack{\{M' \mid \ C \in T_{M'},\ \text{not } \lim_R(T_M, T_{M'}),\\ \text{and not } \lim_{\mathsf{Inv}(R)}(T_{M'}, T_M)\}}} \hat{x}_{M'} \geq 1,$$

and $(A2)$ is satisfied. □

Lemma 7 *If the equation system $\mathcal{E}_{C_0,\mathcal{T}}$ has an admissible solution, then C_0 is finitely satisfiable w.r.t. $\mathcal{T}$.*

Proof. Let $\{\hat{x}_M \mid M \text{ a mosaic}\}$ be an admissible solution of $\mathcal{E}_{C_0,\mathcal{T}}$. In what follows, we construct a finite interpretation $\mathcal{I}$ from this solution. For each mosaic M, fix a set $\hat{M}$ such that $\#\hat{M} = \hat{x}_M$. Moreover, set

$$P = \max\{n \mid (\geqslant n\ R\ C) \in \mathsf{cl}(C_0, \mathcal{T}) \text{ or } (\leqslant n\ R\ C) \in \mathsf{cl}(C_0, \mathcal{T})\}.$$

We define

$$\Delta^{\mathcal{I}} = \bigcup \hat{M} \times \{0, \dots, P-1\}.$$

In the following, we write $m(e) = M$ if $e \in \hat{M}$ and $t(e) = T$ if $e \in \hat{M}$ for some mosaic M with $T_M = T$. For each concept name $A \in \mathbf{C}$, we put

$$A^{\mathcal{I}} := \{(e,i) \in \Delta^{\mathcal{I}} \mid A \in t(e)\}.$$

Role names $R \in \mathbf{R}$ are harder to deal with. We start with defining some auxiliary notions. For each role $R \in \mathsf{rol}(C_o, \mathcal{T})$ and each pair of types $T, T' \in \mathsf{type}(C_0, \mathcal{T})$ such that $\lim_R(T,T')$ but not $\lim_{\mathsf{Inv}(R)}(T', T)$, fix a relation

$$\gamma^R_{T,T'} \subseteq \bigcup_{\{M \mid T_M = T\}} (\hat{M} \times \{0, \dots, P-1\}) \quad \times \bigcup_{\{M \mid T_M = T'\}} (\hat{M} \times \{0, \dots, P-1\})$$

such that

1. for each (e,i) with $t(e) = T$, we have
$$\#\{((e,i),(e',j)) \in \gamma^R_{T,T'}\} \geq L_{m(e)}(R,T');$$
2. for each (e,i) with $t(e) = T'$, we have
$$\#\{((e',j),(e,i)) \in \gamma^R_{T,T'}\} = E_{m(e')}(\mathsf{Inv}(R),T).$$

Let us show that such a relation exists: by (E2), there exists a mapping
$$f : \bigcup_{\{M \mid T_M = T\}} \hat{M} \quad \times \bigcup_{\{M \mid T_M = T'\}} \hat{M} \quad \rightarrow \quad \mathbb{N}$$
such that

a) for each e with $t(e) = T$, we have
$$\sum_{\{e' \mid t(e') = T'\}} f(e,e') = L_{m(e)}(R,T');$$
b) for each e' with $t(e') = T'$, we have
$$\sum_{\{e \mid t(e) = T\}} f(e,e') \leq E_{m(e')}(\mathsf{Inv}(R),T).$$

Using f, we define a relation
$$r \subseteq \bigcup_{\{M \mid T_M = T\}} (\hat{M} \times \{0,\ldots,P-1\}) \quad \times \bigcup_{\{M \mid T_M = T'\}} (\hat{M} \times \{0,\ldots,P-1\})$$
such that

c) for each (e,i) with $t(e) = T$, we have
$$\#\{((e,i),(e',j)) \in r\} = L_{m(e)}(R,T');$$
d) for each (e,i) with $t(e) = T'$, we have
$$\#\{((e',j),(e,i)) \in r\} \leq E_{m(e')}(\mathsf{Inv}(R),T).$$

More precisely, this is done by setting
$$r = \{((e,i),(e',i')) \mid i' = i + k \mod P \text{ for some } k \text{ with } 1 \leq k \leq f(e,e')\}.$$

Using the facts that f satisfies a) and b) and that $f(e,e') \leq P$ for any e,e', it is readily checked that r satisfies c) and d).[2] Finally, we can augment r to $\gamma^R_{T,T'}$ by performing, for each $(e,i) \in \mathsf{ran}(r)$ with
$$k := \#\{((e',j),(e,i)) \in r\} < E_{m(e')}(\mathsf{Inv}(R),T), \qquad (*)$$

[2]To see that $f(e,e') \leq P$ for any e,e', one may use Properties 2 to 4 of mosaics to show that, for any M, R, and T, we have $L_M(R,T) \leq P$ and $E_M(R,T) \leq P$.

the following step:

Obviously, $(*)$ implies $E_{m(e)}(\mathsf{Inv}(R), T) > 0$. Hence by (A1) there exists a mosaic M such that $\hat{M} \neq \emptyset$ and $T_M = T$. Fix an $e' \in \hat{M}$. Since $P \geq E_{m(e')}(\mathsf{Inv}(R), T)$, we may fix a set $X \subseteq \{0, \dots, P-1\}$ such that

$$\#X = E_{m(e')}(\mathsf{Inv}(R), T) - k$$

and $((e', j), (e, i)) \notin r$ for each $j \in X$. We augment r with the set $\{((e', j), (e, i)) \mid j \in X\}$.

We have now finished the construction of $\gamma^R_{T,T'}$. As an abbreviation, for each role R we define

$$\Gamma^R = \bigcup_{\substack{\{T,T' \in \mathsf{type}(C_0,\mathcal{T}) \mid \\ \lim_R(T,T') \text{ and not } \lim_{R^-}(T',T)\}}} \gamma^R_{T,T'}.$$

One more relation needs to be defined before the interpretation of role names can be given: for each role name R and each pair of types $T, T' \in \mathsf{type}(C_0, \mathcal{T})$ such that $\lim_R(T, T')$ and $\lim_{R^-}(T', T)$, fix a relation

$$\lambda^R_{T,T'} \subseteq \bigcup_{\{M \mid T_M = T\}} (\hat{M} \times \{0, \dots, P-1\}) \quad \times \bigcup_{\{M \mid T_M = T'\}} (\hat{M} \times \{0, \dots, P-1\})$$

such that

1. for each (e, i) with $t(e) = T$, we have

$$\#\{((e, i), (e', j)) \in \lambda^R_{T,T'}\} = E_{m(e)}(R, T');$$

2. for each (e, i) with $t(e) = T'$, we have

$$\#\{((e', j), (e, i)) \in \lambda^R_{T,T'}\} = E_{m(e')}(\mathsf{Inv}(R), T).$$

The exact construction is omitted since it is very similar to those of $\gamma^R_{T,T'}$: first construct an appropriate function f and then turn it into a relation r which can immediately be used as $\lambda^R_{T,T'}$ (the additional augmentation step of the construction of $\gamma^R_{T,T'}$ need not be applied).

As an abbreviation, for each role name R we define

$$\Lambda^R = \bigcup_{\substack{\{T,T' \in \mathsf{type}(C_0,\mathcal{T}) \mid \\ \lim_R(T,T') \text{ and } \lim_{R^-}(T',T)\}}} \lambda^R_{T,T'}.$$

Finally, for each role name R, set

$$\Omega^R := \{\langle (e, i), (e', i') \rangle \mid \text{ not } \lim_R(t(e), t(e')), \text{ and not } \lim_{R^-}(t(e'), t(e))\}$$

We are now ready to define the interpretation $R^\mathcal{I}$ of role names:

$$R^\mathcal{I} := \Omega^R \cup \Gamma^R \cup (\Gamma^{R^-})^\smile \cup \Lambda^R,$$

where $r^{\smile}$ denotes the converse of the relation r.

We now show that $\mathcal{I}$ is a model of C_0 w.r.t. $\mathcal{T}$. We first establish a technical claim:

Claim 1: For all $(e,i) \in \Delta^{\mathcal{I}}$, roles R, and types T' with $\lim_{\mathsf{Inv}(R)}(T', t(e))$, we have $\#\{(e',i') \mid ((e,i),(e',i')) \in R^{\mathcal{I}} \text{ and } t(e') = T'\} = E_{m(e)}(R, T')$.

Proof: In the following, a "witness" is an element $(e',i') \in \Delta^{\mathcal{I}}$ such that $((e,i),(e',i')) \in R^{\mathcal{I}}$ and $t(e') = T'$. We show that there are exactly $E_{m(e)}(R,T')$ witnesses by a case distinction.

- *Not* $\lim_R(t(e), T')$ and R is a role name. It is readily checked that then witnesses are added to $R^{\mathcal{I}}$ only through the $(\Gamma^{R^-})^{\smile}$ component, more precisely through the relation $\gamma^{R^-}_{(T',t(e))}$. By Property 2 of this relation, the number of witnesses added in this way is precisely $E_{m(e)}(R,T')$ as desired.
- *Not* $\lim_R(t(e), T')$ and $R = S^-$ for some role name S. Then pairs $((e',i'),(e,i))$ with $t(e') = T'$ are added to $S^{\mathcal{I}}$ only through the Γ^S component, more precisely through the relation $\gamma^S_{(T',t(e))}$. Again by Property 2 of this relation and since $S = \mathsf{Inv}(R)$, the number of tuples added in this way is precisely $E_{m(e)}(R,T')$ as desired.
- $\lim_R(t(e), T')$ and R is a role name. It is readily checked that witnesses are added to $R^{\mathcal{I}}$ only through the Λ^R component, more precisely through the relation $\lambda^R_{(T',t(e))}$. By Property 1 of this relation, the number of witnesses added in this way is precisely $E_{m(e)}(R,T')$ as desired.
- $\lim_R(t(e), T')$ and $R = S^-$ for some role name S. Then pairs $((e',i'),(e,i))$ with $t(e') = T'$ are added to $S^{\mathcal{I}}$ only through the Λ^S component, more precisely through the relation $\lambda^S_{(T',t(e))}$. By Property 2 of this relation, the number of tuples added in this way is precisely $E_{m(e)}(R,T')$ as desired.

Using Claim 1 just established, we can now prove another claim which is central for showing that $\mathcal{I}$ is a model of the input concept C_0 and the input TBox $\mathcal{T}$:

Claim 2: If $C \in t(e)$, then $(e,i) \in C^{\mathcal{I}}$ for each $i < P$.

The proof is by structural induction. Fix an $(e,i) \in \Delta^{\mathcal{I}}$ such that $C \in t(e)$.

- C is a concept name. Then $e \in C^{\mathcal{I}}$ follows from the definition of $\mathcal{I}$.
- $C = \neg D$. Since every concept in $\mathsf{cl}(C_0, \mathcal{T})$ is in NNF, D is a concept name. If $\neg D \in T$, then $D \notin T$ by definition of types. Thus $e \in (\neg D)^{\mathcal{I}}$ by definition of $\mathcal{I}$.
- $C = (\leqslant n\ R\ D)$. A "witness" for C is an element $(e',i') \in \Delta^{\mathcal{I}}$ such that $((e,i),(e',i')) \in R^{\mathcal{I}}$ and $D \in t(e')$. We show that there exist at most n witnesses.

By definition, we have $\lim_{\mathsf{Inv}(R)}(t(e'), t(e))$ for every witness (e', i'). Claim 1 thus yields an exact bound

$$m = \sum_{\{T \mid D \in T\}} E_{m(e)}(R, T)$$

for the number of witnesses. By Property 3 of mosaics, we have $n \geq m$.

- $C = (\geqslant n\ R\ D)$. Again, a "witness" for C is an element $(e', i') \in \Delta^{\mathcal{I}}$ such that $((e, i), (e', i')) \in R^{\mathcal{I}}$ and $D \in t(e')$. We need to show that there exist at least n witnesses.

  Firstly, for each type T with $\lim_R(t(e), T)$ and not $\lim_{\mathsf{Inv}(R)}(T, t(e))$, there are at least $L_{m(e)}(R, T)$ witnesses (e', i) with $t(e') = T$: if R is a role name, then the fact that $\Gamma^R \subseteq R^{\mathcal{I}}$ and Property 1 of $\gamma^R_{t(e),T}$ yield the desired result. Similarly, if $R = S^-$ for some role name S, then the facts that $(\Gamma^{S^-})^\smile = (\Gamma^R)^\smile \subseteq S^{\mathcal{I}}$ and $R^{\mathcal{I}} = (S^{\mathcal{I}})^\smile$ together with Property 1 of $\gamma^R_{t(e),T}$ yield the desired result.

  Together with Claim 1 and Properties 1 and 2 of mosaics, we thus have a lower bound

$$m = \sum_{\{T \mid D \in T\}} L_{m(e)}(R, T) + \sum_{\{T \mid D \in T\}} E_{m(e)}(R, T)$$

  on the number of witnesses. If $m \geq n$, then we are done. Otherwise, (A2) ensures that there exists a mosaic M such that $D \in T_M$, not $\lim_R(t(e), T_M)$, not $\lim_{\mathsf{Inv}(R)}(T_M, t(e))$, and $\hat{M} \neq \emptyset$. Since $\Omega^R \subseteq R^{\mathcal{I}}$, this yields

$$\#\{((e, i), (e', i')) \in R^{\mathcal{I}} \mid m(e') = M\} = P.$$

  Now $P \geq n$ by definition of P and we are done.

- $C = D \sqcap E$ and $C = D \sqcup E$. For this case, the claim follows immediately from the definition of types and the induction hypothesis.

As a consequence, $\mathcal{I}$ is a model of C_0 and $\mathcal{T} = \{\top \doteq C_{\mathcal{T}}\}$: by Equation (E1) and due to the fact that $\hat{x}_M > 0$ implies $\#\hat{M} > 0$, there is a mosaic M such that $C_0 \in T_M$ and $\#\hat{M} > 0$. Fix an $e \in \hat{M}$. Claim 2 implies that $(e, i) \in C_0^{\mathcal{I}}$ for $i < P$ and thus $\mathcal{I}$ is a model of C_0. Moreover, by definition of types, we have $C_{\mathcal{T}} \in T_M$ for each mosaic M. This fact together with Claim 2 implies that $\mathcal{I}$ is a model of $\mathcal{T}$. □

Since the number of mosaics is exponential in the size of C_0 and $\mathcal{T}$, the size of $\mathcal{E}_{C_0,\mathcal{T}}$ and of the admissibility condition is also exponential in the size of C_0 and $\mathcal{T}$. To prove an ExpTime upper bound for the finite satisfiability of $\mathcal{ALCQI}$-concepts, it thus remains to show that the existence of an admissible solution for the equation systems $\mathcal{E}_{C_0,\mathcal{T}}$ can be decided in deterministic polynomial time. Before we actually do this, we first fix some notation.

We assume linear inequalities to be of the form $\Sigma_i c_i x_i \geq b$. Such an inequality is called *positive* if $b \geq 0$. A system of linear inequalities is described by a tuple $(V, \mathcal{E})$, where V is a set of variables and $\mathcal{E}$ the set of inequalities. Such a system is called *simple* if all inequalities are positive and all coefficients are (possibly negative) integers.

A *side condition* for an inequality system $(V, \mathcal{E})$ is a constraint of the form

$$x > 0 \Longrightarrow x_1 + \cdots + x_\ell > 0, \text{ where } x, x_1, \ldots x_\ell \in V.$$

Let $(V, \mathcal{E})$ be an inequality system and I a set of side conditions for $(V, \mathcal{E})$. We say that $(V, \mathcal{E})$ admits an *I-admissible solution* if it admits a solution satisfying all constraints from I.

It is not hard to check that the inequality systems from Definition 5 are simple and that the conditions (A1) and (A2) can be polynomially transformed into side conditions:

- (E1) is already simple,
- (E2) can obviously transformed into $\sum \ldots - \sum \ldots \geq 0$,
- the equality (E3) is transformed into two inequalities of the form $\sum \ldots - \sum \ldots \geq 0$,
- each implication due to (A1) is transformed into polynomially many by using a separate side condition for each summand appearing in the premise (this works since we are interested in non-negative solutions only). Next, the coefficients on the left-hand sides of the premise are then omitted by dropping those side-conditions whose coefficient is zero and replacing all other coefficients with 1.
- (A2) is already in the form of a side condition.

In the following, we prove that the existence of a non-negative, integer, and I-admissible solution for a simple system of inequalities $(V, \mathcal{E})$ and a set of side conditions I can be decided in deterministic polynomial time. In the proof, we use a lemma that was established by Calvanese in [5] and builds on results of Papadimitriou [11]. We state this lemma for the sake of completeness.

Lemma 8 *Let $(V, \mathcal{E})$ be a system of $m = \#\mathcal{E}$ linear inequalities in $n = \#V$ variables, in which all coefficients and constants are from the interval $[-a; a]$ of integers. Then, if $(V, \mathcal{E})$ has a solution in $\mathbb{N}^n$, it also has one in $\{0, 1, \ldots, H(V, \mathcal{E})\}^n$, where $H(V, \mathcal{E}) = (n + m)(ma)^{2m+1}$.*

We can now establish the PTime upper bound.

Lemma 9 *Let $(V, \mathcal{E})$ be a simple equation system and I a set of side conditions for $(V, \mathcal{E})$. Then the existence of an integer, non-negative, and I-admissible solution for $(V, \mathcal{E})$ can be decided in (deterministic) time polynomial in $\#V + \#\mathcal{E} + \#I$.*

Proof. For a positive integer k, we use $\mathcal{E}_I(k)$ to denote the set of inequalities

$$\{x \leq k \cdot (x_1 + \cdots + x_k) \mid x > 0 \Longrightarrow x_1 + \cdots + x_k > 0 \in I\}.$$

It is readily checked that every non-negative solution of $(V, \mathcal{E} \cup \mathcal{E}_I(k))$ is a (non-negative and) I-admissible solution of $(V, \mathcal{E})$. We prove the following claim:

Claim: There is an integer $k_\mathcal{E}$ exponential in $\#V + \#\mathcal{E} + \#I$ such that $(V, \mathcal{E})$ admits a non-negative, integer, and I-admissible solution iff $(V, \mathcal{E} \cup \mathcal{E}_I(k_\mathcal{E}))$ admits a non-negative (rational) solution.

Proof: Let $n = \#V$, $m = \#\mathcal{E}$, and $r = \#I$. Then we choose

$$k_\mathcal{E} = (n + m + r)(m + r)^{2(m+r)+1}.$$

It remains to show that $k_\mathcal{E}$ is as required:

For the "if" direction, let $\overline{S}$ be a non-negative solution of $(V, \mathcal{E} \cup \mathcal{E}_I(k_\mathcal{E}))$. As noted above, $\overline{S}$ is also a (non-negative and) I-admissible solution of $(V, \mathcal{E})$. Since all inequations in $(V, \mathcal{E})$ are positive, we can convert $\overline{S}$ into an integer solution by multiplying $\overline{S}$ with the smallest common multiplier of the denominators in $\overline{S}$.

Now for the "only if" direction: assume that there exists an integer, non-negative, and I-admissible solution $\overline{S}$ of $(V, \mathcal{E})$. It is readily checked that this implies the existence of a set $P \subseteq V$ such that $\overline{S}$ is also an (integer and non-negative) solution of the system $(V, \mathcal{E} \cup \mathcal{E}_P)$, where

$$\begin{aligned} \mathcal{E}_P &= \{x > 0 \mid x \in P\} \cup \\ &\quad \{x_1 + \cdots + x_j > 0 \mid x \in P \text{ and } x > 0 \Longrightarrow x_1 + \cdots + x_j > 0 \in I\} \cup \\ &\quad \{x = 0 \mid x \in V \setminus P\}. \end{aligned}$$

By Lemma 8, the existence of $\overline{S}$ implies the existence of a non-negative integer solution $\overline{S}'$ of $(V, \mathcal{E} \cup \mathcal{E}_P)$ which is bounded by $h_P = H(V, \mathcal{E} \cup \mathcal{E}_P)$. It is easily seen that the solution $\overline{S}'$ is also an (integer and non-negative) solution of $(V, \mathcal{E} \cup \mathcal{E}_I(n))$ for any $n \geq h_P$. It remains to note that, since $\mathcal{E}_P$ contains one inequality for each variable in V and at most one inequality for each implication in I, we have $h_P \leq k_\mathcal{E}$.

In view of the claim just established, it is now easy to show that the existence of a non-negative integer and I-admissible solution for a simple system of inequalities $(V, \mathcal{E})$ and a set of side conditions I can be decided in deterministic polynomial time: we may clearly view $(V, \mathcal{E} \cup \mathcal{E}_I(k_\mathcal{E}))$ as a linear programming problem. Since $k_\mathcal{E}$ is exponential in $\#V + \#\mathcal{E} + \#I$, the binary representation of $k_\mathcal{E}$ is polynomial in $\#V + \#\mathcal{E} + \#I$. Thus, the existence of a rational (non-negative) solution for $(V, \mathcal{E} \cup \mathcal{E}_I(k_\mathcal{E}))$ can be checked in (deterministic) time polynomial in $\#V + \#\mathcal{E} + \#I$ [14]. □

Putting together Lemmas 6, 7, and 9, we obtain the ExpTime upper bound. The corresponding lower bound is a consequence of the ExpTime-hardness of unrestricted satisfiability of $\mathcal{ALC}$ w.r.t. TBoxes [7; 12; 13] and the fact that this DL has the finite model property.

Theorem 10 *Finite satisfiability of $\mathcal{ALCQI}$-concepts w.r.t. TBoxes is* ExpTime-*complete if numbers are coded in unary.*

If numbers in number restrictions are coded binarily, the algorithm developed in this section does no longer yield an ExpTime upper bound: in this case, the number of mosaics is double exponential in the size of the input concept and TBox. Since it is not clear whether and how the presented algorithm can be modified in order to yield an ExpTime upper bound for the case of binary coding, we resort to a different approach to attacking this problem: in the next section, we reduce finite $\mathcal{ALCQI}$-satisfiability to the finite satisfiability of $\mathcal{ALCFI}$-concepts. Since the employed reduction is polynomial, in this way we obtain an ExpTime upper bound for the finite satisfiability of $\mathcal{ALCQI}$-concepts w.r.t. TBoxes, even if numbers are coded in binary.

4 Binary Coding of Numbers

In this section, we prove that finite $\mathcal{ALCQI}$-concept satisfiability w.r.t. TBoxes is decidable in ExpTime even if numbers are coded in binary. The proof is by a reduction to finite $\mathcal{ALCFI}$-concept satisfiability w.r.t. TBoxes. Since, in the case of $\mathcal{ALCFI}$, the size of numbers appearing in number restrictions is constant (independently of the coding), the results presented in the previous section imply that finite $\mathcal{ALCFI}$-concept satisfiability w.r.t. TBoxes is ExpTime-complete. Thus, this logic is a suitable target for reduction. In contrast to existing reductions of $\mathcal{ALCQI}$ to $\mathcal{ALCFI}$, which only work in the case of potentially infinite models (such as the one presented in [6]), we have to take special care to deal with finite models.

Before we go into technical details, let us describe the intuition behind the reduction. The general idea is to replace counting via qualified number restrictions with counting via concept names: to count up to a number n, we reserve concept names $B_0, \dots, B_{\lceil \log(n) \rceil}$ representing the bits of numbers between 0 and n. For the actual counting, we can then use well-known (propositional logic) formulas that encode incrementation. But how can we use this approach to count the number of role successors? Intuitively, we rearrange the successors of each domain element in a way that allows to replace qualifying number restrictions with the functionality of roles provided by $\mathcal{ALCFI}$ and counting via concept names. Consider, for example, the domain element x and its R-successors displayed on the left-hand side of Figure 2. Ignoring the "direct" R-successors of x on the right-hand side for a moment, it is obvious that the R-successors are rearranged along a path that is built using an auxiliary role L_R. Employing the $(\leqslant 1\ R\ \top)$ constructor of $\mathcal{ALCFI}$, each node on this path has precisely one

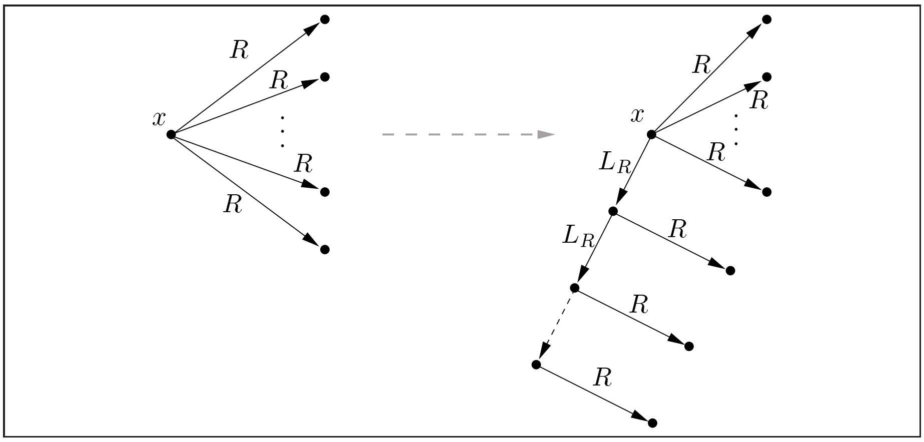

Figure 2: Representing role successor relationships.

L_R-predecessor and at most one L_R-successor. The counting via concept names is then performed along the domain elements on L_R-paths.

However, we cannot gather all original R-successors of x on the L_R-path. The reason for this is as follows: assume we are at some domain element on the L_R-path descending from x and move along this domain element's outgoing R-edge. Then we reach either a "real" domain element or arrive on an $\mathsf{Inv}(R)$-path. If the latter is the case, we have to ensure that, moving up the $\mathsf{Inv}(R)$-path, we will finally reach a "real" domain element. To do this, we count the lengths of auxiliary paths via concept names:[3] once we have moved up to node 0 of the path, its predecessor must be "real". Since, however, we do not know how many R-successors an object had in the original model, we do not know how many bits to reserve for this counting. The solution is to gather only those R-successors of x on the R-path which are constrained by a $(\leqslant n\ R\ C)$ concept applying to x or which are witnesses for a $(\geqslant n\ R\ C)$ concept applying to x—this helps since the number of such domain elements is known in advance. All other domain elements remain "direct" successors of x.

Fix an $\mathcal{ALCQI}$-concept C and an $\mathcal{ALCQI}$-TBox $\mathcal{T}$ whose finite satisfiability is to be decided. In the following, we use $\mathsf{cnam}(C,\mathcal{T})$ to denote the set of concept names appearing in C and $\mathcal{T}$, $\mathsf{rnam}(C,\mathcal{T})$ to denote the set of role names appearing in C and $\mathcal{T}$, and, as above, $\mathsf{rol}(C,\mathcal{T})$ to denote the set

$$\mathsf{rnam}(C,\mathcal{T}) \cup \{R^- \mid R \in \mathsf{rnam}(C,\mathcal{T})\}.$$

W.l.o.g., we assume C and $\mathcal{T}$ to be in NNF. In order to translate C and $\mathcal{T}$ to $\mathcal{ALCFI}$, we need to introduce some additional concept and role names:

1. a fresh (i.e., not appearing in C or $\mathcal{T}$) concept name Real;

[3]this counter is a different one than the ones mentioned above

2. for each $R \in \mathsf{rnam}(C, \mathcal{T})$, a fresh concept name H_R and a fresh role name L_R;

3. for each concept $D \in \mathsf{cl}(C, \mathcal{T})$ of the form $(\bowtie n\ R\ C)$, where $\bowtie$ is used as a placeholder for $\geqslant$ or $\leqslant$, we reserve a fresh concept name X_D;

4. again for each concept $D \in \mathsf{cl}(C, \mathcal{T})$ of the form $(\bowtie n\ R\ C)$, we reserve additional fresh concept names $B^{\bowtie n}_{C,R,0}, \ldots, B^{\bowtie n}_{C,R,k}$, where $k = \lceil \log(n+1) \rceil$;

5. for each role $R \in \mathsf{rol}(C, \mathcal{T})$, we reserve fresh concept names $B_{R,0}, \ldots, B_{R,k}$, where $k = \lceil \log(\mathsf{depth}_R) \rceil$ and

$$\mathsf{depth}_R = \sum_{(\bowtie\ n\ R\ C) \in \mathsf{cl}(C,\mathcal{T})} n;$$

The concept names $B_{R,i}$ are used to count the length of auxiliary L_R paths as described above. The concept names $B^{\bowtie n}_{C,R,0}$ are also used for counting: for an $\mathcal{ALCQI}$-concept $(\bowtie\ n\ R\ C)$, they are used to count the "occurrence" of R-successors in C along the L_R path, and will thus replace the $\mathcal{ALCQI}$-concept $(\bowtie n\ R\ C)$. Note that the number of newly introduced concept and role names is polynomial in the size of C and $\mathcal{T}$. We will use $\overline{B^{\bowtie n}_{C,R}}$ to refer to the number encoded by the concept names

$$B^{\bowtie n}_{C,R,0}, \ldots, B^{\bowtie n}_{C,R,\lceil \log(n+1) \rceil}$$

and $\overline{B_R}$ to refer to the number encoded by the concept names

$$B_{R,0}, \ldots, B_{R,\lceil \log(\mathsf{depth}_R) \rceil}.$$

Moreover, we will use the following abbreviations:

- $(\overline{B_R} = i)$ to denote the $\mathcal{ALCFI}$-concept expressing that $\overline{B_R}$ equals i (and similar for $\overline{B^{\bowtie n}_{C,R}} = i$ and the comparisons "$<$" and "$>$");

- $\mathsf{incr}(\overline{B_R}, S)$ to denote the $\mathcal{ALCFI}$-concept expressing that, for all S-successors, the number $\overline{B_R}$ is incremented by 1 modulo depth_R (and similar for $\mathsf{incr}(\overline{B^{\bowtie n}_{C,R}}, S)$). More precisely, these concepts are defined as follows (we use concepts $C \to D$ as an abbreviation for $\neg C \sqcup D$):

$$\begin{aligned}
&(B_{R,0} \to \forall S.\neg B_{R,0}) \sqcap (\neg B_{R,0} \to \forall S.B_{R,0}) \sqcap \\
&\bigsqcap_{k=1..n} \big(\bigsqcap_{j=0..k-1} B_{R,j} \big) \to \big((B_{R,k} \to \forall S.\neg B_{R,k}) \sqcap (\neg B_{R,k} \to \forall S.B_{R,k}) \big) \sqcap \\
&\bigsqcap_{k=1..n} \big(\bigsqcup_{j=0..k-1} \neg B_{R,j} \big) \to \big((B_{R,k} \to \forall S.B_{R,k}) \sqcap (\neg B_{R,k} \to \forall S.\neg B_{R,k}) \big).
\end{aligned}$$

We can now inductively define a translation $\gamma(C)$ of the concept C into a Boolean formula (which is also an $\mathcal{ALCFI}$-concept):

$$\begin{aligned}
\gamma(A) &:= A \\
\gamma(\neg C) &:= \neg\gamma(C) \\
\gamma(C \sqcap D) &:= \gamma(C) \sqcap \gamma(D) \\
\gamma(C \sqcup D) &:= \gamma(C) \sqcup \gamma(D) \\
\gamma(\geqslant n\ R\ C) &:= X_{(\geqslant n\ R\ C)} \\
\gamma(\leqslant n\ R\ C) &:= X_{(\leqslant n\ R\ C)}
\end{aligned}$$

Now set $\sigma(C) := \gamma(C) \sqcap \mathsf{Real}$ and, for $\mathcal{T} = \{\top \doteq C_\mathcal{T}\}$,

$$\sigma(\mathcal{T}) := \{\top \doteq \mathsf{Real} \rightarrow \gamma(C_\mathcal{T})\} \cup \mathsf{Aux}(C, \mathcal{T})$$

where the TBox $\mathsf{Aux}(C, \mathcal{T})$ is defined in Figure 3 in which we use $D \sqsubseteq E$ as abbreviation for $\top \doteq D \rightarrow E$, and in which all $\bigsqcup_{\dots}$, $\bigsqcap_{\dots}$ range over all concepts in $\mathsf{cl}(C, \mathcal{T})$ of the form specified. In what follows, we will use Ei to refer to the i'th concept equation and E$i.j$ to refer to its j'th line.

Equations E1, E2, and E3 ensure the behaviour sketched above of Real, H_R, and the counting concepts B_R and $B^{\bowtie n}_{D,R}$. Equation E5 ensures that the counting concepts $B^{\bowtie n}_{D,R}$ are updated correctly along an L_R path. To guarantee that a "real" element d satisfies "number restrictions" $X_{(\bowtie\, n\; R\; D)}$, E4 ensures that we see enough R-successors in D for at least restrictions $(\geqslant n\; R\; D)$ along an L_R path starting at d, whereas E6 guarantees that we do not see too many such successors along an L_R path for at most restrictions $(\leqslant n\; R\; D)$.

Lemma 11 *C is finitely satisfiable w.r.t. $\mathcal{T}$ iff $\sigma(C)$ is finitely satisfiable w.r.t. $\sigma(\mathcal{T})$.*

Proof. Let us start with the "if" direction. Hence, assume that $\sigma(C)$ is finitely satisfiable w.r.t. $\sigma(\mathcal{T})$. The proof strategy is to take a finite model of $\sigma(C)$ and $\sigma(\mathcal{T})$ and transform it into a finite model of C and $\mathcal{T}$. However, we cannot take an arbitrary model for this purpose, but need to select a special, so-called singular one: let $\mathcal{I}$ be a model of $\sigma(C)$ and $\sigma(\mathcal{T})$. For each domain element $d \in \mathsf{Real}^\mathcal{I}$ and each $R \in \mathsf{rol}(C, \mathcal{T})$, we inductively define a sequence of domain elements $h_0^{d,R}, \dots, h_{\ell_{d,R}}^{d,R}$ as follows:

- set $h_0^{d,R} = d$;
- set $h_{i+1}^{d,R}$ to the L_R-successor of $h_i^{d,R}$ (which is unique due to E1.3) if it exists. Otherwise, $\ell_{d,R} = i$.

The constructed sequence is finite due to the use of the $\overline{B_R}$ counter in E2.2, E3.3, and E3.5. The model $\mathcal{I}$ is called *singular* if, for all roles $R \in \mathsf{rol}(C, \mathcal{T})$, nodes $d \in \mathsf{Real}^\mathcal{I}$, and $i < j \leq \ell_{d,R}$, we have

$$\{e \mid (h_i^{d,R}, e) \in R^\mathcal{I}\} \cap \{e \mid (h_j^{d,R}, e) \in R^\mathcal{I}\} = \emptyset.$$

Claim 1. If $\sigma(C)$ is finitely satisfiable w.r.t. $\sigma(\mathcal{T})$, then there is a finite, singular model of $\sigma(C)$ and $\sigma(\mathcal{T})$.

Proof: Let $\mathcal{I}$ be a finite model for $\sigma(C)$ and $\sigma(\mathcal{T})$. Fix an injective mapping δ from $\Delta^\mathcal{I}$ to $\{0, \dots, (\#\Delta^\mathcal{I} - 1)\}$. Then we construct a new (finite) interpretation $\mathcal{J}$ by copying $\mathcal{I}$ sufficiently often and "bending R edges" from one copy of $\mathcal{I}$ into others. More precisely, $\mathcal{J}$ is defined as follows:

- $\Delta^\mathcal{J} := \{\langle d, i\rangle \mid d \in \Delta^\mathcal{I} \text{ and } i < \#\Delta^\mathcal{I}\}$;
- $A^\mathcal{J} := \{\langle d, i\rangle \in \Delta^\mathcal{J} \mid d \in A^\mathcal{I}\}$ for all concept names A;

$$
\begin{aligned}
\top \;\doteq\; & \bigsqcap_{R\in\mathsf{rol}(C,\mathcal{T})} \forall R.(\mathsf{Real} \sqcup H_{\mathsf{Inv}(R)}) \sqcap \\
& \qquad \forall L_R.H_R \sqcap \\
& \qquad (\leqslant 1\; L_R\; \top) \sqcap \\
& \bigsqcap_{(\bowtie\; n\; R\; D)} \big(X_{(\bowtie\; n\; R\; D)} \leftrightarrow \forall L_R.X_{(\bowtie\; n\; R\; D)}\big) \sqcap \\
& \bigsqcap_{R\in\mathsf{rol}(C,\mathcal{T})} \bigsqcap_{A\in\mathsf{cnam}(C,\mathcal{T})} (A \leftrightarrow \forall L_R.A) \sqcap \\
& \bigsqcap_{D} \neg\gamma(D) \to \gamma(\dot{\neg}(D))
\end{aligned}
$$

$$
\begin{aligned}
\mathsf{Real} \;\sqsubseteq\; & \bigsqcap_{R\in\mathsf{rol}(C,\mathcal{T})} \neg H_R \sqcap \\
& \qquad \forall L_R.(\overline{B_R} = 0) \sqcap \\
& \qquad (\leqslant 0\; L_R^-\; \top) \sqcap \\
& \bigsqcap_{(\bowtie\; n\; R\; D)} \big(X_{(\bowtie\; n\; R\; D)} \to \forall L_R.(\overline{B_{D,R}^{\bowtie n}} = 0)\big) \sqcap \\
& \bigsqcap_{(\leqslant\; n\; R\; D)} \big(X_{(\leqslant n\; R\; D)} \to \forall R.\neg\gamma(D)\big) \sqcap \\
& \bigsqcap_{\substack{(\geqslant\; n\; R\; D)\\ \text{with } n>0}} \big(X_{(\geqslant n\; R\; D)} \to \exists L_R.\top\big)
\end{aligned}
$$

$$
\begin{aligned}
H_R \;\sqsubseteq\; & (=1\; R\; \top) \sqcap \\
& (=1\; L_R^-\; \top) \sqcap \\
& \mathsf{incr}(\overline{B_R}, L_R) \sqcap \\
& (\overline{B_R} = 0) \to \exists L_R^-.\mathsf{Real} \sqcap \\
& (\overline{B_R} = (\mathsf{depth}_R - 1)) \to (\leqslant 0\; L_R\; \top)
\end{aligned}
$$

$$
H_R \sqsubseteq \bigsqcap_{(\geqslant n\; R\; D)} (X_{(\geqslant\; n\; R\; D)} \sqcap \overline{B_{D,R}^{\geqslant n}} < n \sqcap \forall R.\neg\gamma(D) \sqcap \forall L_R.\bot) \to \bot)
$$

$$
H_R \sqsubseteq \bigsqcap_{(\bowtie\; n\; R\; D)} (\exists R.\gamma(D) \to \mathsf{incr}(\overline{B_{D,R}^{\bowtie n}}, L_R))
$$

$$
H_R \sqsubseteq \bigsqcap_{(\leqslant\; n\; R\; D)} \big((X_{(\leqslant\; n\; R\; D)} \sqcap \overline{B_{D,R}^{\leqslant n}} = n \sqcap \exists R.\gamma(D)) \to \bot\big)
$$

Figure 3: The TBox $\mathsf{Aux}(C, \mathcal{T})$.

- $L_R^{\mathcal{J}} := \{(\langle d, i\rangle, \langle e, i\rangle) \in \Delta^{\mathcal{J}} \times \Delta^{\mathcal{J}} \mid (d, e) \in L_R^{\mathcal{I}}\}$ for all role names L_R with $R \in \mathsf{rol}(C, \mathcal{T})$;
- $R^{\mathcal{J}} := \{(\langle d, i\rangle, \langle e, (\delta(d) + i \bmod \#\Delta^{\mathcal{I}})\rangle) \mid (d, e) \in R^{\mathcal{I}}\}$ for all role names R appearing in C or $\mathcal{T}$.

It is straightforward to check that $\mathcal{J}$ is a singular model for $\sigma(C)$ and $\sigma(\mathcal{T})$, which finishes the proof of Claim 1.

Before we continue, let us state an important property of the sequences of domain elements $h_0^{d,R}, \ldots, h_{\ell_{d,R}}^{d,R}$:

($*$) Let $d, e \in \mathsf{Real}^{\mathcal{I}}$ such that $d \neq e$ and $R \in \mathsf{rol}(C, \mathcal{T})$. Then, for all $i \leq \ell_{d,R}$, and $j \leq \ell_{e,R}$, we have $h_i^{d,R} \neq h_j^{e,R}$. This is an easy consequence of the choice of the elements $h^{d,R}$ and $h^{e,R}$ together with E2.3 and E3.2.

Now let $\mathcal{I}$ be a singular, finite model for $\sigma(C)$ and $\sigma(\mathcal{T})$ and fix, for each $d \in \mathsf{Real}^{\mathcal{I}}$ and $R \in \mathsf{rol}(C, \mathcal{T})$, a sequence of domain elements $h_0^{d,R}, \ldots, h_{\ell_{d,R}}^{d,R}$ as above. We use $\mathcal{I}$ to define an interpretation $\mathcal{J}$ as follows:

$$\begin{array}{lcl}
\Delta^{\mathcal{J}} & := & \mathsf{Real}^{\mathcal{I}} \\
A^{\mathcal{J}} & := & A^{\mathcal{I}} \cap \mathsf{Real}^{\mathcal{I}} \\
R^{\mathcal{J}} & := & \{(d, e) \in \Delta^{\mathcal{J}} \times \Delta^{\mathcal{J}} \mid \exists i \leq \ell_{d,R},\ k \leq \ell_{e,\mathsf{Inv}(R)} : (h_i^{d,R}, h_j^{e,\mathsf{Inv}(R)}) \in R^{\mathcal{I}}\}
\end{array}$$

It remains to establish the following claim:

Claim 2. For all $d \in \Delta^{\mathcal{J}}$ and $D \in \mathsf{cl}(C, \mathcal{T})$, $d \in \gamma(D)^{\mathcal{I}}$ implies $d \in D^{\mathcal{J}}$.

For assume that Claim 2 is true. Since $\mathcal{J}$ is a model of $\sigma(C)$, by definition of σ there exists a $d \in (\gamma(C) \sqcap \mathsf{Real})^{\mathcal{I}}$. Clearly we have $d \in \Delta^{\mathcal{J}}$ and thus Claim 2 yields $d \in C^{\mathcal{J}}$. Hence, $\mathcal{J}$ is a model of C. By definition of $\sigma(\mathcal{T})$ and the semantics, we have $\mathsf{Real}^{\mathcal{I}} = (\gamma(C_{\mathcal{T}}) \sqcap \mathsf{Real})^{\mathcal{I}}$. Together with Claim 2 and definition of $\mathcal{J}$, we obtain $\Delta^{\mathcal{J}} = C_{\mathcal{T}}^{\mathcal{J}}$ and thus $\mathcal{J}$ is a model of $\mathcal{T}$.

We prove Claim 2 by induction on the *norm* $||\cdot||$ of concepts D which is defined inductively as follows:

$$\begin{array}{rcl cl}
||A|| & := & ||\neg A|| & := & 0 \text{ for } A \text{ concept name} \\
||C_1 \sqcap C_2|| & := & ||C_1 \sqcup C_2|| & := & 1 + ||C_1|| + ||C_2|| \\
||(\geqslant n\ R\ D)|| & := & ||(\leqslant n\ R\ D)|| & := & 1 + ||D||
\end{array}$$

Let $d \in \Delta^{\mathcal{J}} \cap \gamma(D)^{\mathcal{I}}$ for some $D \in \mathsf{cl}(C, \mathcal{T})$. Then $d \in \mathsf{Real}^{\mathcal{I}}$. Since C and $\mathcal{T}$ are in NNF, D is also in NNF. We only treat the interesting cases:

- Let $D = (\geqslant n\ R\ E)$ and $d \in \gamma(D)^{\mathcal{I}} = (X_{(\geqslant n\ R\ E)})^{\mathcal{I}}$. By E1.4, we have $h_i^{d,R} \in (X_{(\geqslant n\ R\ E)})^{\mathcal{I}}$ for $1 \leq i \leq \ell_{d,R}$. Hence, by exploiting the counter $\overline{B_{E,R}^{\geqslant n}}$ and its use in E2.4, E2.6, E4, and E5, it is straightforward to show that there exist a subset $I \subseteq \{1, \ldots, \ell_{d,R}\}$ of cardinality at least n such that, for each $i \in I$, there exists an $e_i \in \Delta^{\mathcal{I}}$ such that $(h_i^{d,R}, e_i) \in R^{\mathcal{I}}$ and

$e_i \in \gamma(E)^{\mathcal{I}}$. Due to singularity, we have that $i \neq j$ implies $e_i \neq e_j$ for all $i, j \in I$. By E1.1, we have $e_i \in \mathsf{Real}^{\mathcal{I}}$ or $e_i \in H_{\mathsf{Inv}(R)}$ for all $i \in I$. Using the counter $\overline{B_{\mathsf{Inv}(R)}}$ and E3.2, E3.3, E3.4, it is thus readily checked that, for each $i \in I$, there exists an $f_i \in \Delta^{\mathcal{I}}$ such that $f_i \in \mathsf{Real}^{\mathcal{I}}$ and e_i appears among the $h_0^{f_i,\mathsf{Inv}(R)}, \ldots, h_{\ell_{f_i,\mathsf{Inv}(R)}}^{f_i,\mathsf{Inv}(R)}$. By Property $(*)$, $i \neq j$ implies $f_i \neq f_j$ for all $i, j \in I$. By definition of $\mathcal{J}$, we have $(d, f_i) \in R^{\mathcal{J}}$ for each $i \in I$:

- if R is a role name, then this is an immediate consequence of the definition of $\mathcal{J}$;
- if $R = S^-$ for some role name S, then $(f_i, d) \in S^{\mathcal{J}}$ by definition of $\mathcal{J}$. The semantics yields $(d, f_i) \in R^{\mathcal{J}}$.

It thus remains to verify that $f_i \in E^{\mathcal{J}}$: clearly, $\gamma(E)$ is a Boolean formula over the set of concept names

$$\mathsf{cnam}(C, \mathcal{T}) \cup \{X_F \mid F = (\bowtie n\ R\ F') \in \mathsf{cl}(C, \mathcal{T})\}.$$

Since $e_i \in \gamma(E)^{\mathcal{I}}$, E1.4 and E1.5 thus yield $f_i \in \gamma(E)^{\mathcal{I}}$ for each $i \in I$. Since $f_i \in \mathsf{Real}^{\mathcal{I}}$, it remains to apply the induction hypothesis.

- Let $D = (\leqslant n\ R\ E)$ and $d \in \gamma(D)^{\mathcal{I}} = (X_{(\leqslant n\ R\ E)})^{\mathcal{I}}$. Assume that there exists a subset $W \subseteq \Delta^{\mathcal{J}}$ of cardinality greater than n such that, for each $e \in W$, we have $(d, e) \in R^{\mathcal{J}}$ and $e \in E^{\mathcal{J}}$. By definition of $\mathcal{J}$, this implies that, for each $e \in W$, there are $s_e \leq \ell_{d,R}$ and $t_e \leq \ell_{e,R}$ such that $(h_{s_e}^{d,R}, h_{t_e}^{e,\mathsf{Inv}(R)}) \in R^{\mathcal{I}}$:

  - if R is a role name, then this is an immediate consequence of the definition of $\mathcal{J}$;
  - if $R = S^-$ for some role name S, then $(d, e) \in R^{\mathcal{I}}$ implies $(e, d) \in S^{\mathcal{I}}$. By definition of $\mathcal{J}$, this means that there are $s_e \leq \ell_{d,R}$ and $t_e \leq \ell_{e,R}$ such that $(h_{t_e}^{e,S}, h_{s_e}^{d,R}) \in S^{\mathcal{I}}$. By semantics and since $S = \mathsf{Inv}(R)$, we obtain $(h_{s_e}^{d,R}, h_{t_e}^{e,\mathsf{Inv}(R)}) \in R^{\mathcal{I}}$.

  We clearly have $W \subseteq \mathsf{Real}^{\mathcal{I}}$. We prove the following three Properties:

  1. $e \neq e'$ implies $h_{s_e}^{d,R} \neq h_{s_{e'}}^{d,R}$ for all $e, e' \in W$. By Property $(*)$, $e \neq e'$ implies $h_{t_e}^{e,\mathsf{Inv}(R)} \neq h_{t_{e'}}^{e',\mathsf{Inv}(R)}$ for all $e, e' \in W$. Thus, E3.1 yields $h_{s_e}^{d,R} \neq h_{s'_e}^{d,R}$ if $e \neq e'$.
  2. $h_{t_e}^{e,\mathsf{Inv}(R)} \in \gamma(E)^{\mathcal{I}}$ for each $e \in W$. Suppose that $e \notin \gamma(E)^{\mathcal{I}}$. Then $e \in (\neg\gamma(E))^{\mathcal{I}}$ and, by E1.6, $e \in \gamma(\dot{\neg}E)^{\mathcal{I}}$. Since $e \in \mathsf{Real}^{\mathcal{I}}$ and we are performing induction over the norm of concepts rather than standard structural induction, the induction hypothesis yields $e \in (\dot{\neg}E)^{\mathcal{J}}$, a contradiction to $e \in E^{\mathcal{J}}$. Thus, $e \in \gamma(E)^{\mathcal{I}}$. Since $\gamma(E)$ is a Boolean formula, it follows from E1.4 and E1.5 that $h_{t_e}^{e,\mathsf{Inv}(R)} \in \gamma(E)$.

3. $s_e \neq 0$ for all $e \in W$. For assume that $s_e = 0$. Then $h^{d,R}_{s_e} = d$. By E2.5 and since $d \in (X_{(\leqslant n\ R\ E)})^{\mathcal{I}}$ and $(d, h^{e,\mathsf{Inv}(R)}_{t_e}) \in R^{\mathcal{I}}$, this yields $h^{e,\mathsf{Inv}(R)}_{t_e} \in (\neg(\gamma(E)))^{\mathcal{I}}$ in contradiction to Property 2.

Properties 1 to 3 imply the existence of a subset $I \subseteq \{1, \ldots, \ell_{d,R}\}$ of cardinality greater than n such that, for each $i \in I$, there exists an $e \in \Delta^{\mathcal{I}}$ with $(h^{d,R}_i, e) \in R^{\mathcal{I}}$ and $e \in \gamma(E)^{\mathcal{I}}$. Exploiting the counter $\overline{B^{\leqslant n}_{E,R}}$ and its use in E2.4, E5, and E6, it is readily checked that this is a contradiction to $\mathcal{I}$ being a model for $\mathsf{Aux}(C, \mathcal{T})$.

Now for the "only if" direction: let $\mathcal{I}$ be a finite model of C and $\mathcal{T}$. For each $d \in \Delta^{\mathcal{I}}$ and each $R \in \mathsf{rol}(C, \mathcal{T})$, fix a subset $W_{d,R} \subseteq \Delta^{\mathcal{I}}$ of cardinality at most depth_R such that the following conditions are satisfied:

1. $(d, e) \in R^{\mathcal{I}}$ for all $e \in W_{d,R}$;
2. for all $(\geqslant n\ R\ D) \in \mathsf{cl}(C, \mathcal{T})$ with $d \in (\geqslant n\ R\ D)^{\mathcal{I}}$, we have
$$\#\{e \in W_{d,R} \mid e \in D^{\mathcal{I}}\} \geq n;$$
3. for all $(\leqslant n\ R\ D) \in \mathsf{cl}(C, \mathcal{T})$ with $d \in (\leqslant n\ R\ D)^{\mathcal{I}}$, we have
$$\{e \in \Delta^{\mathcal{I}} \mid (d, e) \in R^{\mathcal{I}} \text{ and } e \in D^{\mathcal{I}}\} \subseteq W_{d,R};$$

Using the semantics and the definition of depth_R, it is easy to show that such subsets indeed exist. Next, fix a linear ordering on $W_{d,R}$, i.e., an injective mapping $\nu_{d,R} : W_{d,R} \longrightarrow \{0, \ldots, \#W_{d,R} - 1\}$. We use these mappings to define a finite model $\mathcal{J}$ of $\sigma(C)$ w.r.t. $\sigma(\mathcal{T})$ as follows:

- $\Delta^{\mathcal{J}} = \Delta^{\mathcal{I}} \cup \{x_{d,R,e} \mid d \in \Delta^{\mathcal{I}},\ R \in \mathsf{rol}(C, \mathcal{T}),\ \text{and } e \in W_{d,R}\}$;
- $A^{\mathcal{J}} = A^{\mathcal{I}} \cup \{x_{d,R,e} \mid d \in A^{\mathcal{I}},\ R \in \mathsf{rol}(C, \mathcal{T}), \text{ and } e \in W_{d,R}\}$ for all $A \in \mathsf{cnam}(C, \mathcal{T})$;
- $X^{\mathcal{J}}_{(\bowtie n\ R\ D)} = (\bowtie n\ R\ D)^{\mathcal{I}} \cup \{x_{d,R,e} \mid d \in (\bowtie n\ R\ D)^{\mathcal{I}} \text{ and } e \in W_{d,R}\}$ for all $(\bowtie n\ R\ D) \in \mathsf{cl}(C, \mathcal{T})$;
- $\mathsf{Real}^{\mathcal{J}} = \Delta^{\mathcal{I}}$;
- $H^{\mathcal{J}}_R = \{x_{d,R,e} \mid d \in \Delta^{\mathcal{I}} \text{ and } e \in W_{d,R}\}$ for all $R \in \mathsf{rol}(C, \mathcal{T})$;
- For each $R \in \mathsf{rol}(C, \mathcal{T})$, the counter $\overline{B_R}$ is defined as follows: $\overline{B_R} = 0$ for all instances of $\mathsf{Real}^{\mathcal{J}}$; for the instances of $H^{\mathcal{J}}_R$, we define $\overline{B_R}$ as follows:
$$\overline{B_R} = i \text{ for those } x_{d,R,e} \in H^{\mathcal{J}}_R \text{ with } \nu_{d,R}(e) = i;$$

- For each concept $D \in \mathsf{cl}(C,\mathcal{T})$ with D of the form $(\bowtie n\ R\ D)$, the counter $\overline{B^{\bowtie n}_{D,R}}$ is defined as follows: $\overline{B^{\bowtie n}_{D,R}} = 0$ for all instances of $\mathsf{Real}^{\mathcal{J}}$; for instances $x_{d,R,e}$ of $H^{\mathcal{J}}_R$, we set

$$\overline{B^{\bowtie n}_{D,R}} = \#\{e' \in W_{d,R} \mid \nu_{d,R}(e') < \nu_{d,R}(e) \text{ and } e' \in D^{\mathcal{I}}\};$$

- $$\begin{aligned} R^{\mathcal{I}} = \; & \{(x_{d,R,e}, x_{e,\mathsf{Inv}(R),d}) \mid d,e \in \Delta^{\mathcal{I}} \text{ with } e \in W_{d,R} \text{ and } d \in W_{e,\mathsf{Inv}(R)}\} \cup \\ & \{(x_{d,R,e}, e) \mid d,e \in \Delta^{\mathcal{I}} \text{ with } e \in W_{d,R} \text{ and } d \notin W_{e,\mathsf{Inv}(R)}\} \cup \\ & \{(d,x) \mid d \in \mathsf{Real}^{\mathcal{J}},\ (d,e) \in R^{\mathcal{I}} \text{ and } x = x_{e,\mathsf{Inv}(R),d} \text{ or} \\ & \qquad x = e \text{ and } d \notin W_{e,\mathsf{Inv}(R)}\} \end{aligned}$$
  for all $R \in \mathsf{rnam}(C,\mathcal{T})$;

- $$\begin{aligned} L_R = \; & \{(d, x_{d,R,e}) \mid d \in \Delta^{\mathcal{I}},\ e \in W_{d,R},\ \text{and } \nu_{d,R}(e) = 0\} \cup \\ & \{(x_{d,R,e}, x_{d,R,e'}) \mid d \in \Delta^{\mathcal{I}},\ e,e' \in W_{d,R},\ \text{and } \nu_{d,R}(e') = \nu_{d,R}(e)+1\}. \end{aligned}$$

Since the translation $\sigma(C)$ of an $\mathcal{ALCQI}$-concept C is a Boolean formula, it is trivial to prove the following claim by structural induction (using the definition of $\mathcal{J}$):

Claim 3. For all $d \in \Delta^{\mathcal{I}}$ and $D \in \mathsf{cl}(C,\mathcal{T})$, $d \in D^{\mathcal{I}}$ implies $d \in \gamma(D)^{\mathcal{J}}$.

Since $\mathcal{I}$ is a model of C, Claim 3 clearly implies that there is a $d \in \Delta^{\mathcal{I}}$ such that $d \in \gamma(C)^{\mathcal{J}}$. By definition of $\mathsf{Real}^{\mathcal{J}}$, we thus have $d \in \sigma(C)^{\mathcal{J}}$ and thus $\mathcal{J}$ is a model of $\sigma(C)$. Moreover, also by Claim 3 $\mathcal{J}$ is a model of the TBox $\{\gamma(D) \doteq \gamma(E) \mid D \doteq E \in \mathcal{T}\}$. Since it is tedious but straightforward to verify that $\mathcal{J}$ is also a model of the TBox $\mathsf{Aux}(C,\mathcal{T})$ (details are left to the reader), $\mathcal{J}$ is thus a model of $\sigma(\mathcal{T})$. ☐

Taking together Theorem 10, which implies that finite satisfiability of $\mathcal{ALCFI}$-concepts w.r.t. TBoxes is in ExpTime, and Lemma 11, we obtain the following theorem:

Theorem 12 *Finite satisfiability of $\mathcal{ALCQI}$-concepts w.r.t. TBoxes is* ExpTime-*complete if numbers are coded in binary.*

5 ABox Consistency

In this section, we extend the complexity bounds obtained in Sections 3 and 4 to a more general reasoning task: finite $\mathcal{ALCQI}$-ABox consistency. As noted in the introduction, ABoxes can be understood as describing a"snapshot" of the world. We should like to note that (finite) $\mathcal{ALCQI}$-ABox consistency has important applications: whereas finite $\mathcal{ALCQI}$-concept satisfiability algorithms can be used to decide the consistency of conceptual database models and infer implicit IS-A relationships as described in the introduction, $\mathcal{ALCQI}$-ABox consistency can be used as the core component of algorithms deciding containment of conjunctive queries w.r.t. conceptual database models—a task that DLs have succesfully been used for and that calls for finite model reasoning [2; 10].

Definition 13 (ABox) *Let O be a countable infinite set of* object names. *An* ABox assertion *is an expression of the form $a : C$ or $(a, b) : R$, where a and b are object names, C is a concept name, and R a role name. An* ABox *is a finite set of ABox assertions.*

Interpretations $\mathcal{I}$ are extended to ABoxes as follows: additionally, the interpretation function $\cdot^{\mathcal{I}}$ maps each object name to an element of $\Delta^{\mathcal{I}}$ such that $a \neq b$ implies $a^{\mathcal{I}} \neq b^{\mathcal{I}}$ for all $a, b \in O$ (the so-called unique name assumption*). An interpretation $\mathcal{I}$* satisfies *an assertion $a : C$ if $a^{\mathcal{I}} \in C^{\mathcal{I}}$ and an assertion $(a, b) : R$ if $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$. It is a* model *for an ABox $\mathcal{A}$ if it satisfies all assertions in $\mathcal{A}$. An ABox is called* finitely consistent *if it has a finite model.* ♢

In the following, we will polynomially reduce finite $\mathcal{ALCQI}$-ABox consistency to finite $\mathcal{ALCQI}$-concept satisfiability. Thus, we prove that $\mathcal{ALCQI}$-ABox consistency is ExpTime-complete independently of the way in which numbers are coded. We start with fixing some notation.

Let $\mathcal{A}$ be an ABox and $\mathcal{T}$ a TBox. For each object name a used in $\mathcal{A}$, $\mathsf{refl}_{\mathcal{A}}(a)$ denotes the set of role names R such that

$$\{(a, a) : R, (a, a) : R^-\} \cap \mathcal{A} \neq \emptyset.$$

For each object a and role $R \in \mathsf{rol}(\mathcal{A}, \mathcal{T})$, $N_{\mathcal{A}}(a, R)$ denotes the set of object names b such that $b \neq a$ and

$$\{(a, b) : R,\ (b, a) : \mathsf{Inv}(R)\} \cap \mathcal{A} \neq \emptyset.$$

Moreover, we use $\mathsf{cl}(\mathcal{A}, \mathcal{T})$ to denote the smallest set containing all sub-concepts of concepts appearing in $\mathcal{A}$ and $\mathcal{T}$ that is closed under $\dot\neg$. It can easily be shown that the cardinality of $\mathsf{cl}(\mathcal{A}, \mathcal{T})$ is linear in the size of $\mathcal{A}$ and $\mathcal{T}$. We use $\mathsf{rol}(\mathcal{A}, \mathcal{T})$ to denote the set of all roles (i.e., role names or inverses of role names) used in $\mathcal{A}$ or $\mathcal{T}$.

Definition 14 (Type) *A* type *T for an ABox $\mathcal{A}$ and a TBox $\mathcal{T}$ is defined as in Definition 3 where $\mathsf{cl}(C_0, \mathcal{T})$ is replaced with $\mathsf{cl}(\mathsf{A}, \mathcal{T})$.* ♢

In the following, we will sometimes identify types T with the conjunction $\bigsqcap_{C \in T} C$ and write, e.g., $d \in T^{\mathcal{I}}$ for $d \in (\bigsqcap_{C \in T} C)^{\mathcal{I}}$. Again, the number of types for an ABox $\mathcal{A}$ and a TBox $\mathcal{T}$ is exponential in the size of $\mathcal{A}$ and $\mathcal{T}$. The central notion in the reduction of finite $\mathcal{ALCQI}$-ABox consistency to finite $\mathcal{ALCQI}$-concept satisfiability is that of a reduction candidate:

Definition 15 (Reduction Candidate) *Let $\mathcal{A}$ be an ABox and $\mathcal{T}$ a TBox. A* reduction candidate *for $\mathcal{A}$ and $\mathcal{T}$ is a function t that maps each object name a appearing in $\mathcal{A}$ to a type $t(a)$ for $\mathcal{A}$ and $\mathcal{T}$ such that $a : C \in \mathcal{A}$ implies $C \in t(a)$.*

Let t be a reduction candidate for $\mathcal{A}$ and $\mathcal{T}$. For each object a, role $R \in \mathsf{rol}(\mathcal{A}, \mathcal{T})$, and type $T \in \mathsf{ran}(t)$ we use $\#_t^{\mathcal{A}}(a, R, T)$ to denote the number of objects b such that $b \in N_{\mathcal{A}}(a, R)$ and $t(b) = T$.

Now, for each object a used in $\mathcal{A}$, we define a reduction concept $C_t^{\mathcal{A}}(a)$ *as follows:*

$$\begin{aligned} C_t^{\mathcal{A}}(a) \quad := \quad & t(a) \sqcap X \sqcap \bigsqcap_{R \in \mathsf{refl}_{\mathcal{A}}(a)} \exists R.(t(a) \sqcap X) \sqcap \\ & \bigsqcap_{R \in \mathsf{rol}(\mathcal{A},\mathcal{T})} \bigsqcap_{T \in \mathsf{ran}(t)} (\geqslant \#_t^{\mathcal{A}}(a, R, T)\ R\ (T \sqcap \neg X)), \end{aligned}$$

where X is a concept name not used in $\mathcal{A}$ and $\mathcal{T}$. The reduction candidate t is called realizable *iff, for every object a used in $\mathcal{A}$, the reduction concept $C_t^{\mathcal{A}}(a)$ is finitely satisfiable w.r.t. $\mathcal{T}$.* $\diamondsuit$

First we establish a technical lemma.

Lemma 16 *Let $\mathcal{A}$ be an ABox, $\mathcal{T}$ a TBox, t a reduction candidate for $\mathcal{A}$ and $\mathcal{T}$, and a an object name used in $\mathcal{A}$. If the reduction concept $C_t^{\mathcal{A}}(a)$ is finitely satisfiable w.r.t. $\mathcal{T}$, then there exists a finite model $\mathcal{I}$ of $\mathcal{T}$ and $C_t^{\mathcal{A}}(a)$ and some $d \in (C_t^{\mathcal{A}}(a))^{\mathcal{I}}$ such that, for all roles R, $(a, a) : R \in \mathcal{A}$ implies $(d, d) \in R^{\mathcal{I}}$.*

Proof. Let $\mathcal{I}$ be a model of $C_t^{\mathcal{A}}(a)$ and $\mathcal{T}$ and let $d \in (C_t^{\mathcal{A}}(a))^{\mathcal{I}}$. We construct a new interpretation $\mathcal{I}''$ in two steps:

1. Define a new interpretation $\mathcal{I}'$ as follows:

   - $\Delta^{\mathcal{I}'} = \Delta^{\mathcal{I}} \times \{0, 1\}$;
   - $A^{\mathcal{I}'} = \{(d, i) \in \Delta^{\mathcal{I}'} \mid d \in A^{\mathcal{I}}\}$ for all concept names A;
   - $R^{\mathcal{I}'} = \{((d, i), (e, j)) \mid (d, e) \in R^{\mathcal{I}} \text{ and } i \neq j\}$ for all role names R.

   Again, using structural induction, it is readily checked that, for each $d \in \Delta^{\mathcal{I}}$ and $C \in \mathsf{cl}(\mathcal{A}, \mathcal{T})$, $d \in C^{\mathcal{I}}$ implies $(d, i) \in C^{\mathcal{I}'}$ for $i \in \{0, 1\}$. Thus $(d, 0) \in (C_t^{\mathcal{A}}(a))^{\mathcal{I}'}$ (the same holds for $(d, 1)$) and $\mathcal{I}'$ is a model of $\mathcal{T}$. Moreover, $\mathcal{I}'$ clearly satisfies the following property: for all roles $R \in \mathsf{rol}(\mathcal{A}, \mathcal{T})$ and $d \in \Delta^{\mathcal{I}'}$, we have $(d, d) \notin R^{\mathcal{I}'}$.

2. We now construct the interpretation $\mathcal{I}''$ from $\mathcal{I}'$. Since the inner structure of elements from $\Delta^{\mathcal{I}'}$ is not important, we henceforth refer to $(d, 0)$ as d'. For each role name $R \in \mathsf{refl}_{\mathcal{A}}(a)$, fix a domain element $e_R \in \Delta^{\mathcal{I}'}$ such that $(d', e_R) \in R^{\mathcal{I}'}$ and $e_R \in t(a)^{\mathcal{I}'}$. Such domain elements exist since $C_t^{\mathcal{A}}(a)$ contains the conjunct $\bigsqcap_{R \in \mathsf{refl}_{\mathcal{A}}(a)} \exists R.(t(a) \sqcap X)$. The interpretation $\mathcal{I}''$ is now defined as follows:

   - $\Delta^{\mathcal{I}''} = \Delta^{\mathcal{I}'}$;
   - $A^{\mathcal{I}''} = A^{\mathcal{I}'}$ for all concept names A;
   - $R^{\mathcal{I}''} = R^{\mathcal{I}'}$ for all role names $R \notin \mathsf{refl}_{\mathcal{A}}(a)$;
   - $R^{\mathcal{I}''} = (R^{\mathcal{I}'} \setminus \{(d', e_R)\}) \cup \{(d', d'), \}$. for all role names $R \in \mathsf{refl}_{\mathcal{A}}(a)$.

Using structural induction, it is not hard to check that, for each $d \in \Delta^{\mathcal{I}'}$ and $C \in \mathsf{cl}(\mathcal{A},\mathcal{T})$, $d \in C^{\mathcal{I}'}$ implies $d \in C^{\mathcal{I}''}$. Thus, $d' \in (C_t^{\mathcal{A}}(a))^{\mathcal{I}''}$ and $\mathcal{I}''$ is a model of $\mathcal{T}$. Moreover, $(a,a) : R \in \mathcal{A}$ implies $(d',d') \in R^{\mathcal{I}''}$: this is true by definition of $\mathsf{refl}_{\mathcal{A}}$ and $\mathcal{I}''$ if R is a role name. If $R = S^-$ for some role name S, then $(a,a) : R \in \mathcal{A}$ implies $S \in \mathsf{refl}_{\mathcal{A}}$. Thus $(d',d') \in S^{\mathcal{I}''}$ by definition of $\mathcal{I}''$. By semantics, we obtain $(d',d') \in S^{\mathcal{I}''}$ as required. □

The following lemma describes the relationship between ABoxes and reduction candidates.

Lemma 17 *Let $\mathcal{A}$ be an ABox and $\mathcal{T}$ a TBox. $\mathcal{A}$ is finitely consistent w.r.t. $\mathcal{T}$ iff there exists a realizable reduction candidate for $\mathcal{A}$ and $\mathcal{T}$.*

Proof. For the "only if" direction, let $\mathcal{I}$ be a model of $\mathcal{A}$ and $\mathcal{T}$. We construct a reduction candidate t as follows:

$$\text{for each object } a \text{ in } \mathcal{A}, \text{ set } t(a) = \{D \in \mathsf{cl}(\mathcal{A},\mathcal{T}) \mid a^{\mathcal{I}} \in D^{\mathcal{I}}\}.$$

It remains to prove that t is realizable. Let a be an object in $\mathcal{A}$. We construct a model $\mathcal{I}''$ of $C_t^{\mathcal{A}}(a)$ from $\mathcal{I}$ in two steps as follows: first, construct $\mathcal{I}'$ from $\mathcal{I}$ as in the proof of Lemma 16 and set

$$X^{\mathcal{I}'} = \{(d,0) \in \Delta^{\mathcal{I}'} \mid d \in \Delta^{\mathcal{I}}\}.$$

Then $\mathcal{I}''$ is obtained from $\mathcal{I}'$ by "bending some R edges", everything else is unchanged: for each role name $R \in \mathsf{refl}\mathcal{A}(a)$, set

$$\begin{aligned} R^{\mathcal{I}''} = \; & R^{\mathcal{I}'} \setminus \{((a^{\mathcal{I}},0),(a^{\mathcal{I}},1)),((a^{\mathcal{I}},1),(a^{\mathcal{I}},0))\} \cup \\ & \{((a^{\mathcal{I}},0),(a^{\mathcal{I}},0)),((a^{\mathcal{I}},1),(a^{\mathcal{I}},1))\}. \end{aligned}$$

It can be easily verified that $\mathcal{I}''$ is indeed a finite model of $C_t^{\mathcal{A}}(a)$ w.r.t. $\mathcal{T}$.

For the "if" direction, assume that there exists a realizable reduction candidate t for $\mathcal{A}$ and $\mathcal{T}$. This implies that, for each object name a used in $\mathcal{A}$, there is a finite model $\mathcal{I}_a$ of $C_t^{\mathcal{A}}(a)$ and $\mathcal{T}$. For each such model $\mathcal{I}_a$, fix a domain element $d_a \in \Delta^{\mathcal{I}_a}$ such that $d_a \in (C_t^{\mathcal{A}}(a))^{\mathcal{I}_a}$. By Lemma 16, we may w.l.o.g. assume that, for all object names a used in $\mathcal{A}$ and roles R, $(a,a) : R \in \mathcal{A}$ implies $(d_a,d_a) \in R^{\mathcal{I}_a}$. Moreover, we assume that $a \neq b$ implies $\Delta^{\mathcal{I}_a} \cap \Delta^{\mathcal{I}_b} = \emptyset$.

In the following, we use the models $\mathcal{I}_a$ to construct a (finite) model $\mathcal{I}$ of $\mathcal{A}$ and $\mathcal{T}$. First fix, for each object name a used in $\mathcal{A}$ and each role $R \in \mathsf{rol}(\mathcal{A},\mathcal{T})$, an injective function η_a from $N_{\mathcal{A}}(a,R)$ to $\Delta^{\mathcal{I}_a}$ such that, for all $b \in N_{\mathcal{A}}(a,R)$, we have $(d_a,\eta_a(b)) \in R^{\mathcal{I}_a}$ and $\eta_a(b) \in (t(b) \sqcap \neg X)^{\mathcal{I}_a}$. Such functions do clearly exist due to the conjunct $\bigsqcap_{R \in \mathsf{rol}(\mathcal{A},\mathcal{T})} \bigsqcap_{T \in \mathsf{ran}(t)} (\geqslant \#_t^{\mathcal{A}}(a,R,T)\ R\ (T \sqcap \neg X))$ of $C_t^{\mathcal{A}}(a)$. Then define the interpretation $\mathcal{I}$ as follows:

- $\Delta^{\mathcal{I}} := \bigcup_{a \text{ used in } \mathcal{A}} \Delta^{\mathcal{I}_a}$;
- $A^{\mathcal{I}} := \bigcup_{a \text{ used in } \mathcal{A}} A^{\mathcal{I}_a}$ for all concept names A;

- $R^{\mathcal{I}} := \bigcup_{a \text{ used in } \mathcal{A}} \big((R^{\mathcal{I}_a} \setminus \bigcup_{b \in N_{\mathcal{A}}(a,R)} \{(d_a, \eta_a(b))\}) \cup \bigcup_{b \in N_{\mathcal{A}}(a,R)} \{(d_a, d_b)\}\big)$ for all role names R;
- $a^{\mathcal{I}} := d_a$ for each object name a used in $\mathcal{A}$.

It is straightforward to prove the following claim using structural induction:

Claim: for each object name a used in $\mathcal{A}$, $d \in \Delta^{\mathcal{I}_a}$, and $C \in \mathsf{cl}(\mathcal{A}, \mathcal{T})$, $d \in C^{\mathcal{I}_a}$ implies $d \in C^{\mathcal{I}}$.

It is thus readily checked that $\mathcal{I}$ is indeed a (finite) model of $\mathcal{A}$ and $\mathcal{T}$:

1. Let $a : C \in \mathcal{A}$. Then the claim together with $d_a \in (C_t^{\mathcal{A}}(a))^{\mathcal{I}_a}$ yields $a^{\mathcal{I}} = d_a \in C^{\mathcal{I}}$ since $t(a)$ is a conjunct of $C_t^{\mathcal{A}}(a)$ and $a : C \in \mathcal{A}$ implies $C \in t(a)$.
2. Let $(a, a) : R \in \mathcal{A}$. Since $a^{\mathcal{I}} = d_a$, we have $(a^{\mathcal{I}}, a^{\mathcal{I}}) \in R^{\mathcal{I}_a}$ by choice of $\mathcal{I}_a$. Since $d_a \in X^{\mathcal{I}_a}$ by definition of $C_t^{\mathcal{A}}(a)$ and, for each $b \in N_{\mathcal{A}}(a, R)$, $\eta_a(b) \in (\neg X)^{\mathcal{I}_a}$ by definition of η_b, we have $a^{\mathcal{I}} \notin \mathsf{dom}(\eta_a)$. Thus, $(a^{\mathcal{I}}, a^{\mathcal{I}}) \in R^{\mathcal{I}_a}$ implies $(a^{\mathcal{I}}, a^{\mathcal{I}}) \in R^{\mathcal{I}}$, both if R is a role name and if R is the inverse of a role name.
3. Let $(a, b) : R \in \mathcal{A}$ where $a \neq b$. If R is a role name, then $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$ by definition of $\mathcal{I}$. If $R = S^-$ for some role name S, then we have $a \in N_{\mathcal{A}}(b, S)$. Thus, $(b^{\mathcal{I}}, a^{\mathcal{I}}) \in S^{\mathcal{I}}$ by definition of $\mathcal{I}$ implying $(a^{\mathcal{I}}, b^{\mathcal{I}}) \in S^{\mathcal{I}}$ by the semantics.
4. Finally, the claim together with the fact that, for each object name a used in $\mathcal{A}$, $\mathcal{I}_a$ is a model of $\mathcal{T}$ clearly implies that $\mathcal{I}$ is also a model of $\mathcal{T}$.

□

It is now easy to establish a tight complexity bound for finite $\mathcal{ALCQI}$-ABox consistency.

Theorem 18 *Finite $\mathcal{ALCQI}$-ABox consistency w.r.t. TBoxes is* ExpTime-*complete if numbers are coded in binary.*

Proof. Let $\mathcal{A}$ be an ABox and $\mathcal{T}$ a TBox. Since the number of types for $\mathcal{A}$ and $\mathcal{T}$ is exponential in the size of $\mathcal{A}$ and $\mathcal{T}$ and the number of object names used in $\mathcal{A}$ is linear in the size of $\mathcal{A}$, the number of reduction candidates for $\mathcal{A}$ and $\mathcal{T}$ is exponential in the size of $\mathcal{A}$ and $\mathcal{T}$. Thus, to decide finite consistency of $\mathcal{A}$ w.r.t. $\mathcal{T}$, we may simply enumerate all reduction types for $\mathcal{A}$ and $\mathcal{T}$ and check them for realizability: by Lemma 17, $\mathcal{A}$ is finitely consistent w.r.t. $\mathcal{T}$ if we find a realizable reduction type. Since the size of the reduction concepts is clearly polynomial in the size of $\mathcal{A}$ and $\mathcal{T}$, by Theorem 12 the resulting algorithm can be executed in deterministic time exponential in $\mathcal{A}$ and $\mathcal{T}$. □

Note that our choice of the unique name assumption is *not* crucial for this result: if we want to decide finite consistency of an ABox $\mathcal{A}$ without the unique name assumption, we may use the following approach: enumerate all possible partitionings of the object names used in $\mathcal{A}$. For each partitioning, choose a representative for each partition and then replace each object name with the representative of its partition. Obviously, the ABox $\mathcal{A}$ is finitely consistent *without* the unique name assumption if and only if any of the resulting ABoxes is finitely consistent *with* the unique name assumption. Clearly, this yields an ExpTime upper bound for finite ABox consistency without the unique name assumption.

6 Outlook

In this paper, we have determined finite model reasoning in the description logic $\mathcal{ALCQI}$ to be ExpTime-complete. This shows that reasoning w.r.t. finite models is not harder than reasoning w.r.t. arbitrary models, which is known to be also ExpTime-complete [6]. We hope that, ultimately, this research will lead to the development of finite model reasoning systems that behave equally well as existing DL reasoners doing reasoning w.r.t. arbitrary models. Note, however, that the current algorithm is *best-case* ExpTime since it constructs an exponentially large equation system. It can thus not be expected to have an acceptable runtime behaviour if implemented in a naive way. Nevertheless, we believe that the use of equation systems and linear programming is indispensable for finite model reasoning in $\mathcal{ALCQI}$. Thus, efforts to obtain efficient reasoners should perhaps concentrate on methods to avoid best-case exponentiality such as on-the-fly construction of equation systems. Moreover, the reductions presented in Section 4 and 5 can also not be expected to exhibit an acceptable run-time behaviour and it would thus be interesting to try to replace them by more "direct" methods.

Theoretically, there exist at least two interesting directions in which the presented research can be continued: first, while finite $\mathcal{ALCQI}$-concept satisfiabiltiy w.r.t. TBoxes is sufficient for reasoning about conceptual database models as described in the introduction, finite $\mathcal{ALCQI}$-ABox consistency it is not yet sufficient for deciding the containment of conjunctive queries w.r.t. a given conceptual model—an intermediate reduction step is required. It would thus be interesting to analyze the complexity of query containment in finite models. We believe that it is possible to obtain an ExpTime upper bound by building on the results presented in Section 5. Secondly, it would be interesting to extend $\mathcal{ALCQI}$ with nominals, i.e. with concept names interpreted as singleton sets. Finite and standard reasoning in the resulting DL $\mathcal{ALCQOI}$ is known to be NExpTime-hard [16]. An extension in this direction is rather challenging since the results established in this paper crucially rely on the fact that adding disjoint copies of a model preserves the model's properties. Unfortunately, in the presence of nominals, this is no longer true.

References

[1] Franz Baader, Diego Calvanese, Deborah McGuinness, Daniele Nardi, and Peter F. Patel-Schneider, editors. *The Description Logic Handbook: Theory, Implementation, and Applications.* Cambridge University Press, 2002. To appear.

[2] D. Calvanese, G. De Giacomo, and M. Lenzerini. On the decidability of query containment under constraints. In *Proc. of PODS-98*, pages 149–158, 1998.

[3] D. Calvanese, M. Lenzerini, and D. Nardi. Description logics for conceptual data modeling. In Jan Chomicki and Gnter Saake, editors, *Logics for Databases and Information Systems*, pages 229–263. Kluwer Academic Publisher, 1998.

[4] Diego Calvanese. Finite model reasoning in description logics. In *Proc. of KR-96.* Morgan Kaufmann, 1996.

[5] Diego Calvanese. *Unrestricted and Finite Model Reasoning in Class-Based Representation Formalisms.* PhD thesis, Dip. di Inf. e Sist., Univ. di Roma "La Sapienza", 1996.

[6] G. De Giacomo and M. Lenzerini. Tbox and Abox reasoning in expressive description logics. In *Proc. of KR-96*, pages 316–327. Morgan Kaufmann, 1996.

[7] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *J. of Computer and System Science*, 18:194–211, 1979.

[8] E. Franconi and G. Ng. The i.com tool for intelligent conceptual modelling. In *Working Notes of the ECAI2000 Workshop KRDB2000*, 2000.

[9] I. Horrocks and U. Sattler. A description logic with transitive and inverse roles and role hierarchies. *J. of Logic and Computation*, 1999.

[10] I. Horrocks, U. Sattler, S. Tessaris, and S. Tobies. How to decide query containment under constraints using a description logic. In A. Voronkov, editor, *Proc. of LPAR 2000*, number 1955 in LNAI. Springer-Verlag, 2000.

[11] C. H. Papadimitriou. On the complexity of integer programming. *J. of the ACM*, 28(2):765–769, 1981.

[12] V. R. Pratt. Models of program logics. In *Proceedings of the 20th Annual Symposium on Foundations of Computer Science*, San Juan, Puerto Rico, 1979.

[13] K. Schild. A correspondence theory for terminological logics: Preliminary report. In *Proc. of IJCAI-91*, pages 466–471, Sydney, 1991.

[14] A. Schrijver. *Theory of linear and integer programming.* Wiley, 1986.

[15] Bernhard Thalheim. Fundamentals of cardinality constraints. In *Proceedings of the Conference on Entity-Relationship-Approaches 1992 (ER92)*, number 645 in LNCS, pages 7–23. Springer Verlag, 1992.

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