On the Complexity of Horn Description Logics

Markus Krötzsch, Sebastian Rudolph, and Pascal Hitzler

Institute AIFB, Universität Karlsruhe, Germany
{mak|sru|phi}@aifb.uni-karlsruhe.de

Abstract. Horn-$\mathcal{SHIQ}$ has been identified as a fragment of the description logic $\mathcal{SHIQ}$ for which inferencing is in PTime with respect to the size of the ABox. This enables reasoning with larger ABoxes in situations where the TBox is static, and represents one approach towards tractable description logic reasoning.
In this paper, we show that reasoning in Horn-$\mathcal{SHIQ}$, in spite of its low data-complexity, is ExpTime-hard with respect to the overall size of the knowledge base. While this result is not unexpected, the proof is not a mere modification of existing reductions since it has to account for the restrictions of Hornness. We establish the result for Horn-$\mathcal{FLE}$, showing that Hornness does not simplify TBox reasoning even for very restricted description logics. Moreover, we derive a context-free grammar that defines Horn-$\mathcal{SHIQ}$ in a simpler and more intuitive way than existing characterisations.

1 Introduction

The development of description logics (DLs) has been dominated by the desire to obtain powerful yet decidable formalisms for specifying knowledge. High complexity of reasoning was considered inevitable for obtaining practically useful logics, and highly efficient algorithms were developed to still solve arising reasoning problems. Accordingly, the DL-based flavours *OWL Lite* and *OWL DL* of the Web Ontology Language [1] are already ExpTime- and NExpTime-complete, respectively.

However, reasoning in those expressive logics remains intractable, and even modern, optimised algorithms are of limited scalability. This has triggered the renewed investigation of description logics with tractable inference problems, and various such fragments have been proposed.[1] Typically, these logics aim at low complexity of reasoning with respect to the size of the entire knowledge base (the so-called *combined complexity*), as in the case of $\mathcal{EL}^{++}$ [2] and *DL-Lite* [3], which are both polynomial in this sense. Alternatively, one can consider the complexity with respect to the number of simple assertions within the knowledge base, which is also known as the *data complexity*. This has led to the investigation of Horn-$\mathcal{SHIQ}$ as an expressive fragment of the description logic $\mathcal{SHIQ}$ [4] that is known to be of polynomial data complexity [5].

Horn-$\mathcal{SHIQ}$ supports all logical operators of $\mathcal{SHIQ}$ but syntactically restricts their use in various ways. This leads to the aforementioned low data complexity, but it also results in a rather involved description of the syntax, which merely supplies a criterion for verifying Hornness of some given knowledge base. Extending prior results

[1] See http://owl-workshop.man.ac.uk/Tractable.html for an overview.

Table 1. Concept constructors in $\mathcal{SHIQ}$. Semantics refers to an interpretation $\mathcal{I}$ with domain $\mathcal{D}$.

Name	Syntax	Semantics
top	$\top$	$\mathcal{D}$
bottom	$\bot$	$\emptyset$
negation	$\neg C$	$\mathcal{D} \setminus C^{\mathcal{I}}$
conjunction	$C \sqcap D$	$C^{\mathcal{I}} \cap D^{\mathcal{I}}$
disjunction	$C \sqcup D$	$C^{\mathcal{I}} \cup D^{\mathcal{I}}$
univ. restriction	$\forall R.C$	$\{x \in \mathcal{D} \mid (x, y) \in R^{\mathcal{I}}$ implies $y \in C^{\mathcal{I}}\}$
exist. restriction	$\exists R.C$	$\{x \in \mathcal{D} \mid$ for some $y \in \mathcal{D},\ (x, y) \in R^{\mathcal{I}}$ and $y \in C^{\mathcal{I}}\}$
qualified number	$\leq n\, R.C$	$\{x \in \mathcal{D} \mid \#\{y \in \mathcal{D} \mid (x, y) \in R^{\mathcal{I}}$ and $y \in C^{\mathcal{I}}\} \leq n\}$
restriction	$\geq n\, R.C$	$\{x \in \mathcal{D} \mid \#\{y \in \mathcal{D} \mid (x, y) \in R^{\mathcal{I}}$ and $y \in C^{\mathcal{I}}\} \geq n\}$

for Horn-$\mathcal{ALCHIQ}$ [6], we offer a new constructive definition describing the syntax of Horn-$\mathcal{SHIQ}$ with a simple context-free grammar.

But the main open problem that we address is the question for the combined complexity of Horn description logics. Since $\mathcal{SHIQ}$ is in ExpTime [7], the same holds for Horn-$\mathcal{SHIQ}$, but it is not known whether this upper bound is tight. We settle this question by showing that even small fragments of Horn-$\mathcal{SHIQ}$ are ExpTime-hard, and thus demonstrate that Hornness often does not simplify the complexity of TBox reasoning. The presented complexity proof is not a mere corollary of existing results, but employs a novel, self-contained reduction of the halting problem for polynomially space-bounded alternating Turing machines. This makes it a simple alternative for showing known ExpTime-hardness results for logics like $\mathcal{ALC}$ or $\mathcal{EL}$ with functional roles.

After a short introduction to the relevant description logics in Sect. 2, we present a simple description of Horn-$\mathcal{SHIQ}$ and other Horn-DLs in Sect. 3. Our main results regarding the ExpTime-complexity of Horn-DLs are shown in Sect. 4. In Sect. 5, we discuss our results and open questions for future research.

2 Preliminaries

We briefly repeat some basic definitions of DLs and introduce our notation.

Definition 1. *A knowledge base of the description logic $\mathcal{SHIQ}$ is based on a set* N_R *of* role names, *a set* N_C *of* concept names, *and a set* N_I *of* individual names. *The set of* $\mathcal{SHIQ}$ (abstract) roles *is* $\mathsf{N}_R \cup \{R^- \mid R \in \mathsf{N}_R\}$, *and we set* $\mathsf{Inv}(R) = R^-$ *and* $\mathsf{Inv}(R^-) = R$. *In the following, we leave this vocabulary implicit and assume that A, B are concept names, a, b are individual names, and R, S are abstract roles.*

A $\mathcal{SHIQ}$ knowledge base consists of three finite sets of axioms that are referred to as RBox, TBox, *and* ABox. *A $\mathcal{SHIQ}$ RBox may contain axioms of the form $S \sqsubseteq R$ iff it also contains* $\mathsf{Inv}(R) \sqsubseteq \mathsf{Inv}(S)$, *and axioms of the form* $\mathsf{Trans}(R)$ *iff it also contains* $\mathsf{Trans}(\mathsf{Inv}(R))$. *By $\sqsubseteq^*$ we denote the reflexive-transitive closure of $\sqsubseteq$. A role R is* transitive *whenever there is a role S such that* $\mathsf{Trans}(S)$, *$R \sqsubseteq^* S$ and $S \sqsubseteq^* R$. R is* simple *if it has no transitive subroles, i.e., if $S \sqsubseteq^* R$ implies that S is not transitive. Roles that are not simple are also called* complex.

Table 2. Definition of $\mathsf{clos}(\mathcal{KB})$. $\mathsf{NNF}(C)$ denotes the negation normal form of some concept C. For details see [8].

- If $C \sqsubseteq D \in \mathcal{KB}$, then $\mathsf{NNF}(\neg C \sqcup D) \in \mathsf{clos}(\mathcal{KB})$,
- If $C(a) \in \mathcal{KB}$, then $\mathsf{NNF}(C) \in \mathsf{clos}(\mathcal{KB})$,
- If $C \in \mathsf{clos}(\mathcal{KB})$ and D is a subconcept of C, then $D \in \mathsf{clos}(\mathcal{KB})$,
- If $\leq n\, R.C \in \mathsf{clos}(\mathcal{KB})$, then $\mathsf{NNF}(\neg C) \in \mathsf{clos}(\mathcal{KB})$,
- If $\forall R.C \in \mathsf{clos}(\mathcal{KB})$, $S \sqsubseteq^* R$, and $\mathrm{Trans}(S) \in \mathcal{KB}$, then $\forall S.C \in \mathsf{clos}(\mathcal{KB})$.

A $\mathcal{SHIQ}$ TBox consists of axioms of the form $C \sqsubseteq D$, where C and D are concept expressions *constructed from concept names by the operators shown in Table 1. A $\mathcal{SHIQ}$ ABox consists of axioms of the form $A(a)$, $\neg A(a)$, $R(a, b)$, $\neg S(a, b)$, $a \approx b$, and $a \not\approx b$, where S is a simple role.*

The above definition is fairly standard, with some minor exceptions. First, we allow for negated simple role assertions within ABoxes. This is known to not make the logic more complex or even undecidable, see [8] for some discussion. Second, we restrict ABox concept statements to possibly negated atomic concepts. Our ABoxes thus are *extensionally reduced*, but it is known that this does not restrict the expressivity of the logic since complex ABox statements can easily be moved into the TBox by introducing auxiliary concept names. Third, we do not explicitly consider concept/role equivalence $\equiv$, since it can be modelled via mutual concept/role inclusions.

We adhere to the common model-theoretic semantics for $\mathcal{SHIQ}$ with general concept inclusion, which we will not repeat here (see, e.g., [8] for details). Table 1 recalls the semantics of concept operators in $\mathcal{SHIQ}$.

We will consider various fragments of $\mathcal{SHIQ}$ below. A $\mathcal{SHIQ}$ knowledge base is in $\mathcal{ALCHIQ}$ if it contains no transitivity axioms. It is in $\mathcal{FLE}$ if the RBox is empty and only $\forall$, $\exists$, $\sqcap$, and $\top$ are used within the TBox. The fragment of $\mathcal{FLE}$ without $\exists$ ($\forall$) is called $\mathcal{FL}_0$ ($\mathcal{EL}$).

3 A simple description of Horn-$\mathcal{SHIQ}$

The data complexity of a description logic inference task is the complexity of inferencing with respect to the size of its (extensionally reduced) ABox. In [5], Horn-$\mathcal{SHIQ}$ was introduced as a particular fragment of the description logic $\mathcal{SHIQ}$ that is distinguished by its low PTime *data complexity*. While the exposition in [5] involved various recursively defined auxiliary functions, we present a simpler definition that extends the definition of Horn-$\mathcal{ALCHIQ}$ given in [6].

The original definition of Horn-$\mathcal{SHIQ}$ involves a preprocessing step for eliminating transitivity axioms by transforming a $\mathcal{SHIQ}$ knowledge base into an equisatisfiable $\mathcal{ALCHIQ}$ knowledge base. For showing that our following definition of Horn-$\mathcal{SHIQ}$ is correct, we first briefly repeat this transformation procedure.

For a $\mathcal{SHIQ}$ knowledge base $\mathcal{KB}$, a set of concept terms $\mathsf{clos}(\mathcal{KB})$ is defined recursively as shown in Table 2. Now $\mathcal{KB}$ is transformed into an $\mathcal{ALCHIQ}$ knowledge base $\Omega(\mathcal{KB})$ by

– eliminating all transitivity axioms Trans(S), and by
– adding the axiom $\forall R.C \sqsubseteq \forall S.(\forall S.C)$, for every concept $\forall R.C \in \mathsf{clos}(\mathcal{KB})$ and role S, such that $S \sqsubseteq^* R$ and Trans(S) $\in \mathcal{KB}$.

It was shown in [8] that $\mathcal{KB}$ is satisfiable iff $\Omega(\mathcal{KB})$ is satisfiable. A similar reduction was already introduced in [7, Chapter 6], but we focus on the transformation used for defining Horn-$\mathcal{SHIQ}$. Based on the prior definition of Horn-$\mathcal{ALCHIQ}$, a Horn-$\mathcal{SHIQ}$ knowledge base in [5] was defined as a $\mathcal{SHIQ}$ knowledge base $\mathcal{KB}$ for which $\Omega(\mathcal{KB})$ is in Horn-$\mathcal{ALCHIQ}$. We are now ready to provide a simpler formulation.

Proposition 1. *We say that a $\mathcal{SHIQ}$ axiom $C \sqsubseteq D$ is* Horn *if the concept expression $\neg C \sqcup D$ has the form $\mathbf{C}_1^+$ as defined by the context-free grammar in Table 3.*

A $\mathcal{SHIQ}$ knowledge base with an extensionally reduced ABox is in Horn-$\mathcal{SHIQ}$ iff all of its TBox axioms are Horn.

Proof. In [6] it was already shown that a knowledge base is in Horn-$\mathcal{ALCHIQ}$ iff its TBox consists of $\mathcal{ALCHIQ}$-axioms that are Horn in the above sense. Here we only show that the components with complex roles account for the additional axioms that can be constructed in Horn-$\mathcal{SHIQ}$. This is achieved by analysing the axioms that are introduced by the above transformation. Indeed, axioms of the form $\forall R.C \sqsubseteq \forall S.(\forall S.C)$ might fail to be Horn since they correspond to expressions $\exists R.\neg C \sqcup \forall S.(\forall S.C)$. The latter are generally not Horn, since disjunctions in $\mathbf{C}_1^+$ must have the form $\mathbf{C}_0^+ \sqcup \mathbf{C}_1^+$. Since $\exists R.\neg C$ cannot be of the form $\mathbf{C}_0^+$, this requires that $\forall S.(\forall S.C)$ is in $\mathbf{C}_0^+$. But this can only be the case if C is in $\mathbf{C}_0^+$ as well. $\exists R.\neg C$ in this case also is in $\mathbf{C}_1^+$, since $\mathbf{C}_0^- \subseteq \mathbf{C}_1^+$ and $\mathbf{C}_0^+ \subseteq \mathbf{C}_1^-$. This can be shown by an easy induction over the structure of $\mathbf{C}_0^-$ which we omit here (the base case is $\mathbf{A}$; the mutual dependency between $\mathbf{C}_0^+$ and $\mathbf{C}_0^-$ is not problematic during the induction steps).

We thus have described the axioms that can be introduced without problems during transitivity elimination. A closer look at the elimination procedure reveals that the introduction of axioms depends on the existence of formulae of the form $\forall R.C \in \mathsf{clos}(\mathcal{KB})$, where R has a transitive subrole, i.e. R is not simple. We must ensure that C is in $\mathbf{C}_0^+$ in this case. The last two lines of Table 2 obviously cannot directly contribute to the inclusion of formulae $\forall R.C$ in $\mathsf{clos}(\mathcal{KB})$ (unless another problematic axiom is already present). Moreover, since we restrict to extensionally reduced ABoxes, the second line is not relevant either. Consequently, a formula $\forall R.C$ is in $\mathsf{clos}(\mathcal{KB})$ iff it is a subconcept of the negation normal form of some concept $\neg D \sqcup E$ with $D \sqsubseteq E \in \mathcal{KB}$.

Now consider a $\mathcal{SHIQ}$ knowledge base $\mathcal{KB}$ which has a TBox in Horn-$\mathcal{ALCHIQ}$ when ignoring any transitivity axioms. From the above considerations we conclude: $\mathcal{KB}$ is in Horn-$\mathcal{SHIQ}$ iff, for every TBox axiom $D \sqsubseteq E$, every non-simple role R, and every subconcept $\forall R.C$ of $\mathsf{NNF}(\neg D \sqcup E)$, we find that C is in $\mathbf{C}_0^+$. For subconcepts of positive polarity, this is exactly captured by the distinction between $\forall \mathbf{S}.\mathbf{C}_1^+$ and $\forall \mathbf{R}.\mathbf{C}_0^+$ in the definition of $\mathbf{C}_1^+$. Subconcepts of the form $\mathbf{C}_1^-$ have negative polarity in the constructed axiom, so the dual descriptions $\exists \mathbf{S}.\mathbf{C}_1^-$ and $\exists \mathbf{R}.\mathbf{C}_0^-$ characterise the required restrictions. Clearly, no further restrictions are required, and the given restrictions cannot be relaxed without introducing non-Horn axioms during the elimination procedure. □

The advantage of the above definition, besides its simplicity and brevity, is that it provides a local criterion for checking Hornness by investigating the structure of single

Table 3. A grammar for defining Horn-$\mathcal{SHIQ}$. **A**, **R**, and **S** denote the sets of all concept names, role names, and simple role names, respectively. The presentation is slightly simplified by exploiting associativity and commutativity of $\sqcap$ and $\sqcup$, and by omitting $\geq 1\, R.C$ if $\exists R.C$ is present. The grammar for Horn-$\mathcal{ALCHIQ}$ [6] is obtained for the special case that all roles are simple.

$$\begin{aligned}
\mathbf{C}_1^+ &::= \top \mid \bot \mid \neg\mathbf{C}_1^- \mid \mathbf{C}_1^+ \sqcap \mathbf{C}_1^+ \mid \mathbf{C}_0^+ \sqcup \mathbf{C}_1^+ \mid \exists\mathbf{R}.\mathbf{C}_1^+ \mid \forall\mathbf{S}.\mathbf{C}_1^+ \mid \forall\mathbf{R}.\mathbf{C}_0^+ \mid \geq n\, \mathbf{R}.\mathbf{C}_1^+ \mid \leq 1\, \mathbf{R}.\mathbf{C}_0^- \mid \mathbf{A} \\
\mathbf{C}_1^- &::= \top \mid \bot \mid \neg\mathbf{C}_1^+ \mid \mathbf{C}_0^- \sqcap \mathbf{C}_1^- \mid \mathbf{C}_1^- \sqcup \mathbf{C}_1^- \mid \exists\mathbf{S}.\mathbf{C}_1^- \mid \exists\mathbf{R}.\mathbf{C}_0^- \mid \forall\mathbf{R}.\mathbf{C}_1^- \mid \geq 2\, \mathbf{R}.\mathbf{C}_0^- \mid \leq n\, \mathbf{R}.\mathbf{C}_1^+ \mid \mathbf{A} \\
\mathbf{C}_0^+ &::= \top \mid \bot \mid \neg\mathbf{C}_0^- \mid \mathbf{C}_0^+ \sqcap \mathbf{C}_0^+ \mid \mathbf{C}_0^+ \sqcup \mathbf{C}_0^+ \mid \forall\mathbf{R}.\mathbf{C}_0^+ \\
\mathbf{C}_0^- &::= \top \mid \bot \mid \neg\mathbf{C}_0^+ \mid \mathbf{C}_0^- \sqcap \mathbf{C}_0^- \mid \mathbf{C}_0^- \sqcup \mathbf{C}_0^- \mid \exists\mathbf{R}.\mathbf{C}_0^- \mid \mathbf{A}
\end{aligned}$$

axioms. The original definition hides this locality by relying on a transitivity elimination procedure that operates on the whole knowledge base. We adopt the definition of Proposition 1 to characterise the Horn-version of fragments of $\mathcal{SHIQ}$, such as Horn-$\mathcal{FLE}$, as well. Note that Horn-$\mathcal{SHIQ}$ includes all of $\mathcal{EL}$, i.e. Horn-$\mathcal{EL}$ is just $\mathcal{EL}$.

4 Complexity of Horn-$\mathcal{SHIQ}$

To show that Horn-$\mathcal{SHIQ}$ is ExpTime-complete, note that inclusion in ExpTime is obvious since it is a fragment of $\mathcal{SHIQ}$ which is also in ExpTime [7]. To show hardness of the satisfiability problem, we show that even the smaller fragment Horn-$\mathcal{FLE}$ is ExpTime-hard. We establish a polynomial reduction of reasoning in this logic to the halting problem of polynomially space-bounded alternating Turing machines.

4.1 Alternating Turing machines

Definition 2. *An* alternating Turing machine *(ATM) $\mathcal{M}$ is a tuple (Q, Σ, Δ, q_0) where*

- *$Q = U \dot\cup E$ is the disjoint union of a finite set of* universal states *U and a finite set of* existential states *E,*
- *Σ is a finite* alphabet *that includes a* blank *symbol $\square$,*
- *$\Delta \subseteq (Q \times \Sigma) \times (Q \times \Sigma \times \{l, r\})$ is a* transition relation*, and*
- *$q_0 \in Q$ is the* initial state.

A (universal/existential) configuration *of $\mathcal{M}$ is a word $\alpha \in \Sigma^* Q \Sigma^*$ ($\Sigma^* U \Sigma^*$/$\Sigma^* E \Sigma^*$). A configuration α' is a* successor *of a configuration α if one of the following holds:*

1. *$\alpha = w_l q \sigma \sigma_r w_r$, $\alpha' = w_l \sigma' q' \sigma_r w_r$, and $(q, \sigma, q', \sigma', r) \in \Delta$,*
2. *$\alpha = w_l q \sigma$, $\alpha' = w_l \sigma' q' \square$, and $(q, \sigma, q', \sigma', r) \in \Delta$,*
3. *$\alpha = w_l \sigma_l q \sigma w_r$, $\alpha' = w_l q' \sigma_l \sigma' w_r$, and $(q, \sigma, q', \sigma', l) \in \Delta$,*

where $q \in Q$ and $\sigma, \sigma', \sigma_l, \sigma_r \in \Sigma$ as well as $w_l, w_r \in \Sigma^*$. Given some natural number s, the possible transitions in space *s are defined by additionally requiring that $|\alpha'| \leq s + 1$.*

The set of accepting configurations *is the least set which satisfies the following conditions. A configuration α is accepting iff*

- *α is a universal configuration and all its successor configurations are accepting, or*

– *α is an existential configuration and at least one of its successor configurations is accepting.*

Note that universal configurations without any successors here play the rôle of accepting final configurations, and thus form the basis for the recursive definition above.

$\mathcal{M}$ accepts *a given word $w \in \Sigma^*$ (in space s) iff the configuration $q_0 w$ is accepting (when restricting to transitions in space s).*

This definition is inspired by the complexity classes NP and co-NP, which are characterised by non-deterministic Turing machines that accept an input if either at least one or all possible runs lead to an accepting state. An ATM can switch between these two modes and indeed turns out to be more powerful than classical Turing machines of either kind. In particular, ATMs can solve ExpTime problems in polynomial space [9].

Definition 3. *A language L is accepted by a polynomially space-bounded ATM iff there is a polynomial p such that, for every word $w \in \Sigma^*$, $w \in L$ iff w is accepted in space $p(|w|)$.*

Fact 1. *The complexity class* APSpace *of languages accepted by polynomially space-bounded ATMs coincides with the complexity class* ExpTime*.*

We thus can show ExpTime-hardness of Horn-$\mathcal{SHIQ}$ by polynomially reducing the halting problem of ATMs with a polynomially bounded storage space to inferencing in Horn-$\mathcal{SHIQ}$. In the following, we exclusively deal with polynomially space-bounded ATMs, and so we omit additions such as "in space s" when clear from the context.

4.2 Simulating ATMs in Horn-$\mathcal{FLE}$

In the following, we consider a fixed ATM $\mathcal{M}$ denoted as in Definition 2, and a polynomial p that defines a bound for the required space. For any word $w \in \Sigma^*$, we construct a Horn-$\mathcal{FLE}$ knowledge base $K_{\mathcal{M},w}$ and show that acceptance of w by the ATM $\mathcal{M}$ can be decided by inferencing over this knowledge base.

In detail, $K_{\mathcal{M},w}$ depends on $\mathcal{M}$ and $p(|w|)$, and has an empty ABox.[2] Acceptance of w by the ATM is reduced to checking concept subsumption, where one of the involved concepts directly depends on w. Intuitively, the elements of an interpretation domain of $K_{\mathcal{M},w}$ represent possible configurations of $\mathcal{M}$, encoded by the following concept names:

– A_q for $q \in Q$: the ATM is in state q,
– H_i for $i = 0, \ldots, p(|w|) - 1$: the ATM is at position i on the storage tape,
– $C_{\sigma,i}$ with $\sigma \in \Sigma$ and $i = 0, \ldots, p(|w|) - 1$: position i on the storage tape contains symbol σ,
– A: the ATM accepts this configuration.

This approach is pretty standard, and it is not too hard to axiomatise a successor relation S and appropriate acceptance conditions in $\mathcal{ALC}$ (see, e.g., [10]). But this reduction is not applicable in Horn-$\mathcal{SHIQ}$, and it is not trivial to modify it accordingly.

[2] The RBox is empty for $\mathcal{FLE}$ anyway.

Table 4. Knowledge base $K_{\mathcal{M},w}$ simulating a polynomially space-bounded ATM. The rules are instantiated for all $q, q' \in Q$, $\sigma, \sigma' \in \Sigma$, $i, j \in \{0, \ldots, p(|w|) - 1\}$, and $\delta \in \Delta$.

(1) Left and right transition rules:

$A_q \sqcap H_i \sqcap C_{\sigma,i} \sqsubseteq \exists S_\delta.(A_{q'} \sqcap H_{i+1} \sqcap C_{\sigma',i})$ with $\delta = (q, \sigma, q', \sigma', r)$, $i < p(|w|) - 1$

$A_q \sqcap H_i \sqcap C_{\sigma,i} \sqsubseteq \exists S_\delta.(A_{q'} \sqcap H_{i-1} \sqcap C_{\sigma',i})$ with $\delta = (q, \sigma, q', \sigma', l)$, $i > 0$

(2) Memory:

$H_j \sqcap C_{\sigma,i} \sqsubseteq \forall S_\delta.C_{\sigma,i} \quad i \neq j$

(3) Existential acceptance:

$A_q \sqcap \exists S_\delta.A \sqsubseteq A$ for all $q \in E$

(4) Universal acceptance:

$A_q \sqcap H_i \sqcap C_{\sigma,i} \sqcap \bigsqcap_{\delta \in \tilde{\Delta}} (\exists S_\delta.A) \sqsubseteq A$ $\quad q \in U$, $x \in \{r \mid i < p(|w|) - 1\} \cup \{l \mid i > 0\}$

$\tilde{\Delta} = \{(q, \sigma, q', \sigma', x) \in \Delta\}$

One problem that we encounter is that the acceptance condition of exitential states is a (non-Horn) disjunction over possible successor configurations. To overcome this, we encode individual transitions by using a distinguished successor relation for each translation in Δ. This allows us to explicitly state which conditions must hold for a particular successor without requiring disjunction. For the acceptance condition, we use a recursive formulation as employed in Definition 2. In this way, acceptance is propagated backwards from the final accepting configurations.

In the case of $\mathcal{ALC}$, acceptance of the ATM is reduced to concept satisfiability, i.e. one checks whether an accepting initial configuration can exist. This requires that acceptance is faithfully propagated to successor states, so that any model of the initial concept encodes a valid traces of the ATM. Axiomatising this requires many exclusive disjunctions, such as "The ATM always is in *exactly* one of its states H_i." Since it is not clear how to model this in a Horn-DL, we take a dual approach: reducing acceptance to concept subsumption, we require the initial state to be accepting in *all* possible models. We therefore may focus on the task of propagating properties to successor configurations, while not taking care of disallowing additional statements to hold. Our encoding ensures that, whenever the inital configuration is not accepting, there is at least one "minimal" model that reflects this.

After this informal introduction, consider the knowledge base $K_{\mathcal{M},w}$ given in Table 4. The roles S_δ, $\delta \in \Delta$, describe a configuration's successors using the translation δ. The initial configuration for word w is described by the concept expression I_w:

$$I_w := A_{q_0} \sqcap H_0 \sqcap C_{\sigma_0,0} \sqcap \ldots \sqcap C_{\sigma_{|w|-1},|w|-1} \sqcap C_{\square,|w|} \sqcap \ldots \sqcap C_{\square,p(|w|)-1},$$

where σ_i denotes the symbol at the ith position of w. We will show that checking whether the initial configuration is accepting is equivalent to cheking whether $I_w \sqsubseteq A$ follows from $K_{\mathcal{M},w}$. The following is obvious from the characterisation given in Table 3.

Lemma 1. *$K_{\mathcal{M},w}$ and $I_w \sqsubseteq A$ are in Horn-$\mathcal{FLE}$.*

Next we need to investigate the relationship between elements of an interpretation that satisfies $K_{\mathcal{M},w}$ and configurations of $\mathcal{M}$. Given an interpretation $\mathcal{I}$ of $K_{\mathcal{M},w}$, we say that an element e of the domain of $\mathcal{I}$ *represents* a configuration $\sigma_1 \ldots \sigma_{i-1} q \sigma_i \ldots \sigma_m$ if $e \in A_q^{\mathcal{I}}$, $e \in H_i^{\mathcal{I}}$, and, for every $j \in \{0, \ldots, p(|w|) - 1\}$, $e \in C_{\sigma,j}^{\mathcal{I}}$ whenever

$$j \leq m \text{ and } \sigma = \sigma_m \qquad \text{or} \qquad j > m \text{ and } \sigma = \square.$$

Note that we do not require uniqueness of the above, so that a single element might in fact represent more than one configuration. As we will see below, this does not affect our results. If e represents a configuration as above, we will also say that e has state q, position i, symbol σ_j at position j etc.

Lemma 2. *Consider some interpretation $\mathcal{I}$ that satisfies $K_{\mathcal{M},w}$. If some element e of $\mathcal{I}$ represents a configuration α and some transition δ is applicable to α, then e has an $S_\delta^\mathcal{I}$-successor that represents the (unique) result of applying δ to α.*

Proof. Consider an element e, state α, and transition δ as in the claim. Then one of the axioms (1) applies, and e must also have an $S_\delta^\mathcal{I}$-successor. This successor represents the correct state, position, and symbol at position i of e, again by the axioms (1). By axiom (2), symbols at all other positions are also represented by all $S_\delta^\mathcal{I}$-successors of e. □

Lemma 3. *A word w is accepted by $\mathcal{M}$ iff $I_w \sqsubseteq A$ is a consequence of $K_{\mathcal{M},w}$.*

Proof. Consider an arbitrary interpretation $\mathcal{I}$ that satisfies $K_{\mathcal{M},w}$. We first show that, if any element e of $\mathcal{I}$ represents an accepting configuration α, then $e \in A^\mathcal{I}$.

We use an inductive argument along the recursive definition of acceptance. If α is a universal configuration then all successors of α are accepting, too. By Lemma 2, for any δ-successor α' of α there is a corresponding $S_\delta^\mathcal{I}$-successor e' of e. By the induction hypothesis for α', e' is in $A^\mathcal{I}$. Since this holds for all δ-successors of α, axiom (4) implies $e \in A^\mathcal{I}$. Especially, this argument covers the base case where α has no successors.

If α is an existential configuration, then there is some accepting δ-successor α' of α. Again by Lemma 2, there is an $S_\delta^\mathcal{I}$-successor e' of e that represents α', and $e' \in A^\mathcal{I}$ by the induction hypothesis. Hence axiom (3) applies and also conclude $e \in A^\mathcal{I}$.

Since all elements in $I_w^\mathcal{I}$ represent the initial configuration of the ATM, this shows that $I_w^\mathcal{I} \subseteq A^\mathcal{I}$ whenever the initial configuration is accepting.

It remains to show the converse: if the initial configuration is not accepting, there is some interpretation $\mathcal{I}$ such that $I_w^\mathcal{I} \nsubseteq A^\mathcal{I}$. To this end, we define a canonical interpretation M of $K_{\mathcal{M},w}$ as follows. The domain of M is the set of all configurations of $\mathcal{M}$ that have size $p(|w|) + 1$ (i.e. that encode a tape of length $p(|w|)$, possibly with trailing blanks). The interpretations for the concepts A_q, H_i, and $C_{\sigma,i}$ are defined as expected so that every configuration represents itself but no other configuration. Especially, I_w^M is the singleton set containing the initial configuration. Given two configurations α and α', and a transition δ, we define $(\alpha, \alpha') \in S_\delta^M$ iff there is a transition δ from α to α'. A^M is defined to be the set of accepting configurations.

By checking the individual axioms of Table 4, it is easy to see that M satisfies $K_{\mathcal{M},w}$. Now if the initial configuration is not accepting, $I_w^M \nsubseteq A^M$ by construction. Thus M is a counterexample for $I_w \sqsubseteq A$ which thus is not a logical consequence. □

We can summarise our results as follows.

Theorem 1. *Checking concept subsumption in Horn-$\mathcal{FLE}$ is* ExpTime*-complete.*

Proof. Inclusion is obvious as Horn-$\mathcal{FLE}$ is a fragment of $\mathcal{ALC}$, which is in ExpTime. Regarding hardness, Lemma 3 shows that the word problem for polynomially space-bounded ATMs can be reduced to checking concept subsumption in $K_{\mathcal{M},w}$. By Lemma 1,

$K_{\mathcal{M},w}$ is in Horn-$\mathcal{FLE}$. The reduction is polynomially bounded due to the restricted number of axioms: there are at most $2 \times |Q| \times p(|w|) \times |\Sigma| \times |\Delta|$ axioms of type (1), $p(|w|)^2 \times |\Sigma| \times |\Delta|$ of type (2), $|Q| \times |\Sigma|$ of type (3), and $|Q| \times p(|w|) \times |\Sigma|$ of type (4). □

It is worth to discuss this result. The logic $\mathcal{FL}_0$ which admits only $\top$, $\sqcap$, and $\forall$ is known to be ExpTime-complete already [2]. Since we additionally use $\exists$, it might appear that Theorem 1 is trivial. However, the condition of Hornness severely restricts the use of $\forall$, and indeed we conjecture that Horn-$\mathcal{FL}_0$ actually is in P.

On the other hand, checking concept subsumption in the description logic $\mathcal{EL}$ which allows $\top$, $\sqcap$, and $\exists$ is in P [11]. This shows that the axioms (2) in Table 4 are really necessary. Without them, inferencing for this knowledge base would merely be polynomial.[3] This observation makes the axioms (2) particularly interesting for further study. Especially, we obtain the following corollary.

Theorem 2. *Let $\mathcal{EL}^{\leq 1}$ denote $\mathcal{EL}$ extended with number restrictions of the form $\leq 1\, R.\top$. Horn-$\mathcal{EL}^{\leq 1}$ is* ExpTime*-complete.*

Proof. Indeed, we can replace the axioms (2) in Table 4 with the following statements:

$$\top \sqsubseteq \leq 1\, S_\delta.\top \qquad\qquad H_j \sqcap C_{\sigma,i} \sqcap \exists S_\delta.\top \sqsubseteq \exists S_\delta.C_{\sigma,i} \quad i \neq j$$

It is easy to see that this formulation allows us to establish a result as in Lemma 2, which is the only place where the original axioms (2) had been required. □

ExpTime-completeness of $\mathcal{EL}^{\leq 1}$ was shown in [2], but the above theorem sharpens this result to the Horn case, and provides a more direct proof. Theorems 1 and 2 thus can be viewed as sharpenings of the hardness results on extensions of $\mathcal{EL}$.

5 Discussion and outlook

We have provided simple, self-contained characterisations of both the syntax and complexity of Horn-$\mathcal{SHIQ}$, and we believe that both contribute to an improved understanding of Horn-fragments in description logics. Our results show that, in spite of its positive effect on data complexity, Hornness in many cases cannot alleviate the high complexity of TBox reasoning.

The direct proofs of our results yield further insights regarding the source of the arising complexity. Existential role restrictions generally have the potential to increase the size of the admissible models beyond the number of explicitly given individuals. But as $\mathcal{EL}$ illustrates, existential restrictions alone do not suffice to enforce an exponential number of additional individuals. Indeed, for elements introduced by existential restrictions, one can only conclude logical properties that are directly imposed by the axiom introducing the new element. In contrast, successor elements arising in the above proofs represent arbitrary combinations of certain logical properties (e.g. tape configurations) without having an axiom for each such combination.

The key is that multiple axioms can *independently* propagate properties to the same successor element, and in this way enable an exponential number of combinations of

[3] This also holds for instance classification and satisfiability checking which are decided by checking concept subsumtion in $\mathcal{EL}^{++}$, which is still tractable [2].

such properties. In Theorem 1, independent propagation is achieved by universal quantification. In Theorem 2, restricting the number of overall successors allows us to combine properties within one successor. We conjecture that the interplay between existential and universal/number restrictions is still needed, and that Horn-$\mathcal{FL}_0$ is in P.

Another question is whether unqualified existential restrictions $\exists.\top$ still increase complexity, i.e. whether Horn-$\mathcal{FL}^-$ [12] is ExpTime-hard or not. A positive answer would subsume both the above Theorem 1 and a similar result on $\mathcal{AL}$ as presented in [12, Theorem 3.27]. Since $\mathcal{AL}$ provides atomic negation and universal restrictions, but only unqualified existential restrictions, none of the two results implies the other and it is not obvious how to adjust either of the proofs accordingly.

Finally, though most extensions of $\mathcal{EL}^{++}$ increase the complexity [2], it is still conceivable that this can be prevented in some cases by restricting to Horn-logic. A first candidate for this investigation would be Horn-$\mathcal{ELU}$, which adds (Horn) disjunctions to $\mathcal{EL}$. In general, we think that further research in Horn DLs can contribute to the development of practically meaningful inferencing that is still tractable.

References

1. McGuinness, D., v. Harmelen, F.: OWL Web Ontology Language Overview (2004) W3C Recommendation 10 Feb. 2004, http://www.w3.org/TR/owl-features/.
2. Baader, F., Brandt, S., Lutz, C.: Pushing the EL envelope. In: Proc. 19th Int. Joint Conf. on Artificial Intelligence (IJCAI-05), Edinburgh, UK, Morgan-Kaufmann Publishers (2005)
3. Calvanese, D., de Giacomo, G., Lembo, D., Lenzerini, M., Rosati, R.: Tailoring OWL for data intensive ontologies. In: Proc. 1st Workshop OWL Experiences and Directions. (2005)
4. Horrocks, I., Sattler, U., Tobies, S.: Practical reasoning for very expressive description logics. Journal of the Interest Group in Pure and Applied Logic **8** (2000) 239–264
5. Hustadt, U., Motik, B., Sattler, U.: Data complexity of reasoning in very expressive description logics. In: Proc. 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-05), Edinburgh, UK, Morgan-Kaufmann Publishers (2005) 466–471
6. Krötzsch, M., Hitzler, P., Vrnadecic, D., Sintek, M.: How to reason with OWL in a logic programming system. In: Proc. 2nd Int. Conf. on Rules and Rule Markup Languages for the Semantic Web (RuleML 2006), Athens, Georgia, USA, Springer (2006)
7. Tobies, S.: Complexity Results and Practical Algorithms for Logics in Knowledge Representation. PhD thesis, RWTH Aachen, Germany (2001)
8. Motik, B.: Reasoning in Description Logics using Resolution and Deductive Databases. PhD thesis, Universität Karlsruhe (TH), Germany (2006)
9. Chandra, A.K., Kozen, D.C., Stockmeyer, L.J.: Alternation. J. ACM **28** (1981) 114–133
10. Lutz, C., Sattler, U.: Description logics (2005) Tutorial at the ICCL Summer School 2005, Dresden, Germany. Course material available online.
11. Baader, F.: Terminological cycles in a description logic with existential restrictions. In Gottlob, G., Walsh, T., eds.: Proc. 18th Int. Joint Conf. on Artificial Intelligence (IJCAI-03), Acapulco, Mexico, Morgan-Kaufmann Publishers (2003) 325–330
12. Baader, F., Calvanese, D., McGuinness, D., Nardi, D., Patel-Schneider, P., eds.: The Description Logic Handbook: Theory, Implementation and Applications. Cambridge University Press (2003)